

Electron Neutrino Appearance at the NOvA Experiment

Alexander Booth

Submitted for the degree of Doctor of Philosophy

University of Sussex

October 2021

Declaration

I hereby declare that this thesis has not been and will not be, submitted in whole or in part to another University for the award of any other degree.

Signature:

Alexander Booth

UNIVERSITY OF SUSSEX

ALEXANDER BOOTH

ELECTRON NEUTRINO APPEARANCE AT THE NOvA EXPERIMENT

ABSTRACT

NOvA is a long-baseline neutrino oscillation experiment searching for electron neutrino appearance and muon neutrino disappearance. To do this, NOvA uses the NuMI beam at Fermi National Accelerator Laboratory along with two functionally identical detectors, separated by a baseline of 809 km. A near detector, which is close to the point of neutrino production, provides a measurement of initial beam energy spectra and flavour composition. The spectra are then extrapolated to a far detector and compared to data to look for oscillations. The experiment is able to constrain several parameters of the Pontecorvo-Maki-Nakagawa-Sakata matrix and is sensitive to the neutrino mass hierarchy.

This thesis presents an analysis of 6 years of NuMI data collected by the NOvA far detector corresponding to a 14 ktonne equivalent exposure of 13.60×10^{20} and 12.50×10^{20} protons on target, in neutrino and antineutrino beam modes respectively. Several techniques to develop NOvA's most recent 3-flavour oscillation analysis are explored, focussing on the electron neutrino appearance channel. A new fit of far detector data is performed, probing the parameters governing 3-flavour neutrino oscillations with improved sensitivity equivalent to collecting 5 to 6% more data. NOvA is expected to take data for about 12 years and so this improvement is equivalent to approximately half a year of data taking. Each year of operation of the accelerator complex for NOvA costs tens of millions of dollars, providing strong motivation to maximise the sensitivity of our analysis of the data.

A best fit at $\sin^2 \theta_{23} = 0.46^{+0.04}_{-0.03}$, $\delta_{CP} = (0.12^{+0.21}_{-0.19}) \pi$ and $\Delta m_{32}^2 = (2.41^{+0.06}_{-0.06}) \times 10^{-3} \text{ eV}^2$ in the normal neutrino mass hierarchy and lower octant of θ_{23} is obtained. In the upper octant $\sin^2 \theta_{23} = 0.57^{+0.02}_{-0.04}$. Two points of maximal CP violation, $\delta_{CP} = 3\pi/2$ in the inverted mass hierarchy and $\delta_{CP} = \pi/2$ in the normal mass hierarchy, are excluded at 3.4σ and 2.4σ respectively. Maximal muon-tau mixing, the oscillation point at which $\theta_{23} = \pi/4$, is allowed at the $1\text{-}\sigma$ level.

Acknowledgements

The 15 months leading up to the submission of this thesis were some of the most challenging I have ever experienced. Thanks to COVID-19, it is 15 months that I hope to never have to repeat. I feel immense gratitude to all of my family and friends, both in and out of the physics world, who guided me through this time.

Firstly, I would like to give a special thanks my supervisor, Jeff Hartnell. Your insight and attention to detail have been invaluable in my learning of neutrino physics and my development as a researcher more generally. I would also like to say thank you to Michelle Stancari; without you, I would not have thought about pursuing a PhD in experimental particle physics.

Thank you to my friends and housemates at Fermilab: Karl Warburton, Mike Wallbank, Tyler Alion and Tristan Blackburn, for showing me the ropes, introducing me to the delights of Chicagoland and the many road trips. Dom Brailsford, for the Rapid City speak easy and vim config. Diana Mendez, for all of the answers to my many questions.

Thank you to my friends and housemates at CERN: Mario and Olivia Grandi, for all of the dinners and baking exchanges. Ioannis Iiotidis, for the many, many beers in Russia and the rescue in Geneva. All those in the UK DAQ office, for the long lunch breaks at Restaurant 3. Celio Moura and Roman Berner, for the many, many caipirinhas in Sao Paulo.

Thank you to all of the residents of the physics Airbnb: Luke Vinton, your hummus will never be forgotten. Iker de Icaza-Astiz, for the Taco Tuesdays, sea swims and weekend tours by bike. Pierre Lasorak and Jigmet Yangchan, for the wine, the curries, the badminton and the highest stakes rock-paper-scissors game I am likely to ever have.

Thank you to my fellow Brightonians: Aran Borkum, for the early morning paddle-

boards. Georgia Chisnall and Nicolo Tuccori, for the many beach beers. Brett Mayes, for the afternoon in Hollywood. Michal Rigan, Charlie Mills and Yibing Zhang, for all of the fun of 4C10.

Finally, I'd like to say thank you to my family, Craig, Helen and Alice Booth, who have always given to me their unconditional support. Last but very not least, I'd like to give a special mention to Brianna Mills. Her unwavering support and strict enforcement of work-life balance have kept me on the right track these past 10 years. Thank you.

Preface

An overview of the content of each chapter in the thesis is given below, with original contributions from the author highlighted accordingly. The general goals of this thesis were suggested by my supervisor, Professor Jeffrey John Hartnell, who is an analysis coordinator for the Neutrinos from the Main Injector (NuMI) Off-Axis Neutrino Experiment (NOvA) experiment. Note that the bibliography contains many references to documentation which is only available to members of the NOvA collaboration. The inclusion of these references is to ensure that the original work of others is suitably recognised. Where an internal document is referred to, every effort has been made to ensure that the associated material in the thesis is self explanatory.

Chapter 1 introduces and motivates the subject of the thesis.

Chapter 2 presents the theoretical background relevant to neutrino oscillations and summarises the current status of the field. Particular emphasis is given to how NOvA, as a long-baseline experiment, probes the oscillation parameters. The chapter draws from a number of sources, following a comprehensive review of associated literature.

Chapter 3 describes the major components of the NOvA experimental set-up. Sections cover the neutrino source, near and far detectors, and a brief summary of simulation. The chapter is based on the work of other members of the NOvA collaboration.

Chapter 4 introduces various aspects of the 3-flavour oscillation analysis methodology including energy reconstruction, event selection and the extrapolation of information from the Near Detector (ND) to the Far Detector (FD). The chapter discusses primarily the work of the NOvA collaboration at large, with my own contributions appropriately highlighted.

Chapter 5 describes the development of a technique to quantify the uncertainty in ND event rate normalisation due to potential deficiencies in the simulation of the environment, particularly in relation to the pile-up of events. The systematic was first introduced by another collaborator. I implemented it into the most recent analysis and the various improvements presented in the chapter are my own work.

Chapter 6 gives an overview of changes and additions made to the electron neutrino appearance analysis, with the goal of increasing NOvA's sensitivity to the neutrino mass hierarchy and the Charge-Parity (CP) violating phase, δ_{CP} . The chapter uses NOvA's standard 3-flavour oscillation analysis as a base from which to build. All improvements are my own work.

Chapter 7 presents the results of a 3-flavour fit to neutrino and antineutrino beam data to measure oscillation parameters using the improved analysis. Code written by other members of the NOvA collaboration was adapted by myself to produce some of the plots which are shown. Aside from this, the chapter is based entirely on my own work.

Chapter 8 summarises the thesis and presents a self-reflection of the research project.

Contents

List of Tables	xiii
List of Figures	xx
1 Introduction	1
2 Neutrino Physics	3
2.1 First Detection and Evidence of Neutrino Oscillations	3
2.2 Neutrino Interactions	6
2.2.1 The Weak Force	6
2.2.2 Interactions in the NOvA Detectors	7
2.3 Oscillations	8
2.3.1 Time Evolution	10
2.3.2 Designing an Oscillation Experiment	13
2.3.3 Oscillations in Matter	15
2.4 Oscillations at NOvA	17
2.4.1 Muon Neutrino Disappearance	18
2.4.2 Electron Neutrino Appearance	19
2.5 Experimental Status	21
2.5.1 Measurements in the Solar Sector	22
2.5.2 Measurements in the Reactor Sector	22
2.5.3 Measurements in the Atmospheric Sector	26
2.5.4 Constraints on δ_{CP}	27
2.5.5 Latest 3-flavour Results from NOvA	29
3 The NOvA Experiment	34
3.1 The NuMI Beam	34
3.1.1 Primary Proton Beam	35

3.1.2	Secondary Hadron Beam	36
3.1.3	Decay Pipe & Neutrino Beam	37
3.1.4	Off-axis Design	38
3.2	The NOvA Detectors	39
3.2.1	The Cell	42
3.2.2	Liquid Scintillator	42
3.2.3	Optical Fibre	42
3.2.4	Photodetectors & Front End Boards	44
3.2.5	Data Concentrator Modules & Triggering	44
3.2.6	Detector Assembly	45
3.2.7	Near Detector	45
3.2.8	Far Detector	48
3.3	Simulation	49
3.3.1	Beam Simulation	49
3.3.2	Interactions & Cross-sections	51
3.3.3	Detector Simulation	52
4	Analysis Methodology	53
4.1	Low-level Software	53
4.1.1	Event Reconstruction	53
4.1.2	Energy Reconstruction	62
4.1.3	Particle Identification	66
4.2	High Level Software	70
4.3	Data Quality & Event Selection	72
4.3.1	Data & Spill Quality	72
4.3.2	Event Quality	73
4.3.3	Containment	73
4.3.4	Particle Identification (PID)	74
4.4	Analysis Binning at the FD	78
4.4.1	ν_μ -CC Binning	78
4.4.2	ν_e -CC Binning	79
4.5	Predictions at the FD	82
4.5.1	Decomposition & Extrapolation	83
4.5.2	Estimation of Cosmic Induced Events	91
4.6	Systematic Uncertainties	93

4.6.1	Flux	94
4.6.2	Cross-section	95
4.6.3	Detector Response and Calibration	96
4.6.4	Other Uncertainties	99
4.7	Oscillation Model Fit to Data	101
4.7.1	Treatment of Oscillation Parameters	104
4.7.2	Treatment of Systematic Uncertainties	104
5	Context Effects on Near Detector Selection Efficiency	106
5.1	Single Neutrino Overlay Methodology	107
5.2	Extracting Selection Efficiency Difference	109
5.2.1	Pre-selection Cuts	109
5.2.2	Slicing Efficiency	110
5.2.3	3-Flavour Cut Flow	113
5.3	Results	115
5.4	Implementation as a Systematic	118
5.5	Energy Resolution	118
6	Analysis Improvements	124
6.1	Studies at the FD	124
6.1.1	Searching for Gains	126
6.1.2	Optimisation of Bin Edges	134
6.1.3	Optimisation by Eye	136
6.2	Studies at the ND	143
6.2.1	Simulation of Rock Events	151
6.3	Fake Data Sensitivities	156
6.3.1	The Effect of Systematic Uncertainty on Sensitivities	156
6.3.2	$\Delta\chi^2$ Studies	158
7	Results	166
7.1	Studies at the ND	166
7.1.1	Decomposition	166
7.1.2	Reconstructed Variables at the ND	176
7.2	Analysis of FD Data	176
7.2.1	Candidate Events & New Best Fit	176
7.2.2	Confidence Intervals	181

CONTENTS	xi
7.2.3 Reconstructed Variables at the FD	189
7.3 Systematic Uncertainties	189
8 Conclusion	198
Acronyms	201
Bibliography	205
A Neutrino Physics Supplementary Material	216
A.1 Explicit Calculation of 3-flavour Transition Probability	216
B Analysis Bin Mapping	219
C Supplementary Data to Simulation Comparison Plots	221
C.1 ND	221
C.2 FD	240

List of Tables

2.1	3-flavour oscillation parameters from a fit to 2020 global data by the NuFit collaboration.	33
3.1	Most common decay modes for charged pions and kaons.	38
3.2	Chemical composition of the NOvA scintillator.	43
4.1	Muon neutrino candidate cut flow for the 2020 3-flavour analysis.	75
4.2	Muon antineutrino candidate cut flow for the 2020 3-flavour analysis.	76
4.3	Electron neutrino candidate cut flow for the 2020 3-flavour analysis.	77
4.4	Electron antineutrino candidate cut flow for the 2020 3-flavour analysis.	77
4.5	Systematic parameters: flux.	94
4.6	Systematic parameters: neutrino cross-section.	97
4.7	Systematic parameters: detector response and calibration.	99
4.8	Systematic parameters: other systematics.	102
5.1	Cut flow table for overlaid single ν_μ -CC events.	116
5.2	Cut flow table for overlaid single ν_e -CC events.	117
5.3	Summary of the single overlaid neutrino selection efficiency difference for the 2020 analysis.	117
6.1	$\Delta\kappa_{CP}^2$ values for the 2020 3-flavour analysis.	126
6.2	Summary of $\Delta\kappa_{CP}^2$ following changes to energy binning at the FD in Forward Horn Current (FHC).	129
6.3	Summary of $\Delta\kappa_{CP}^2$ following changes to energy binning at the FD in Reverse Horn Current (RHC).	130
6.4	Summary of $\Delta\kappa_{CP}^2$ following changes to Particle Identification (PID) binning and cuts at the FD in FHC.	132

6.5	Summary of $\Delta\kappa_{CP}^2$ following changes to PID binning and cuts at the FD in RHC.	133
6.6	A summary of the most optimal selection bin edges.	136
6.7	Selection bin boundaries in FHC.	138
6.8	Selection bin boundaries in RHC.	141
6.9	A summary of sensitivities when using the thesis axes compared to the 2020 axes. FD only. Statistics only.	143
6.10	A summary of sensitivities when using the thesis axes compared to the 2020 axes. With all systematics and using full extrapolation.	158
7.1	Decomposition results for the electron-like ND sample in FHC.	168
7.2	Decomposition results for the electron-like ND sample in RHC.	169
7.3	Simulated breakdown of predicted events at the FD in FHC with extrapolation at the 2020 3-flavour best fit.	172
7.4	Simulated breakdown of predicted events at the FD in FHC without extrapolation at the 2020 3-flavour best fit.	172
7.5	Simulated breakdown of predicted events at the FD in RHC with extrapolation at the 2020 3-flavour best fit.	174
7.6	Simulated breakdown of predicted events at the FD in RHC without extrapolation at the 2020 3-flavour best fit.	174
7.7	Predicted number of FD events in FHC at the thesis best fit compared to data.	177
7.8	Predicted number of FD events in RHC at the thesis best fit compared to data.	179
7.9	Summary of the $1\text{-}\sigma$ Gaussian allowed intervals for each freely fitted oscillation parameter.	189
B.1	A map between two analysis binnings.	220

List of Figures

2.1	Example vertices of the weak interaction.	7
2.2	Muon neutrino interaction cross-section as a function of energy.	9
2.3	Diagrams showing examples of the processes contributing to the total Charged Current (CC) cross-section.	9
2.4	Appearance probability in the two neutrino flavour picture.	14
2.5	The interaction of neutrinos in matter.	16
2.6	A schematic to explain the neutrino mass hierarchy problem.	19
2.7	Electron neutrino appearance probability at NOvA as a function of energy in vacuum.	20
2.8	Electron neutrino appearance probability at NOvA as a function of energy in matter.	21
2.9	Electron antineutrino survival probability at the Kamioka Liquid-scintillator Antineutrino Detector (KamLAND) and recent contours in the solar parameter space.	23
2.10	Measurements in the reactor sector by Daya Bay.	25
2.11	Recent results in the atmospheric parameter space.	28
2.12	Asimov sensitivity to CP violation at Deep Underground Neutrino Experiment (DUNE).	29
2.13	The number of Protons on Target (POT) as a function of time for the NOvA 2020 dataset.	31
2.14	Muon neutrino candidates in the NOvA 2020 analysis.	31
2.15	Electron neutrino candidates in the NOvA 2020 analysis.	32
2.16	Results of the 2020 electron neutrino appearance analysis at NOvA.	32
3.1	Schematic drawing of the NuMI beam facility.	35
3.2	The Fermi National Accelerator Laboratory (Fermilab) accelerator complex.	36
3.3	Possible trajectories of particles incident on the magnetic focussing horns.	37

3.4	Neutrino energy as a function of parent pion energy at different angles with respect to the beam axis.	39
3.5	Reconstructed energy spectra for ν_μ -CC in the ND and FD constructed at different angles with respect to the beam axis.	40
3.6	A scaled depiction of the NOvA detectors.	41
3.7	A single NOvA rigid Polyvinyl Chloride (PVC) extrusion.	41
3.8	Emission and absorption spectra for the NOvA optical fibre.	43
3.9	A plan view of the ND cavern with respect to the Main Injector Neutrino Oscillation Search (MINOS) installations and beamline.	46
3.10	Recent photographs of the NOvA ND.	47
3.11	Photographs of the FD hall.	48
3.12	Major components of the neutrino flux at the FD and ND in both beam modes.	50
3.13	Data-Monte Carlo (MC) comparison of reconstructed visible hadronic energy for events passing the ν_μ -CC selection.	52
4.1	Establishing central hits with TDSlicer.	56
4.2	Diagrammatic explanation of Prim's algorithm.	56
4.3	Neutrino event topologies in the ND.	57
4.4	Electron neutrino reconstruction chain flowchart.	60
4.5	The Kalman track algorithm at work.	61
4.6	Tools for the calorimetric energy calibration.	63
4.7	Relative calibration procedure.	65
4.8	Spine based fits used in the ν_μ -CC and $\bar{\nu}_\mu$ -CC energy estimator.	67
4.9	A heuristic explanation of deep learning at NOvA.	69
4.10	True breakdown of interaction type for the ν_e cosmic rejection classification score.	71
4.11	Diagrammatic explanation of the 2020 analysis cut flows.	76
4.12	Disappearance analysis FD binning scheme.	79
4.13	Muon neutrino energy estimator resolution for the four hadronic energy quartiles.	80
4.14	ν_μ -CC predicted spectrum at the FD split by hadronic energy quartiles. . .	81
4.15	Electron neutrino analysis binning for the 2020 3-flavour analysis.	82
4.16	A schematic to explain the extrapolation procedure in the disappearance analysis.	84

4.17	Distributions of reconstructed muon transverse momentum at the ND and FD.	85
4.18	Muon and electron neutrino flux at the ND, split by ancestry.	89
4.19	Beam Electron Neutrino (BEN) contained and uncontained samples at the ND.	89
4.20	The distribution of transverse and forward momentum for ancestor pions. .	90
4.21	The distribution of number of Michel electrons in the ν_e like ND sample. . .	92
5.1	Fully reconstructed spill in data, before and after the single simulated neutrino is overlaid.	107
5.2	Fully reconstructed spill in MC, before and after the single simulated neutrino is overlaid.	108
5.3	A visual summary of pre-selection cuts.	111
5.4	True energy distribution of overlaid neutrinos.	112
5.5	The distribution of ‘slice efficiency’ for all overlaid neutrinos.	113
5.6	Application of ν_μ -CC analysis cuts to the sliced samples of single neutrinos sequentially.	114
5.7	Selection efficiency as a function of key variables.	117
5.8	The distribution of fractional residuals in the ν_μ -CC energy estimator for selected overlaid muon neutrinos in neutrino mode data and MC.	119
5.9	The distribution of fractional residuals in the ν_e -CC energy estimator for selected overlaid electron neutrinos in neutrino mode data and MC.	120
5.10	The distribution of fractional residuals in the ν_μ -CC energy estimator for selected overlaid muon neutrinos in neutrino mode data and MC as a function of true neutrino energy.	122
5.11	The distribution of fractional residuals in the ν_e -CC energy estimator for selected overlaid electron neutrinos in neutrino mode data and MC as a function of true neutrino energy.	123
6.1	Summary of data equivalent sensitivity improvement following changes to energy binning at the FD.	128
6.2	Summary of data equivalent sensitivity improvement following changes to energy binning at the FD.	131
6.3	FD event distribution with the most optimal selection bin edges.	137
6.4	The contribution of each analysis bin to the total $\Delta\kappa_{CP}^2$	139

6.5	ν_e -CC appearance reconstructed energy prediction at the FD at the 2020 3-flavour best fit, FHC without extrapolation.	140
6.6	ν_e -CC appearance reconstructed energy prediction at the FD at the 2020 3-flavour best fit, RHC without extrapolation.	142
6.7	Reconstructed neutrino energy distributions at the ND, FHC.	145
6.8	Reconstructed neutrino energy distributions at the ND, RHC.	146
6.9	ν_e -CC appearance Convolutional Visual Network (CVN) A prediction at the FD at the 2020 3-flavour best fit, FHC without extrapolation.	147
6.10	$\bar{\nu}_e$ -CC appearance CVN A prediction at the FD at the 2020 3-flavour best fit, RHC without extrapolation.	148
6.11	ν_e -CC appearance CVN B prediction at the FD at the 2020 3-flavour best fit, FHC without extrapolation.	149
6.12	$\bar{\nu}_e$ -CC appearance CVN B prediction at the FD at the 2020 3-flavour best fit, RHC without extrapolation.	150
6.13	Reconstructed neutrino energy distributions at the ND, FHC. Rock excluded from decomposition weights.	154
6.14	Reconstructed neutrino energy distributions at the ND, RHC. Rock excluded from decomposition weights.	155
6.15	Rock contributions to the electron neutrino-like sample at the ND in FHC.	156
6.16	Rock contributions to the electron neutrino-like sample at the ND in FHC II.	157
6.17	The effect of systematic parameters on the FD prediction.	159
6.18	A comparison of $\Delta\kappa_{CP}^2$ with and without systematic uncertainties.	160
6.19	Fake data sensitivities of the new versus 2020 analysis at the 2020 3-flavour best fit.	162
6.20	One Dimensional (1D) fake data sensitivities of the new versus 2020 analysis at the 2020 3-flavour best fit.	163
6.21	Estimates of the contribution of each category of systematic uncertainty to the total uncertainty on each of the oscillation parameters.	165
7.1	Reconstructed neutrino energy distributions at the ND, FHC.	170
7.2	Reconstructed neutrino energy distributions at the ND, RHC.	171
7.3	ν_e -CC appearance reconstructed energy prediction at the FD at the 2020 3-flavour best fit, FHC with extrapolation.	173
7.4	ν_e -CC appearance reconstructed energy prediction at the FD at the 2020 3-flavour best fit, RHC with extrapolation.	175

7.5	ν_e -CC appearance candidate reconstructed energy spectra at the FD, FHC.	178
7.6	$\bar{\nu}_e$ -CC appearance candidate reconstructed energy spectra at the FD, FHC.	180
7.7	Bi-event plot at the thesis best fit prediction.	182
7.8	$\sin^2 \theta_{23}$ versus δ_{CP} Gaussian contours in the normal hierarchy.	183
7.9	$\sin^2 \theta_{23}$ versus δ_{CP} Gaussian contours in the inverted hierarchy.	183
7.10	Δm_{32}^2 versus $\sin^2 \theta_{23}$ Gaussian contours in the normal hierarchy.	184
7.11	Δm_{32}^2 versus $\sin^2 \theta_{23}$ Gaussian contours in the inverted hierarchy.	184
7.12	1D Gaussian contours for δ_{CP} in each hierarchy and octant combination. . .	186
7.13	1D Gaussian contours for $\sin^2 \theta_{23}$ in each hierarchy.	187
7.14	1D Gaussian contours for $ \Delta m_{32}^2 $ in each hierarchy and octant combination.	188
7.15	Summary of systematic pulls at the thesis best fit.	191
7.16	Estimates of the contribution of each category of systematic uncertainty to the total uncertainty on δ_{CP}	192
7.17	Estimates of the contribution of each category of systematic uncertainty to the total uncertainty on $\sin^2 \theta_{23}$	193
7.18	Estimates of the contribution of each category of systematic uncertainty to the total uncertainty on Δm_{32}^2	194
7.19	Estimates of the contribution of each systematic uncertainty to the total uncertainty on δ_{CP}	195
7.20	Estimates of the contribution of each systematic uncertainty to the total uncertainty on $\sin^2 \theta_{23}$	196
7.21	Estimates of the contribution of each systematic uncertainty to the total uncertainty on Δm_{32}^2	197
C.1	Reconstructed electro-magnetic energy distributions at the ND, FHC. . . .	222
C.2	Reconstructed electro-magnetic energy distributions at the ND, RHC. . . .	223
C.3	Reconstructed hadronic energy distributions at the ND, FHC.	224
C.4	Reconstructed hadronic energy distributions at the ND, RHC.	225
C.5	Reconstructed hadronic energy fraction distributions at the ND, FHC. . . .	226
C.6	Reconstructed hadronic energy fraction distributions at the ND, RHC. . . .	227
C.7	Electron neutrino CVN A score distributions at the ND, FHC.	228
C.8	Electron neutrino CVN A score distributions at the ND, RHC.	229
C.9	Reconstructed vertex X distributions at the ND, FHC.	230
C.10	Reconstructed vertex X distributions at the ND, RHC.	231
C.11	Reconstructed vertex Y distributions at the ND, FHC.	232

C.12 Reconstructed vertex Y distributions at the ND, RHC.	233
C.13 Reconstructed vertex Z distributions at the ND, FHC.	234
C.14 Reconstructed vertex Z distributions at the ND, RHC.	235
C.15 Number of reconstructed showers in a slice distributions at the ND, FHC. .	236
C.16 Number of reconstructed showers in a slice distributions at the ND, RHC. .	237
C.17 Number of hits in the primary shower distributions at the ND, FHC.	238
C.18 Number of hits in the primary shower distributions at the ND, RHC.	239
C.19 ν_e -CC appearance candidate reconstructed electro-magnetic energy spectra at the FD, FHC.	241
C.20 $\bar{\nu}_e$ -CC appearance candidate reconstructed electro-magnetic energy spectra at the FD, RHC.	242
C.21 ν_e -CC appearance candidate reconstructed hadronic energy spectra at the FD, FHC.	243
C.22 $\bar{\nu}_e$ -CC appearance candidate reconstructed hadronic energy spectra at the FD, RHC.	244
C.23 ν_e -CC appearance candidate reconstructed hadronic energy fraction spectra at the FD, FHC.	245
C.24 $\bar{\nu}_e$ -CC appearance candidate reconstructed hadronic energy fraction spectra at the FD, RHC.	246
C.25 ν_e -CC appearance candidate CVN A score spectra at the FD, FHC.	247
C.26 $\bar{\nu}_e$ -CC appearance candidate CVN A score spectra at the FD, RHC.	248
C.27 ν_e -CC appearance candidate CVN B score spectra at the FD, FHC.	249
C.28 $\bar{\nu}_e$ -CC appearance candidate CVN B score spectra at the FD, RHC.	250
C.29 ν_e -CC appearance candidate cosmic rejection Boosted Decision Tree (BDT) score spectra at the FD, FHC.	251
C.30 $\bar{\nu}_e$ -CC appearance candidate cosmic rejection BDT score spectra at the FD, RHC.	252
C.31 ν_e -CC appearance candidate reconstructed vertex X position spectra at the FD, FHC.	253
C.32 $\bar{\nu}_e$ -CC appearance candidate reconstructed vertex X position spectra at the FD, RHC.	254
C.33 ν_e -CC appearance candidate reconstructed vertex Y position spectra at the FD, FHC.	255

C.34 $\bar{\nu}_e$ -CC appearance candidate reconstructed vertex Y position spectra at the FD, RHC.	256
C.35 ν_e -CC appearance candidate reconstructed vertex Z position spectra at the FD, FHC.	257
C.36 $\bar{\nu}_e$ -CC appearance candidate reconstructed vertex Z position spectra at the FD, RHC.	258
C.37 ν_e -CC appearance candidate number of reconstructed showers spectra at the FD, FHC.	259
C.38 $\bar{\nu}_e$ -CC appearance candidate number of reconstructed showers spectra at the FD, RHC.	260
C.39 ν_e -CC appearance candidate reconstructed calorimetric energy of the primary shower spectra at the FD, FHC.	261
C.40 $\bar{\nu}_e$ -CC appearance candidate reconstructed calorimetric energy of the primary shower spectra at the FD, RHC.	262
C.41 ν_e -CC appearance candidate number of hits in the primary shower spectra at the FD, FHC.	263
C.42 $\bar{\nu}_e$ -CC appearance candidate number of hits in the primary shower spectra at the FD, RHC.	264
C.43 ν_e -CC appearance candidate reconstructed transverse momentum fraction of the primary shower spectra at the FD, FHC.	265
C.44 $\bar{\nu}_e$ -CC appearance candidate reconstructed transverse momentum fraction of the primary shower spectra at the FD, RHC.	266

Chapter 1

Introduction

Neutrino physics has a rich history stretching back almost 100 years. First proposed as a solution to an apparent observation of energy and angular momentum conservation violation, the neutrino has subsequently been extensively studied using a variety of experimental techniques. Since 1998, it has been known that the assumption made in the Standard Model of Particle Physics (SM) that neutrinos are massless is incorrect. Despite the neutrino's elusive nature, initially thought to be so weakly interacting that they could never be detected, several measurements of different neutrino sources have demonstrated that neutrinos oscillate between flavours. The existence of this oscillation showed that neutrinos do in fact have a (small) non-zero mass.

Some of the parameters which govern neutrino oscillations are now relatively well constrained, but there are still many unanswered questions. It is not yet known which neutrino mass eigenstate is the heaviest or whether a fundamental symmetry between μ and τ flavours exists. Perhaps the biggest question of all is whether neutrinos are a source of Charge-Parity (CP) violation in the lepton sector. Finding the solution could help provide an explanation to the observed asymmetry of matter and antimatter present in the universe. Unsurprisingly, neutrinos and neutrino oscillations continue to be the focus of substantial research efforts around the globe.

Neutrinos from the Main Injector (NuMI) Off-Axis Neutrino Experiment (NOvA) is a currently operating long-baseline neutrino oscillation experiment searching for electron neutrino and antineutrino appearance and muon neutrino and antineutrino disappearance. To do this, NOvA uses the NuMI beam at Fermi National Accelerator Laboratory (Fermilab) along with two functionally identical detectors, separated by 809 km. A Near

Detector (ND), which is close to the point of neutrino production, provides a measurement of initial beam spectra and flavour composition. The spectra are then extrapolated to the Far Detector (FD) to look for oscillations. The experiment is able to constrain several parameters of the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, the mass squared splitting and is sensitive to the neutrino mass hierarchy.

In this thesis, an analysis of NuMI data corresponding to a 14 ktonne equivalent exposure of 13.60×10^{20} and 12.50×10^{20} protons on target collected by the NOvA far detector, in neutrino and antineutrino beam modes respectively is presented. With a particular focus on the electron neutrino appearance channel, several techniques to improve NOvA's most recent 3-flavour oscillation analysis are explored. A new fit of FD data is performed, probing the parameters governing 3-flavour neutrino oscillations with improved sensitivity equivalent to collecting 5 to 6% more data compared to the most recent public result [1]. There is strong motivation to maximise the sensitivity of our analysis of the data. Each year of operation of the accelerator complex for NOvA costs tens of millions of dollars. NOvA is expected to take data for about 12 years and so this improvement is equivalent to approximately half a year of data taking.

The structure of the thesis is outlined in the Preface.

Chapter 2

Neutrino Physics

The story of the neutrino starts in the year 1914, when James Chadwick demonstrated that the energy spectrum of the electron emitted in nuclear β decay was continuous [2]. At the time it was assumed that the decay was two body, taking the form $N \rightarrow N' + e$ where N is some nucleus and e is an electron. Therefore, Chadwick's observation implied a violation of the principle of conservation of energy. As a solution to this problem, Wolfgang Pauli postulated in 1930 that the spectrum could be explained by the existence of a neutral, weakly interacting fermion emitted in addition to the electron [3]. He called this new particle the 'neutron' and proposed that it should have a mass similar to that of the electron. Three years after Pauli's initial idea, Fermi incorporated the particle into a new model for β decay, he renamed the particle the 'neutrino' [4].

A summary of the neutrino's first detection and evidence of its oscillation is given here. The remainder of the chapter goes on to give a comprehensive summary of the theory most relevant to the thesis, with particular emphasis given to 3-flavour neutrino oscillations. Sections 2.2 and 2.3 introduce neutrino interactions and neutrino oscillations respectively. Section 2.4 focuses on oscillations at NOvA. Finally, Section 2.5 presents the current global experimental picture of the field.

2.1 First Detection and Evidence of Neutrino Oscillations

It wasn't until 1956, almost 30 years after the existence of the neutrino was first postulated by Pauli, that the first detection was made by Fred Reines and Clyde Cowan [5,6]. Their method was to place large tanks of water and liquid scintillator close to the core of a nuclear reactor, using the huge flux of electron antineutrinos to observe the inverse β

decay process

$$p + \bar{\nu}_e \rightarrow n + e^+. \quad (2.1)$$

The water was doped with cadmium chloride, a neutron absorber. The combination of two prompt gammas from the positron annihilation and the later, single gamma from the capture of the neutron, was used to infer a signal. For this work, Reines was awarded the Nobel Prize in Physics in 1995.

Following the detection of the electron neutrino, it took just six years for the muon neutrino to be discovered using the world's first accelerator-based neutrino beam at Brookhaven National Laboratory [7]. The general scheme of colliding protons with a target to produce secondary hadrons and subsequently allowing them to decay to neutrino-charged lepton pairs, is a technique still used today. The third generation neutrino, the tau, was observed directly at Fermilab by the Direct Observation of the Nu Tau (DONUT) collaboration in 2000 [8]. The collaboration's first result was based on just four events on a background of 0.34 ± 0.05 . Although a significance of around 3.5σ , below the normal 5σ threshold, the result was generally accepted by the community given that the particle was expected to be there. In a final publication in 2007, the collaboration reported a measurement of the tau neutrino Charged Current (CC) cross-section relative to the electron and muon neutrino cross-sections [9]. The values obtained were consistent with the SM expectation.

Despite the discovery of the tau neutrino completing the SM in 2000¹, the phenomena that would become known as ‘neutrino oscillations’ were first suggested by Bruno Pontecorvo in 1957 [10]. He postulated transitions between neutrinos and antineutrinos in a scheme analogous to K^0/\bar{K}^0 oscillation. Although nothing initially came of his proposal, the idea was revisited, first by Maki, Nakagawa and Sakata. They assumed that the electron and muon neutrinos were mixed states of two mass eigenstates [11]. Pontecorvo then developed this idea to formulate the first intuitive understanding of two flavour neutrino oscillations in 1968 [12].

The first experimental observation of neutrino oscillations came in the same year, when Ray Davis set up an experiment consisting of a tank of chlorine cleaning solution, underground at the Homestake mine in South Dakota². The solution was capable of

¹Excluding the Higgs boson.

²Acceptance that it was in fact neutrino oscillations that were being observed did not come until many years later.

neutrino capture via

$$\nu_e + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar} + e^-. \quad (2.2)$$

Davis would extract the argon atoms and count them, giving a direct measurement of the solar neutrino flux [13]. Electron neutrinos are produced in large numbers in nuclear reactions in the Sun, but the flux measured by Davis was around one third of the number predicted in calculations by John Bahcall in the standard solar model [14]. This became known as the solar neutrino problem and was believed to be the result of flaws in Davis' methods or in theory. However, subsequent experiments using different technologies confirmed the deficit. The water Cherenkov detector Kamiokande [15] and the gallium based experiments SAGE [16] and GALLEX [17], all reported lower than theoretically predicted rates of solar neutrinos in 1989, 1991 and 1992 respectively.

Around the same time, Kamiokande made one of the first measurements of neutrino oscillations using atmospheric neutrinos [18]. Atmospheric neutrinos are produced when primary cosmic rays (usually protons), collide with the Earth's atmosphere and the products (mostly pions) decay. In a naive calculation, the expected ratio of muon neutrinos to electron neutrinos is two to one. Kamiokande found that the value of this ratio depended on the neutrino's zenith angle - roughly two to one for neutrinos coming from directly overhead, dropping off to approximately one to one for up-going neutrinos. This demonstrated that neutrinos travelling a greater distance after production, through the Earth on their way to the detector, were 'disappearing.' This observation could be explained by oscillations. In 1998 Kamiokande's upgrade, Super-Kamiokande, reported an energy dependent asymmetry in the muon neutrino rate with a significance greater than 6σ [19]. Therefore, this data provided conclusive evidence of neutrino oscillations.

In 2002, the Sudbury Neutrino Observatory (SNO) experiment resolved the solar neutrino problem by simultaneously measuring the electron and total solar neutrino fluxes. The flux of electron neutrinos showed the same deficit that had been observed by previous experiments and the total flux was in agreement with the standard solar model [20].

For the achievements of the Super-Kamiokande and SNO collaborations, Takaaki Kajita and Arthur McDonald respectively were awarded the Nobel Prize in Physics in 2015 'for the discovery of neutrino oscillations, which shows that neutrinos have mass' [21]. Since their discovery, the study of neutrino oscillations has progressed into the measure-

ment of the parameters governing oscillations and establishing the neutrino mass hierarchy. In addition, several fundamental questions remain: whether neutrinos are Dirac or Majorana, the origin of their mass, the number of neutrino flavours and others. Finding answers to these questions provides exciting prospects for the future of the field.

2.2 Neutrino Interactions

2.2.1 The Weak Force

In the SM, neutrinos interact via the weak force only. This force couples leptons differing by one unit of electric charge in Charged Current (CC) interactions, mediated by the massive, spin-1 W^\pm bosons. Examining the form of the 4-vector current

$$J^k \sim \bar{u}(p') \frac{1}{2} \gamma^k (1 - \gamma^5) u(p) \quad k \in [0, 3], \quad (2.3)$$

corresponding to such vertices, highlights two key features of the interaction [22]. Here, $u(p)$ and $u(p')$ are Dirac spinors representing incoming and outgoing particles with momenta p and p' respectively. γ^l are the Dirac γ matrices. An example diagram of such a vertex is shown in Figure 2.1a. In the case that neutrinos are Dirac particles, the structure of the vertex factor is vector minus axial vector (V-A). Therefore, weak interactions do not conserve parity, making them unique in the SM³. An additional important consequence of the vertex is that only left-handed chiral *particle* states and right-handed chiral *anti-particle* states participate in the weak interaction⁴. This follows immediately from the inclusion of $\frac{1}{2}(1 - \gamma^5)$ in the vertex factor, which is the left-handed chiral projection operator, P_L . This projects out left-handed particle states and right-handed antiparticle states. By construction, for a right-handed particle or left-handed antiparticle state v , $P_L v = 0$.

³Physical quantities are often classified by their dimension and behaviour under parity transformation (spatial inversion through the origin). Vector quantities change sign under parity transformation, 3-momentum for example, $\vec{p} \rightarrow -\vec{p}$. Conversely, axial vector quantities (vectors that can be expressed as a cross-product) do not change sign under parity transformation, angular momentum for example, $\vec{L} \rightarrow \vec{L}$. There are also pseudoscalars, single valued quantities that change sign under parity due to them being expressible as a scalar product of a vector and axial vector. An important example is the normalised projection of spin onto 3-momentum, known as helicity $h \equiv \frac{\vec{S} \cdot \vec{p}}{|\vec{p}|}$ [22].

⁴Chirality has no obvious physical interpretation, loosely it can be understood as how a spinor transforms under a Lorentz transformation. For a given boost, a left-handed state ‘rotates’ one way and a right-handed state the other. It can also be interpreted as the helicity of a particle in the massless limit. Mathematically, the chiral states are eigenstates of γ^5 .




Figure 2.1: Diagrams showing examples of CC (**left**) and Neutral Current (NC) (**right**) weak vertices. λ is a charged lepton and ν_λ is a neutrino of corresponding flavour, there is no flavour change at the vertex. Both interactions violate parity, but there is only a transfer of charge in CC. Taken from [23].

Unification of the weak force with Quantum Electrodynamics (QED) to become the $SU(2) \times U(1)$ electroweak theory gives rise to the Neutral Current (NC) interaction, mediated by the massive (but uncharged), spin-1 Z boson. An example diagram of a NC vertex is shown in Figure 2.1b. Both the CC and NC channels are key in NOvA's oscillation analyses, whether that be as signal or background. In the context of a neutrino's interaction with the detector mass, the channels that NOvA uses are of the following form:

$$\begin{aligned} \bar{\nu}_\lambda^{(-)} + X \rightarrow \lambda^\pm + X' & \qquad \bar{\nu}_\lambda^{(-)} + X \rightarrow \bar{\nu}_\lambda^{(-)} + X', \end{aligned} \quad (2.4)$$

where X is the detector mass and \pm in the final state refers to an antineutrino or neutrino in the initial state respectively. The key difference between CC on the left and NC on the right is that CC has a charged lepton in the final state. Identifying the flavour of this lepton gives information about the flavour of the incoming neutrino. The identity of the other final state particles X' , depends on the exact identity of X . X could be, for example, a proton or the entire nucleus. The various possible initial and final states are discussed in the next section.

2.2.2 Interactions in the NOvA Detectors

The final states of CC and NC interactions in the NOvA detectors depend on how the initial state neutrino interacts with the detector mass. These possibilities have different cross-sections, all of which contribute to the total cross-section. The dependencies of the cross-sections on neutrino energy in NOvA's region of interest are shown in Figure 2.2. There are four main CC processes, Quasi-elastic (QE), Meson Exchange Current (MEC), Baryon Resonance Production (RES) and Deep Inelastic Scattering (DIS), described as

follows:

- QE: A neutrino (antineutrino) of a given flavour interacts with a single neutron (proton), producing a negatively (positively) charged lepton of the same flavour. See Figure 2.3a.
- MEC: A neutrino (antineutrino) interacts with a correlated pair of nucleons within a nucleus, leading to multiple nucleons in the final state. See Figure 2.3b.
- RES: The exchanged W boson has enough energy to create a nuclear resonance which subsequently undergoes a strong decay, typically emitting a pion and a nucleon from the nucleus. See Figure 2.3c.
- DIS: The W boson has enough energy to interact with a single quark, breaking up the nucleon containing that quark. This results in a hadronisation, appearing experimentally as a number of strongly interacting particles. See Figure 2.3d.

The simulation of the different processes discussed here, is covered in detail in Chapter 3.

2.3 Oscillations

It has been experimentally determined that neutrinos oscillate. In the most simple terms, there is a non-zero probability that a neutrino created in one flavour at one point in spacetime will be measured at a later spacetime point with a different flavour. In the context of quantum mechanics, this experimental fact implies that the neutrino's flavour eigenstates (eigenstates of the weak interaction) are not the same as the neutrino's mass eigenstates (states diagonalising the free particle Hamiltonian). These two bases are in fact rotated relative to each other. Consequently, each of the flavour states can be written as a linear combination of the mass states and vice-versa. This relationship can be summarised as follows in matrix form

$$|\nu_\lambda\rangle = \sum_{j=1}^N U_{\lambda j}^* |\nu_j\rangle. \quad (2.5)$$

In this expression, $|\nu_\lambda\rangle$ and $|\nu_j\rangle$ are the neutrino flavour and mass eigenstates respectively and N corresponds to the number of active neutrino flavours. In the case of 3-flavour oscillations, $\lambda \in [e, \mu, \tau]$ and $N = 3$. Additionally, U is a 3×3 unitary ($U^{-1} = U^\dagger$) matrix known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, after the three people who first introduced it formally and Pontecorvo, acknowledged for predicting the




Figure 2.2: Muon neutrino cross-sections as a function of neutrino energy for a carbon target as predicted by GENIE [24]. Made for this thesis by P. Lasorak, a fellow NOvA collaborator.




Figure 2.3: Diagrams showing examples of the processes contributing to the total CC cross-section. Taken and adapted from [25].

existence of oscillations [11, 12]. Unlike the analogous mixing matrix in the quark sector, the PMNS matrix is quite off diagonal. This means that each flavour state is relatively strongly coupled with each mass state.

In general, nine parameters are required to completely define this type of matrix: three angles and six complex phases. However, not all of these phases are physically relevant, three of them can be absorbed into the definitions of the particles themselves [22]. Following this rewriting, the PMNS matrix can then be expressed as the product of four matrices (Equation 2.6) [26].

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\frac{\alpha_{21}}{2}} & 0 \\ 0 & 0 & e^{i\frac{\alpha_{31}}{2}} \end{pmatrix} \quad (2.6)$$

$$s_{ab} := \sin \theta_{ab} \quad , \quad c_{ab} := \cos \theta_{ab} \quad , \quad a, b \in [1, 3], \quad a \neq b. \quad (2.7)$$

This representation of the leptonic mixing matrix is perhaps the most intuitive. Loosely speaking, each individual matrix describes a ‘sector’, grouping terms which are traditionally probed by particular type of experimental set up. At zeroth order, ‘set up’ in neutrino physics refers to an experiment’s L/E value, or the ratio of a neutrino’s propagation distance and energy. Consequently from left to right, the matrices represent the atmospheric, reactor and solar sectors respectively. A summary of measurements in each of these sectors is given in Sections 2.5.3, 2.5.2 and 2.5.1 respectively. The rightmost matrix is only not equal to the identity when α_{21} and α_{31} are non-zero, that is to say, in the case that neutrinos are Majorana rather than Dirac. As will be shown in Equation 2.25, oscillation experiments are not sensitive to these phases. For this reason, this component of the PMNS matrix will be ignored, leaving three angles (θ_{12} , θ_{13} , θ_{23}) and one complex phase (δ). δ is often written δ_{CP} , as it describes to what extent neutrinos violate CP symmetry.

2.3.1 Time Evolution

It is necessary to derive an expression to describe the probability of neutrinos changing flavour as they propagate to obtain an observable that can be studied experimentally. There are several ways to do this. The most complete, but still non-relativistic, method is by treating each flavour eigenstate as a superposition of ‘source’ wave packets, one for each mass eigenstate. However, the derivation which follows is the most simplistic, considering the neutrinos to be plane waves. Certain assumptions made in this description

are not necessary when using the wave packet method. Moreover, the important concept of coherence is lost. Nevertheless, the fundamental result is the same for the purposes of this thesis. A number of sources were used in parallel to ensure that the following mathematical picture is as transparent as possible [27, 28].

Consider first a neutrino with definite flavour state $|\nu_\lambda\rangle$ produced at spacetime point $X_0 = (0, \vec{0})$. The neutrino then propagates in the absence of any external potentials. The mass states $|\nu_j(X)\rangle$, which are eigenstates of the free particle Hamiltonian, therefore, evolve according to the time-dependent Schrödinger equation (Equation 2.8)

$$i\partial_t|\nu_j(X)\rangle = -\frac{1}{2m_j}\nabla^2|\nu_j(X)\rangle. \quad (2.8)$$

The solutions are plane waves of the form

$$|\nu_j(X)\rangle = e^{-iP_j \cdot X}|\nu_j(X_0)\rangle, \quad (2.9)$$

where X is the (later) spacetime point of detection and P_j is the 4-momentum of the j 'th mass state such that

$$P_j \cdot X = E_j t - \vec{p}_j \cdot \vec{x}. \quad (2.10)$$

Equation 2.5 can then be inverted to write the mass states as a linear combination of the flavour states

$$|\nu_j\rangle = \sum_{\eta \in [e, \mu, \tau]} U_{\eta j} |\nu_\eta\rangle. \quad (2.11)$$

Substituting firstly, this expression into the plane wave solutions (Equation 2.9) and secondly, these solutions back into Equation 2.5, leaves the following expression

$$|\nu_\lambda(X)\rangle = \sum_{\eta \in [e, \mu, \tau]} \sum_{j=1}^N U_{\lambda j}^* e^{-iP_j \cdot X} U_{\eta j} |\nu_\eta\rangle. \quad (2.12)$$

This shows that for an arbitrary spacetime point X , the flavour state is a superposition of the states $|\nu_\eta\rangle$. Therefore, the neutrino produced at $X = X_0$ is now a composite state of different flavoured neutrinos. The transition amplitude for detecting a neutrino in a flavour state λ' at spacetime point X , given that it was produced in flavour state λ at X_0 is given by the inner product

$$\mathcal{A}(\nu_\lambda \rightarrow \nu_{\lambda'}) = \langle \nu_{\lambda'} | \nu_\lambda(X) \rangle \quad (2.13)$$

$$= \sum_{\eta \in [e, \mu, \tau]} \sum_{j=1}^N U_{\lambda j}^* e^{-iP_j \cdot X} U_{\eta j} \delta_{\lambda' \eta} \quad (2.14)$$

$$= \sum_{j=1}^N U_{\lambda j}^* U_{\lambda' j} e^{-iP_j \cdot X}. \quad (2.15)$$

The probability of observing such a transition is then simply the coherent sum

$$\mathcal{P}(\nu_\lambda \rightarrow \nu_{\lambda'}) = |\mathcal{A}(\nu_\lambda \rightarrow \nu_{\lambda'})|^2 \quad (2.16)$$

$$= \sum_{j=1}^N U_{\lambda j}^* U_{\lambda' j} e^{-i P_j \cdot X} \sum_{k=1}^N U_{\lambda k} U_{\lambda' k}^* e^{i P_k \cdot X} \quad (2.17)$$

$$= \sum_{j,k=1}^N U_{\lambda j}^* U_{\lambda' j} U_{\lambda k} U_{\lambda' k}^* e^{-i(P_j - P_k) \cdot X}. \quad (2.18)$$

Recalling Equation 2.10, the exponential phase can be expanded and simplified

$$(P_j - P_k) \cdot X = (E_j - E_k)t - (\vec{p}_j - \vec{p}_k) \cdot \vec{x} \quad (2.19)$$

$$= \left(\sqrt{m_j^2 + |\vec{p}_j|^2} - \sqrt{m_k^2 + |\vec{p}_k|^2} \right) t - (\vec{p}_j - \vec{p}_k) \cdot \vec{x} \quad (2.20)$$

$$\approx |\vec{p}| \left(\sqrt{1 + \frac{m_j^2}{|\vec{p}|^2}} - \sqrt{1 + \frac{m_k^2}{|\vec{p}|^2}} \right) t \quad (2.21)$$

$$\approx |\vec{p}| \left(\left(1 + \frac{m_j^2}{2|\vec{p}|^2} \right) - \left(1 + \frac{m_k^2}{2|\vec{p}|^2} \right) \right) L \quad (2.22)$$

$$= \frac{L}{2E} (m_j^2 - m_k^2) := \frac{\Delta m_{jk}^2 L}{2E}, \quad (2.23)$$

where the following assumptions were made:

- All mass states have the same spatial momentum $\vec{p} = \vec{p}_j = \vec{p}_k$.
- Neutrinos are ultra-relativistic ($m \ll E$), therefore, $|\vec{p}| \sim E$ and $t \sim L$ where L is the distance propagated by the neutrinos between production and detection. This quantity is often referred to as the ‘baseline’ in oscillation experiments. Terms $\mathcal{O}(m^4)$ and higher are neglected.

Putting all of these pieces together gives the following expression for the oscillation probability

$$\mathcal{P}(\nu_\lambda \rightarrow \nu_{\lambda'}) = \sum_{j,k=1}^N U_{\lambda j}^* U_{\lambda' j} U_{\lambda k} U_{\lambda' k}^* e^{-i \frac{\Delta m_{jk}^2 L}{2E}}. \quad (2.24)$$

This can be written in the more intuitive and familiar form for $N = 3$ (three neutrinos)

$$\begin{aligned} \mathcal{P} \left(\overset{(-)}{\nu_\lambda} \rightarrow \overset{(-)}{\nu_{\lambda'}} \right) &= \delta_{\lambda\lambda'} - 4 \sum_{j < k} \text{Re} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin^2 \left(\frac{\Delta m_{jk}^2 L}{4E} \right) \\ &\quad \pm 2 \sum_{j < k} \text{Im} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin \left(\frac{\Delta m_{jk}^2 L}{2E} \right), \end{aligned} \quad (2.25)$$

where $\delta_{\lambda\lambda'}$ is the Kronecker delta and the sign of the imaginary component is positive / negative for antineutrinos / neutrinos respectively. This follows immediately by recognising

that a neutrino state with negative helicity (left-handed) is related to an antineutrino state with positive helicity (right-handed) by CP transformation. Applying CP to the expression for probability amounts to taking the complex conjugate, flipping the sign of the imaginary term only. Other noteworthy properties:

- Dependence on terms of the form $U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}$ means that the Majorana phases of the PMNS matrix cancel. Oscillation experiments, therefore, have no sensitivity to these phases.
- In addition to the PMNS parameters, there is a dependence on the difference in the square of the neutrino masses $(\Delta m_{jk}^2 := m_j^2 - m_k^2)$. This implies that for the probability of transition to be non-zero, at least one of the neutrinos must be massive and their masses should be distinct.
- Considering the probability to be a function of L (or L/E) implies that the probability varies sinusoidally. The transition probability is, therefore, different at different points along the flight of the neutrino, repeating after some period. This is what is described as ‘oscillations’.
- In the case that $\lambda \neq \lambda'$, the equation gives what is known as the ‘appearance’ probability. Conversely, in the case that $\lambda = \lambda'$, the equation gives what is known as the ‘survival’ or ‘disappearance’ probability. Since the products $U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}$ are real this case, the final, imaginary term is zero. As a result, measurements using the disappearance channel yield limited information about the (complex) CP violating phase, δ_{CP} .

See Appendix A.1 for an explicit derivation of Equation 2.25 from the transition amplitude (Equation 2.13).

2.3.2 Designing an Oscillation Experiment

From Section 2.3.1, it is clear that in the 3-flavour picture, both appearance and disappearance probabilities are relatively complicated functions of the PMNS parameters, mass splittings, propagation distance and energy. To illustrate how an experiment might be designed to probe these parameters, it is useful to consider the approximation that there are two neutrino flavours. In this simplified picture, the neutrino mixing matrix becomes a real, Two Dimensional (2D) rotation matrix.

$$U_{N=2} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \quad (2.26)$$

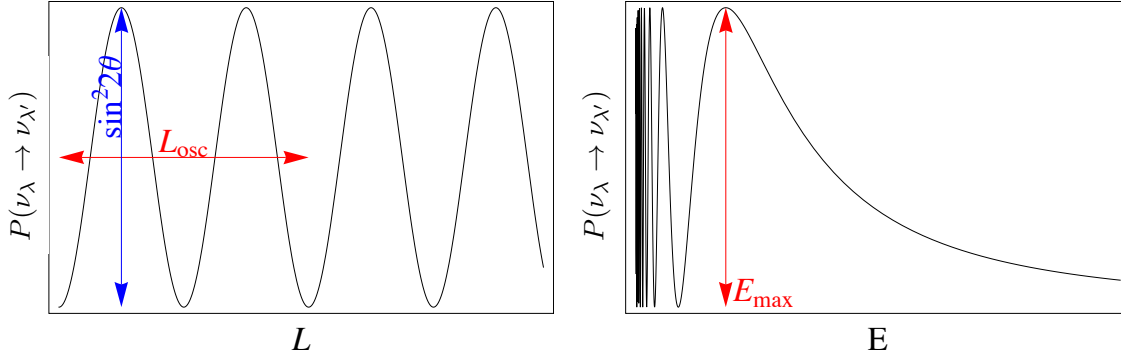


Figure 2.4: Appearance probability as a function of propagation distance (**left**) and neutrino energy (**right**), in the two neutrino flavour picture. Moving to increasingly higher energies, the appearance probability tends to zero. Taken from [28].

Moving to more tangible units by reintroducing factors of \hbar and c , the oscillation probabilities are given by [27]

$$P(\nu_\lambda \rightarrow \nu_{\lambda'}) = \sin^2(2\theta) \sin^2\left(1.27 \frac{\Delta m^2 (\text{eV}^2) L (\text{km})}{E (\text{GeV})}\right), \quad \lambda \neq \lambda' \quad (2.27)$$

$$P(\nu_\lambda \rightarrow \nu_\lambda) = 1 - P(\nu_\lambda \rightarrow \nu_{\lambda'}), \quad (2.28)$$

where $\Delta m^2 \equiv m_2^2 - m_1^2$. The appearance probability (Equation 2.27) is plotted in Figure 2.4 as both a function of propagation length (with fixed energy) and as a function of energy (with fixed propagation length). It oscillates as a sine function squared, where the amplitude is driven by the mixing angle θ . Conversely, the frequency is driven by the difference in squared masses Δm^2 and can be extracted by considering the period of oscillation at a fixed energy, or oscillation length, defined to be

$$L_{\text{osc}} (\text{km}) = \pi \frac{E (\text{GeV})}{1.27 \Delta m^2 (\text{eV}^2)}. \quad (2.29)$$

With this in mind, an experiment with a monochromatic neutrino source can be imagined, where the oscillation probability is sampled at different propagation distances or ‘baselines’.

In practise, due to the large active masses required by neutrino detectors, it is the probability as a function of neutrino energy at a fixed baseline which is usually measured experimentally. The position of the first oscillation maximum from right to left, marked in the right of Figure 2.4, is given by

$$E_{\text{max}} = 1.27 \frac{\Delta m^2 (\text{eV}^2) L (\text{km})}{\pi/2}. \quad (2.30)$$

This expression can be reached by differentiating the probability with respect to energy and finding the root corresponding to the largest value of E . Therefore, the mass splitting in the case of two neutrinos is directly proportional to E_{\max} .

As alluded to in Section 2.3, the key to probing different components of the PMNS matrix is designing experiments with optimised values neutrino energy and baseline. More specifically, it is beneficial for their ratio to be the same order as the square mass splitting, that is to say $E/L \sim \Delta m^2$. If $E/L \gg \Delta m^2$, then $P(\nu_\lambda \rightarrow \nu_{\lambda'}) \sim \sin^2(2\theta) (\Delta m^2)^2 (L/E)$. At leading order, therefore, as $E \rightarrow \infty$ and assuming a fixed baseline, $P \rightarrow 0$ and there are no oscillations. At the opposite limit, $E/L \ll \Delta m^2$, oscillations become too ‘fast’, and energy resolution becomes a limiting factor. Note also in this limit that it is possible for oscillations occur over scales of L which are smaller than that of a detector. This would be a problem for experiments like NOvA which currently measure using a single baseline.

2.3.3 Oscillations in Matter

Up to this point, it has been assumed that the neutrino mass states are eigenstates of the free particle Hamiltonian. In practice, neutrinos often oscillate while passing through matter, not vacuum, when on route from source to detector. In the case of NOvA, the neutrinos travel through the Earth for approximately 800 km. Largely, the Earth is, of course, composed of non-exotic matter such as protons, neutrons and electrons, as opposed to, for example, μ or τ leptons. For this reason, electron neutrinos and antineutrinos have an additional coherent forward scattering process with the electrons in matter via the W boson. These interactions, shown in Figure 2.5, change the energy levels of the mass eigenstates, thereby affecting oscillation probabilities.

This behaviour was predicted in the 1970’s and the mathematics describing it was developed in the 1980’s - it is known as the Mikheyev-Smirnov-Wolfenstein (MSW) or ‘matter’ effect [29, 30]. The matter electrons contribute an additional potential term

$$V_e = \pm \sqrt{2} G_F N_e, \quad (2.31)$$

where G_F is Fermi’s constant, N_e is the electron number density, the positive sign is for neutrinos and the negative sign for antineutrinos. To demonstrate the effect of this potential, the two neutrino approximation can once again be considered [23, 31]. As with the vacuum case, the time dependent Schrödinger equation can be used to describe the time




Figure 2.5: Diagrams showing CC (**left** and **middle**) and NC (**right**) coherent scattering of neutrinos or antineutrinos on electrons. The NC channel happens equally for all neutrino flavours whereas the CC scattering is unique to the electron flavour. Taken from [23].

evolution of the neutrino flavour states but with the additional potential V_e .

$$i \frac{d}{dt} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \left[U_{N=2} \begin{pmatrix} \frac{m_1^2}{2E} & 0 \\ 0 & \frac{m_2^2}{2E} \end{pmatrix} U_{N=2}^\dagger + \begin{pmatrix} V_e & 0 \\ 0 & 0 \end{pmatrix} \right] \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} \quad (2.32)$$

$$= \frac{\Delta m^2}{4E} \begin{pmatrix} -\cos 2\theta \pm 4EV_e & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = H_{M_f} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}. \quad (2.33)$$

The first term on the right hand side of Equation 2.32 is the free particle Hamiltonian, rotated into the flavour basis by sandwiching it between $U_{N=2}$ and $U_{N=2}^\dagger$, where $U_{N=2}$ is given by Equation 2.26⁵. Multiplying out this product and adding the external potential results in Equation 2.33, an effective Hamiltonian in the flavour basis, H_{M_f} describing neutrino propagation in matter.

H_{M_f} must now be diagonalised to find the energy levels $E_{M_f,\eta}$ and their corresponding flavour eigenstates in matter. This can be done by a unitary transformation, parametrised by some angle θ_M , namely

$$H_{M_f,\text{diag}} \equiv \begin{pmatrix} E_{M_f,1} & 0 \\ 0 & E_{M_f,2} \end{pmatrix} \quad (2.34)$$

$$= U_{M,N=2}^\dagger H_{M_f} U_{M,N=2} \\ = \begin{pmatrix} \cos(\theta_M) & -\sin(\theta_M) \\ \sin(\theta_M) & \cos(\theta_M) \end{pmatrix} H_{M_f} \begin{pmatrix} \cos(\theta_M) & \sin(\theta_M) \\ -\sin(\theta_M) & \cos(\theta_M) \end{pmatrix}. \quad (2.35)$$

⁵The free particle Hamiltonian's eigenvalues $\left(\frac{m_j^2}{2E}\right)$ come simply from taking the ultra-relativistic limit of the plane wave solution, $E_j = \sqrt{p^2 + m_j^2} \approx p + \frac{m_j^2}{2p} \approx p + \frac{m_j^2}{2E}$, $p \sim E$. p is a term common to all states and, therefore, enters as a common phase, dropping out when computing probabilities and therefore irrelevant to oscillations. Hence, it can be ignored.

By directly comparing the matrix entries of Equation 2.34 and Equation 2.35, it can be shown that for $E_{Mf,j} = \mp \frac{\Delta m_M^2}{2E}$, negative for $j = 1$, positive for $j = 2$,

$$\sin(2\theta_M) \equiv \frac{\sin(2\theta)}{A_M} \quad (2.36)$$

$$\Delta m_M^2 \equiv \Delta m^2 A_M \quad (2.37)$$

$$A_M = \sqrt{\left(\cos(2\theta) \mp \frac{2EV_e}{\Delta m^2}\right)^2 + \sin^2(2\theta)}. \quad (2.38)$$

The calculation to reach this point is non-trivial but written in this form allows for the most intuitive evaluation of the result. It can be seen that both the Hamiltonian and mixing matrix take the same form as in vacuum oscillations, but with modified effective masses. This implies that the oscillation probabilities also take the same form but with modified parameters $\theta \leftrightarrow \theta_M$ and $\Delta m^2 \leftrightarrow \Delta m_M^2$. The \pm sign in the matter potential (Equation 2.31) appears as a \mp in the expression for A_M , implying different effective parameters for neutrinos and antineutrinos for the same matter. Note that in the limit that $N_e \rightarrow 0$, representing matter becoming less dense and tending toward vacuum, $A_M \rightarrow 1$ meaning $\theta_M \rightarrow \theta$ and $\Delta m_M^2 \rightarrow \Delta m^2$, as expected.

Although deriving the analogous result for three neutrino flavours is far more complex, the take-away message is the same. Propagation in matter changes the effective mixing angles and mass splittings, leading to different appearance and disappearance probabilities for neutrinos compared to antineutrinos. This apparent ‘fake CP’ violation can pose problems for oscillation experiments if matter effects are not correctly taken into account. Largely, however, the matter effect is of great use to experiments. In the case of NOvA, it helps to determine the neutrino mass hierarchy. This will be discussed in detail in Section 2.4.

2.4 Oscillations at NOvA

In its 3-flavour long-baseline oscillation physics program, NOvA has the following goals:

1. Measure the atmospheric mass splitting, Δm_{32}^2 .
2. Determine the value (and octant) of θ_{23} . The lower octant of θ_{23} is defined to be the region where $\theta_{23} < \pi/4$ whereas $\theta_{23} > \pi/4$ defines the upper octant.
3. Constrain the values of the CP violating phase, δ_{CP} .
4. Determine the neutrino mass hierarchy.

Items 1. and 2. can be achieved (except for the octant determination) via the disappearance channel $\mathcal{P}(\nu_\mu \rightarrow \nu_\mu)$ alone. Items 3. and 4., however, require the appearance channel $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$. Exactly how these measurements are made is the subject of the next sections.

2.4.1 Muon Neutrino Disappearance

Defining $\Delta_{ij} \equiv \Delta m_{ij}^2 L/4E$, the relevant disappearance probability for NOvA in vacuum is given by

$$\begin{aligned}
 \mathcal{P}(\nu_\mu \rightarrow \nu_\mu) &= 1 - 4 \sum_{j < k} U_{\mu j} U_{\mu j}^* U_{\mu k}^* U_{\mu k} \sin^2(\Delta_{jk}) \\
 &= 1 - 4 (|U_{\mu 1}|^2 |U_{\mu 2}|^2 \sin^2(\Delta_{12}) + |U_{\mu 2}|^2 |U_{\mu 3}|^2 \sin^2(\Delta_{23}) + |U_{\mu 1}|^2 |U_{\mu 3}|^2 \sin^2(\Delta_{13})) \\
 &\approx 1 - 4 (c_{23}^2 s_{12}^2 c_{23}^2 c_{12}^2 \sin^2(\Delta_{12}) + c_{23}^2 c_{12}^2 s_{23}^2 \sin^2(\Delta_{23}) + c_{23}^2 s_{12}^2 s_{23}^2 \sin^2(\Delta_{13})) \\
 &\approx 1 - 4 \cos^2 \theta_{23} \sin^2 \theta_{23} (\cos^2 \theta_{12} + \sin^2 \theta_{12}) \sin^2(\Delta_{23}) \\
 &= 1 - (2 \cos \theta_{23} \sin \theta_{23})^2 \sin^2(\Delta_{23}) \\
 &= 1 - \sin^2 2\theta_{23} \sin^2(\Delta_{23}). \tag{2.39}
 \end{aligned}$$

Two simplifications based on current experimental knowledge have been used to reach Equation 2.39. Firstly, θ_{13} is small compared to the other mixing angles, therefore, $\sin^2 \theta_{13} \sim 0$ and $\cos \theta_{13} \sim 1$. Secondly, $|\Delta m_{21}^2|$ is known to be a factor of 30 smaller than $|\Delta m_{31}^2|$. By the cyclic relation $|\Delta m_{31}^2| = |\Delta m_{32}^2| \pm |\Delta m_{21}^2|$ (\pm depending on the neutrino mass hierarchy), $|\Delta m_{32}^2| \approx |\Delta m_{31}^2|$. This is known as the One Mass Scale Dominance (OMSD) approximation. Therefore, at leading order and as highlighted in the discussion of the two flavour approximation (Section 2.3.2), the amplitude of oscillation as a function of energy directly relates to $\sin^2 2\theta_{23}$. Additionally, the position of the first oscillation maximum directly relates to Δm_{32}^2 .

A limitation of using the disappearance channel to measure θ_{23} is that the probability is symmetric about $\theta_{23} = \pi/4$, or, invariant under transformations $\theta_{23} \rightarrow \pi/2 - \theta_{23}$ [26]. In practise, therefore, NOvA has no sensitivity to the octant of θ_{23} in this channel. The case that $\theta_{23} = \pi/4$ is known as ‘maximal mixing’. It is defined in this way since it describes the situation where the third mass eigenstate ν_3 , is made up of equal amounts of ν_μ and ν_τ ⁶. Whether nature has chosen maximal mixing is of great interest to the neutrino community, potentially pointing to a fundamental symmetry in the lepton sector. This

⁶ $\theta_{23} = \pi/4 \implies U_{\mu 3} = U_{\tau 3} = \frac{1}{2} \cos \theta_{13}$.

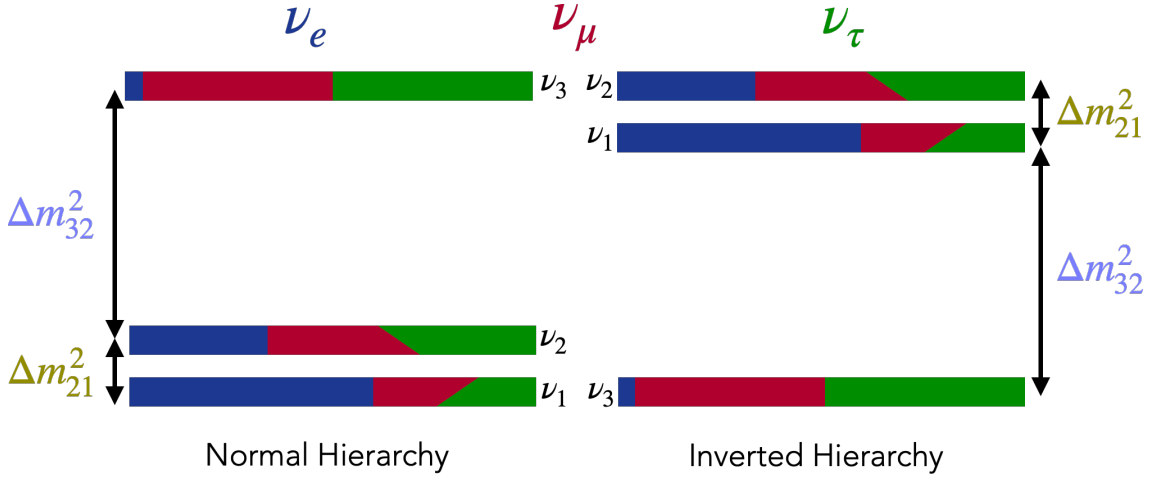


Figure 2.6: A schematic describing the two neutrino mass hierarchy scenarios. It is not currently known experimentally whether the third neutrino mass eigenstate is the heaviest or the lightest. The colour breakdown of each mass state represents their rough flavour composition. Diagonal boundaries between colours represent varying values of δ_{CP} .

situation is distinct from maximum *disappearance* or, the value of θ_{23} which leads to the largest amplitude of oscillation. For muon neutrinos at NOvA, maximum disappearance happens at around $\theta_{23} = 0.514$ due to higher order terms and small contribution of the matter effect. The point of maximum disappearance for muon antineutrinos is distinct from that of neutrinos, again due to contributions from the matter effect.

A final interesting observation is that these oscillations depend on the sine squared of the atmospheric mass squared difference. For this reason, it is not possible to establish the sign of Δm_{32}^2 . Therefore, it is not known whether the third mass eigenstate is heavier or lighter the other two. This is known as the neutrino mass hierarchy problem, explained diagrammatically in Figure 2.6. If an experiment were sensitive enough to measure simultaneously and to the necessary precision Δm_{32}^2 and Δm_{31}^2 , the hierarchy could be established. Currently, no experiment probes the hierarchy in this way. NOvA's sensitivity is in the electron neutrino appearance channel, discussed next.

2.4.2 Electron Neutrino Appearance

Defining again $\Delta_{ij} \equiv \Delta m_{ij}^2 L / 4E$, the relevant appearance probability for NOvA in matter is given by [32]

$$\mathcal{P}(\nu_{\mu}^{(-)} \rightarrow \nu_e^{(-)}) = \mathcal{P}_{\text{atm}} + \mathcal{P}_{\text{sol}} + 2\sqrt{\mathcal{P}_{\text{atm}}\mathcal{P}_{\text{sol}}}(\cos\delta \cos\Delta_{32}^{(+)} - \sin\delta \sin\Delta_{32}), \quad (2.40)$$




Figure 2.7: Electron neutrino appearance probability in vacuum for muon neutrinos (**left**) and muon antineutrinos (**right**), for a given set of oscillation parameters as indicated in the legend. The solid blue curves show the case that CP is conserved, or $\delta_{CP} = 0, \pi$. The green and red curves correspond to maximum CP violation, $\delta_{CP} = 3\pi/2$ and $\delta_{CP} = \pi/2$ respectively. Taken from [33].

where

$$\sqrt{\mathcal{P}_{\text{atm}}} \equiv \sin \theta_{23} \sin 2\theta_{13} \sin \left(\frac{\sin(\Delta_{31} - aL)}{\Delta_{31} - aL} \right) \Delta_{31}, \quad (2.41)$$

$$\sqrt{\mathcal{P}_{\text{sol}}} \equiv \cos \theta_{23} \sin 2\theta_{12} \sin \left(\frac{\sin(aL)}{aL} \right) \Delta_{21}. \quad (2.42)$$

Like the disappearance case, OMSD has been applied and higher order terms in $\sin^2 \theta_{13}$ have been dropped. Additionally, $a \equiv \pm G_F N_e / \sqrt{2}$ where the positive sign is for neutrinos and the negative sign for antineutrinos. For perspective, $|a|$ for the Earth is roughly $1/3500 \text{ km}$.

At NOvA's baseline, the \mathcal{P}_{sol} term is negligible since it depends on higher order terms in the solar mass splitting. The next-to-leading order cross term (the term picking up contributions from both \mathcal{P}_{atm} and \mathcal{P}_{sol}) is the 'CP violating piece'. If δ is equal to 0 or π , Equation 2.40 becomes the same for neutrinos and antineutrino, implying that CP is conserved. If δ takes any other value, this implies CP is violated and why δ is often written δ_{CP} . Figures 2.7a and 2.7b show that although a small effect, NOvA has some sensitivity to δ_{CP} via the number of appearance ν_e 's and $\bar{\nu}_e$'s. At the first appearance maximum the value of δ_{CP} determines whether there is: an enhancement of neutrinos and suppression of antineutrinos ($\delta_{CP} = 3\pi/2$), an enhancement of antineutrinos and suppression of neutrinos ($\delta_{CP} = \pi/2$), equal numbers of both $\delta_{CP} = 0, \pi$.




Figure 2.8: Electron neutrino appearance probability in vacuum compared to matter for muon neutrinos (**left**) and muon antineutrinos (**right**), for a given set of oscillation parameters indicated in the legend. The black curves show the case of vacuum, the blue and red curves represent the normal and inverted hierarchies in matter respectively. The shaded areas are drawn out by varying values of δ_{CP} . Taken from [33].

As discussed in Section 2.4.1, the electron neutrino appearance channel gives NOvA its sensitivity to the neutrino mass hierarchy. Oscillations through matter bring in a dependence on the hierarchy. For electron neutrinos in the normal hierarchy, there is an enhancement of the appearance probability compared to vacuum. In the inverted hierarchy, there is a suppression. This behaviour, and the corresponding behaviour for electron antineutrinos is shown in Figure 2.6.

2.5 Experimental Status

As alluded to in Section 2.1, the experimental study of neutrino oscillations already has a rich history. Five of the key parameters governing oscillations are relatively well measured: the mixing angles (θ_{12} , θ_{13} , θ_{23}) and mass splittings (Δm_{21}^2 , Δm_{32}^2). The sign of Δm_{32}^2 and the octant of θ_{23} , however, are yet to be determined. The complex phase δ_{CP} is the least well constrained. NOvA, amongst other running (Tokai to Kamioka (T2K) [34]) and future (Deep Underground Neutrino Experiment (DUNE) [35]) accelerator based experiments, is working to improve the measurement of this parameter, the mass hierarchy and the octant.

This section is not intended to be a comprehensive historical record but highlights the most recent and the most precise measurements in each of the ‘sectors’ (see Section 2.3). The latest public 3-flavour oscillation results from NOvA are discussed in Section 2.5.5.

For completeness, Table 2.1 summarises the latest oscillation fits to global data from the NuFit collaboration [36].

2.5.1 Measurements in the Solar Sector

The ‘solar’ sector refers to the parameters θ_{12} and Δm_{21}^2 . Historically probed using neutrinos produced in nuclear reactions in the Sun with experiments such as SNO [20] and Super-Kamiokande [19], the current best measurement of the mass splitting comes from the Kamioka Liquid-scintillator Antineutrino Detector (KamLAND) [37]. KamLAND uses electron antineutrinos with energies $\mathcal{O}(1 - 10)$ MeV from an array of 53 commercial nuclear reactors. The detector, at a flux weighted average distance of 180 m from the reactors, is a 13 m spherical balloon filled with 1 ktonne of liquid scintillator at a depth of 2700 m water equivalent. The electron antineutrino survival probability as a function of the effective baseline, divided by neutrino energy is shown in Figure 2.9a. The oscillatory nature of the data is clearly visible. The blue histogram shows a 3-flavour oscillation fit to KamLAND data alone, corresponding to best fit values of $\sin^2 \theta_{12} = 0.316_{-0.026}^{+0.034}$ and $\Delta m_{21}^2 = 7.54_{-0.18}^{+0.19} \times 10^{-5} \text{ eV}^2$. The KamLAND measurement of the mass splitting, however, is in tension at about the $1.4\text{-}\sigma$ level with a joint fit of Super-Kamiokande and SNO data. A comparison of each experiment’s contours, presented at The XXIX International Conference on Neutrino Physics and Astrophysics (Neutrino 2020), is shown in Figure 2.9b.

2.5.2 Measurements in the Reactor Sector

The ‘reactor’ sector refers to the parameters θ_{13} and Δm_{31}^2 . Despite its small size, θ_{13} is currently the most precisely known mixing angle. Three experiments, all proposed in the mid-2000s, have contributed to this effort using inverse β decay to measure the flux of electron antineutrinos from fission processes inside nuclear reactors. In general, electron antineutrino appearance probabilities are measured as a function of energy and fitted to a 3-flavour model. The appearance formula is dependent on both θ_{13} and Δm_{31}^2 (and Δm_{32}^2) at leading order. For the mass splittings, reactor experiments use an approximation which replaces the part of the formula dependent on Δm_{31}^2 and Δm_{32}^2 with an *effective* mass splitting, Δm_{ee}^2 ⁷. This parameter can then either be treated as a free or be constrained

⁷Details of this procedure are given in [39]. In short, defining a Δm_{ee}^2 removes dependence on the mass hierarchy. It is not, however, a fundamental parameter. Therefore, a transformation must be made to recover a physically meaningful quantity, reintroducing the hierarchy.

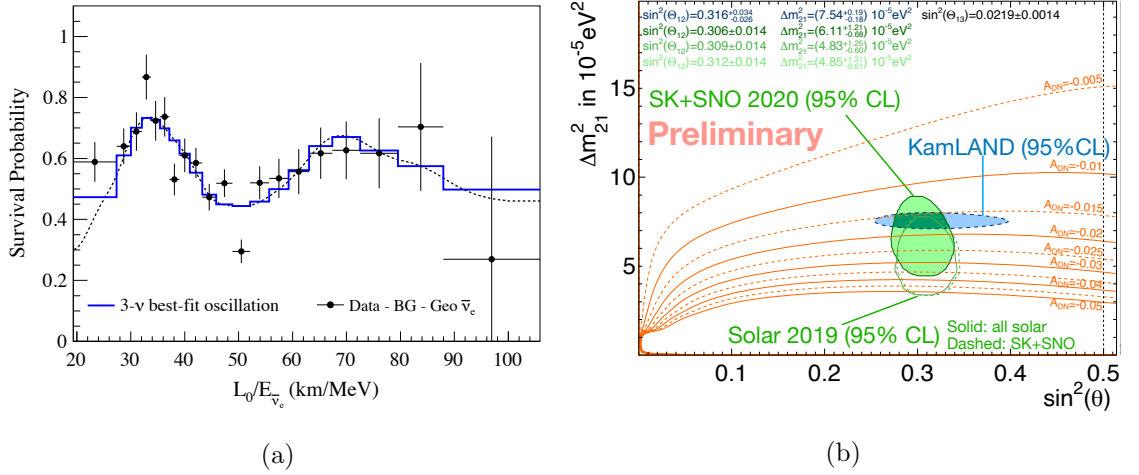


Figure 2.9: **Left:** Election antineutrino survival probability as a function of effective baseline over neutrino energy for the KamLAND experiment. The data (black) is fitted with a 3-flavour oscillation model with no external constraints (blue). Taken from [37]. **Right:** A comparison of contours in the Δm_{21}^2 and θ_{12} space for KamLAND, Super-Kamiokande and SNO. Some very mild tension between the Δm_{21}^2 best fit values hailing from reactor data and solar data, can be seen. A_{DN} is a metric to quantify the amount of asymmetry in the number of neutrino events observed in the day compared to the night. Taken from [38].

by measurements from the atmospheric sector.

Early indications of a small, but non-zero value of θ_{13} came from the Double-Chooz collaboration in 2011 [40]⁸. The experiment consists of one (eventually two) detector(s) placed at 1050 m (400 m) from a pair of nuclear reactors in France. The neutrino target, 10 m³ of gadolinium doped liquid scintillator, is beneath a hill offering 300 m water equivalent of shielding from cosmic rays. In their latest publication in April 2020, the data prefers a best fit of $\sin^2 2\theta_{13} = 0.105 \pm 0.014$ in the 3-flavour model [41].

Shortly after the initial indications from Double-Chooz, the Daya Bay collaboration reported an analysis disfavouring the $\theta_{13} = 0$ hypothesis at 5.2σ [42]. At the time of this result, the experiment consisted of six gadolinium doped liquid scintillator detectors, each with a target mass of 20 tonnes. They were clustered in three groups (experiment halls) at effective baselines of 512 m, 561 m and 1579 m, from six reactors. A single effective baseline for each experiment hall is established by translating the flux from all six reactors. Unlike the first Double-Chooz result which used the Main Injector Neutrino Oscillation Search (MINOS) measurement of the atmospheric mass splitting to constrain its fit [43], Daya Bay freely fits for the effective mass splitting Δm_{ee}^2 . Figure 2.10a shows the ratio of the measured number of events and the no oscillations expectation for each experiment hall. The χ^2 profile for $\sin^2 2\theta_{13}$ is inset in the top right, showing the exclusion of zero at great than $5\text{-}\sigma$. Another reactor experiment, Reactor Experiment for Neutrino Oscillation (RENO) (situated in South Korea), confirmed this result less than a year later at $4.9\text{-}\sigma$ [44]. Daya Bay's latest result was published in 2018 [45]. A best fit of $\sin^2 2\theta_{13} = 0.0856 \pm 0.0029$ and $\Delta m_{32}^2 = 2.471_{-0.070}^{+0.068}(-2.575)_{-0.070}^{+0.068} \times 10^{-3} \text{ eV}^2$ assuming the normal (inverted) hierarchy, was found for the 1958 days of data. The corresponding contours are shown in Figure 2.10b. These measurements of θ_{13} and Δm_{32}^2 are currently world leading.

A non-zero value of θ_{13} is key to measuring the CP violating phase δ_{CP} via long-baseline accelerator neutrino oscillation experiments. This is clear from the \mathcal{P}_{sol} term of Equation 2.40. By the same reasoning, experiments like NOvA do in fact have some sensitivity to θ_{13} . However, the behaviour of the probability when changing θ_{13} is completely degenerate with the behaviour when changing θ_{23} . Thus, the sensitivity in this channel is

⁸The first indications came from T2K. In the same year, they reported electron neutrino appearance from a muon neutrino beam with a significance of $2.5\text{-}\sigma$ [34].

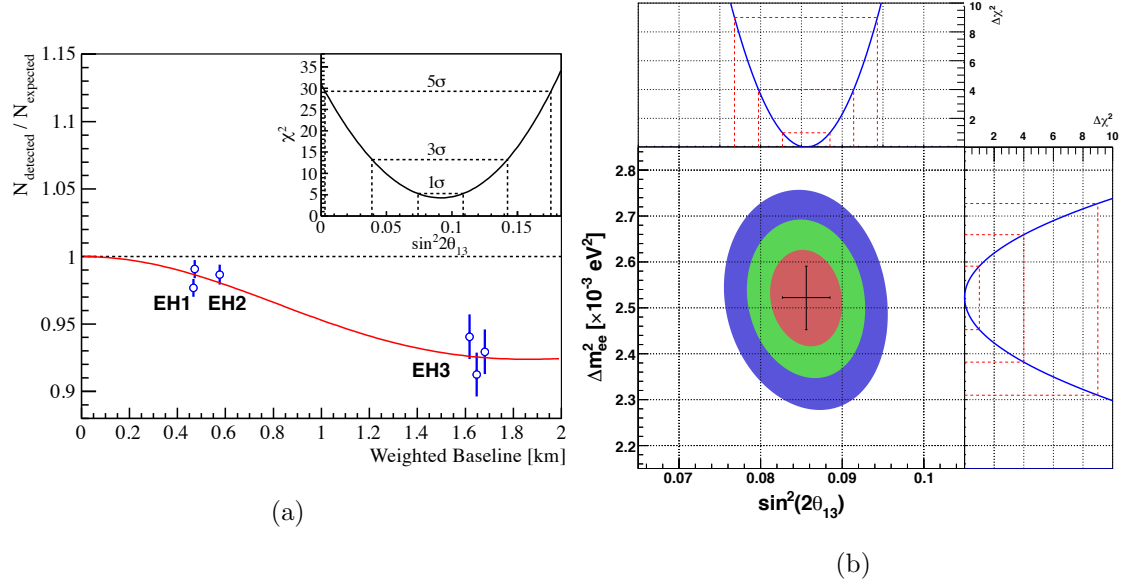


Figure 2.10: **Left:** Ratio of the number of observed inverse β decay events and no oscillations prediction for Daya Bay's first analysis in 2012. Each blue point represents data from a single detector and the red curve is the best fit. Error bars on the data show both the statistical and total uncorrelated systematic error. The inset plot shows the χ^2 profile for $\sin^2 2\theta_{13}$. Zero is clearly excluded at greater than 5- σ . Taken from [42]. **Right:** The 1-, 2- and 3- σ confidence intervals in θ_{13} and Δm_{ee}^2 space for Daya Bay's most recent publication. The panels on the top and right of the figure show the One Dimensional (1D) $\Delta\chi^2$ for each parameter. The best fit and corresponding 1D uncertainties are given by the black cross. Taken from [45].

much less than that achieved in the electron antineutrino disappearance channel. Long-baseline experiments measuring $\mathcal{P}(\nu_\mu \rightarrow \nu_e)$, therefore, usually constrain θ_{13} in their fits using reactor measurements.

2.5.3 Measurements in the Atmospheric Sector

The ‘atmospheric’ sector refers to the parameters θ_{23} and Δm_{32}^2 . It has been and continues to be, one of the most studied. The importance of the parameters to the neutrino mass hierarchy and potential μ/τ symmetry has already been discussed in Section 2.4.1. The IceCube, MINOS, Super-Kamiokande and T2K collaborations have all released results in this sector since 2019 [38, 46–48], as has NOvA [1]. NOvA’s result is discussed in detail in Section 2.5.5, a brief summary of the results from the other collaborations is given here.

- **IceCube:** The IceCube experiment detects atmospheric and astrophysical neutrinos via their interaction with approximately one cubic kilometre of instrumented ice, 1.5 km below the surface, close to the South Pole. To probe the atmospheric mixing parameters, tau neutrino appearance from atmospheric muon neutrinos ($\nu_\mu \rightarrow \nu_\tau$) is used. Neutrinos with reconstructed energies between 5.6 GeV and 56 GeV are taken forward into the final selection. In their last analysis published in 2019, the normal hierarchy and upper octant were preferred with a 2D best fit at $\sin^2 \theta_{23} = 0.58_{-0.13}^{+0.04}$ and $\Delta m_{32}^2 = 2.55_{-0.11}^{+0.12} \times 10^{-3} \text{ eV}^2$.
- **MINOS:** At Neutrino 2020, MINOS and its successor MINOS+ presented its final oscillation analysis, combining beam and atmospheric data from ten years of running. The experiment was made up of two detectors, a near and a far, separated by a baseline of 735 km. They were functionally identical, constructed from alternating planes of plastic scintillator strips and steel. They were also magnetised and placed on-axis with respect to the muon neutrino enhanced beam with an energy peaked around 3 and 7 GeV for MINOS and MINOS+ respectively. Muon neutrino disappearance and electron neutrino appearance was used to probe the atmospheric parameters. The normal ordering and lower octant were found to be slightly preferred, with best fit values $\sin^2 \theta_{23} = 0.43_{-0.04}^{+0.20}$ and $\Delta m_{32}^2 = 2.40_{-0.09}^{+0.08} \times 10^{-3} \text{ eV}^2$. Maximal mixing was disfavoured at around the 1- σ level.
- **Super-Kamiokande:** Super-Kamiokande’s importance to the understanding of solar neutrinos has already been highlighted. It also definitively demonstrated the oscillation of atmospheric neutrinos and, therefore, showed that neutrinos have mass. The

detector has a fiducial mass of 22.5 ktonnes and is composed of a cylindrical tank of water. Cherenkov light from final state particles of neutrino interactions is picked up by approximately 11000 Photo-multiplier Tubes (PMTs). Different particles in the medium produce light signals of distinct topologies and this property can be used to identify particle type. The light can also be used to determine the direction of the incoming neutrino, allowing the rate of neutrinos of different flavours to be measured as a function of angle (baseline). In the absence of oscillations, the rate is roughly isotropic. A 3-flavour oscillation model can be applied to explain angular dependence in the data. Super-Kamiokande's most recent result was released at Neutrino 2020, the normal hierarchy was preferred and the best fit was found to be at $\sin^2 \theta_{23} = 0.44^{+0.05}_{-0.02}$ and $\Delta m_{32}^2 = 2.40^{+0.11}_{-0.12} \times 10^{-3} \text{ eV}^2$.

- T2K: NOvA's main 'competitor', T2K is a long-baseline experiment in Japan. A complex of three near detectors sits 280 m downstream of a muon neutrino enhanced beam that has an energy peaked at 600 MeV. The Super-Kamiokande detector functions as a far detector, 295 km further downstream and 2.5 degrees off-axis. Muon neutrino disappearance and electron neutrino appearance are used to probe the atmospheric parameters. In April 2020, a T2K article focussing on CP violation was published [49]. This is briefly discussed in Section 2.5.4. Three months later at Neutrino 2020, a more general result was presented. The data preferred the normal hierarchy and upper octant, with best fit at $\sin^2 \theta_{23} = 0.55^{+0.02}_{-0.05}$ and $\Delta m_{32}^2 = 2.49^{+0.06}_{-0.08} \times 10^{-3} \text{ eV}^2$.

Figure 2.11 shows the most recent contours from all of the experiments discussed in this section. Assuming the normal neutrino mass hierarchy, there is general agreement at the 90% confidence level. With the exception of Super-Kamiokande, contours are largely symmetric about maximal mixing ($\sin^2 \theta_{23} = 0.5$).

2.5.4 Constraints on δ_{CP}

The parameter quantifying to what extent neutrino oscillations violate CP symmetry, δ_{CP} , is currently the field's least well measured. MINOS, NOvA and T2K all provide weak constraints, generally in slight tension with each other.

In a paper published in April 2020, T2K observed data explainable by a sizeable increase in the neutrino oscillation probability [49]. This led to the exclusion at $3\text{-}\sigma$ of 46% of δ_{CP} values, corresponding to those giving an enhancement of the antineutrino




Figure 2.11: Contours in θ_{23} and Δm_{32}^2 space, from the most recent atmospheric sector analyses. Assuming the normal mass hierarchy, the 90% confidence intervals are shown; each collaboration is represented by its own colour and is labelled. Taken from [50].




Figure 2.12: Asimov sensitivity to CP violation at DUNE, as a function of the true value of δ_{CP} , for ten years of exposure. The green shaded area represents the effect of changing θ_{23} within the 3- σ NuFit range of values [36]. Similar treatment of Δm_{32}^2 and θ_{13} yield smaller variations in the sensitivity compared to the central value evaluation. Taken from [51].

oscillation probability. With the addition of more data, at Neutrino 2020 T2K reported a similar preference but with weaker limits, excluding 35% at 3- σ [48].

Future accelerator based oscillation experiments with longer baselines, larger detectors and more powerful beams hope to make a more precise measurement of δ_{CP} . Figure 2.12 shows the Asimov sensitivity of DUNE [51], a planned liquid argon based experiment, as a function of the true value of δ_{CP} . If the normal hierarchy is assumed and with 10 years of running, DUNE has the ability to exclude the CP conserved case at 5- σ for greater than 50% of δ_{CP} 's range.

2.5.5 Latest 3-flavour Results from NOvA

This section will briefly summarise the results of the latest (2020) 3-flavour analysis of NOvA's data, presented at Neutrino 2020 [1]. Chapter 6 discusses improvements made to

this analysis and Chapter 7 presents the results of those improvements. It is important to note that the same dataset and simulation are used in the version presented here and the version presented in later chapters.

Figure 2.13 shows the Protons on Target (POT) accumulating against time for the dataset used in the 2020 analysis. At NOvA, the beam mode where the number of neutrinos in the beam is enhanced is known as Forward Horn Current (FHC) mode. The corresponding antineutrino mode is known as Reverse Horn Current (RHC) (see Section 3.1.2 for details). The total FHC and RHC exposures are 13.60×10^{20} and 12.50×10^{20} POT respectively, collected between 2014 and 2020. 3-flavour oscillation results are extracted from a simultaneous fit of data in both beam modes, in the muon neutrino disappearance and electron neutrino appearance channels.

Figures 2.14a and 2.14b show the distributions of reconstructed neutrino energy at the FD for selected muon neutrinos and antineutrinos in FHC and RHC respectively. There were 211 candidates in FHC (on a background of 8.2 events) and 105 candidates in RHC mode (on a background of 2.1 events). Analogous spectra for electron neutrinos are shown in Figure 2.15⁹. The number of appearance electron neutrinos is less than the number of surviving muon neutrinos. 82 and 33 candidates were observed in FHC and RHC respectively.

The resulting confidence intervals in $\sin^2 \theta_{23}$ and δ_{CP} space are shown in Figure 2.16a for both the normal and inverted hierarchies. The best fit at $\sin^2 \theta_{23} = 0.57_{-0.03}^{+0.04}$ and $\Delta m_{32}^2 = 2.41 \pm 0.07 \times 10^{-3} \text{ eV}^2$ shows a preference in the data for the normal hierarchy (at $1\text{-}\sigma$) and upper octant (at $1.2\text{-}\sigma$). With regard to δ_{CP} , the best fit is found to be at 0.82π and there is at least one value in each choice of hierarchy and octant that is not disfavoured at greater than $1\text{-}\sigma$. $\delta_{CP} = 3\pi/2$ in the normal hierarchy is excluded at greater than $3\text{-}\sigma$ while $\delta_{CP} = \pi/2$ in the inverted hierarchy is disfavoured at approximately $2\text{-}\sigma$. This is nicely explained in Figure 2.16b by the so-called ‘bi-event’ plot, or, the number of electron neutrino candidates compared to electron antineutrino candidates. No strong asymmetry in the number of each type of event is observed.

⁹The electron neutrino spectra are not directly analogous to the muon neutrino spectra. Disappearance conserves CP, therefore, antineutrinos in the neutrino mode sample are counted as signal and vice versa in antineutrino mode. Appearance does not necessarily conserve CP, therefore, electron (anti)neutrinos in the (neutrino) antineutrino sample are counted as background.




Figure 2.13: The number of POT as a function of time for the NOvA 2020 dataset. Each point represents 24 hours of data taking and the solid lines show the cumulative totals. Orange and blue regions represent neutrino and antineutrino beam running modes respectively. Taken from [52].

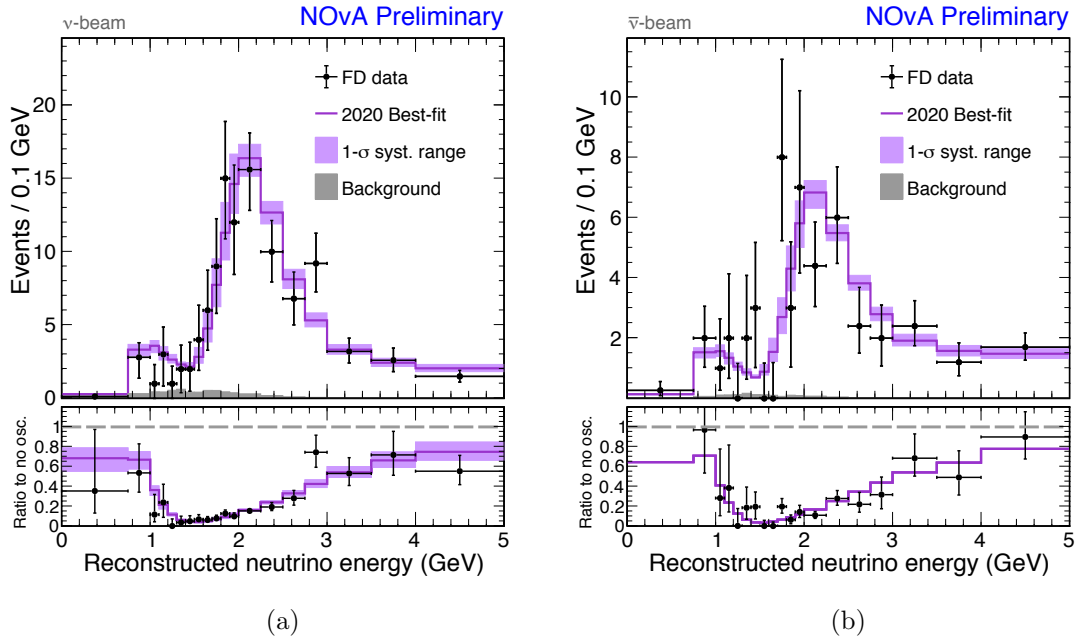


Figure 2.14: The 2020 analysis reconstructed neutrino energy spectra at the FD with muon neutrino selection cuts and energy estimator applied. Neutrino and antineutrino beam modes are shown on the left and right figures respectively. The ratio to no oscillations for both data and the best fit prediction are shown beneath the spectra. Taken from [1].

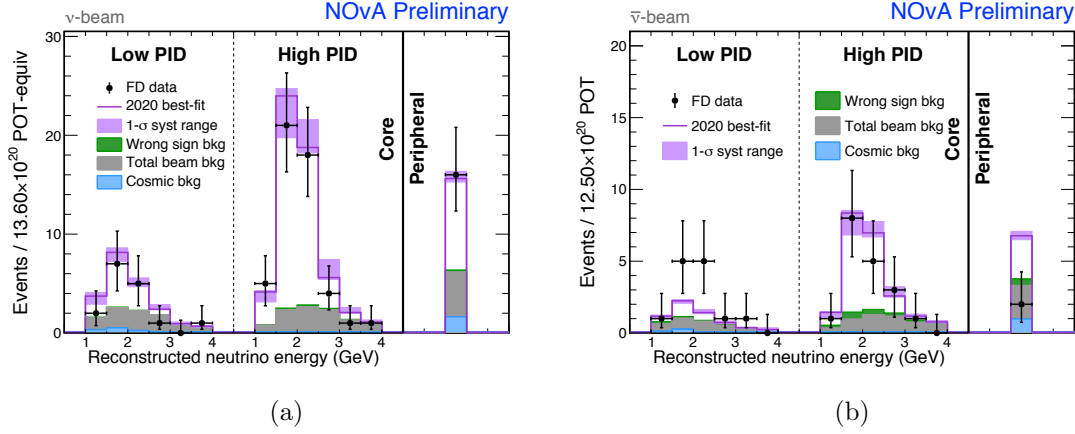


Figure 2.15: The 2020 analysis reconstructed neutrino energy spectra at the FD with electron neutrino selection cuts and energy estimator applied. Neutrino and antineutrino beam modes are shown on the left and right figures respectively. The binning scheme is explained fully in Section 4.4.2. Taken from [1].

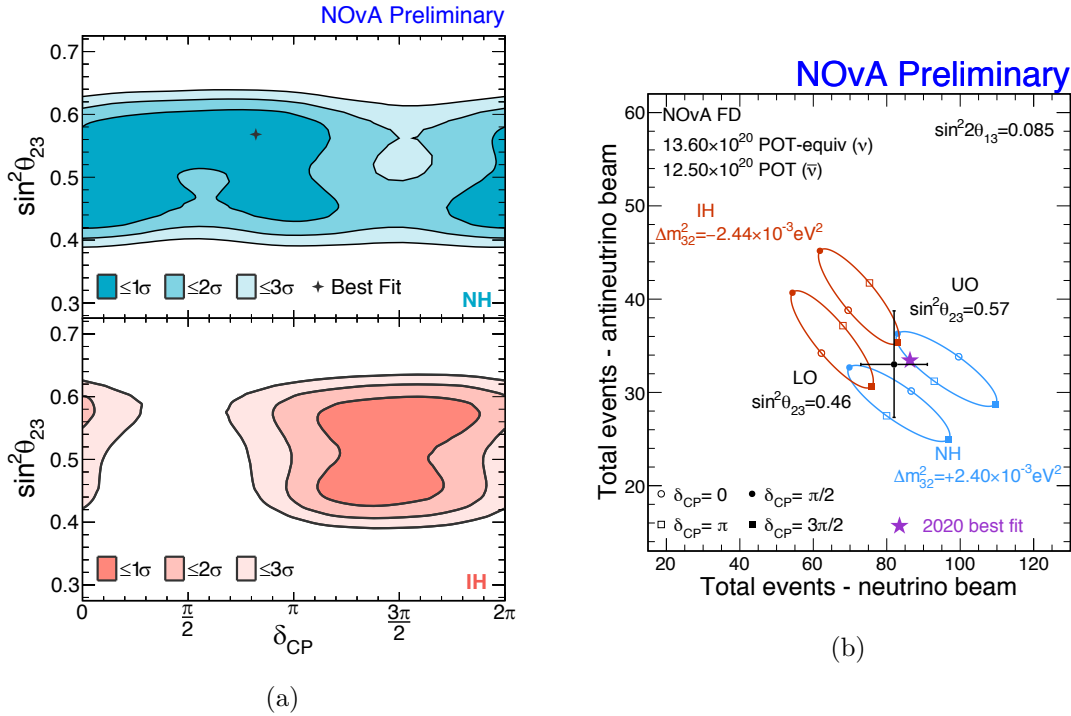


Figure 2.16: **Left:** 1-, 2- and 3- σ contours for the 2020 3-flavour oscillations fit in $\sin^2\theta_{23}$ and δ_{CP} space. **Right:** So-called ‘bi-event’ plot, the number of electron neutrino candidates compared to electron antineutrino candidates is shown, as well as the best fit prediction. For choices of oscillation parameters in the four octant and hierarchy scenarios, an ellipse is traced out by varying the value of δ_{CP} . No strong asymmetry in the number of events was observed. Taken from [1].

		Normal Ordering (best fit)		Inverted Ordering ($\Delta\chi^2 = 2.7$)	
		bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range
without Super-Kamiokande atmospheric data	$\sin^2 \theta_{12}$	$0.304^{+0.013}_{-0.012}$	$0.269 \rightarrow 0.343$	$0.304^{+0.013}_{-0.012}$	$0.269 \rightarrow 0.343$
	$\theta_{12}/^\circ$	$33.44^{+0.78}_{-0.75}$	$31.27 \rightarrow 35.86$	$33.45^{+0.78}_{-0.75}$	$31.27 \rightarrow 35.87$
	$\sin^2 \theta_{23}$	$0.570^{+0.018}_{-0.024}$	$0.407 \rightarrow 0.618$	$0.575^{+0.017}_{-0.021}$	$0.411 \rightarrow 0.621$
	$\theta_{23}/^\circ$	$49.0^{+1.1}_{-1.4}$	$39.6 \rightarrow 51.8$	$49.3^{+1.0}_{-1.2}$	$39.9 \rightarrow 52.0$
	$\sin^2 \theta_{13}$	$0.02221^{+0.00068}_{-0.00062}$	$0.02034 \rightarrow 0.02430$	$0.02240^{+0.00062}_{-0.00062}$	$0.02053 \rightarrow 0.02436$
	$\theta_{13}/^\circ$	$8.57^{+0.13}_{-0.12}$	$8.20 \rightarrow 8.97$	$8.61^{+0.12}_{-0.12}$	$8.24 \rightarrow 8.98$
	$\delta_{CP}/^\circ$	195^{+51}_{-25}	$107 \rightarrow 403$	286^{+27}_{-32}	$192 \rightarrow 360$
	$\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$	$7.42^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.04$	$7.42^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.04$
	$\frac{\Delta m_{3l}^2}{10^{-3} \text{ eV}^2}$	$+2.514^{+0.028}_{-0.027}$	$+2.431 \rightarrow +2.598$	$-2.497^{+0.028}_{-0.028}$	$-2.583 \rightarrow -2.412$
		Normal Ordering (best fit)		Inverted Ordering ($\Delta\chi^2 = 7.1$)	
		bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range
with Super-Kamiokande atmospheric data	$\sin^2 \theta_{12}$	$0.304^{+0.012}_{-0.012}$	$0.269 \rightarrow 0.343$	$0.304^{+0.013}_{-0.012}$	$0.269 \rightarrow 0.343$
	$\theta_{12}/^\circ$	$33.44^{+0.77}_{-0.74}$	$31.27 \rightarrow 35.86$	$33.45^{+0.78}_{-0.75}$	$31.27 \rightarrow 35.87$
	$\sin^2 \theta_{23}$	$0.573^{+0.016}_{-0.020}$	$0.415 \rightarrow 0.616$	$0.575^{+0.016}_{-0.019}$	$0.419 \rightarrow 0.617$
	$\theta_{23}/^\circ$	$49.2^{+0.9}_{-1.2}$	$40.1 \rightarrow 51.7$	$49.3^{+0.9}_{-1.1}$	$40.3 \rightarrow 51.8$
	$\sin^2 \theta_{13}$	$0.02219^{+0.00062}_{-0.00063}$	$0.02032 \rightarrow 0.02410$	$0.02238^{+0.00063}_{-0.00062}$	$0.02052 \rightarrow 0.02428$
	$\theta_{13}/^\circ$	$8.57^{+0.12}_{-0.12}$	$8.20 \rightarrow 8.93$	$8.60^{+0.12}_{-0.12}$	$8.24 \rightarrow 8.96$
	$\delta_{CP}/^\circ$	197^{+27}_{-24}	$120 \rightarrow 369$	282^{+26}_{-30}	$193 \rightarrow 352$
	$\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$	$7.42^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.04$	$7.42^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.04$
	$\frac{\Delta m_{3l}^2}{10^{-3} \text{ eV}^2}$	$+2.517^{+0.026}_{-0.028}$	$+2.435 \rightarrow +2.598$	$-2.498^{+0.028}_{-0.028}$	$-2.581 \rightarrow -2.414$

Table 2.1: 3-flavour oscillation parameters from a fit to 2020 global data by the NuFit collaboration. The lower and upper sections show results obtained with and without atmospheric data from Super-Kamiokande included^A. Best fits in the normal and inverted hierarchy scenarios are included in the left and right columns respectively. Taken from [36].

^AThere was not enough information available to the NuFit collaboration to make an independent analysis of Super-Kamiokande atmospheric data comparable in detail to that performed by the collaboration. The Super-Kamiokande data is included in the global fit via a tabulated χ^2 map.

Chapter 3

The NOvA Experiment

NOvA is a two-detector, long-baseline experiment designed primarily to study the PMNS matrix and neutrino mass hierarchy, via muon neutrino disappearance and electron neutrino appearance in a muon neutrino beam. The two detectors, a near and a far, are functionally equivalent and used to measure the neutrino energy and flavour composition of the NuMI beam at two distances from the target. Both detectors are located off-axis at 14 mrad and their almost identical construction allows for partial cancellation of systematic uncertainties correlated between the them. The 300 tonne Near Detector (ND) is located 100 m underground, 1.015 km from the point at which the protons hit the target at Fermilab, allowing the flux to be sampled before oscillations have taken place. The much larger, 14 kilo-tonne Far Detector (FD) operates on the surface near Ash River, Minnesota, a further 809 km from the target. This detector is used to search for oscillation phenomena. The choice of baseline and off-axis position means that the FD experiences a narrow energy flux at the first oscillation maximum, enhancing sensitivity to oscillations.

This chapter gives an overview of the NOvA experimental set-up. Section 3.1 discusses the neutrino source and Section 3.2 explains the near and far detectors. The operational design differs very little from the one initially described in the NOvA Technical Design Report (TDR) [53], where additional information on many of the topics discussed in this chapter can be found.

3.1 The NuMI Beam

Neutrinos from the Main Injector (NuMI) [54] is the name given to the facility providing the flux of neutrinos or antineutrinos that NOvA uses for its oscillation and cross-section




Figure 3.1: Schematic drawing of the NuMI beam facility showing the target hall, decay pipe, hadron monitor, absorber and muon monitors. Taken from [54].

analyses. A schematic drawing of the facility is shown in Figure 3.1, displaying all of its key components. It is just one part of a complex of accelerators and other machinery located at Fermi National Accelerator Laboratory (Fermilab) in the US mid-west, shown in Figure 3.2. Protons originating from an ion source are gradually accelerated through the complex until they reach 120 GeV. These primary particles are then directed to impinge on the NuMI target, this is typically done in batches of $\sim 10^{13}$ particles over $10 \mu\text{s}$ at intervals of approximately 1.33 s. This is known as a ‘spill’ and creates a directed cascade of secondary hadrons, mostly pions, which go on to decay to neutrino-charged lepton pairs. The next sections walk through this process in more detail.

3.1.1 Primary Proton Beam

The 120 GeV protons directed at the NuMI target start as a beam of 35 keV hydrogen ions, generated by an ion source. These ions are fed into Fermilab’s 150 m long Linear Accelerator (LINAC), accelerating them to 400 MeV, before their electrons are stripped using a thin gold foil and they are passed to the Booster. The Booster is a synchrotron with a diameter of 152 m and accelerates the bare protons to 8 GeV. From here they are passed into the fixed field Recycler ring where a technique known as ‘slip-stacking’ is used. Slip-stacking is key to giving NuMI the ability to run at higher intensities and involves grouping together 2 batches of protons from the Booster. Due to the much larger size of the Recycler compared to the Booster, it can accommodate up to six of these double bunches. This is known as the ‘6+6’ configuration. As a final step, the slip-stacked protons are passed to the Main Injector (MI) where they are accelerated to their final momentum of 120 GeV. While this is being done, in parallel, the next batches of protons are being stacked in the recycler.




Figure 3.2: The Fermilab accelerator complex, showing the path of protons from ion source to NuMI target [55].

3.1.2 Secondary Hadron Beam

Once at 120 GeV, the protons are extracted and directed to impinge on the 0.95 m long graphite target in the NuMI beam facility (see Figure 3.1). The collisions between the accelerated protons and carbon atoms of the target create a spray of hadrons, mostly pions and kaons, which are focussed to form a secondary beam.

The focussing is done by a pair of magnetic horns placed down-stream of the target, each composed of an outer and parabolically shaped inner conductor. When pulsed with a high current, a magnetic field with a strength inversely proportional to the radial distance from the central axis is created. The shape of the inner conductor results in the path length of a charged particle moving through the conductor to be proportional to the square of its radial entry distance. Putting these two pieces of information together under the thin lens approximation, the result is a lens with a focal length proportional to a particle's incident momentum. The sign of the current can be flipped to focus or deflect a certain particle charge. The modes where negatively and positively charged particles are focussed are known as Reverse Horn Current (RHC) and Forward Horn Current (FHC) modes respectively. Figure 3.3 shows possible trajectories of particles incident on the horns. It




Figure 3.3: Possible trajectories of particles incident on the magnetic focussing horns. The use of two horns ensures that any particle under or over focussed by the first horn, is re-focussed by the second. Particles which do not enter the region between the inner and outer conductors, those with low transverse momentum, are undeflected. Taken from [54].

was demonstrated that the use of two horns over one increases focussing efficiency by 50%. Additionally, the relative positions of each of the horns and target can be varied, to tune the energy of the beam. The further the most down-stream horn is away, the higher the beam energy.

3.1.3 Decay Pipe & Neutrino Beam

After passing the second horn, the focussed hadrons enter the 675 m long decay pipe. The decay pipe is filled with helium which, due to its long interaction length, gives the secondary mesons time to decay in flight to tertiary mesons, charged leptons and neutrinos. At this point, the vast majority of particles in the beam are pions with momenta of order GeV, although there is some contribution from kaons. The most relevant decay modes in each horn current are summarised in Table 3.1. The meson's daughters are Lorentz boosted down the pipe towards the Absorber, designed to prevent any unwanted particles from travelling further. Around 80% of unwanted beam which reaches the Absorber consists of primary protons which didn't interact, the remainder is mostly secondary protons and mesons that did not decay. A small fraction (4%) is made up of electrons, neutrons and gammas. Approximately 240 m of rock after the Absorber is used to range out any remaining muons, some however propagate further and decay themselves via (in RHC mode) $\mu^- \rightarrow \nu_\mu + e^- + \bar{\nu}_e$. This is one way that an electron neutrino component is introduced to the beam, constituting an irreducible background to the ν_e appearance

Horn Current	Decay Mode	Branching Ratio
FHC	$\pi^+ \rightarrow \mu^+ + \nu_\mu$	0.9999
	$K^+ \rightarrow \mu^+ + \nu_\mu$	0.6355
RHC	$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$	0.9999
	$K^- \rightarrow \mu^- + \bar{\nu}_\mu$	0.6355

Table 3.1: Most common decay modes for charged pions and kaons. Branching ratios taken from [26].

analysis. The same background can also arise from the decays of K^0 and K^\pm via, for example,

$$K_L^0 \rightarrow \pi^- + e^+ + \nu_e \quad (3.1)$$

$$K^+ \rightarrow \pi^0 + e^+ + \nu_e. \quad (3.2)$$

The energy spectrum of these electron neutrinos has a peak at a higher energy than those hailing from pions, around 7.5 GeV at NOvA. Charged and neutral kaons are also the dominant parent of muon neutrinos at higher energy. These neutrinos have an energy spectrum which peaks around 12 GeV.

3.1.4 Off-axis Design

A key feature of NOvA is that both detectors sit 14.6 mrad off the axis of the NuMI beam. This is a technique which exploits the kinematics of a two body decay, the dominant decay mode in the creation of a neutrino beam. Taking the specific example of $\pi^+ \rightarrow \mu^+ + \nu_\mu$, in the rest frame of the parent pion, the decay is isotropic. In the case of beam however, in the laboratory frame, the pions decay in flight and the daughters are boosted. The neutrino energy, therefore, becomes a function of the pion energy and the angle of the neutrino with respect to the boost direction, described by the following equation:

$$E_\nu = E_\pi \left(1 - \frac{m_\mu^2}{m_\pi^2} \right) \frac{1}{1 + \theta^2 \gamma^2}, \quad (3.3)$$

E_π , m_π and γ are the energy, rest mass and Lorentz factor respectively of the parent pion. Additionally, m_μ is the rest mass of the muon and θ is the (small) angle between the pion and neutrino directions. The equation can be derived by considering conservation of 4-momentum between parent and daughter particles.




Figure 3.4: Neutrino energy as a function of parent pion energy at different angles with respect to the beam axis [56].

This relationship highlights that the muon neutrino energy spectrum that the detectors see can be tuned. Figures 3.4 and 3.5 show this effect. The reconstructed neutrino energy spectra for ν_μ -CC events at the Near and Far detectors is plotted for different off-axis angles. Although moving further off-axis reduces the overall flux, it becomes more narrowly peaked around the first oscillation maximum (~ 2 GeV). This is a consequence of neutrino energy becoming less dependent on pion energy as the detectors are moved further away from the beam axis.

Another benefit of the detectors being off-axis is enhanced background rejection, particularly in the case of NC events where there is no outgoing charged lepton. In such an interaction, the neutrino carries much of the energy away, causing a tendency in the visible energy to “feed down” to lower energies. With a wide beam, there are more NC events from higher energy neutrinos which have the ability to feed down into the signal region. An effect which is greatly reduced when the beam is narrowed.

3.2 The NOvA Detectors

The primary goal of the NOvA detectors is to identify muons and electrons produced in CC neutrino interactions and measure their energy. The two detectors are liquid scintillator based tracking calorimeters, that are functionally identical. Their main distinctions are their size and location with respect to the surface. The ND is 300 tonne and 105 m




Figure 3.5: Reconstructed energy spectra for ν_μ -CC in the ND (**left**) and FD (**right**) constructed at different angles with respect to the beam axis [57].

underground, in contrast to the FD which is 14 kilotonne and on the surface.

To distinguish muons from electrons, and both of these from backgrounds such as NC events, the detectors are segmented and make use of low- Z materials (primarily carbon). The segmentation is done via the use of many Polyvinyl Chloride (PVC) cells filled with liquid scintillator, each read out by their own wavelength shifting fibre. The array of cells must be sufficiently granular to separate using topology, the showering particles, such as electrons, from non-showering particles, such as muons and pions. Using a low Z material specifically aids with electron-photon separation, an important concept for reducing the amount of NC background in the electron neutrino appearance sample. Often in NC events, a π_0 particle will be produced. This neutral particle deposits no observable energy in the NOvA detectors until it decays, most commonly to two photons. The two single photons are often difficult to resolve, producing the same signature as an electron in the detector. In a true electron neutrino CC interaction, the electron is produced exactly at the primary vertex of the interaction. A π_0 , although produced at the vertex, will travel on average some distance determined by the interaction length of the material before decaying. The vertex at which the two photons are produced is, therefore, displaced from the primary vertex. The long interaction length of carbon ensures that the displaced vertex from a $\pi^0 \rightarrow \gamma\gamma$ decay is large enough to be resolved.

A diagram of the two detectors and the fundamental detection principle is shown in Figure 3.6. Common features, key to their construction and design are explained in more detail in the next sections. Sections 3.2.7 and 3.2.8 then expand on their differences.




Figure 3.6: A scaled depiction of the NOvA detectors, including a comparison to the average person. A ‘zoom’ showing the alternating orientation of cell planes is also included, along with a schematic drawing of a wavelength shifting fibre looped in a single cell. The detection principle is also illustrated. A charged particle passes through the cell, causing the liquid to produce scintillation light. This light bounces around the cell until captured by the fibre, shifted and transported to be collected and amplified by an Avalanche Photodiode (APD). The image on the far right shows the pixelated readout of a single APD.




Figure 3.7: A single NOvA rigid PVC extrusion, 16 cells wide. Taken from [53].

3.2.1 The Cell

The fundamental building block of the NOvA detectors is the cell: rectangular structures of rigid PVC, filled with liquid scintillator and containing a looped, wavelength shifting fibre. Light is collected by the fibre which then shifts and transports it to be collected. Cells are extruded in groups of 16 with a cross-sectional area of 3.8×5.9 cm, as depicted in Figure 3.7. The extrusions are much longer in the FD (15.6 m) than the ND (3.9 m), leading to greater attenuation of light along the fibres in the FD. To ensure that as much of the light as possible is captured by the fibre, the PVC was chosen to be highly reflective. Since light reflects eight times on average before hitting the fibre, an increase in reflectivity of 1% results in an $\sim 8\%$ increase in the amount of light seen by the fibre.

Two groups of these 16 cells are glued together to form a 32 cell ‘module’. These plane-like structures are placed normally to the beam direction, with successive layers turned by 90 degrees with respect to one another. This formation, shown in Figure 3.6, enables NOvA to do Three Dimensional (3D) tracking of particles.

3.2.2 Liquid Scintillator

PVC constitutes approximately 35% of the total detector mass. The other 65% is active material, the liquid scintillator held within the cells. Its composition is detailed in Table 3.2. Aside from the mineral oil used as a solvent, the largest component is pseudocumene - the primary scintillant, which emits light in the range 360 - 390 nm (ultra-violet). Wavelength shifters are included to shift this light to 400 - 450 nm (violet / blue), to line-up with the absorption spectrum of the wavelength shifting fibre. Furthermore, a small amount of tocopherol is added as an antioxidant. This is important since oxygen contamination of scintillator can lead to a degradation of light output. Finally, to avoid charge build up and possible sparking while filling the PVC cells, small amounts of antistatic agent are added. This addition has no effect on the light yield. In total, approximately three million gallons of the scintillator mixture are required to fill the FD and thirty-thousand gallons are required to fill the ND.

3.2.3 Optical Fibre

Each cell contains a looped, wavelength shifting fibre designed to capture light and transport it via total internal reflection to the pixelated array of an APD. A looped fibre can yield almost a factor of four more light than a single fibre with a non-reflecting end, since

Component	Purpose	Mass Fraction
Mineral Oil	Solvent	95.8%
Pseudocumene	Scintillant	4.1%
2,5-Diphenyloxazole (PPO)	Wavelength shifter	0.091%
bis-MSB	Wavelength shifter	0.0013%
Stadis-425	Antistatic agent	0.0003%
Tocopherol	Antioxidant	0.0010%

Table 3.2: Chemical composition of the NOvA scintillator [53].




Figure 3.8: Emission and absorption spectra for the NOvA optical fibre. The wavelengths of light emitted by the scintillator are shifted upwards. Due to an overlap in the spectra, light below 495 nm is severely attenuated, most notably the largest emission peak at 475 nm. Taken from [53].

light is able to travel in either direction down it. The fibre has a polystyrene core, coated with two additional layers of material with decreasing refractive index to maximise total internal reflection. The total diameter of the fibre is 0.7 mm, with the double cladding constituting around 12% of that.

The APDs used in NOvA have a higher quantum efficiency at longer wavelengths. The polystyrene of the core is, therefore, mixed with a dye which shifts the 400 - 450 nm light emitted by the scintillator further toward green, 490 - 550 nm. Due to an overlap in the emission and absorption spectra for the dye, shown in Figure 3.8, wavelengths of light below around 500 nm are severely attenuated.

3.2.4 Photodetectors & Front End Boards

NOvA uses APDs manufactured by Hamamatsu for its photodetection, chosen for their high quantum efficiency in the 520 - 550 nm range of light exiting the fibre ends. They are packaged in arrays of 32 pixels for a one-to-one mapping with the number of cells in a module (see Section 3.2.1). Each individual pixel interfaces with both ends of a cell's single fibre, as shown in the right side of Figure 3.6.

APDs convert light to an electric signal using the photoelectric effect. When photons are absorbed, electron-hole pairs are made. An electric field separates these two charge carriers and if the potential difference is sufficiently high, their velocity will increase to a level where they create additional electron-hole pairs by impact ionisation. These electron-hole pairs are then themselves accelerated and the process is repeated, amplifying the photocurrent. To reduce noise generated by thermal creation of electron-hole pairs, thermal-electric coolers are used to maintain an operating temperature of -15°C . This is key for NOvA to be sensitive to faint light created at the far end of a FD cell.

Electric signals from the avalanche multiplication are read by a Front End Board (FEB) containing an Application Specific Integrated Circuit (ASIC), Analogue to Digital Converter (ADC) and Field Programmable Gate Array (FPGA). This is where the data acquisition process begins. The rate at which a FEB samples an APD is detector dependent, discussed further in Section 3.2.7. Shaping is performed first by the ASIC, giving pulses characteristic rise and fall times. Shaped, analogue signals are then digitised by the ADC and passed onto the FPGA for processing, where the times and amplitudes of the signals are extracted. Any 'hits' with an amplitude above a pre-programmed threshold are sent on for further processing.

3.2.5 Data Concentrator Modules & Triggering

Following low level processing by the FEBs, the APD hits are sent to a farm of Data Concentrator Module (DCM)s. Up to 64 APDs feed into a single DCM. Their purpose is to integrate all information they receive in a $50\text{ }\mu\text{s}$ snapshot in time and produce a data packet (microslice). This data packet is then sent to a buffer farm via Ethernet where it is kept while a triggering decision is made. There are two main types of trigger:

- **Clock Based**, data around a particular point in time is stored. An example is the NuMI trigger which stores $550\text{ }\mu\text{s}$ of data symmetrically about the $10\text{ }\mu\text{s}$ beam spill.

A Master Timing Distribution Unit (MTDU) connects to an external Global Positioning System (GPS) antenna, driving a master clock known as ‘NOvA Time’. The MTDU drives multiple slave TDUs, each of which synchronise a group of DCMs back to the master clock. The MTDU is also connected to the accelerator network. In the case of the NuMI trigger, each time there is a NuMI spill, the MTDU converts the accelerator timestamp to NOvA time. Data in the buffer nodes with a timestamp around the spill time is readout.

A related trigger type is a ‘signal’ based trigger. This is when an external signal to store data for an specified duration is received. An example is the cosmic pulse trigger which stores 550 μs of FD data at 10 Hz. The data is used for detector calibration and estimating cosmic background in the various analyses.

- **Data Driven**, data is read out because specific conditions are met. An example would be the supernova trigger. This trigger looks for multiple small clusters of hits and their distribution in time. The length of data read out (in time) is determined by the time series of the small clusters and the maximum amount of storage on the buffer nodes.

3.2.6 Detector Assembly

The NOvA detectors are built from the PVC extrusions described in Section 3.2.1. Extrusion modules (two PVC extrusions glued together) were constructed at the Module Factory at the University of Minnesota, Minneapolis. Here, the modules were tested and assembled with other detector components. A completed extrusion module features an end plate stuck across the entire assembly, with centre and side plates added for protection against leakage. The top is sealed with a manifold cover through which the fibres are direct to the ‘snout’. The snout routes the fibres to the APD pixel array and provides the fill and vent ports that allow the module to be filled with liquid scintillator. An electronics box, supported by the snout, contains the APD and FEB. Once complete and leak tested, the modules were shipped to either the ND or FD sites.

3.2.7 Near Detector

The NOvA ND is located in a purpose-built cavern, 105 m below ground and 1.015 km downstream of the NuMI target at Fermilab. The cavern is accessible from the shaft used




Figure 3.9: A plan view of the ND cavern with respect to the MINOS installations and beamline. Taken from [53].

by the MINOS experiment [58]. As discussed in Section 3.1.4, the ND sits 14.6 mrad away from the central axis of the beam. Furthermore, since the detector is on a level grade, the NuMI beam enters at a slight downward angle of 3° . Figure 3.9 shows the position of the cavern with respect to the MINOS hall and the axis of the beam.

Images of the ND from the front and back, highlighting all of the main feature are shown in Figure 3.10. It can be seen that the detector has two distinct regions. The most upstream is fully active, constructed from 192 planes, each three extrusion modules, $3.9\text{ m} \times 4.9\text{ m}$, wide (high). The most downstream region is known as the ‘muon catcher’. It has a similar structure to the upstream section but includes 10 layers of 10 cm thick steal placed between each of the 11 pairs of active planes. Since the steal planes were recycled from the smaller NOvA prototype Near Detector on the Surface (NDOS), the muon catcher is $2/3$ the height of the rest of the detector. Its purpose is to contain muons that would otherwise escape through the back. The total ND structure contains 20192 readout channels in 214 planes for a combined length of 15.8 m in the beam direction and 300 tonne total mass.

The primary purpose of the ND is to measure the flux of the NuMI beam before oscillations take place. It is, therefore, much closer to the NuMI target than the FD and experiences a much larger flux. During a typical spill, there are between five and ten neutrino interactions in the ND. To avoid pile-up of physics events, greater timing resolution is required compared to the FD. This is achieved by increasing the rate at which the FEBs sample the APDs by a factor of four (every 125 ns rather than every 500 ns). The result of this increase is a timing resolution in the ND of 5 ns, compared to 15 ns in the FD. The ND has been taking data since August 2014.




Figure 3.10: Recent photographs of the NOvA ND, including a person for perspective. The front of the detector is shown on the **left** image, looking downstream. The muon catcher is shown in the foreground of the **right** image. The electronics boxes containing the FEBS and DCMs are clearly visible on the detector top [55].



Figure 3.11: Photographs of the FD hall. The detector was assembled in 1 kilo-tonne blocks of planes horizontally, before being lifted and pivoted by 90° into position (**left**). The complete detector is shown on the **right** [55].

3.2.8 Far Detector

The NOvA FD is constructed from 896 planes normal to the beam direction and contains a total of 344064 channels. The result is an active area of $15.6\text{ m} \times 15.6\text{ m}$ facing the beam, a length of 60 m in the beam direction and total mass of 14 kilo-tonne. Approximately 65% of that mass is active (scintillating). Images taken during and following the assembly of the FD are shown in Figure 3.11.

Like the ND, the FD sits 14.6 mrad away from the central axis of the beam but, in contrast, is on the surface. The surface position presents a very different challenge for data taking and offline analysis, compared to the ND. To mitigate the number of cosmic rays entering the detector, it is shielded by 4 ft (122 cm) of concrete and 0.5 ft (15 cm) of barite. This overburden provides approximately 14 radiation lengths of protection, reducing the number of cosmic ray induced events in the detector significantly. Even so, the FD still experiences cosmic ray events at roughly 150 kHz. A small fraction of this cosmic data is recorded and used to calibrate the detectors (see Section 4.1.2). It is also used to estimate the number of cosmic rays contributing to the oscillation analyses' backgrounds. In particular, primary cosmic ray photon events can produce electromagnetic showers in the detector, mimicking ν_e -CC interactions. Furthermore, cosmic ray neutrons can interact to produce pions. Such events can be mistaken for ν_μ -CC interactions. In addition to the overburden, several offline techniques are used reduce such backgrounds, explained further

in Section 4.1.3. The FD has been taking data since October 2013.

3.3 Simulation

NOvA uses a wide range of tools to produce accurate simulations of several distinct physical processes. The extraction of oscillation parameters relies on predictions of the beam flux, interaction of neutrinos and the response of the detector to particles produced in these interactions. Without robust modelling and analysis, interpretation of experimental data is not possible.

3.3.1 Beam Simulation

NOvA's model of the neutrino flux hails from the simulation of hadron production in the NuMI target, followed by the subsequent transport and decay of these hadrons. A Geant4 [59] based simulation featuring an accurate description of beamline geometry (including focussing) and material composition, known as G4NuMI [60], is used for this purpose. A 120 GeV beam of protons with a Gaussian profile in the transverse plane is used as input to the simulation. The simulation chain ends at the point when a neutrino is produced¹.

To improve the prediction of hadron production, corrections to the nominal neutrino flux are made based on constraints from external thin and thick target experiments such as NA49 [61]. The corrections are introduced by applying weights to the neutrino yield, based on the hadronic interaction history of each individual neutrino². Weights are calculated by the Package to Predict the Flux (PPFX), developed by the MINERvA collaboration for the NuMI beam [62]. The package also estimates uncertainties on the spectrum based on the external data using a multi-universe technique (Section 4.6.1). Figure 3.12 shows the neutrino composition of the flux after PPFX corrections have been applied, in both FHC and RHC modes.

¹Muons which are produced in conjunction with the neutrinos are tracked further.

²Truth information about hadronic interactions which eventually lead to a neutrino are tabulated at the generation stage. These tables contain details of the kinematics of all interactions and the amount of material traversed by every particle.




Figure 3.12: Major components of the neutrino flux at the FD (**top**) and ND (**bottom**) in FHC (**left**) and RHC (**right**) modes, after corrections calculated using external hadron production data are applied. Taken from [63].

3.3.2 Interactions & Cross-sections

Monte Carlo (MC) neutrino generators use the meson decay points and momenta information from G4NuMI (Section 3.3.1) to generate neutrinos and simulate their interactions with the detectors and surrounding material. For the analysis presented in this thesis, NOvA uses the GENIE generator, version 3 [24]. This version allows for the user to choose and customise the underlying physics models. Weights to the GENIE simulation, motivated by both our own ND and external data, as well as external theory, are applied. The Comprehensive Model Configuration (CMC) that NOvA uses is N18_10j_02_11a, based on the out of the box CMCs G18_10b_02_11a and G18_10j_00_000 with the following changes:

- In the Charged Current Quasi-elastic (CCQE) channel, a sophisticated nuclear physics model from Nieves *et al.*, involving a local Fermi gas model with corrections from the Random Phase Approximation (RPA) is used [64]. Systematic uncertainties come from the ‘z-expansion’ fit of the axial form factor from Meyer *et al.* [65]. This involves replacing the traditional dipole structure of the CCQE form factor with the analytic expression predicted by QCD and making a fit to neutrino-deuteron scattering data.
- For CC MEC mediated events, where the neutrino interacts with a correlated pair of nucleons, the Valencia model is used [66]. A fit of this model to NOvA ND data in the 2D reconstructed visible hadronic energy and three momentum transfer space is made to improve data-MC agreement.
- For RES events, the Berger-Sehgal model [67] with parameters tuned by the GENIE collaboration to better agree with neutrino-deuteron scattering data is used.
- For DIS, the Bodek-Yang model with parameters tuned by NOvA to neutrino-deuteron scattering data is used [68].
- The GENIE hN model for Final State Interactions (FSI) is used, with parameters tuned by NOvA to better match pion-carbon scattering data.

Figure 3.13 shows the data and MC comparison of the reconstructed visible hadronic energy in the ND of events passing the ν_μ -CC selection, with and without NOvA’s adjustment to the cross-section model. Better data-MC agreement is achieved, mostly driven by the changes made in relation to MEC events.




Figure 3.13: Data-MC comparison of reconstructed visible hadronic energy in FHC (**left**) and RHC (**right**) beam modes, for events passing the ν_μ -CC selection. The stacked histograms represent the breakdown of MC by true interaction mode. The black solid histogram shows the total simulation after tuning. The corresponding modes in default GENIE are shown by the dashed black histograms. Taken from [70].

An important background for NOvA, most notably at the FD, arises from interactions of cosmic rays. To accurately reflect cosmic activity in simulation, cosmic ray data is ‘overlaid’ into the simulation of beam events. Cosmic rays for detector calibration are simulated using the Cosmic Ray Shower Library (CRY) generator [69].

3.3.3 Detector Simulation

Following the production of particles by a neutrino or cosmic ray interaction, governed by the generators (Section 3.3.2), propagation and energy deposition in the detector materials is simulated by GEANT4 [59]. Evaluation of light production and transport through fibres, modelling of the electronic noise of APD, and the response of FEBs is handled by NOvA specific software modules.

Chapter 4

Analysis Methodology

Performing an oscillation analysis of data from a two-detector, long-baseline experiment involves several steps. Deposits of energy or ‘hits’ in the detectors are clustered to isolate physics events and reconstructed to establish higher level properties (Section 4.1). Based on these properties, cuts are applied to create signal-rich subsets which can then be binned further as a function of, for example, reconstructed energy (Sections 4.3 and 4.4). These spectra can then be tested against a 3-flavour neutrino model dependent on both the oscillation parameters and terms capturing systematic uncertainties (Sections 4.6 and 4.7). In a process called extrapolation, the nominal FD MC is tweaked based on measurements of the neutrino flux at the ND (Section 4.5). The Common Analysis Format Analysis (CAFAna) framework (Section 4.2) provides a standardised and efficient way to tackle many of these steps [71].

This chapter provides a detailed overview of the process to go from raw hits in the NOvA detectors to constraints on the 3-flavour neutrino oscillation parameters, focussing on the electron neutrino appearance channel. The chapter discusses primarily the work of the NOvA collaboration at large, with my own contributions appropriately highlighted.

4.1 Low-level Software

4.1.1 Event Reconstruction

Hit Finding

The search for muon and electron neutrino candidates starts with scintillation light producing hits in the detectors’ APDs that are digitised to give ADC counts as a function of time. A threshold determining whether a particular ADC time bin is interesting is

defined by a technique called Dual Correlated Sampling (DCS). The difference in ADC counts between time bins x and $x - 3$ must be greater than some value set individually for each APD pixel. Using DCS protects against the baseline wandering. Each hit above threshold is recorded in a data product known as a *cell hit*, carrying information about the cell, plane, time and size of the energy deposition. The time is calibrated to remove any differences attributable to different hits being processed by different hardware. After this calibration, the time of a given hit with respect to the data trigger, and to every other hit in the acquisition is known. The cell provides a transverse position and the plane a longitudinal position which, along with the calibrated time, can be used to cluster hits occurring spatially and temporally close to each other. The amount of energy deposited, characterised by the ADC charge value, is also stored along with a corresponding number of photo-electrons (PE) produced based on average APD response¹. The cell hits are now known as *cal hits* and are sorted and saved according to their *header* information². As a result, cal hits from different runs, subruns, and triggers are written to separate files³.

Clustering

NOvA uses the clustering algorithm **TDSlicer** to associate cal hits which occur close together in time and space [72]. Due to the alternating plane configuration of both detectors, explained in Section 3.2.1, each cal hit carries xzt or yzt (*cell, plane, time*) information. Clustering is, therefore, initially done in 3D. The resulting groups are then merged to form a set of Four Dimensional (4D) clusters in $xyzt$, known as *slices*. Hits that are not associated with a slice are labelled as noise and added to the ‘noise slice’ for the event which is carried forward with through the rest of the analysis chain so that it is available to analysers.

The process by which 3D clusters are established is itself two step. First, to establish ‘centres’ of clustering in time, the work of Rodriguez and Liao is used [73]. For each given

¹The conversion from ADC to PE is a simple scaling of the ADC by a single number, different for the ND and FD.

²*Cal* refers to the fact that the hits have been time calibrated. No energy calibration has been carried out at this point.

³The lowest level simulation files are designed to have exactly the same format as these raw data files (with the addition of associated ‘truth’ information or, the information with which the events were generated). This allows reconstruction algorithms applied to both data and MC further down the processing chain to be identical.

hit, a metric quantifying the density (ρ) of hits around that given hit is calculated. An isolation score (δ) for each hit is then evaluated, describing the minimum Euclidean distance to any other hit with higher density. Hits considered ‘central’ are chosen based on their density and isolation, with differing cuts for the ND and FD. Figure 4.1 shows the distribution of density and isolation for a typical clustering problem. From these central hits, the full 3D clusters are created using Prim’s algorithm [74], building them by starting with the central hits and successively adding the hit outside the cluster that is nearest to some hit inside the cluster. There is a cut applied to the distance between two hits, beyond which the hits will not be clustered. Figure 4.2 shows a diagrammatic representation of the algorithm. Circles A, B, C and D can be interpreted as hits and the labelled lines connecting them as the distances between them. Starting Prim’s algorithm at hit A, moving from left to right in the Figure, D is first added to the cluster since it is the closest connected hit to A. Next, B or C can be added. B is added to the cluster via its connection with D. Its connection with A could equally have been chosen. At this point, connecting A to B is no longer valid since that would link two hits already in the cluster. Finally C is added.

Once slices have been established, they are treated as self-contained interactions. All further reconstruction is applied independently to the group of hits forming each individual slice. Due to the very different topologies of ν_μ and ν_e interactions in the NOvA detectors (see Figure 4.3), two separate reconstruction chains are applied to optimise the identification and characterisation of each type. These are discussed next, focussing mainly on ν_e given its relevance to this thesis.

ν_e Reconstruction

The first stage in the electron neutrino reconstruction chain is to identify the global event vertex. A **Hough transform** is used to identify prominent features in the slice, characterised by straight lines [77]. Each 2D point (hit) has an associated Hough space (r, θ) , or, the set of all straight lines passing through that point parametrised by $r_i = x \cos \theta_i + y \sin \theta_i$. The parameters r and θ represent the length and angle with respect to the x axis of the shortest path between the origin and the line in the Hough space⁴. In this example, the set of all straight lines for a given point traces out a sinusoidal curve in the (r, θ) space. The curve is unique to each point. If the curves of several points intersect at the same values of r_i and θ_i , it implies those points are co-linear. In practise, the classic al-

⁴Equivalently, the Hough space could be all (m, c) for lines through a point parametrised by $y = mx + c$. The polar form, however, is more robust against ‘vertical’ lines.




Figure 4.1: The distribution of isolation score δ versus density ρ (**right**) for a typical clustering problem (**left**). The numbers correspond to the density orderings of the data points. Points in the centre of clusters, those with local or global maximal density, tend to have high isolation score. In this plot, the global maximum point 1, is given an isolation score equal to the largest possible distance any two points can have (by virtue of the fact there is no higher density point). The local maximum in the red cluster, point 10, by virtue of the fact there is no higher density point in the red cluster, must have an isolation ‘partner’ in the blue cluster, hence the large distance and large isolation of this point. Points around the maxima score low on isolation due to their proximity to the maxima but have high density. Points 26, 27 and 28 on the other hand have too great an isolation to be associated to any cluster but are also not dense enough to form their own cluster. Therefore, they are treated as noise. Taken from [73].




Figure 4.2: A schematic explanation of a simple application of Prim's algorithm. Adapted from [75].

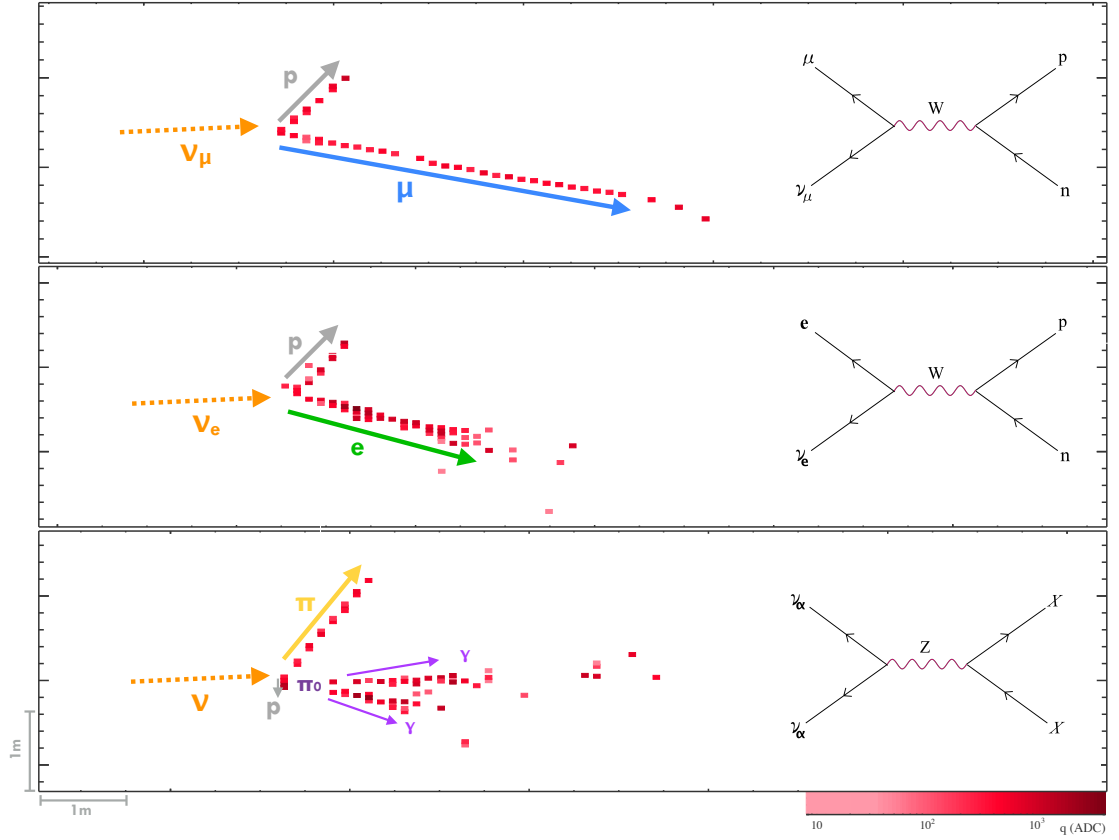


Figure 4.3: Three distinct event topologies seen in NOvA ND data. A ν_μ -CC event (top) is characterised by a long, straight track from the outgoing muon and a shorter, ‘stubby’ track left by the final state proton. In ν_e -CC events (middle), the muon track is replaced with a shower-like object. Both of these event types can be distinguished from a NC event (bottom) by searching for electromagnetic showers displaced from the primary vertex, indicating the presence of a neutral pion. Each event type presents its own unique challenge for reconstruction. Taken from [76].

gorithm calculates r_i and θ_i for all pairs of hits and bins in (r, θ) space, creating a Hough map. Peaks in this distribution correspond to lines which feature in the Hough spaces of several hits and are, therefore, most likely to correspond to lines connecting several hits. NOvA uses a modified version of the classic algorithm [78] which is more robust against falsely assigning lines, accounting for the limits of detector resolution by means of a Gaussian smeared weight. Valid lines for each slice are chosen by taking points in the Hough map where the sum of the weights exceeds some threshold. The algorithm has a tendency to produce too many valid lines. To minimise this, hits associated to the most dominant line (the Hough map peak) are removed and the algorithm re-run. This process is repeated until a predetermined maximum number of valid lines is reached, or, there are no points in the Hough map exceeding the threshold. This iterative process helps to ensure that smaller lines are not missed due to the presence of larger, more dominant ones.

The guidelines laid out by the Hough transform are used as seeds in the **Elastic Arms** algorithm [79]. The algorithm determines the event’s global vertex, assuming that all of the activity in the slice has a common origin. Starting with the intersection points and directions of the Hough lines as seeds, a series of vectors or “arms” pointing away from the intersections points are constructed. For each choice of vertex and arm direction, the value of an energy cost function is calculated. In NOvA, the cost function has three terms. The first captures how well the arms describe the hits and the second penalises for hits that are not assigned to any arm. The third term is a experiment specific, penalising any arm which starts at a distance greater than the mean free path for pair production ($7/9$ of a radiation length) from the vertex. This allows for physically meaningful displaced vertices, common in NC interactions. The optimal vertex is the configuration of arms and vertex position which minimises the cost function.

Once the global vertex has been established, the **Fuzzy k-means** algorithm is used to construct prong objects [80]. Prongs are collections of hits associated with a single particle candidate emerging from the vertex. “Fuzzy” highlights that initially the number of prongs is unknown and that a hit can be a member of multiple prongs. Furthermore, there is no requirement that a hit must be part of any prong, ensuring that hits can still be labelled as noise. The extent to which a particular hit belongs to a particular prong is known as ‘membership’ and is calculated based on the distance of the hit to the centre of a given prong. The algorithm takes hits in the 2D plane (each view is treated separately),

drawing a vector from each of them to the global vertex. These vectors are represented by their angle with respect to the beam direction, allowing for a 1D angular density of hits to be calculated⁵. To establish prong centres and the membership of each hit, an iterative process is used. It is first assumed that all hits belong to one prong, centred on the region of highest density. This position is then corrected, based on the angular separation of each hit from the centre of this prong, until the amount of correction between subsequent steps falls below a certain threshold. Any hits with a membership score of less than some threshold are used to make a new angular distribution of hits, the centre of which is taken as the first iteration point for a second prong. This process is repeated until all hits belong to at least one prong, or, a maximum number of prongs is reached. The final step of Fuzzy-k is to match 2D prongs between views to make 3D prongs. The cumulative energy as a function of distance from the start of each 2D prong is calculated. Counterpart 2D prongs should have a very similar energy profile in either view. Every profile in the XZ view is subtracted from every profile YZ view and the absolute value taken. Pairs with smallest maximal value in this metric are taken as matched [81].

The ν_e reconstruction chain is summarised in the form of a flowchart in Figure 4.4. The final step is to do particle identification, discussed in Section 4.1.3.

ν_μ Reconstruction

The slice level muon neutrino reconstruction is far less complex than that for electron neutrinos, with only one major stage in the processing. ν_μ -CC events are often dominated by the outgoing muon which, unlike showering electrons, produce a sharp, straight track. Any change in trajectory is likely due multiple scattering at small angles along the track. NOvA's tracking algorithm, **Kalman track** [82], is designed for reconstructing such topologies and is based on a Kalman filter routine [83]. Starting with the most upstream hits and working forward, hits are iteratively accepted or rejected in developing tracks based on how the inclusion of that hit affects the track's χ^2 . The calculation of the χ^2 is weighted in a way that favours adjacent hits with small angular deviations. This process is done in the XZ and YZ views separately, as displayed in Figure 4.5. Tracks in the two views are then matched to form 3D objects based on the Z position of the tracks' start and end points.

⁵Based on a hit's distance from the vertex, an associated angular uncertainty due to multiple scattering is assigned. The scale is determined empirically from a simulated sample of muons and electrons between 1 and 2 GeV.

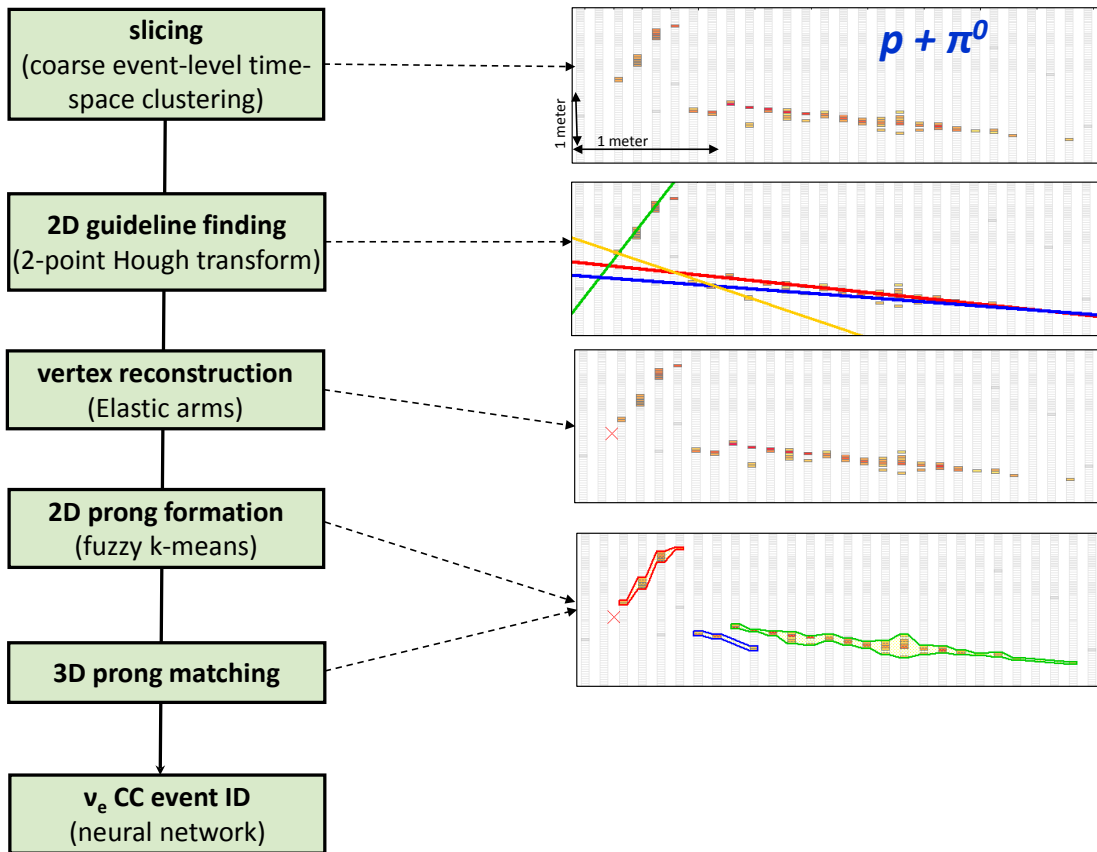


Figure 4.4: The reconstruction chain for the electron neutrino analysis, from the isolation of individual physics events (slicing) to particle identification. Taken from [81].




Figure 4.5: An example of a simulated ν_μ -CC event at the FD reconstructed using the Kalman track algorithm. The top and bottom planes show the XZ and YZ views respectively. The different colours highlight hits associated with separately established tracks in the event. Taken from [82].

A related algorithm is **window tracker**, used to reconstruct both cosmic rays which stop in, and pass straight through the detectors [84]. It works under the assumption that cosmic ray muons follow a straight line trajectory within short portions of Z values. Starting with a set number of planes covering some narrow (user specified) range of Z at the most upstream side of a slice, all hits are compared to a signal (energy) weighted average position. Any hits more than approximately 12 cells away from this position are dropped and those that remain are checked for collinearity. Any point contributing more than 2 units of χ^2 to a straight line fit through the points are removed. A new fit is performed on those that are left and the line is extrapolated into the next window along in Z. The hits in that window are then checked for collinearity and the process repeats. As with Kalman track, a set of 2D tracks are formed in each of the two views and are eventually matched to form 3D tracks.

4.1.2 Energy Reconstruction

As discussed in Chapter 2, long-baseline neutrino experiments traditionally probe the oscillation parameters by measuring the number of neutrino events at a fixed distance as a function of energy. It is, therefore, vital that the true energy of an incident neutrino can be reconstructed. This is done using the charge depositions made by the final state particles. Following the formation of slices (see Section 4.1.1), hits within these slices must be calibrated against standard candles to ensure that their translation to energy units is consistent across detector elements and over time. Following calibration of the hits, analysis level energy estimators are then developed. The estimators for ν_e -CC and ν_μ -CC candidate events are distinct due to the different way in which muons interact with the detectors compared to electrons.

Calibration of Calorimetric Energy

NOvA performs calorimetric energy calibration to allow the deposited energy of any two pieces of recorded activity in the detectors to be meaningfully compared. They must be both comparable to each other and on an absolute scale. It can be summarised in two steps. First, the relative calibration removes any positional dependence in detector response due to, for example, attenuation of light along cell fibres. Next, the absolute calibration assigns an interpretable energy scale (GeV) to the corrected hits from the relative calibration. This process is repeated individually for the different data taking periods, accounting for the drifting of the energy scale due to ageing of the detector as a function of time.

At the FD, both of the calibration stages exploit the abundance and interaction behaviour of cosmic ray muons. Their tracks are easily identified and their energy loss by ionisation is described well by the Bethe-Bloch equation. Hits used in the calibration must belong to tracks identified as cosmic rays that also pass quality cuts affirming they have been well reconstructed. Furthermore, the hits must meet the ‘tricell’ criteria. Namely, a hit is a tricell hit if each neighbouring cell in the same plane of cells also contains a hit. Shown visually in Figure 4.6a, this requirement ensures that for a given hit there is a well defined path length over which the muon deposited energy by ionisation within the central cell. A hit which is not tricell is not guaranteed to have been caused by a muon which passed through both the upper and lower horizontal boundaries of the cell. Without this condition, the path length of the muon within the hit cell cannot be calculated by




Figure 4.6: **Left:** The tricell criterion. The hit in the dark red cell is used in the calibration since hits were also produced in the neighbouring cells. The corresponding path length is obtained using the reconstructed track angle and trigonometry. Taken from [85]. **Right:** The distribution of attenuation corrected PE per unit length as a function of distance from the end of the track, for all tricell hits in selected stopping muon tracks. Each black point shows the mean of a fit made to the peak in each distance slice. Hits in between 100 cm and 200 cm, where the response is most uniform, are used in the absolute calibration to set the energy scale. Taken from [86].

trigonometry.

- **Relative calibration:** As discussed in Section 4.1.1, a single number characterising the average expected APD response is used to convert signals in units of ADC to units of PE. For a given true amount of energy deposited, this observed number of PE is a function of true position in the detectors. More specifically, the APD readout sits at one end of each of the cells. Hits at the furthest point from the readout are greatly attenuated when travelling along the fibre compared to those close by. Therefore, given two equal energy deposits at different points along a cell, the one occurring closest to the readout will be incorrectly interpreted as having a higher energy. The relative calibration accounts for this. A fit to the distribution of mean recorded PE per centimetre from tricell hits in each cell as a function of distance from the readout is performed. In these fits, there are seven free parameters divided amongst an exponential component and a polynomial component. The polynomial component accounts for the different way in which light is reflected toward the ends of the cells compared to the bulk, given the non-reflective nature of the end caps. Plots showing an example fit and the resulting correction are given in Figure 4.7.

The different behaviour of light deposited at the ends of the cells can be seen in the drop in PE per centimetre in data near the edges of the plot. This corrected number of PE, independent of position, is passed onto the absolute calibration.

- **Absolute calibration:** The absolute calibration aims to take the position independent number of PE and convert it to the physically meaningful unit of energy, GeV. To do this, tricell hits from cosmic ray muons which stop inside the detector are used⁶. Figure 4.6b shows the attenuation corrected number of PE deposited per unit path length for such hits, as a function of their distance from the end of their track. This relationship is well described by the Bethe-Bloch equation. In the range of distances between 100 cm and 200 cm from the track end, known as the track window, the detector response is found to be uniform within 1.8% [87]. This track window sits on top of the minimum ionising region of the muon, allowing for a comparison against the known minimum deposition rate in MeV/cm for muons in NOvA's scintillator. Specifically, the calorimetric energy scale is calculated as the mean detector response to a stopping muon tricell hit within the track window in truth (in MeV/cm) divided by the reconstructed mean detector response to a stopping muon tricell hit within the track window (in corrected PE/cm). This number is calculated per detector, per data taking period, for both reconstructed data and reconstructed simulation.

Discrepancies exist between the energy scales calculated for the data and for the MC used in the analysis presented in this thesis [88]. All are below 5% and lead to the inclusion of systematic uncertainties that will be discussed in Section 4.6.

ν_e Energy Reconstruction

For ν_e -CC and $\bar{\nu}_e$ -CC candidate events, the total energy of the incident neutrino $E_{\nu/(\bar{\nu})}$ is calculated using the following quadratic expression

$$E_{\nu/(\bar{\nu})} = \alpha_0 (\alpha_1 E_{\text{EM}} + \alpha_2 E_{\text{Had}} + \alpha_3 E_{\text{EM}}^2 + \alpha_4 E_{\text{Had}}^2). \quad (4.1)$$

E_{EM} is the calorimetric energy of the hits assigned to the reconstructed electromagnetic activity. Prong Convolutional Visual Network (CVN) (see Section 4.1.3) is used to differentiate between activity which is electromagnetic and activity which is hadronic. In

⁶A sample of stopping muons can be isolated by simply looking for tracks where one end is contained within the detector or, by looking for the presence of a Michel electron.




Figure 4.7: **Left:** The mean number of PE per unit path length for tricell hits in cosmic ray data, as a function of position along a FD cell with respect to its centre. The best fit attenuation curve is shown in blue. **Right:** The mean of the ratio of reconstructed and true energy for hits as a function of position along a FD cell with respect to its centre. The red and blue data points show the means before and after relative calibration respectively. Taken from [85].

this context, calorimetric energy is the sum of calibrated hits, scaled additionally to account for dead material and detector thresholds⁷. E_{Had} is the calorimetric energy of any hits remaining in the slice which are not attributed to the electromagnetic activity. The parameters α_i are obtained in a procedure which minimises the variance of reconstructed neutrino energy from true neutrino energy in simulation. The true neutrino energy is weighted so that a fit to it sees a beam flux which is flat in energy, ensuring that the estimator is not biased toward the expected beam peak. The FHC energy resolution, 10.3%, can be read off this plot. The corresponding resolution in RHC is 9.1% [89]. Using a quadratic function improves the energy resolution with respect to using a linear function by approximately 2%. Moving to a cubic fit showed limited improvement on the quadratic fit [90].

ν_μ Energy Reconstruction

For ν_μ -CC and $\bar{\nu}_\mu$ -CC candidate events, the total energy of the incident neutrino $E_{\nu/(\bar{\nu})}$ is calculated in the following way

$$E_{\nu/(\bar{\nu})} = E_{\mu/(\bar{\mu})} + E_{\text{Had}}. \quad (4.2)$$

⁷The scaling factor is obtained from a linear fit of true neutrino energy and visible energy

$E_{\mu/(\bar{\mu})}$ is the energy of the outgoing (anti)muon, estimated from the length of its track. E_{Had} refers to the system's hadronic energy and is estimated by summing the visible energy of all hits not attributed to the muon. This visible energy is then corrected using a spline-based fit obtained from simulation. Figure 4.8b and Figure 4.8d show these piece-wise functions used to estimate the conversion from reconstructed visible energy to hadronic energy, for FHC and RHC respectively. The functions to go from reconstructed track length to true muon energy, used to obtain $E_{\mu/(\bar{\mu})}$, are similarly shown in Figure 4.8a and Figure 4.8c. The smearing matrices corresponding to hadronic energy are clearly less diagonal than those corresponding to muon energy. As a result, the hadronic energy resolution is approximately 30%, much lower than the 4% achieved for the muon. In the latest 3-flavour analysis, the muon neutrino and antineutrino energy resolutions were 9% and 8% respectively at the FD [91]. The difference between the two beam modes is expected, given that on average in FHC, the fraction of energy given to the hadronic shower is greater.

4.1.3 Particle Identification

Performing an oscillation analysis relies on the experiment's ability to observe change in neutrino flavour. To do this slices must be classified according to whether they are, for example, ν_{μ} -CC, ν_e -CC or NC. NOvA uses a number of particle identification techniques and the most important for the 3-flavour analysis are discussed next.

Convolutional Visual Network (CVN)

The CVN [92] is a classification technique belonging to the Convolutional Neural Network (CNN) family of deep learning methods. The general use of CVNs at NOvA is beneficial in two primary ways. Firstly, it allows the analyses to be decoupled from reconstruction inefficiencies such as the misplacing of a vertex. This is because the event slices (discussed throughout this section) are fed into the network as images or 'pixel maps', where each pixel has an intensity proportional to deposited charge. Secondly, the networks extract their own features from these images that they deem important, rather than have a human explicitly decide. This allows for the exploitation of more complex relationships between variables that humans may otherwise miss. 256 such key features are extracted via the application of a series of linear operations on the original image and placed into 'feature maps'. A heuristic explanation of this process is shown in Figure 4.9a.




Figure 4.8: Piece-wise linear fits of the density distributions of true muon energy against reconstructed muon track length (**left**) and true neutrino energy minus reconstructed muon energy against reconstructed visible hadronic energy (**right**) in simulation. The dashed vertical red lines denote boundaries between regions which are fitted to a single straight line. The solid red lines show the fit results, used as scale factors in the ν_μ -CC (**top**) and $\bar{\nu}_\mu$ -CC (**bottom**) energy estimators. The apparent ‘ridge’ around 1 GeV in the hadronic energy plots is a binning effect. Bins above 1 GeV are larger and, therefore, contain more statistics. Taken from [91].

The output of the network is a score between zero and one which is an estimate of how likely a given event is to be truly ν_μ -CC, ν_e -CC, ν_τ -CC, NC or cosmic activity. NOvA's CVN is trained separately in FHC and RHC, using simulated beam events and cosmic ray data. Intrinsic beam ν_e -CC events are treated no differently to appearance ν_e -CC events. Improvements in sensitivity to oscillation parameters presented in this thesis involve binning the ν_e appearance data according to electron neutrino CVN (CVNe) score. Figure 4.9b shows the distribution of CVNe broken down by true interaction type in FD simulation.

In addition to this event level CVN, there is also ProngCVN which attempts to classify individual reconstructed prongs within a slice [94]. The architecture of the network is analogous to event level CVN, with the additional evaluation of XY and YZ pixel maps containing individual prongs. These individual prongs are determined by NOvA's traditional reconstruction meaning the classifier is not entirely decoupled from reconstruction inefficiencies. The output scores of the network correspond to how likely a particular prong is to have originated from an electron, photon, muon, pion or proton. It is currently used by the electron neutrino energy estimator to determine which groups of hits constitute electromagnetic activity [95].

Reconstructed Muon Identifier (ReMId)

Reconstructed Muon Identifier (ReMId) is a Boosted Decision Tree (BDT) used to identify how muon-like each track in a slice is [96]. It is trained on simulated muon tracks from ν_μ -CC events and NC events, representing signal and background respectively. Scores are assigned based on four reconstructed variables provided by the Kalman track algorithm (see Section 4.1.1): deposited energy per unit length log-likelihood, scattering log-likelihood, track length and the fraction of planes overlapping with any hadronic activity. The deposited energy per unit length (dE/dx) is different for, for example, pions and muons⁸. Calculating a log-likelihood provides a metric for comparing each track against the muon energy profile hypothesis. Similarly, scattering log-likelihood captures to what extent each track deviates from a straight line. The score of the highest scoring track for a given slice is taken forward and used as a classification probability describing how likely that event is to be ν_μ -CC.

⁸In addition to ionisation, pions can lose energy via hadronic scattering.




Figure 4.9: **Left:** A heuristic explanation of CVN. Using a series of image processing transformations applied to heat maps representing deposited charge, 256 features key to classification of interaction type are extracted. These 256 features are used to build feature maps. Each feature can be loosely interpreted as a variable or a strength of response to a certain condition in the underlying image. For example, the green, blue and purple features could be seen to as responses to muon tracks, electron showers and hadronic activity respectively. Since this feature map is extracted from a ν_μ -CC event, a stronger response in the green and purple features compared to blue is expected. Taken from [93]. **Right:** The distribution of CVNe in FD simulation broken down by constituent true interaction type. High values of CVNe are overwhelmingly dominated by appearance electron neutrinos. The second largest component are from electron neutrinos arising from the beam. This component is the main (and irreducible background) to the electron neutrino appearance analysis. These events are, of course, electron neutrinos, they just did not ‘start life as’ muon neutrinos. Taken from [92].

Cosmic Rejection BDTs

Due to the surface position of the FD, the background which requires the highest level of suppression comes from cosmic rays. Both the ν_e and ν_μ analyses have their own, separately trained BDTs which assign a score to each slice quantifying how likely the event is to not be a cosmic interaction. For training samples, standard simulated beam events of the appropriate flavour are used for signal, and cosmic data is used for background. A wide range of kinematic variables are used as inputs to the BDTs for both analyses. For ν_μ there are seven such variables [97]: cosine of the angle of the primary track with respect to the beam, primary track length, largest vertical position of the track end points, cosine of the angle of the primary track with respect to vertical, fraction of hits in the slice belonging to the primary track, shortest distance from the primary track to the detector walls and normalised transverse momentum. Note that the BDT does not take any particle identification scores as input, but events used in the training must pass a loose pre-selection: event CVN muon neutrino score must be greater than 0.4 and there must be a track with a ReMid score of greater than 0.4.

The cosmic rejection BDT for ν_e is more complex [98]. Separate BDTs are trained for the core and peripheral samples (introduced in Section 4.3), using eleven and six input variables respectively. The core BDT's variables are similar to those used for ν_μ , but using the shower rather than the track objects. The peripheral BDT's variables were chosen for their separation power among non-fiducial events: normalised X component of reconstructed momentum, normalised Y component of reconstructed momentum, minimum distance of any prong to a detector wall which is not the top, vertex X position, vertex Y position and vertex Z position. As with the muon neutrino BDT, a loose pre-selection is applied to obtain the training sample. Event CVN electron neutrino score must be greater than 0.5. BDT performance is very similar in FHC and RHC, with almost complete background rejection for 90% signal efficiency in both cases. Figure 4.10 shows the classification score for simulated beam events and cosmic data, for each of the separately trained ν_e BDTs. The higher the classification score, the fewer the cosmic ray events.

4.2 High Level Software

Once the raw data and simulation files have been fully reconstructed, the slices are saved as Common Analysis Format (CAF) files [71]. These files are specially organised ROOT




Figure 4.10: The ν_e BDT classification score for simulated beam events and cosmic data, for each of the separately trained samples. **Clockwise from top left:** FHC core, FHC peripheral, RHC peripheral and RHC core. Cosmic, beam background and appearance electron neutrino events are shown by the black, green and blue histograms respectively. The red histograms shows all events. Note that at low BDT scores, cosmic events dominate entirely. Taken from [98].

TTrees [99] which can be easily read by the CAFAna framework. All remaining steps of the oscillation analysis presented in this thesis are performed using this framework. For speed, everything in CAFAna is histogram based. Each event in each CAF file is accessed once, producing histograms (for example, energy spectra) which all subsequent analysis is done in terms of. The software provides a wide range of classes enabling efficient plotting and fitting of oscillated nominal and systematically shifted neutrino spectra, as well as Near to Far Detector extrapolation. For the fit to data itself, CAFAna utilises Minuit [100], the numerical minimisation package.

4.3 Data Quality & Event Selection

NOvA’s 3-flavour analysis relies on fitting FD spectra of ν_μ -CC and ν_e -CC candidates with an oscillation model. Sensitivity to the model is greatly improved when these samples have a good level of signal purity and efficiency. The following sections discuss how the quality of data is assured and how neutrino candidates are isolated and eventually selected.

4.3.1 Data & Spill Quality

Before any analysis specific treatment is applied at the level of slices, there are several stages of processing to ensure that any data used in the measurement of oscillation parameters is not affected by any electronics or data acquisition issues. Firstly, “Good runs”, a set of criteria that determine whether detector and beam conditions were within standard operating ranges during the period that the data were recorded, are applied [101]. Checks are made that, for example, the beam spill occurred during the expected time window, hit rates within each FEB are not too high and all sections of the detector are switched on. Secondly, cuts are made at the *spill* level. This removes short periods of detector or beam failures from runs that would otherwise be marked good. Spills where the beam position, beam width or POT are outside acceptable ranges are removed. Similarly removed are any spills where one or more DCMs feature no hits. The amount of data retained following the application of good runs and spill cuts has increased over time. Measured as a percentage of total POT at the FD, it has climbed from approximately 75% in data taking period one to greater than 99% most recently [101].

As described in Section 3.2, following a beam trigger at the FD, 550 μs of data are read out, with the beam spill occurring between 218 and 228 μs in that window. Only data falling between 217 and 229 μs is taken forward into the oscillation analysis.

4.3.2 Event Quality

The purpose of event quality cuts is to remove any slices which have failed to meet basic reconstruction standards. This includes, for example, ensuring the presence of a vertex or putting a requirement on the number of hits [102]. The exact cuts vary according to detector and are analysis dependent. Furthermore, in the ν_e -CC analysis there are two distinct groups of FD events making up what are known as the core and peripheral samples. The distinction between the two is discussed again in detail in Section 4.4.2. Roughly speaking, however, the peripheral sample is made up of events which are deemed to be highly electron-like by event CVN but fail the containment cuts applied to formulate the core sample. The core sample itself is additionally split into two regions based on electron neutrino CVN score. The core and peripheral samples are completely orthogonal.

- ν_μ -CC: At both the ND and FD, the reconstructed neutrino energy must be less than 5 GeV and there must be at least one Kalman track with an assigned ReMId value (Section 4.1.3). A slice must also contain more than 20 total hits and register activity in more than four contiguous detector planes.
- ν_e -CC: There must be a reconstructed vertex and at least one prong. Additionally, at the FD the reconstructed neutrino energy must be between 1 and 4 GeV, the longest prong must be between 100 and 500 cm and there must be between 30 and 150 total hits in a slice. At the ND, the energy and hit requirements are loosened to the ranges 0 to 4.5 GeV and 20 to 200 hits respectively. For the peripheral sample, the hit and prong length requirements are dropped and the 0 to 4.5 GeV reconstructed energy range is adopted.

4.3.3 Containment

Containment cuts are primarily implemented to ensure that for a given interaction, all of the energy is deposited within the active volume of the detector. This is key to correctly determining the total energy of the incident neutrino. Furthermore, they also help to reject interactions which originated outside the detector (rock events at the ND, cosmic events at the FD). Various metrics to tune these cuts were used in the 2020 3-flavour analysis. The ratio of signal to background, the level of agreement between data and simulation, the overall energy resolution and the fraction of total energy escaping were all considered [102].

- ν_μ -CC at the ND: All activity must be within the ranges $-180 < X, Y < 180$ cm

and $40 < Z < 1525$ cm. Also, only the primary track can enter the muon catcher and the event vertex cannot be located there. Any events with tracks that cross the air gap due to the differing muon catcher height are additionally rejected. Finally, all tracks are extrapolated both forwards and backwards from their end and start points respectively. The projected track must cross at least five planes from the end point and ten planes from the start point before reaching any detector edge to be accepted.

- ν_μ -CC at the FD: Firstly, events with activity closer than 60, 12, 16, 12, 18 and 18 cm from the top, bottom, east, west, front and back detector sides respectively are rejected. Secondly, there must be no hits in the front two or back three detector planes. Finally, a similar track projection cut to that used in the ND is applied. The projection must cover six and seven planes from the end and start of all tracks respectively.
- ν_e -CC at the ND: There are requirements that both the event vertex and all prongs be sufficiently far from the detector edges. Specifically, the entire event must fall within the region $-150 < X < 170$ cm, $-170 < Y < 150$ cm and $150 < Z < 1225$ cm. No part of the event must enter the muon catcher.
- ν_e -CC at the FD: Events with activity closer than 63 cm from the detector top, 12 cm from the bottom, east or west and 18 cm from the front or back are rejected.

4.3.4 Particle Identification (PID)

Particle identification cuts are largely the same in both detectors since ND samples are used to correct the FD predictions (see Section 4.5). As discussed, NOvA's primary method of event flavour classification is to use the output score of event CVN (see Section 4.1.3). For the ν_μ -CC sample, slices with a muon neutrino identification score of greater than 0.8 and containing a Kalman track with a ReMId score of greater than 0.3 are accepted. For the ν_e -CC core sample, slices with an electron neutrino identification score of greater than 0.97 are accepted for the 'high Particle Identification (PID)' bin. Similarly, the corresponding 'low PID' bin accepts events with scores of greater than 0.84 and 0.85 in FHC and RHC respectively.

Additional measures are taken at the FD to ensure the cosmic background is reduced as much as possible. This is done by means of a BDT trained to reject cosmic interactions

Cut Level	ν_μ	$\bar{\nu}_\mu$	Total Beam Bkg	NC	Cosmics	Effic	Purity
NoCut	602	43.1	124	563	2.52×10^6		
Quality	569	41.0	121	490	4.72×10^5	0.946	1.20×10^{-3}
Containment	249	14.1	83.9	336	6.99×10^3	0.413	0.033
Cosmic Rej.	239	13.8	78.2	293	1.54×10^3	0.397	0.111
Full PID	201	12.5	0.717	2.51	6.02	0.334	0.956

Table 4.1: The simulated number of predicted ν_μ -CC candidates at the FD at each stage of the 2020 analysis cut flow at the 2020 3-flavour best fit (see Section 2.5.5), broken down by true interaction type in FHC. Efficiency is calculated with respect to ‘NoCut’, where no slice level cuts have been applied. Taken from [104].

(see Section 4.1.3). In the ν_μ -CC and ν_e -CC analyses, events scoring less than 0.45 (0.45) and 0.49 (0.47) respectively in FHC (RHC) in their specific BDTs are removed. Finally, ν_e -CC candidates must also pass a ‘nearest slice’ cut. This cut looks for cosmogenic-like activity close in space or time to the candidate in question. It is designed to minimise susceptibility to reconstruction failures whereby a single background event is mistakenly sliced into two separate events, one of which looks like signal [95].

As aforementioned, there is a completely independent second sample in the ν_e analysis known as the ‘peripheral.’ A peripheral candidate must pass a more stringent PID criteria than a core candidate. For a CVN score CVNe, and cosmic rejection score CRe, the following must be met. In FHC: CVNe > 0.97 and CRe > 0.60, or, CVNe > 0.995 and CRe > 0.56. In RHC: CVNe > 0.97 and CRe > 0.61, or, CVNe > 0.995 and CRe > 0.57.

Figure 4.11 summarises the cut flow described above for ν_μ -CC and ν_e -CC in diagrammatic form. The number of simulated events at the FD after each major cut category is applied is shown in Tables 4.1 - 4.4. Events are scaled to the 2020 3-flavour analysis exposure.

A large part of this thesis explores to what extent true ν_e -CC events with a lower CVN score can be used improve NOvA’s sensitivity to the oscillation parameters. A full discussion of this begins in Chapter 6.

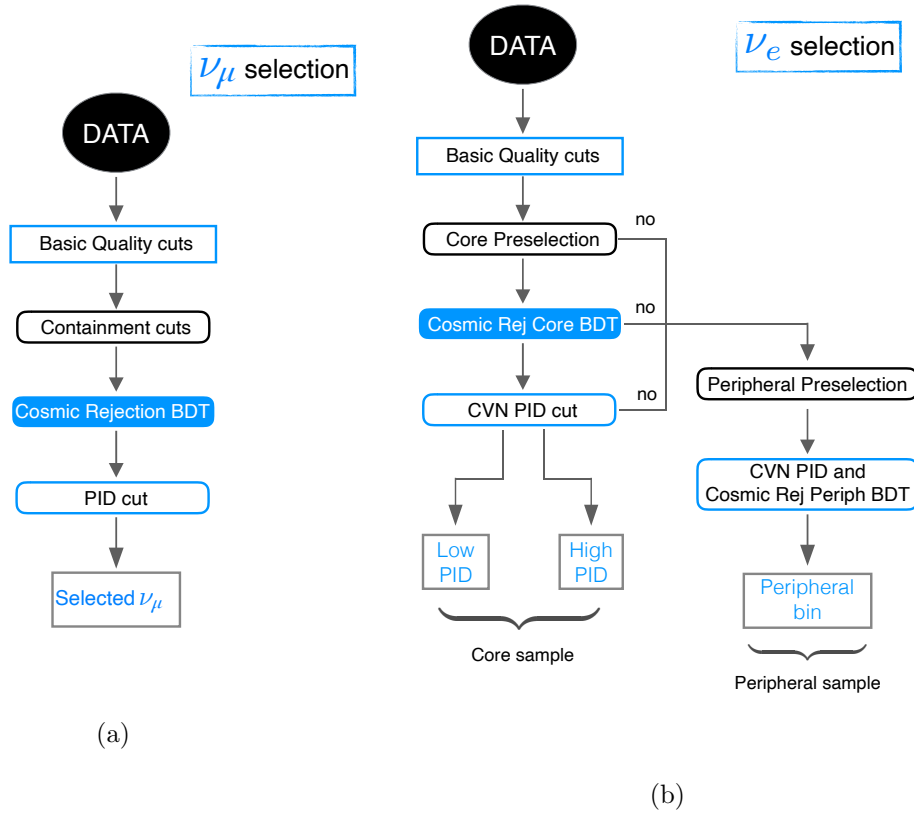


Figure 4.11: A diagrammatic representation of the two FD cut flows. **Left:** ν_μ -CC . **Right:** ν_e -CC . Taken from [103].

Cut Level	$\bar{\nu}_\mu$	ν_μ	Total Beam Bkg	NC	Cosmics	Effic	Purity
NoCut	193	111	45.8	225	1.63×10^6		
Quality	186	106	45.0	195	2.99×10^5	0.964	6.21×10^{-4}
Containment	86.7	35.5	32.0	132	4.38×10^3	0.449	0.019
Cosmic Rej.	84.9	32.8	28.9	107	635	0.440	0.099
Full PID	79.5	25.7	0.336	0.790	1.65	0.412	0.966

Table 4.2: The simulated number of predicted $\bar{\nu}_\mu$ -CC candidates at the FD at each stage of the 2020 analysis cut flow at the 2020 3-flavour best fit (see Section 2.5.5), broken down by true interaction type in RHC. Efficiency is calculated with respect to ‘NoCut’, where no slice level cuts have been applied. Taken from [104].

Cut Level	ν_e	$\bar{\nu}_e$	Total Beam Bkg	NC	Cosmics	Effic	Purity
NoCut	96.8	1.84	1.58×10^3	820	2.95×10^6		
Data Qual.	93.2	1.78	1.5×10^3	680	5.12×10^5	0.964	1.81×10^{-4}
Containment	71.4	1.39	560	510	2.58×10^4	0.738	2.65×10^{-3}
Cosmic Rej.	63.5	1.02	91.5	169	353	0.656	0.094
Full PID	52.0	0.859	11.4	4.73	1.51	0.537	0.738
Basic Cuts	38.7	0.689	549	508	4.53×10^5	0.400	8.53×10^{-5}
Full PID	8.63	0.194	3.84	0.253	1.63	0.0892	0.593

Table 4.3: The simulated number of predicted ν_e -CC candidates at the FD at each stage of the 2020 analysis cut flow at the 2020 3-flavour best fit (see Section 2.5.5), broken down by true interaction type in FHC. Efficiency is calculated with respect to ‘NoCut’, where no slice level cuts have been applied. The core sample is shown first and the peripheral beneath. ‘Basic Cuts’ refers to cuts to ensure quality of reconstruction described in Section 4.3.2 while ‘Full PID’ refers to the complete selection including application of CVN. Taken from [104].

Cut Level	$\bar{\nu}_e$	ν_e	Total Beam Bkg	NC	Cosmics	Effic	Purity
NoCut	27.0	5.14	723	328	1.92×10^6		
Data Qual.	26.4	4.92	693	266	3.25×10^5	0.979	8.12×10^{-5}
Containment	20.9	3.72	252	199	1.61×10^4	0.774	1.26×10^{-3}
Cosmic Rej.	19.3	2.57	27.4	57.8	296	0.715	0.048
Full PID	17.3	1.96	5.43	1.6	0.589	0.639	0.643
Basic Cuts	8.55	2.25	259	200	2.88×10^5	0.317	2.97×10^{-5}
Full PID	3.01	0.419	2.11	0.117	0.998	0.112	0.453

Table 4.4: The simulated number of predicted $\bar{\nu}_e$ -CC candidates at the FD at each stage of the 2020 analysis cut flow at the 2020 3-flavour best fit (see Section 2.5.5), broken down by true interaction type in RHC. Efficiency is calculated with respect to ‘NoCut’, where no slice level cuts have been applied. The core sample is shown first and the peripheral beneath. ‘Basic Cuts’ refers to cuts to ensure quality of reconstruction described in Section 4.3.2 while ‘Full PID’ refers to the complete selection including application of CVN. Taken from [104].

4.4 Analysis Binning at the FD

An optimal binning choice at the FD is key to ensuring that the full potential of all available neutrino beam data is realised. Appropriate binning can be used to maximise resolution in certain areas of phase space, as is the focus of the disappearance analysis. It can also be used to maximally separate signal and background, as is the focus of the appearance analysis. Both are discussed next.

4.4.1 ν_μ -CC Binning

As discussed in Section 2.4.1, in the disappearance analysis the parameters $\sin^2 \theta_{23}$ and Δm_{32}^2 dictate the position of the first disappearance maximum in neutrino energy. At NOvA, the energy region around the maximum is known as ‘the dip’ and occurs between 1 and 2 GeV. Moving to higher energies, not only does NOvA have fewer events but oscillation probability is lower. For this reason, in the dip region, reconstructed neutrino energy is binned in intervals of 0.1 GeV. This allows for a more precise evaluation of both where in energy the minimum of the energy spectrum lies and how deep it is. This of course translates into improved sensitivity to the oscillation parameters as compared to coarser binning. Moving to higher energies, the bins get wider until between 4 and 5 GeV there is just a single bin. There are 19 bins of reconstructed neutrino energy in total. An example FD simulated ν_μ -CC spectrum is shown in Figure 4.12a.

As discussed in Section 4.1.2, the muon neutrino energy resolution is driven by the amount of hadronic energy in the event. To obtain an even better resolution of the dip location, the muon neutrino FD sample is divided into four bins according to the fraction of the total reconstructed energy which is hadronic. All simulated events at the FD which pass the full ν_μ -CC selection cuts (see Section 4.3) are binned in neutrino energy and ordered by their hadronic energy fraction. Each neutrino energy bin is then split into four such that each sub-bin contains an equal number of events. Figure 4.12b shows the hadronic energy fraction bin boundaries as a function of neutrino energy. The result is four hadronic energy quartiles where quartile 1 has the best energy resolution ($\sim 7\%$ in FHC) and quartile 4 the worst ($\sim 12\%$ in FHC). Figure 4.13 shows the distribution of the fractional difference between reconstructed and true neutrino energy for all four quartiles in FHC.

A further benefit of binning by hadronic energy fraction is that true muon neutrino




Figure 4.12: **Left:** ν_{μ} -CC predicted spectrum at the FD for the 2020 oscillation analysis best fit point (purple) and under the assumption of no oscillations (red). By comparing the two histograms, the maximum disappearance region, ‘the dip’, is seen to be between 1 and 2 GeV. This range is binned most finely. Bins can be seen to get wider toward higher energy. Taken from [106]. **Right:** The distribution of hadronic energy fraction against reconstructed neutrino energy for muon neutrino candidates passing all selection cuts in FHC. The blue and purple histograms show the boundaries of the four hadronic energy quartiles in FHC and RHC respectively. Taken from [107].

events with less hadronic activity are more easily identified as ν_{μ} -CC. For this reason, background tends to pile up in the 4th quartile, improving sensitivity through better overall signal and background separation. Figure 4.14 shows a prediction of the ν_{μ} -CC spectrum at the FD when split by quartile, using the 2020 3-flavour analysis best fit point. Both the specialised energy binning and the separation into sub-samples of hadronic energy fraction were developed for the 2018 3-flavour analysis [105].

4.4.2 ν_e -CC Binning

The sensitivity to oscillation parameters in the ν_e -CC analysis is driven by the ability to separate signal and background. In addition to binning in equally sized bins of 0.5 GeV in reconstructed neutrino energy, the core event sample is further split into categories of CVNe score. As discussed in Section 4.1.3, CVNe score describes how electron neutrino-like an event is deemed to be by CVN. In the 3-flavour analysis presented at Neutrino 2020 there were two CVNe categories, low PID and high PID. Loosely speaking, a less pure PID bin and a more pure PID bin respectively. Figure 4.15 shows the true composition in each of these analysis bins for FHC and RHC. The high PID bin background is dominated




Figure 4.13: Muon neutrino energy estimator resolution for the four hadronic energy quartiles in FHC. A Gaussian fit to each histogram is shown. The legend includes the mean and Root Mean Square (RMS) of each of the histograms and the mean and standard deviation of each of the best fit curves. In both beam modes, neutrino interactions involving larger hadronic showers (quartile 4) have a lower energy resolution than those where the vast majority of the energy is assigned to the muon (quartile 1). Taken from [91].




Figure 4.14: The ν_μ -CC FD predicted spectra at the 2020 oscillation analysis best fit, split by hadronic energy fraction quartile in FHC. Quartile 1 contains events with the lowest fraction of hadronic energy and, therefore, has the best energy resolution. Quartile 4 contains events with the highest fraction of hadronic energy and, therefore, has the worst energy resolution. Taken from [106].




Figure 4.15: ν_e -CC FD predicted spectra at the 2020 3-flavour analysis best fit point, broken down by interaction type. **Left:** FHC. **Right:** RHC. In both horn currents, the dominant background in the high PID bin is from beam electron neutrinos. Cosmic interactions make up a large part of the background in the peripheral sample. Taken from [108].

by the (irreducible) beam electron neutrinos.

In addition to the core sample, there is the peripheral sample. This sample contains events deemed to be highly electron neutrino-like by CVN but do not pass the FD containment cuts. For this reason, it is less likely that all of the neutrino energy was deposited within the detector. The energy information is, therefore, not used in the fit. Due to the location of the events in this sample, close to the detector edges, the dominant (non-irreducible) background is from cosmic interactions.

One of the analysis improvements presented in this thesis centres around a detailed examination of the ν_e -CC analysis binning. This is discussed in Chapter 6.

4.5 Predictions at the FD

Predicting the FD reconstructed energy spectrum is done in two ways. Primarily, NOvA uses measurements of the unoscillated neutrino energy spectrum at the ND to correct the nominal FD prediction. This process is known as ‘extrapolation’ and is described in the next sections [109]. Although a data driven correction to the FD central value prediction is beneficial, the main advantage of such a technique is the reduction in the impact of systematic uncertainties correlated between the two detectors. It is particularly effective with regard to errors from interaction cross sections and flux, historically some of the largest. Alternatively, it is possible to rely purely on the MC simulation of the FD. This

is the so-called ‘no extrapolation’ method.

4.5.1 Decomposition & Extrapolation

Neutrino events are one of seven possibilities: ν_μ -CC, $\bar{\nu}_\mu$ -CC, ν_e -CC, $\bar{\nu}_e$ -CC, ν_τ -CC, $\bar{\nu}_\tau$ -CC and NC. Each selected sample is a mixture. Furthermore, oscillations can drastically alter the mixtures between ND and FD. To use information from the ND to help predict what would be seen at the FD it is necessary, therefore, to estimate the composition of selected samples at the ND. This process is called ‘decomposition’. The decomposition method varies according to how a particular ND sample is used. Additionally, the exact details of extrapolation vary according to which FD sample is being predicted. Each case will now be discussed in more detail.

$\nu_\mu \rightarrow \nu_\mu$ Signal

The muon neutrino candidate sample at the ND (see Section 4.3) is used to predict the muon neutrino survival signal at the FD. Due to the high purity of this ND sample, 95.2% and 97.2% in FHC and RHC respectively, all of the observed discrepancy between data and MC is assigned to the ν_μ ($\bar{\nu}_\mu$) component [102]. Somewhat by construction (NOvA’s MEC model is tuned to this same data) there is good agreement in this sample. An overall event excess of 1% and deficit of 3% is observed in FHC and RHC respectively.

As discussed, information obtained by making this measurement is carried to the FD via extrapolation. A schematic describing the process by which this is done is shown in Figure 4.16. As explained in [105], background estimated using simulation is first subtracted from the ND data spectrum. The background subtracted reconstructed neutrino spectrum is unfolded to a true neutrino energy spectrum via a smearing matrix obtained from ND simulation. Next, this true ND spectrum is multiplied by a FD over ND event ratio, designed to capture differences in event rate due to, for example, detector acceptance. A true energy FD spectrum has now been obtained and energy dependent oscillation weights can be applied to reflect a set of oscillation parameters. This oscillated spectrum is converted to a FD reconstructed energy spectrum via a smearing matrix, again obtained in simulation. The backgrounds are then added in.

This process is carried out independently in each bin of hadronic energy fraction (see Section 4.4.1) and additionally in three bins of reconstructed muon transverse momentum




Figure 4.16: The process by which ND data is used to correct the predicted FD signal in the muon neutrino disappearance analysis. Taken from [111].

$|\vec{p}_T^\mu|$, calculated as $|\vec{p}_T^\mu| = |\vec{p}_\mu| \sqrt{1 - \cos^2 \theta_\mu}$. θ_μ is the angle of the muon with respect to the beam [110]. Extrapolating in this way increases the robustness of the analysis, reducing the impact of cross section uncertainties on, in particular, the measurement of Δm_{32}^2 . The reason for this is that differences in kinematics of the ND and FD selected ν_μ samples arise from differences in selection efficiency and acceptance. A known origin of these differences is the varying ability of the two detectors to contain muons within the side walls. Figure 4.17 shows how the distributions of $|\vec{p}_T^\mu|$ for selected events at the ND and FD differ. The extrapolation bin containing events with the lowest $|\vec{p}_T^\mu|$, therefore, will have ND and FD samples whose kinematics more closely match, compared with the bin with the highest $|\vec{p}_T^\mu|$.

Unlike the bins of hadronic energy fraction, following extrapolation, the $|\vec{p}_T^\mu|$ bins are summed. That is to say, they are not explicitly used in the fit for oscillations.

$\nu_\mu \rightarrow \nu_\mu$ Background

Backgrounds to the ν_μ appearance signal at the FD come in two forms. First, event candidates due to neutrinos from the beam, are predicted to be small. For this reason, the nominal FD simulation is used to estimate them. Second, there are event candidates arising from cosmic interactions. Quantifying this background is discussed in Section 4.5.2.

$\nu_\mu \rightarrow \nu_e$ Signal

As with the muon neutrino survival signal, the muon neutrino candidate sample at the ND (see Section 4.3) is used to predict the electron neutrino appearance signal at the FD. The same extrapolation procedure as described in Section 4.5.1 is followed, with the exception




Figure 4.17: **Left:** A comparison of the reconstructed muon transverse momentum for selected ν_μ -CC and $\bar{\nu}_\mu$ -CC candidates in simulation at the ND (red) and FD (blue). ND data is shown in black. **Right:** A similar plot but comparing the reconstructed muon transverse momentum for selected ν_μ -CC and $\bar{\nu}_\mu$ -CC candidates in simulation at the ND (red) and the reconstructed electron transverse momentum for selected ν_e -CC and $\bar{\nu}_e$ -CC candidates at the FD (blue). Both taken from [112].

that the FD over ND event ratio and oscillation weights are appropriately replaced. Furthermore, the FD true to reconstructed neutrino energy smearing matrix maps onto the appearance (see Section 4.4.2) rather than the disappearance energy binning.

Unlike the survival signal, the ND sample is not extrapolated as a function of hadronic energy fraction. Events are, however, subdivided into three bins of $|\vec{p}_T^\mu|$ for extrapolation (see Section 4.5.1) and summed before fitting.

As discussed in Section 4.4.2, the appearance analysis makes use of a ‘peripheral’ sample. In the fit for oscillations, this is a single bin with no explicit dependence on reconstructed neutrino energy, other than events must have a reconstructed neutrino energy less than 4.5 GeV. Before these events are summed into a single bin, extrapolation weights which are calculated for the high PID bin are applied to nominal FD MC predictions for the peripheral sample. This is appropriate given that peripheral events also have been deemed highly electron-like by CVN. In the electron neutrino appearance binning developed in this thesis (see Chapter 6), the peripheral sample is split into 3 bins of reconstructed primary shower calorimetric energy for a given event. ‘Primary shower’

refers to the most energetic shower in the event. Given the high CVN score, however, this shower is likely to be the primary lepton. Extrapolation weights from the high PID bin are applied to these 3 bins by assuming simply that lepton energy is, on average, the majority of the total reconstructed energy. In short, extrapolation weights from the 1.5 to 2.0 GeV, 2.5 to 3 GeV and 3.5 to 4 GeV high-high PID bins are applied to the three peripheral bins sequentially respectively.

$\nu_\mu \rightarrow \nu_e$ Background

The electron neutrino candidate sample at the ND (see Section 4.3) is used to predict the background components of the electron neutrino appearance spectra at the FD. There are three major components to this background: electron neutrinos which are intrinsic to the beam or Beam Electron Neutrino (BEN) events, NC events and ν_μ -CC events. Two types of data-driven decomposition, BEN decomposition and Michel decomposition, are used to estimate these components at the ND in FHC, discussed next⁹. In RHC, due to a much larger fraction of wrong-sign events in the sample compared to FHC, the simpler method of *proportional* decomposition is used. Here, any overall discrepancy between data and simulation is split equally among the three background components. Namely, it is assumed that the simulation correctly predicts the relative proportions of event types but not the absolute normalisation in each analysis bin.

Once decomposed, the ND data is extrapolated to the FD. In the extrapolation method used here, correction factors from the ND are calculated purely in the reconstructed energy space, rather than first converting to true energy. Extrapolating in this way reduces dependence on the energy estimator. This is important for these background events, where the electron neutrino energy estimator is expected to perform poorly¹⁰. Oscillation probability weights at the FD are still applied in true energy [113].

BEN decomposition: This technique uses muon neutrino events at the ND to estimate the intrinsic electron neutrino component of the beam, a major (and irreducible) background at the FD for the appearance analysis [114]. Figure 4.18b shows the true neutrino energy for electron neutrino events at the ND in FHC, split by their ancestry. At low energies (less than ~ 3 GeV), the majority of electron neutrinos originate from pions

⁹Other possible oscillation channels, $\nu_e \rightarrow \nu_\mu$ and $\nu \rightarrow \nu_\tau$ make a much smaller contribution and are predicted using the nominal FD MC.

¹⁰Note that it is never possible to reconstruct the true neutrino energy in a NC event.

via muon decay

$$\pi^+ \rightarrow \nu_\mu + \mu^+ \quad , \quad \mu^+ \rightarrow \bar{\nu}_\mu + e^+ + \nu_e. \quad (4.3)$$

Figure 4.18a shows the equivalent plot for muon neutrinos. In the low energy region, the vast majority of muon neutrinos also originate from pions. It can also be seen that at high energies, most electron and muon neutrinos originate from neutral or charged kaons

$$K_L^0 \rightarrow \pi^- + e^+ + \nu_e \quad (4.4)$$

$$K^+ \rightarrow \pi^0 + e^+ + \nu_e. \quad (4.5)$$

By looking at samples of low and high energy muon neutrinos in data, the pion and kaon fluxes respectively in simulation can be corrected. Due to the shared ancestry of the muon and electron neutrinos, these corrected fluxes can then be propagated to the selected electron neutrino component at the ND through ‘pions weights’ and a ‘kaon scale’. The low and high energy muon neutrino samples are also known as the ND contained and un-contained samples respectively. The cut for the contained sample is the ND muon neutrino selection cut described in Section 4.3. The un-contained cut is largely the same except for the requirement that the candidate muon track not be contained. The breakdown of simulation compared to data for these samples is shown in Figure 4.19.

The calculation of the kaon scale S^K , defined simply as the ratio of the number of selected ν_μ -CC events from kaons in data versus simulation, is done iteratively via the equations [114]

$$S_0^K = 1, \quad (4.6)$$

$$S_i^\pi = \frac{\sum_C \left(N_{\nu_\mu}^{\text{data}} - S_{i-1}^K \times N_{\nu_{\mu CC} \leftarrow K}^{MC} - N_{\nu_{\mu CC} \leftarrow \text{other}}^{MC} - N_{\nu_\mu \text{BG}}^{MC} \right)}{\sum_C \left(N_{\nu_{\mu CC} \leftarrow \pi}^{MC} \right)}, \quad (4.7)$$

$$S_i^K = \frac{\sum_U \left(N_{\nu_\mu}^{\text{data}} - S_{i-1}^\pi \times N_{\nu_{\mu CC} \leftarrow \pi}^{MC} - N_{\nu_{\mu CC} \leftarrow \text{other}}^{MC} - N_{\nu_\mu \text{BG}}^{MC} \right)}{\sum_U \left(N_{\nu_{\mu CC} \leftarrow K}^{MC} \right)}. \quad (4.8)$$

C is the set of bins in the contained sample in the region 1 - 3 GeV where muon neutrinos with pion ancestors dominate. U is the set of bins in the un-contained sample in the region 4.5 - 10 GeV where muon neutrinos with kaon ancestors dominate. The pion scale from Equation 4.7, required to make contained events match in data and MC, is used in Equation 4.8 in obtaining the kaon scale. Then, that improved kaon scale is fed back into Equation 4.7. The procedure stops when subsequent calculations of S^K differ by less than 10^{-4} . For the analysis presented in this thesis, an overall increase of 5.8% in the rate of

electron neutrinos with kaon parents relative to simulation is obtained [102].

With regard to correcting the electron neutrino component with pion ancestry, events in the contained sample between 0 and 5 GeV are used. The calculation and application of weights has the ability to be more sophisticated given the reliability of the muon neutrino energy estimator for this subset of events and the direct kinematic relationship between the produced muon neutrinos and their parent pion. Figure 4.20 shows the distribution of ancestor transverse and forward momentum for selected muon and electron neutrinos in each of their respective samples. There is a large amount of overlap in the phase space of pions which go on to produce beam electron neutrinos compared to muon neutrinos.

Firstly, therefore, a set of weights w_{ν_μ} as a function of reconstructed energy E_{ν_μ} is constructed as follows [95]

$$w_{\nu_\mu}(E_{\nu_\mu}) = \frac{N_{\nu_\mu}^{\text{data}}(E_{\nu_\mu}) - N_{\nu_\mu CC \leftarrow \text{other}}^{MC}(E_{\nu_\mu}) - N_{\nu_\mu BG}^{MC}(E_{\nu_\mu})}{N_{\nu_\mu CC \leftarrow \pi}^{MC}(E_{\nu_\mu}) + N_{\nu_\mu CC \leftarrow K}^{MC}(E_{\nu_\mu})}. \quad (4.9)$$

This is the ratio of the (estimated) number of ν_μ -CC events from pions and kaons in data and simulation¹¹. Secondly, the computed ratios $w_{\nu_\mu}(E_{\nu_\mu})$ are propagated to the parent pion level as a function of the pion momentum, (p_T^π, p_z^π) . The result is a weight $w(p_T^\pi, p_z^\pi)$ for each pion (p_T^π, p_z^π) , where the entire energy spectrum for the neutrinos coming from pions which contribute to the same (p_T^π, p_z^π) bin is summed over [115]

$$w(p_T^\pi, p_z^\pi) = \frac{\sum_k w_{\nu_\mu}(E_{\nu_\mu}(k)) \times E_{\nu_\mu}(k)}{\sum_k E_{\nu_\mu}(k)}. \quad (4.10)$$

These weights are then applied to the ν_e -CC events with pion ancestors selected in the electron neutrino ND sample. The corrected ν_e from pion spectrum is then given by summing over (p_T^π, p_z^π) in each analysis bin. The final estimate of intrinsic beam electron neutrinos is the sum of the corrected pion and kaon components and the uncorrected simulation of other ancestors [95].

Michel decomposition: This decomposition technique is employed to correct NC and ν_μ -CC events, using the number of Michel electrons in each selected event [117]. A true ν_μ event will contain a muon which, 92% of the time in pure carbon, will decay to produce a Michel electron via $\mu^- \rightarrow e^- + \nu_\mu$ [118]. Michel electrons can also be produced

¹¹Note that the inclusion of $N_{\nu_\mu CC \leftarrow K}^{MC}$ in the denominator means any difference in data and simulation is split equally among muon neutrinos with pion and kaon ancestors according to their proportions in simulation. Figure 4.18a shows, however, that the kaon contribution in this sample is (negligibly) small.

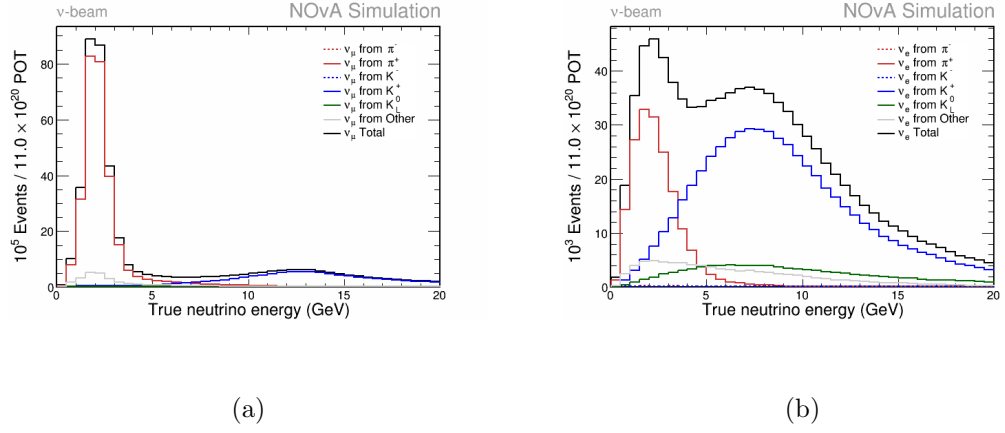


Figure 4.18: **Left:** True neutrino energy of muon neutrino events with a true vertex inside the ND, broken down by the ancestor type in FHC. **Right:** True neutrino energy of electron neutrino events with a true vertex inside the ND, broken down by the ancestor type in FHC. Contained muon neutrino events in the 0 - 5 GeV range are used to calculate the pion weights, while 1 - 3 GeV contained and 4.5 - 10 GeV un-contained muon neutrino events are used for the kaon scale. Both taken from [116].

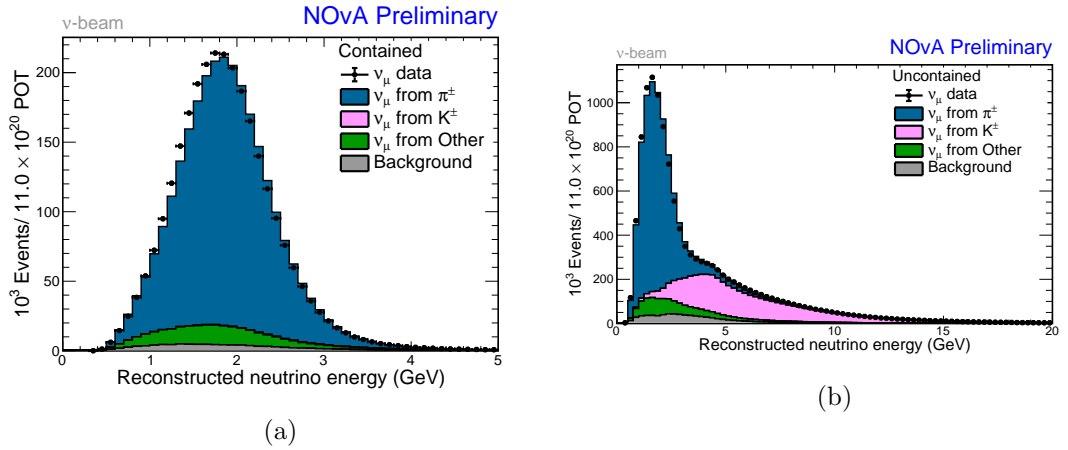


Figure 4.19: **Left:** A breakdown of simulation compared to data for the BEN contained sample. **Right:** A breakdown of simulation compared to data for the BEN un-contained sample. Although not shown on either figure, data and simulation match within systematic uncertainty. Both taken from [116].

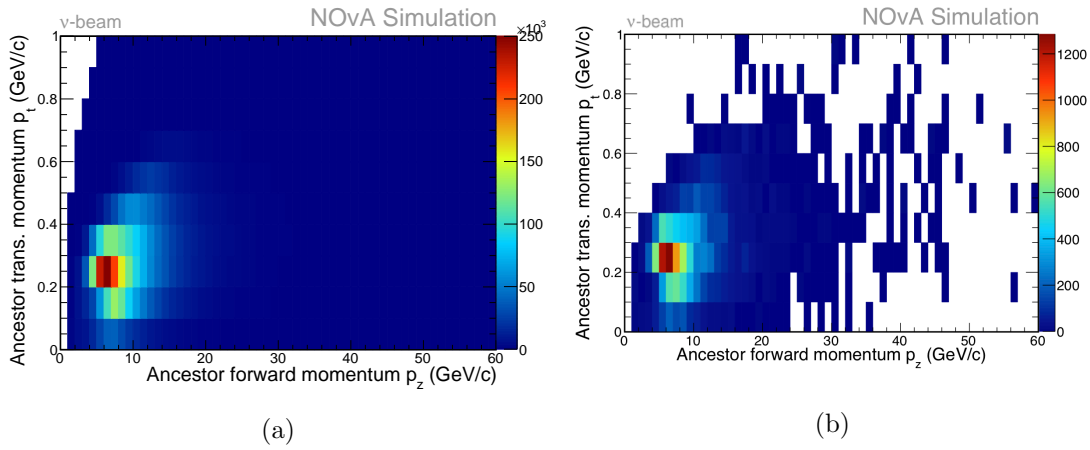


Figure 4.20: **Left:** The distribution of transverse and forward momentum for pions which go onto to produce selected muon neutrinos in the contained BEN sample. **Right:** The distribution of transverse and forward momentum for pions which go onto to produce electron neutrinos which pass the ND electron neutrino selection. A large amount of overlap can be seen between the phase space of the two. Note that this plot was produced for the 2020 3-flavour analysis. The analysis presented in this thesis employs a looser ND electron neutrino selection cut (see Chapter 6). The selection, however, has not changed in the region of phase space where beam electron neutrinos constitute a large fraction of the background. Both taken from [116].

in the hadronic system, usually via $\pi^\pm \rightarrow \mu^\pm \rightarrow e^\pm + \nu + \bar{\nu}$. True ν_μ -CC events will contain on average one extra Michel electron compared to true NC (and beam ν_e) events. It can be seen from Figure 4.21 that the fraction of true muon neutrinos with one or more Michel electrons is greater than for NC events. The number of Michel electrons then can, in principle, be used to estimate the ν_μ -CC and NC contributions in data.

A Michel finding algorithm is run which looks for activity close in space but delayed in time with respect to the neutrino candidate slice [119]. To differentiate between true Michels and background, a log-likelihood based PID known as the Michel electron Identifier (MID) was developed. For a candidate event, MID is the difference in log-likelihoods calculated under the Michel and not a Michel hypotheses based on its calorimetric energy, number of cells containing hits, and the spatial and temporal separation from the parent slice. Distributions analogous to Figure 4.21 are constructed for each analysis bin for events passing a MID cut and the ν_μ -CC and NC events are scaled using a log-likelihood fit to the data in this space [117]. Since there is no separation of beam ν_e and NC events using this method and the ν_e events have already been corrected by BEN decomposition, Michel decomposition holds them fixed.

In the 2020 3-flavour analysis, fairly large corrections of 38.2% and 13.9% were made to the ν_μ -CC and NC populations respectively [102]. These events, however, arise from specific kinematic regions within the possible phase space of these events and are poorly constrained by external data. For the new analysis introduced later in this thesis, there is a discussion of the results of decomposition and final extrapolated FD predictions in Chapter 7.

4.5.2 Estimation of Cosmic Induced Events

In addition to misidentified events from the neutrino beam, a very important source of background to the oscillation analyses at the FD comes from cosmic induced activity. To estimate this background, an approach using data from the cosmic trigger (recorded orthogonally to any beam data) is used. Full ν_μ and ν_e analysis selected cuts (see Section 4.3) are applied to the data and the FD analysis bins are populated. The resulting spectra have some cosmic exposure or ‘livetime’ associated with them. The integrals of these histograms are scaled down to match the livetime corresponding to NuMI beam data. The error on each bin is taken to be the $1\text{-}\sigma$ Poissonian uncertainty before scaling,




Figure 4.21: The number of Michel electrons in each selected event in the 2020 analysis electron neutrino ND sample in simulation. Events are separated by true interaction type and the histograms are not stacked. Data is shown in black. The fraction of true muon neutrinos with one or more Michel electrons is greater than for NC events. Taken from [116].

this is propagated to the fit and treated as a systematic uncertainty [120,121].

As a final addition to the FD spectrum, beam events with a true interaction vertex outside the detector volume passing the electron neutrino FD selection are added. These ‘rock’ events are introduced at the very end since they are simulated in a separate, special sample. Their very small contribution to the analysis presented in this thesis is discussed in Chapter 6.

4.6 Systematic Uncertainties

NOvA accounts for around 60 individual systematic uncertainties in the 3-flavour oscillation analysis. In the software, each systematic is treated as its own ‘universe’ where the distributions (for example, energy spectra) are re-calculated at every step (for example, extrapolation). There are three ways in which the effect of a particular systematic can be captured. Firstly, some systematics such as the flux model and most of the neutrino interaction uncertainties, use functions which re-weight the nominal distributions directly. Others, such as the muon energy scale, necessitate tweaking one or two underlying parameters before generating new distributions. Finally, any systematic which exhibits complex behaviour, such as calibration or detector response, requires new MC files to be made with the earlier stages of the simulation and reconstruction re-performed [102].

These methods are used to obtain ND and FD predictions in a given systematic universe at the ± 1 - and $2\text{-}\sigma$ values of the systematic (recall from Section 4.2 that CAFAna is entirely histogram based). The fit, however, needs systematics that are continuous. Cubic interpolation is used in each analysis bin to obtain the event counts ‘in between’¹². The next sections discuss the different systematic groups in detail and in tabulated form define some shorthand notation used in the later chapters. Where appropriate, shorthand names are also defined in the text using **this font**.

¹²Cubic Hermite splines are commonly used for this kind of interpolation of numerical data. ‘Hermite’ refers to the form with which the polynomials are defined, namely, by their values and first derivatives at the two data end points. Cubic splines are the lowest-order splines which can be both forced to be doubly-differentiable at the end points and have the extrapolation outside of the first and last points be a straight line.

Systematic parameters: flux	
ppfx_hadp_beam_pc00	Flux component 00
ppfx_hadp_beam_pc01	Flux component 01
ppfx_hadp_beam_pc02	Flux component 02
ppfx_hadp_beam_pc03	Flux component 03
ppfx_hadp_beam_pc04	Flux component 04

Table 4.5: Shorthand and corresponding full names of all of the systematic parameters related to flux. Both short and long forms are used later in the thesis.

4.6.1 Flux

Differences in data and simulation, arising from both the mismodelling of hadron production in the graphite target and their subsequent transport, are covered by flux uncertainties. As discussed in Section 3.3.1, a package known as PPFX is used to incorporate external data to constrain the part related to hadron production [62]. The package derives an uncertainty using a multi-verse technique. 100 random scenarios are generated, each representing a world with different proton-target cross-sections. For each of these universes a set of weights, directly analogous to those applied to the central value prediction described in Section 3.3.1, are derived. A covariance matrix is constructed for each universe using interactions corrected by these weights, describing the relationship between bins of true energy. These covariance matrices are summed to form an averaged covariance matrix which can be collapsed via a Principle Component Analysis (PCA) [122]. The PCA allows the hadron production uncertainty to be captured in a much smaller set of uncorrelated weights or ‘principle components’. The five principle components leading to the largest variance in the FD over ND event ratio are taken forward into the fit for oscillations.

Uncertainty in subsequent transport, which arises from unknowns affecting beam focussing such as the target position, the magnetic field in the decay pipe, the amount of current in the horns etcetera, are also included in the PCA [122]. For each of the 100 hadron flux universes described above, 20 sets of beam transport parameters are also thrown. Table 4.5 lists the shorthand and corresponding full names of all of the systematic parameters related to flux.

4.6.2 Cross-section

In total, 62 systematic parameters are used to quantify uncertainties in the model due to neutrino cross-sections and FSI. Most are evaluated using weights provided by GENIE [24], which can be applied to the simulated events based on the neutrino interaction's truth information [95]. Some of these uncertainties are expected, however, to have very little effect on the results but significantly increase run time at the fitting stage. To mitigate the effects on run time, the cross-section systematic parameters are broken down into two groups. The systematics are ranked according to the χ^2 their $\pm 1\text{-}\sigma$ shifts make to the nominal FD prediction in the analysis distributions (see Section 4.4) [102]. All systematics whose effect falls below $\chi^2 = 0.005$ in all distributions are classed as small. In the analysis presented in this thesis, the effect of these systematics is summed in quadrature and treated as one effect in the fit (`SumSmallXSec3Flavor2020`)¹³. There are 38 systematics in the group of *small* cross-section uncertainties, the remaining 24 *large* cross-section uncertainties are treated individually. For perspective, one of the largest cross-section systematics is `radcorrnu` which results in a 1.6% increase or decrease in total number of predicted FD events. Conversely, one of the smallest cross-section uncertainties treated as 'large' is `kZExpEV3Syst2020` (described below) which has approximately an 0.7% affect on the FD prediction.

Some of the key remaining large cross-section uncertainties which differ from the default GENIE adjustable parameters are summarised as follows:

- **CCQE:** NOvA utilises the 'z-expansion' parametrisation of the CCQE form factor to asses the uncertainty on its value (see Section 3.3.2) [65]. GENIE provides five adjustable parameters for the shape and normalisation of the form factor. NOvA retains the standard normalisation parameter, which can be loosely seen as changing the 'traditional' axial mass, adjusting the total number of events by -20% and 15% (`ZNormCCQE`). The four remaining shape parameters are used to create four custom weight functions which correctly enforce correlations between themselves (`kZExpEVnSyst2020`, `n = 1, 2, 3, 4`) [123].
- **MEC:** NOvA employs three uncertainties associated to MEC interactions. These are intended to separately capture dependence on neutrino energy, energy transfer and the composition of final state nucleon pairs based on several experimental and

¹³The same group of cross-section systematics found to be small in the 2020 3-flavour analysis were taken as small in the analysis presented in this thesis.

theoretical models [123]. Since these uncertainties are uncorrelated for neutrino and antineutrinos, there are in fact six MEC systematics included in the fit for oscillations (`MECEnuShapeSyst2020Nu/AntiNu`, `MECShapeSyst2020Nu/AntiNu` and `MECInitStateNPFFracSyst2020Nu/AntiNu`).

- **RES:** Historically, GENIE has over-predicted the number of ν_μ -CC RES events at low values of 4-momentum transfer. However, as described in Section 3.3.2, the version of GENIE used for the 2020 3-flavour analysis uses a new resonance model which results in a much smaller excess in this region. Despite this, there are still small amounts of disagreement between data and simulation, and it is not understood if the new model is in fact the reason for the observed improvements. NOvA, therefore, retains a systematic uncertainty related to this effect (`LowQ2RESSupp2020`).
- **DIS:** Motivated by disagreements in data and simulation in the ND and recent reanalysis of bubble chamber data [124], an uncertainty for non-resonant interactions producing pions is introduced (`DISvnCC1pi_2020` is the systematic uncertainty covering interactions with one pion in the final state). Such interactions with a final state hadronic mass of less than 3 GeV, carry a 50% uncertainty. Above 3 GeV, the uncertainty drops from 50% linearly until hitting 5% at 5 GeV.
- **Other uncertainties:** A 2% uncertainty on the ratio of the ν_e to ν_μ and $\bar{\nu}_e$ to $\bar{\nu}_\mu$ cross-sections is applied to account for radiative corrections (`radcorrnu/radcorrnu`). Similarly, a 2% uncertainty which is fully anticorrelated between ν_e and $\bar{\nu}_e$ is introduced to allow for second class currents (`2ndclasscurr`). Both of these uncertainties were motivated by work done by the T2K collaboration [125].

Table 4.6 lists the shorthand and corresponding full names of all of the systematic parameters related to neutrino cross-section.

4.6.3 Detector Response and Calibration

The evaluation of uncertainties arising from the calibration procedure and the detector response model necessitates the production of several additional MC samples. These datasets, discussed next, require some amount of re-simulation of the nominal ND and FD samples but with the relevant systematic parameters altered [102].

- The detector light model describes the number of photons produced in the scintillator which are collected by the fibre. When a charged particle passes through the NOvA

Systematic parameters: neutrino cross-section	
2ndclasscurr	Second class currents
DISvnCC1pi_2020	DIS ν nCC1pi (DIS events producing 1 pion)
hNFSLFateFracEV1_2020	hN FSI fate fraction eigenvector 1
hNFSLMFP_2020	hN FSI mean free path
LowQ2RESSupp2020	RES low- Q^2 suppression
MECEnuShape2020AntiNu	MEC E_ν shape, antineutrinos
MECEnuShape2020Nu	MEC E_ν shape, neutrinos
MECInitStateNPfrac2020AntiNu	MEC initial state np fraction, antineutrinos
MECInitStateNPfrac2020Nu	MEC initial state np fraction, neutrinos
MECShape2020AntiNu	MEC 2020 ($q_0, \vec{q} $) response, antineutrinos
MECShape2020Nu	MEC 2020 ($q_0, \vec{q} $) response, neutrinos
MaCCRES	MaCCRES (Nuclear form factor parameter.)
MaNCRES	MaNCRES (Nuclear form factor parameter.)
MvCCRES	MvCCRES (Nuclear form factor parameter.)
MvNCRES	MvNCRES (Nuclear form factor parameter.)
NuTauScale	ν_τ Scale
radcornue	Radiative corrections for ν_e
radcornuebar	Radiative corrections for $\bar{\nu}_e$
RPAShapeenh2020	RPA shape: higher- Q^2 enhancement (2020)
RPAShapesupp2020	RPA shape: low- Q^2 suppression (2020)
SumSmallXSec3Flavor2020	Summed small cross-section systematic
ZExpAxialFFSyst2020_EV1	CCQE z-exp eigenvector shift 1
ZExpAxialFFSyst2020_EV2	CCQE z-exp eigenvector shift 2
ZExpAxialFFSyst2020_EV3	CCQE z-exp eigenvector shift 3
ZExpAxialFFSyst2020_EV4	CCQE z-exp eigenvector shift 4
ZNormCCQE	CCQE z-exp normalisation

Table 4.6: Shorthand and corresponding full names of all of the systematic parameters related to neutrino cross-section. Both short and long forms are used later in the thesis.

scintillator, it produces both scintillation and Cherenkov photons. Motivated by a discrepancy between data and simulation in proton dE/dx of around 5% (compared to 1.5% for muons) [126], three new samples were produced with changes to this model. In one, the factor determining the level of Cherenkov photons produced was increased to bring the proton disagreement to less than 1% while leaving the muon response unchanged. Since this factor is a property of the materials used by NOvA, which is common to both detectors, this systematic applies to both the ND and FD (**Cherenkov**). The other two samples investigate the overall number of photons collected by the fibre, there is one for the ND and one for the FD¹⁴. There are shifts of $\pm 10\%$ in both X and Y views at the ND. In the FD the shifts are $\pm 16\%$ in the X view and an anticorrelated $\mp 6\%$ shift in the Y view. In both of these samples the absolute calibration constants are counter adjusted to preserve energy response, which is calibrated using muons. These samples are useful for evaluating the effect of the hits that miss or pass the threshold level (**Light_Level_ND/Light_Level_FD**) [23].

- In both the ND and FD, there is a downward trend of approximately 0.24% per year in the number of reconstructed hits in candidate events [127]. This is evidence that the scintillator and/or the detector electronics are ageing. A special ‘drift’ sample, where the light model is decreased in line with the observed downward trend as a function of time, was produced to account for this (**CalibDrift**). Absolute calibration constants are appropriately shifted to preserve muon response.
- As with the light model uncertainties, the overall energy response uncertainty is determined by discrepancies in data versus simulation for candidate protons in the ND. This disagreement at the 5% level is greater than for any other particle or standard candle metrics such as Michel electron distributions and π^0 mass peaks. To evaluate this systematic (**Calibration**), nominal simulated samples are re-used but the absolute calibration constants are turned up and down by 5% before reconstruction. This procedure ensures that the energy of the reconstructed hits changes, but the number of reconstructed hits remains the same compared to nominal. Additionally, to account for differences in reconstructed and true energies for events close to the detector edges, there is a ‘calibration shape’ uncertainty (**CalibShape**) [128].

Table 4.7 lists the shorthand and corresponding full names of all of the systematic parameters related to detector response and calibration.

¹⁴The overall light yield is view and detector dependent due to geometry.

Systematic parameters: detector response and calibration	
CalibDrift	Calibration drift
CalibShape	Calibration shape
Calibration	Absolute calibration
Cherenkov	Cherenkov light
Light_Level_FD	FD light level
Light_Level_ND	ND light level
RelativeCalib	Relative calibration

Table 4.7: Shorthand and corresponding full names of all of the systematic parameters related to detector response and calibration. Both short and long forms are used later in the thesis.

4.6.4 Other Uncertainties

- ND-FD Acceptance:** As discussed in Chapter 4.5, the technique of extrapolation is used to carry information from measurements made in the ND to predicted FD spectra. This helps to reduce the impact of particular systematic uncertainties on the analysis, for example, from neutrino cross-sections. Although the two detectors are ‘functionally identical’, they do differ, particularly in their sizes. This leads to ND and FD ‘equivalent’ samples having different underlying kinematic distributions due to their differing abilities to contain all final state particles. The correction factors that are applied at the FD, therefore, come from a somewhat incorrect region of phase space. Techniques such as extrapolating in bins of muon transverse momentum (see Section 4.5.1) help to counter these effects. For the prediction of electron neutrino appearance signal, however, these kinematic difference can be particularly acute¹⁵. To quantify these differences, an ‘acceptance’ systematic is derived ($\text{accept_signal_ptTextrap_FHC_2020/accept_signal_ptTextrap_RHC_2020}$) [129]. The size of the systematic is estimated by re-weighting the muon neutrino samples in both the ND data and MC to match the FD electron neutrino samples in the following kinematic spaces: true four momentum transfer squared, reconstructed angle of the primary lepton with respect to the beam, total reconstructed transverse

¹⁵Recall that the electron neutrino candidate sample at the FD is extrapolated from the muon neutrino candidate sample at the ND. Differences in the kinematics of the accepted samples arise, therefore, due to the differing containment of muons at the ND and electrons at the FD and differing PID performance.

momentum of the event. The idea being that the kinematics of the FD electron neutrino selected sample are replicated in the ND muon neutrino candidate sample. The re-weighted ND spectra are then passed through the extrapolation procedure and the resulting FD distributions are compared to the nominal. The largest difference between shifted and nominal distributions is taken as the systematic's $1\text{-}\sigma$. These are maximally 0.4% in both FHC and RHC.

- Muon energy scale:** There are five separate systematic parameters designed to characterise the uncertainty associated with using reconstructed muon length in the detectors to determine muon energy [130]. Three of these systematics (`UnCorrND` / `FD` / `MuCatMuEScaleSyst2020`) deal with the associated sources of uncertainty which are uncorrelated between detectors, with the muon catcher and ND treated separately. The size of these systematics are 0.15%, 0.13% and 0.48% in the FD, ND and muon catcher respectively. All three are dominated by uncertainties in ‘mass accounting’. This includes but is not limited to, for example, uncertainty in the volume/density of detector components such as the scintillator, steel or PVC. A fourth systematic parameter quantifies the uncertainties which are correlated between detectors (`CorrMuEScaleSyst2020`), amounting to 0.7%. This uncertainty is dominated by uncertainties in the Fermi density effect and differences in muon range as calculated by GEANT and calculated by others [131]. The final systematic parameter accounts for so called ‘neutron pileup’ in the ND and muon catcher (`PileupMuESyst2020`). Over the course of a $10\text{ }\mu\text{s}$ spill, many neutrino induced neutrons can build up in the ND, given that the neutron capture time is around $50\text{ }\mu\text{s}$. Upon recapture, however, they produce a hit in the detector which can be tagged onto the beginning or end of a true muon track. This leads to muon neutrino events which occur later in the spill having a higher reconstructed muon energy on average.
- Neutron scale:** Discrepancies in data versus simulation in neutron rich samples (using a sub-sample of $\bar{\nu}_\mu$ -CC candidates) have been observed. In particular, simulation shows an excess of neutron candidate prongs at low energy (below 20 MeV) and a deficit at medium energy (above 20 MeV and below 150 MeV) [132]. This motivated an uncertainty which scales the contribution of energy from simulated neutrons with a true visible energy below 20 MeV to the total reconstructed hadronic energy (`NeutronEvisPrimariesSyst2018`). This has the effect of shifting neutrons from the region of excess to the region of deficit.

- **Normalisation:** Systematic uncertainties falling under the umbrella of normalisation can be broken down into two parts. First, there is the ‘correlated normalisation’ (`NormHornCorr`) which has a value of 0.82%. This is constructed from a sum in quadrature of uncertainties surrounding detector masses and the so called ‘GENIE-GEANT mismatch problem’ observed in NOvA code¹⁶ [133]. Second, uncertainties in POT are summed in quadrature with uncertainties resulting from differences in event selection efficiency in data and simulation (separately for FHC and RHC). The latter uncertainties are the subject of Chapter 5. The resulting scales of these uncertainties (`NormFHC2020`/`NormRHC2020`) are 0.72% and 0.90% in FHC and RHC respectively.
- **ND rock scale:** For the analysis presented in this thesis, a new systematic to capture the uncertainty in the number events originating in the rock surrounding the ND was introduced (`NDNueRockEventSyst`). The necessity for and development of this systematic is discussed in more detail in Chapter 6. In short, a 100% uncertainty is applied to any ND event with a true vertex outside the detector. This systematic is ‘ ν_e specific’, so only the electron neutrino appearance signal and background predicted spectra are affected.

Table 4.8 lists the shorthand and corresponding full names of all of the systematic parameters which do not fit into the categories covered by Tables 4.5 - 4.7.

4.7 Oscillation Model Fit to Data

The measurement of oscillation parameters is made by comparing the total FD prediction under a 3-flavour model to data, using a frequentist statistical approach. NOvA assumes that each analysis bin i is described by Poisson statistics. Specifically, in the i ’th analysis bin, the probability P_i of observing O_i events given an expectation of $E_i(\vec{\theta})$ events for a set of oscillation parameters $\vec{\theta}$ is given by [134]

$$P(O_i|E_i(\vec{\theta})) = \frac{E_i^{O_i} e^{-E_i}}{O_i!}. \quad (4.11)$$

¹⁶In the late stages of the 3-flavour analysis, a bug in either GEANT code or the interface of that code with NOvA code was discovered. The bug causes GEANT to drop particles from the list of primary particles it receives from GENIE around 0.6% of the time. It has since been discovered that this bug is fixed in a later version of GEANT4.

Systematic parameters: detector response and calibration	
accept_signalkin_pTextrap_FHC_2020	Acceptance ND to FD kinematics ν_e signal FHC
accept_signalkin_pTextrap_RHC_2020	Acceptance ND to FD kinematics $\bar{\nu}_e$ signal RHC
CorrMuEScaleSyst2020	Muon energy scale correlated
cosmicScale	Cosmic background scale
LeptonAngleSystFDXZ2020	Lepton angle systematic FD XZ
LeptonAngleSystFDYZ2020	Lepton angle systematic FD YZ
LeptonAngleSystNDXZ2020	Lepton angle systematic ND XZ
LeptonAngleSystNDYZ2020	Lepton angle systematic ND YZ
michel_tagging2020	Michel electrons tagging uncertainty
NDNueRockEventSyst	ND rock scale ν_e
NeutronEvisPrimariesSyst2018	Neutron visible energy systematic
NormFHC2020	FHC normalisation
NormHornCorr	Correlated normalisation
NormRHC2020	RHC normalisation
PileupMuESyst2020	ND neutron pile-up
RockScale	FD rock scale
UnCorrFDMuEScaleSyst2020	FD muon energy scale uncorrelated
UnCorrMuCatMuESyst2020	Muon catcher muon energy scale uncorrelated
UnCorrNDMuEScaleSyst2020	ND muon energy scale uncorrelated

Table 4.8: Shorthand and corresponding full names of all of the systematic parameters which do not fit into the categories covered by Tables 4.5 - 4.7. Both short and long forms are used later in the thesis.

The joint probability across all bins is given by the product of the bin-wise probabilities. This function is known as the Poisson likelihood

$$L(\vec{\theta}) = \prod_{i \in \text{bins}} P(O_i | E_i(\vec{\theta})). \quad (4.12)$$

The set of values for the oscillation parameters which maximise the likelihood for a given observation are commonly referred to as the ‘best fit’ values. Maximisation corresponds to finding a stationary point with respect to $\vec{\theta}$, hence a constant with respect to $\vec{\theta}$ can be subtracted freely from the likelihood. The likelihood ratio can, therefore, be defined

$$\lambda(\vec{\theta}) = \prod_{i \in \text{bins}} \frac{E_i^{O_i} e^{-E_i}}{O_i!} - \prod_{i \in \text{bins}} \frac{O_i^{O_i} e^{-O_i}}{O_i!}. \quad (4.13)$$

The likelihood ratio is more useful when rearranged in the following way

$$-2 \ln \lambda(\vec{\theta}) = -2 \ln \left[\prod_{i \in \text{bins}} \frac{E_i^{O_i} e^{-E_i}}{O_i!} \right] + 2 \ln \left[\prod_{i \in \text{bins}} \frac{O_i^{O_i} e^{-O_i}}{O_i!} \right] \quad (4.14)$$

$$= -2 \sum_{i \in \text{bins}} \left[\ln \left(\frac{E_i^{O_i} e^{-E_i}}{O_i!} \right) \right] + 2 \sum_{i \in \text{bins}} \left[\ln \left(\frac{O_i^{O_i} e^{-O_i}}{O_i!} \right) \right] \quad (4.15)$$

$$= -2 \sum_{i \in \text{bins}} [O_i \ln E_i - E_i - \ln O_i!] + 2 \sum_{i \in \text{bins}} [O_i \ln O_i - O_i - \ln O_i!] \quad (4.16)$$

$$= 2 \sum_{i \in \text{bins}} \left[E_i - O_i + O_i \ln \left(\frac{O_i}{E_i} \right) \right]. \quad (4.17)$$

Explicitly

$$-2 \ln \lambda(\vec{\theta}) = 2 \sum_{i \in \text{bins}} \left[E_i(\vec{\theta}) - O_i + O_i \ln \frac{O_i}{E_i(\vec{\theta})} \right]. \quad (4.18)$$

Minimising this function with respect to $\vec{\theta}$ for a dataset O_i is directly analogous to maximising the likelihood given in Equation 4.12 [134]. Formally, the values of the parameters $\vec{\theta}$ which minimise Equation 4.18 such that

$$(-2 \ln \lambda)_{\min} \equiv -2 \ln \lambda(\vec{\theta}_A) < -2 \ln \lambda(\vec{\theta}_B) \quad \forall \quad \vec{\theta}_B \neq \vec{\theta}_A \quad (4.19)$$

are those which most accurately describe the data. As aforementioned, they are known as the ‘best fit’ values [23].

Equation 4.18 can be considered a χ^2 statistic with k degrees of freedom, since $-2 \ln \lambda(\vec{\theta})$ tends to a Gaussian distribution as $k \rightarrow \infty$ ¹⁷. From this point, therefore, $\chi^2(\vec{\theta})$ will be used to denote Equation 4.18. However, the analysis presented in this thesis is not within

¹⁷The general definition of a χ^2 statistic is *any* function of O_i and $E_i(\vec{\theta})$ which asymptotically obeys the ‘classical’ χ^2 distribution [134].

the boundaries of this asymptotic limit. As a result, all confidence intervals shown in this document, which assume Gaussian statistics, are an approximation. Such confidence intervals are obtained by means of a further test statistic $\Delta\chi^2$, defined as

$$\Delta\chi^2(\vec{\theta}) = \chi^2(\vec{\theta}) - \chi_{\min}^2. \quad (4.20)$$

For a $\Delta\chi^2$ with one degree of freedom, for example, the $1\text{-}\sigma$ (“68%”) confidence interval is defined to be the region of phase space where $\Delta\chi^2 < 1$. For the purpose of constructing contours in two dimensions, there are two degrees of freedom. The 1- , 2- and $3\text{-}\sigma$ regions are defined to be where $\Delta\chi^2 < 2.30$, $\Delta\chi^2 < 6.18$ and $\Delta\chi^2 < 11.83$ respectively [26]. Equation 4.20 is also used to quantify the sensitivity of the analysis to a given set of oscillation parameters. Here, so called Asimov ‘fake data’ is constructed at the central value of a simulated FD spectrum, predicted at a given set of oscillation parameters¹⁸

4.7.1 Treatment of Oscillation Parameters

The underlying physics model employed in this thesis uses the full oscillation formulae in matter introduced in Section 2.4. The experiment baseline, $L = 810$ km and density of the Earth, $\rho = 2.84$ g/cm³ are held fixed [135]. The solar oscillation parameters are also held fixed at the values given in the 2019 edition of the Particle Data Group (PDG) [26]: $\Delta m_{21}^2 = 7.53 \times 10^{-5}$ eV² and $\sin^2 2\theta_{12} = 0.851$. As discussed, NOvA has some sensitivity to θ_{13} . This angle, however, is very well measured by other experiments. It is treated, therefore, as a constrained nuisance parameter via the addition of a Gaussian penalty term to the χ^2 statistic. The central value and $1\text{-}\sigma$ uncertainty for $\sin^2 2\theta_{13}$, also taken from the 2019 edition of the PDG, are 0.085 and 0.003 respectively.

The parameters which remain, δ_{CP} , Δm_{32}^2 and $\sin^2 \theta_{23}$ are all completely free in the fit. In order to probe both mass hierarchies, two fits, one seeded with a positive value of Δm_{32}^2 , the other seeded with a negative value of Δm_{32}^2 , are performed. The “best fit” in the hierarchy which is disfavoured corresponds to a local minimum only.

4.7.2 Treatment of Systematic Uncertainties

NOvA incorporates numerous systematic uncertainties into its fit to the data (see Section 4.6). They are treated as uncorrelated and are included in the likelihood fit in the

¹⁸Fake data at NOvA is constructed by taking the central value of a high-statistics simulated FD spectrum, predicted at some set of oscillation parameters and scaled to a particular POT exposure.

form of Gaussian nuisance parameters. The χ^2 test statistic (Equation 4.18) is modified as follows [136]

$$\chi^2(\vec{\theta}) = \min \left(\chi^2(\vec{\theta}, \vec{s}) + \sum_{i \in \text{Systs}} \frac{s_i^2}{\sigma_i^2} \right). \quad (4.21)$$

At each point in oscillation space, therefore, the χ^2 taken forward is the one which is minimised with respect to all systematic parameters i , including the rightmost ‘penalty term’. s_i and σ_i represent the values and 1- σ ranges of the systematic parameters (in units of σ) respectively. As its name suggests, the penalty term acts to increase the total χ^2 according to how far a particular systematic is pulled away from its central value of zero. The larger the pull, the larger the penalty.

Chapter 5

Context Effects on Near Detector Selection Efficiency

The intensity of NuMI beam pulses has a sizeable influence on the efficiency of selecting true ν_μ -CC and ν_e -CC events in the ND. At higher intensities, the density of neutrino interactions is larger. This leads to separate physics events occurring more closely together in both time and space (pile-up), complicating reconstruction. To establish to what extent, if any, selection efficiency in data and simulation are different to one another, a single neutrino overlay procedure is used. This technique, first implemented by a fellow collaborator for the 2018 round of analyses [137], involves inserting individual MC neutrino interactions on top of nominal data and MC beam spills in the ND. The difference in efficiency to select the single interactions in data and MC is evaluated separately for four samples of interest (ν_e -CC, ν_μ -CC) \otimes (FHC, RHC). The results are used as data driven values for part of the systematic uncertainty on the ND normalisation. In the 2018 analysis, only the values for the ν_μ -CC samples were determined. Separate results for ν_e -CC were introduced by myself.

This chapter expands on the procedure described above and was used for the 2020 oscillation analysis. Section 5.1 details the how the relevant data and simulation samples are prepared. Section 5.2 talks about the extraction of the selection efficiency differences and Section 5.3 presents the values obtained. Section 5.4 then explains how those values are used in the fit for oscillations. Section 5.5 shows a study of neutrino energy resolution, which uses the overlaid data sets.

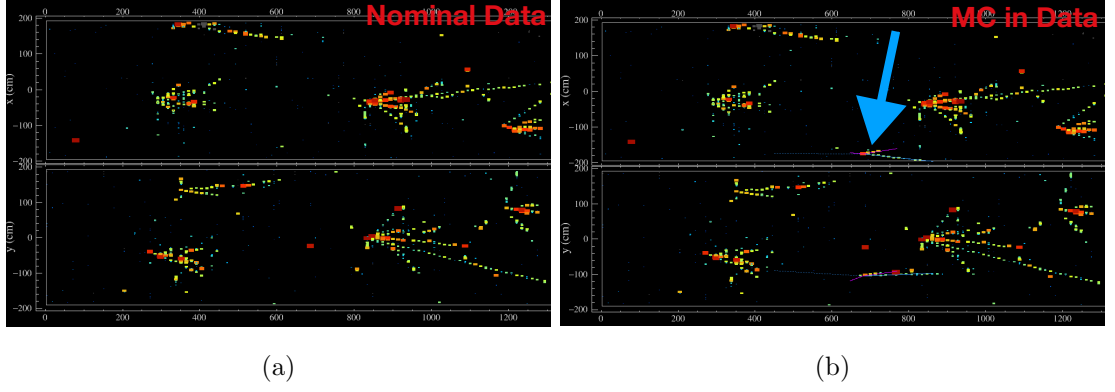


Figure 5.1: Fully reconstructed ND spill in data, before (**left**) and after (**right**) the single simulated neutrino is overlaid. The extra ν_μ is highlighted by the blue arrow. The top and bottom windows in each Figure are birds-eye and side-on views of the detector respectively.

5.1 Single Neutrino Overlay Methodology

Single MC neutrinos are overlaid into a subset of ND spills in data. This gives a ‘true’ number of neutrinos to select. The same overlay and selection procedure in simulated ND spills yields an efficiency directly comparable to the efficiency in data, and the difference is taken as a data-MC ND normalisation difference. Figures 5.1 and 5.2 show example event displays of nominal ND spills, before and after the single neutrino has been overlaid. The evolution of software from the 2018 to the 2020 analysis¹ and the need for a new way to match the beam intensity in simulation to the data, required an in-depth evaluation of how singles events are generated and overlaid. All changes to the method for 2020 are geared toward achieving a more robust evaluation of NOvA’s susceptibility to pile-up and are summarised below.

- In the 2018 analysis, the single neutrino interactions were simulated in the conditions of a particular data run/subrun combination and overlaid onto that specific combination. However, in doing so the ‘random subrun’ procedure for matching the beam intensity in simulation to data was distorted. This led to the need for a complex re-weighting of events, to align the distributions of POT per spill in data and MC a posteriori.

In the 2020 analysis, singles are generated in such a way that any single neutrino can be overlaid onto any nominal spill in data or MC, which are sampled randomly. It is

¹The performance of a newly introduced clustering algorithm, discussed in Section 4.1.1, was of particular interest.




Figure 5.2: Fully reconstructed ND spill in nominal MC, before (**left**) and after (**right**) single simulated neutrino is overlaid. The extra ν_μ is highlighted by the blue arrow.

important to ensure that singles are overlaid into enough spills to both measure the data-MC difference to the required uncertainty, but also to match beam intensity distributions in data and MC.

- In both 2018 and 2020, the overlay procedure has been to add exactly one MC neutrino to each nominal spill, regardless of the intensity of the spill into which it is being overlaid. In 2020, to account for the fact that at higher intensities we expect more neutrino interactions, that is to say, have the rate of overlaid neutrinos per spill vs. intensity match the natural rate of neutrinos per spill vs. intensity, a spill POT dependent weight is applied to the overlaid neutrinos. This is equal to the ratio of the actual POT per spill and a central value of POT per spill (40×10^{12} POT). Events which are selected in higher intensity spills are, therefore, weighted up and those in lower intensity spills are weighted down. For example, a single neutrino being overlaid into spill of 40×10^{12} POT (a high intensity spill) will be assigned a weight of $6/4 = 1.5$. In contrast, a single neutrino being overlaid into spill of 20×10^{12} POT (a low intensity spill) will be assigned a weight of $2/4 = 0.5$.
- Independent evaluation of selection efficiency differences between data and simulation for electron neutrinos and antineutrinos have been introduced. In 2018, the selection efficiency differences for muon neutrinos were applied to electron neutrino candidate distributions. This likely lead to an underestimation of the systematic due to the more complex topology of electron neutrino interactions compared to muon neutrino interactions in NOvA.

Four separate samples ($(\nu_e\text{-CC}, \nu_\mu\text{-CC}) \otimes (\text{FHC}, \text{RHC})$) of single neutrino events are simulated in the ND with the same flux files and GENIE version as the nominal MC. There

is exactly one neutrino in each simulated spill and spills are not run/subrun matched to real data taking conditions. This gives the flexibility to overlay *any* single neutrino into *any* nominal data or MC spill.

The difference in selection efficiency between data and simulation is expected to be small - of order 1%. It is necessary, therefore, to make a sub-percent measurement of the selection efficiency in each case to be able to resolve this difference. This means that after all cuts have been applied and assuming a selection efficiency of 20%, there must be a sample of roughly 200k events. This means that the subset which defines the denominator must contain around 1 million events. However, to get to the sample which constitutes the denominator, a series of pre-selection cuts are made. These are discussed in detail in Section 5.2.1 but in short, also reduce the sample by a factor of four to five. Leaving some room for contingency means that of order 10 million single neutrino events must be generated in each horn current and each neutrino flavour.

The hits from simulated singles are then mixed in to nominal data and MC ND files which have been randomly sampled. These hybrid spills are then reconstructed using the regular reconstruction chain and converted to CAF format. The extraction of the $1\text{-}\sigma$ context systematic can then be extracted using the standard CAFAna analysis framework.

5.2 Extracting Selection Efficiency Difference

5.2.1 Pre-selection Cuts

To extract a reliable efficiency, it is important to have a well defined denominator that both pinpoints the signal which matters for normalisation and is truly comparable between data and MC. This sample is realised by applying a set of pre-selection cuts, designed for the 2018 analysis, to the overlaid neutrinos in the mixed files.

- MC does not simulate various scenarios found in the real data spills - for example, the NuMI trigger still provides data “spills” during beam downtime where the intensity is zero POT. To get around this, the standard spill cuts are applied (see Section 4.3.1).
- To isolate signal, true overlaid ν_μ -CC or $\bar{\nu}_\mu$ -CC events (in the case of ν_μ systematic) and ν_e -CC or $\bar{\nu}_e$ -CC events (in the case of ν_e systematic) are required.
- The true vertex of the neutrino interaction must be inside a fiducial box of 25-

1150 cm in the beam direction and ± 180 cm in the transverse directions². The box, therefore, sits symmetrically in the centre of the detector plane transverse to the beam direction and slightly away from the most upstream plane and muon catcher. This removes events very near the detector edges and those starting in the muon catcher.

A visual summary of the pre-selection cuts is given in Figure 5.3. Note the “NT” stands for “neutrino truth” and refers to the fact that plots are made using truth, not reconstructed, variables. Figures 5.3a to 5.3c showing the x , y and z positions of the true neutrino interaction vertex have some interesting physical features. For a neutrino interaction to be included in one of these plots, it is only necessary for energy from that interaction to be deposited in the detector. The vertex of the interaction can be outside of the detector. The bumps in Figure 5.3a come from interactions with a vertex in the electronics boxes on the top of the detector. The slopes in that same Figure and Figure 5.3b come from the fact that the detector is not centred on the beam. Flux, therefore, is a function of position across the detector. Finally the clear regions of “low” and “high” in Figures 5.3b and Figures 5.3c are caused by the muon catcher. Recall that the muon catcher is approximately two thirds of the height of the rest of the detector and contains different materials.


- Events are also required to have a true energy in the 3-flavour oscillation analysis region of interest, 0 to 7 GeV (0 to 6 GeV for ν_e). These ranges are slightly wider than those of reconstructed neutrino energies used in the analysis to avoid biases at the upper limits.

The sample of events which pass these pre-selection cuts forms the efficiency denominator, corresponding to the row ‘overlaid’ in the cut flow tables shown in Section 5.3. A plot showing the true energy distribution of this sample is shown in Figure 5.4.


5.2.2 Slicing Efficiency

Imperfections in slicing should also be addressed by this analysis. It is not uncommon for the energy deposited by the overlaid neutrino to be split between multiple slices in a spill. A typical physical example causing this is a neutron being produced at a neutrino interaction vertex travelling a large distance in the detector before interacting. This results in a hit in the detector which is spatially and temporally displace from the parent


²The origins of the coordinate systems for each detector are in the centre of the front (most upstream) faces. Both are right-handed with z increasing moving downstream and y increasing moving vertically upwards.




(a) Requiring the true vertex X position to satisfy $|x| \leq 180$ cm removes events starting close to the side walls.



(b) Requiring the true vertex Y position to satisfy $|y| \leq 180$ cm removes events starting close to the top and bottom.



(c) Requiring the true vertex Z position to satisfy $25 \leq z \leq 1150$ cm removes events starting close to the front and in the muon catcher.



(d) To isolate ν_μ , a cut on true PDG code is applied. There is also a cut requiring the interaction to be charged current.

Figure 5.3: In all of the histograms shown, there is an entry per overlaid single neutrino, weighted by $(\text{spill POT}) / (40 \times 10^{12} \text{ POT})$. The black points show singles overlaid in data and the red points show singles overlaid in MC. Each plot represents a cut made based on the truth information of the overlaid neutrino. These cuts, combined with the requirement to be a charged current event and that the corresponding spill passes standard spill cuts, define the pre-selection and form the efficiency denominator.




Figure 5.4: True energy distribution of overlaid muon neutrinos in data and MC in FHC, after pre-selection cuts are applied. This sample forms the efficiency denominator.

interaction. This hit could then be dismissed as noise or, in the busy environment of the ND, be incorporated into a slice with another neutrino interaction. To ensure that a maximum of one event is counted in each spill, a cut on the (poorly named) CAF variable ‘slice efficiency’ is made. This variable captures the fraction of energy in a given slice associated with the neutrino that contributes the most energy to the slice. Putting a cut at 51% ensures only one slice per spill is considered overlaid. Figure 5.5 shows that placing a cut here does not remove important data-MC differences. Between 0.5 and 0.2 the ratio of data and MC is flat. Below 0.2, the energy contribution to the slices is so low that they are highly unlikely to be selected. The slope in the ratio just above 0.5 and the apparent normalisation offset above this will be addressed by this analysis. After imposing this cut, the sample is referred to as ‘sliced’ in the cut flow tables.

It was shown in that although removing this cut introduces interesting effects at the slice level, once a basic quality cut is applied, the number of slices is comparable to when the cut is placed at 51%. Most importantly at selection level, there is no difference in the value of the normalisation systematic. This implies that the slices containing small amounts of the overlaid neutrino’s energy, do not contain any other meaningful physics





Figure 5.5: The distribution of ‘slice efficiency’ for all overlaid neutrinos in data and MC. There is one entry per reconstructed slice. Placing a cut at 51% ensures that one overlaid event can only contribute one slice.

(other neutrino interactions). It is likely that the hits in these slices are from neutrons that are spatially and temporally separated from the parent interaction, making it difficult to associate the activity.


5.2.3 3-Flavour Cut Flow

The difference in efficiencies in data and simulation comes from the different stages of reconstruction and selection for oscillation analysis. After the ‘sliced’ sample has been made, the standard 3-flavour cuts (see Section 4.3) are applied sequentially. In the case of ν_μ -CC - Data Quality (DQ), DQ + containment, DQ + containment + PID (selected). Similarly for ν_e -CC - pre-selected, selected. An example of how the ν_μ -CC cut flow affects the true energy distribution shape can be seen in the area normalised plots in Figure 5.6. The samples of events which pass each of the full selection cuts form the final efficiency numerators. The entire cut flow from ‘overlaid’ to ‘selected’ can be summarised in tables like Table 5.1 which shows the results for ν_μ -CC FHC.


To estimate the uncertainty on the fractional selection difference, the distribution of




(a) Sliced. Eff. diff: 0.054



(b) Quality. Eff. diff: -0.048



(c) Quality and containment. Eff. diff: -0.27



(d) Full selection. Eff. diff: -0.18

Figure 5.6: Applying ν_μ -CC analysis cuts to the sliced samples of single neutrinos sequentially. Fully selected events form the final efficiency numerator. Plots are area normalised so that the shapes of the distributions can be easily compared. Comparisons of the number of events are given in Tables 5.1, 5.2 and 5.3.

error about each binomially distributed efficiency in data, ϵ_D and MC, ϵ_{MC} is approximated to be normal. The individual errors are then summed in quadrature to get the uncertainty on the fractional difference. A walk-through of the calculation is shown below.

- Recall that each overlaid neutrino carries a weight according to the intensity of the spill into which it is being overlaid, discussed in Section 5.1. Under the normal estimation, the variance of the efficiency ϵ with N overlaid events i , each with weight w_i is given as:

$$\sigma^2 = \frac{\epsilon(1-\epsilon)}{\sum_i^N w_i} \quad (5.1)$$

- Variance of the selection efficiency in data:

$$\sigma_{\text{Data}}^2 = \frac{\epsilon_D(1-\epsilon_D)}{\sum_i^{N_D} w_i} \quad (5.2)$$

- Variance of the selection efficiency in MC:

$$\sigma_{\text{MC}}^2 = \frac{\epsilon_{\text{MC}}(1-\epsilon_{\text{MC}})}{\sum_j^{N_{\text{MC}}} w_j} \quad (5.3)$$

- Standard deviation of the difference of selection efficiencies:

$$\sigma_{(\Delta\text{Data-MC})} = \sqrt{\sigma_{\text{Data}}^2 + \sigma_{\text{MC}}^2} \quad (5.4)$$

- Standard deviation of the difference of selection efficiencies as a percentage of the selection efficiency in MC.

$$\sigma_{(\Delta\text{Data-MC},\%)} = \frac{\sigma_{(\Delta\text{Data-MC})}}{\epsilon_{\text{MC}}} \times 100 \quad (5.5)$$

This number is taken as the uncertainty on the 1- σ systematic value for fractional selection efficiency difference.

5.3 Results

The goal of the single neutrino overlay study in the context of the 2020 oscillation analysis was four numbers. An estimation of the uncertainty arising from differences in selection efficiency between data and MC for ν_μ -CC and ν_e -CC events, in both horn currents. To get these numbers, the number of simulated single neutrinos selected when mixed into nominal data and nominal MC is compared. The fractional difference in efficiency between the two is taken as the 1- σ value of the systematic.

	MC in Data		MC in MC		$\frac{(\text{Data}_{\text{eff}} - \text{MC}_{\text{eff}})}{\text{MC}_{\text{eff}}}, \%$
	Events	Eff, %	Events	Eff, %	
Overlaid	2.53×10^6	100	2.64×10^6	100	0
Sliced	2.50×10^6	98.9	2.61×10^6	98.8	0.054
Quality	2.26×10^6	89.5	2.36×10^6	89.5	-0.048
Quality & Con- tainment	526000	20.8	551000	20.8	-0.27
Selected	386000	15.3	404000	15.3	-0.18 ± 0.21

Table 5.1: Cut flow table for overlaid single ν_μ -CC events in FHC. The final difference in efficiency and associated error between data and simulation is in the bottom right cell.

A summary of the four values is given in Table 5.3 and a detailed breakdown as successive cuts are applied is shown for two of the event types: ν_μ -CC in Table 5.1 and ν_e -CC in Table 5.2. In the muon neutrino FHC and RHC modes, the values of the systematic are either smaller (FHC) or similar in magnitude (RHC) to the 2018 analysis. Both are less than 0.5%. In the electron neutrino FHC and RHC samples (for which no value was calculated in 2018), the values of the systematic are close to 0.5% but slightly larger than for the muon neutrinos in the same beam configuration. Three of the results are consistent with there being no difference between data and MC.

The slicing and selection procedure also appears to be robust against the range of spill intensities experienced by the ND so far. This can be seen by looking at the selection efficiency as a function of POT per spill, shown in Figure 5.7a. There is less than 1% variation in efficiency between the lowest and highest spill POT values. Figure 5.7b shows efficiency fall sharply as the true neutrino energy increase. This is expected due to the ND's (in-)ability to contain higher energy events. Importantly, the behaviour is seen in both data and MC.

	MC in Data		MC in MC		$\frac{(\text{Data}_{\text{eff}} - \text{MC}_{\text{eff}})}{\text{MC}_{\text{eff}}}, \%$
	Events	Eff, %	Events	Eff, %	
Overlaid	1.52×10^6	100	1.53×10^6	100	0
Sliced	1.50×10^6	99.0	1.52×10^6	99.0	0.087
Pre-selected	154000	10.2	155000	10.1	0.41
Selected	111000	7.35	112000	7.32	0.36 ± 0.41

Table 5.2: Cut flow table for overlaid single ν_e -CC events in FHC. The final difference in efficiency and associated error between data and simulation is in the bottom right cell.

	Eff. in Data, %	Eff. in MC, %	$\frac{(\text{Data}_{\text{eff}} - \text{MC}_{\text{eff}})}{\text{MC}_{\text{eff}}}, \%$	Systematic
ν_μ FHC	15.26	15.29	-0.18 ± 0.21	0.21%
ν_μ RHC	18.14	18.23	-0.48 ± 0.16	0.48%
ν_e FHC	7.35	7.32	0.36 ± 0.41	0.41%
ν_e RHC	8.91	8.96	-0.53 ± 0.34	0.53%

Table 5.3: The final values of the single overlaid neutrino selection efficiency difference. The values used as a systematics for the 2020 analysis are shown in the final column of each row.

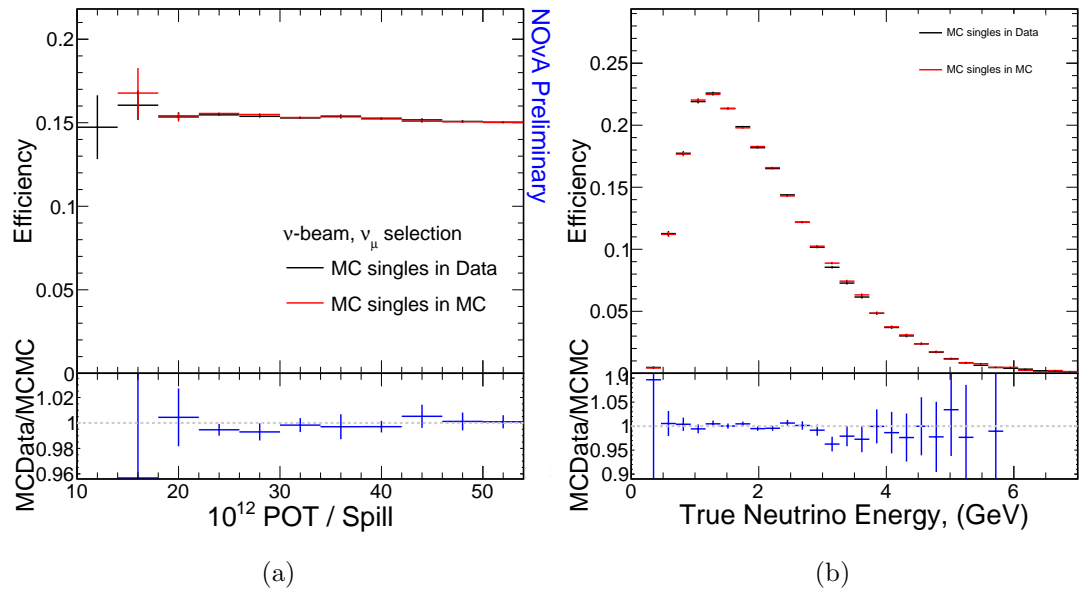


Figure 5.7: Selection efficiency as a function of POT per spill of nominal spill (**left**) and true neutrino energy of overlaid muon neutrinos (**right**).

5.4 Implementation as a Systematic

The four percentages derived from the overlay procedure to capture ND context effects are not implemented into the fit for oscillations as stand-alone systematics. Instead, they are combined with uncertainty on total POT, estimated to be 0.55% [102]. The POT uncertainty is correlated in both detectors and affects all events, but could be separate for each horn current. The context numbers are all independent and in principle could be applied to ND events according to their selected flavour and horn current. In practise due to time constraints, the context numbers are summed in quadrature with the POT uncertainty, according to horn current. The two resulting ‘total’ uncertainties (one for FHC and one for RHC) are applied to all ND events in the appropriate beam mode.

5.5 Energy Resolution

Another interesting study which can be done using the overlay samples described in this chapter examines the neutrino energy resolution. As previously mentioned, the overlaid neutrinos give some level of truth information to refer to in the ‘real’ environment of data. The spread in the distribution of fractional residuals defined in Equation 5.6 is a good metric for the resolution of a given energy estimator

$$R_{E_\nu} = \frac{E_{\nu\text{Reco}} - E_{\nu\text{True}}}{E_{\nu\text{True}}} \quad , \quad \nu \in [\nu_\mu, \nu_e] . \quad (5.6)$$

A wider spread of R_{E_ν} for overlaid neutrinos in data compared to MC may imply some form of oversight in the simulation, as may a difference in the mean. Figures 5.8 and 5.9 show the fractional residuals for overlaid FHC ν_μ -CC and ν_e -CC respectively. It is difficult to extract precise values of energy resolution from these plots. In particular, how much of the tails are included when applying either a Gaussian fit or calculating the RMS greatly affects the value obtained. However, the focus of this study is not to extract absolute values in data and MC, but to compare the two. It can be seen from the ratio plots that in the peak region, both ν_μ -CC and ν_e -CC slightly slope. Gaussian fits in these regions, however, show that the peaks differ by much less than 1% (see legends). The width of these fitted Gaussians confirms that the spread of the distributions in data and MC are very similar, again much less than 1%. Statistics are more limited outside of the fitted regions, but the ratios are largely flat. Calculations of RMS using fractional residuals values $\in [-3, 3]$ brings similar conclusions. The story is the same for RHC events.




Figure 5.8: The distribution of fractional residuals in the ν_μ -CC energy estimator for selected overlaid muon neutrinos in FHC data and MC. Plots are normalised to the number of overlaid single neutrinos in data. Gaussian fits are made around the peaks. Best fit values of the mean and variance are shown in the legend.




Figure 5.9: The distribution of fractional residuals in the ν_e -CC energy estimator for selected overlaid electron neutrinos in FHC data and MC. Plots are normalised to the number of overlaid single neutrinos in data. Gaussian fits are made around the peaks. Best fit values of the mean and variance are shown in the legend.

It is also informative to look at the fractional residuals as a function of true neutrino energy. This is shown in Figures 5.10 and 5.11 for overlaid FHC ν_μ -CC and ν_e -CC respectively. Each data point shows the mean fractional residual in a given bin of true energy and the ‘error bars’ show the raw RMS in that bin. RMS is used here as it is typically one of the quantities which is used by the groups in NOvA which develop the energy estimators and study their resolution. In the lower plots data *minus* MC, rather than the ratio of the two, is shown. The difference in means is more easily interpreted than the ratio given the proximity of both of the means to zero. The error bars show the standard error on the difference between the means. There is good agreement between data and MC in each of the ν_μ -CC and ν_e -CC single neutrino overlay samples, both in per bin central value and RMS. Finally, it is worth noting that both energy estimators become more biased at higher energies, this is due to the fact that they are tuned to the peak (and not the mean) in the 3-flavour oscillation analysis region of interest³.

In summary, the systematic studies presented in this chapter show that pile-up, noise and any other context effects are sufficiently well modelled in the simulation such that no significant difference in selection efficiency or energy reconstruction between ND data and MC are introduced.

³For ν_μ -CC , $0 \text{ GeV} \leq E_{\text{Reco}} < 5 \text{ GeV}$ and for ν_e -CC , $1 \text{ GeV} \leq E_{\text{Reco}} < 4 \text{ GeV}$.




Figure 5.10: The distribution of fractional residuals in the ν_μ -CC energy estimator for selected overlaid muon neutrinos in FHC data and MC as a function of true neutrino energy. The error bars show the raw RMS in each bin. The lower plot shows the distribution of data *minus* MC.




Figure 5.11: The distribution of fractional residuals in the ν_e -CC energy estimator for selected overlaid electron neutrinos in FHC data and MC as a function of true neutrino energy. The error bars show the raw RMS in each bin. The lower plot shows the distribution of data *minus* MC.

Chapter 6

Analysis Improvements

The primary goal of this thesis is to explore ways in which NOvA's existing and relatively mature 3-flavour analysis can be adapted to probe the parameters governing neutrino oscillations with greater sensitivity, particularly δ_{CP} and the neutrino mass hierarchy. This chapter reports on several attempts, both successful and unsuccessful, to improve the standard analysis methodology in the electron neutrino appearance channel. It is found that simultaneous implementation of three changes can achieve a gain in sensitivity equivalent to collecting 5 to 6% more data based on a simple metric. First, two other axes binned as a function of different particle identifiers are introduced. This is in addition to the existing dependence on energy and a single electron neutrino particle identification score. Second, by including regions of phase space with a greater amount of beam and cosmic induced background, more data is made available to the fit for oscillations. A bin edge optimisation procedure is also developed to simultaneously establish the position of bin edges in every variable which maximises the simple metric. Third, a coarsely binned simple energy variable targeting the most energetic shower in an event is introduced to the peripheral events sample. These three changes are combined and a 3-flavour fit to neutrino and antineutrino data is performed. The results of this improved analysis are presented in Chapter 7.

6.1 Studies at the FD

In order to compare different analysis methodologies, it is necessary to first define a metric which quantifies to what extent one is 'better' than another. In the case of this thesis, 'better' corresponds to a greater sensitivity to the neutrino oscillation parameters - specifically δ_{CP} and the neutrino mass hierarchy. To allow quick comparison of different

methodologies, the metric must also be simple. With this in mind, the following metric based on the Poisson log-likelihood ratio [26] was adopted

$$\Delta\kappa_{CP}^2 = 2 \sum_{i \in \text{bins}} \left[E_i^{3\pi/2} - O_i^{\pi/2} + O_i^{\pi/2} \ln \left(\frac{O_i^{\pi/2}}{E_i^{3\pi/2}} \right) \right]. \quad (6.1)$$

Here, $E_i^{3\pi/2}$ is the predicted number of FD events from simulation in analysis bin i in the normal (inverted) neutrino mass hierarchy at the oscillation point

$$\delta_{CP} = \frac{3\pi}{2}, \quad (6.2)$$

$$\sin^2 \theta_{23} = 0.5, \quad (6.3)$$

$$\Delta m_{32}^2 = 2.44 (-2.55) \times 10^{-3} \text{ eV}^2, \quad (6.4)$$

$$\sin^2 2\theta_{13} = 0.085. \quad (6.5)$$

Δm_{32}^2 and $\sin^2 2\theta_{13}$, therefore, are chosen to sit at their 2019 PDG central values [26]. For θ_{23} , maximal mixing is used as the neutral option with respect to the true octant choice. $O_i^{\pi/2}$ corresponds to fake data constructed in analysis bin i at the same oscillation point as $E_i^{3\pi/2}$ but with δ_{CP} set to be maximally opposite

$$\delta_{CP} = \frac{\pi}{2}. \quad (6.6)$$

Fake data at NOvA is constructed by taking the central value of a high-statistics simulated FD spectrum, predicted at some set of oscillation parameters and scaled to a particular POT exposure. Therefore, if nature has chosen $\delta_{CP} = \frac{\pi}{2}$, $\sqrt{\Delta\kappa_{CP}^2}$ describes the significance with which the point $\delta_{CP} = \frac{3\pi}{2}$ can be excluded. Visually, it can be seen as comparing two points on the electron neutrino bi-event plot, Figure 2.16b. Loosely, it quantifies how well a point somewhere in middle, degenerate region (close where current data sits) and a point on the edge, in the region of large asymmetry, can be discerned from one another. $\Delta\kappa_{CP}^2$ can be calculated for two different analysis methodologies, $\Delta\kappa_{CP_A}^2$ and $\Delta\kappa_{CP_B}^2$. Given $\Delta\kappa_{CP_A}^2 < \Delta\kappa_{CP_B}^2$, the fractional difference between the two corresponds directly to the fractional change in exposure that would be necessary to achieve the sensitivity of methodology B using methodology A .

Due to the high number of new analysis methodologies being evaluated, the time taken to construct each methodology computationally must be kept as low as possible. For this reason, in the initial stages of the search for sensitivity improvements when combinatorics are high, only FD simulation is used. Namely, FD predictions are not corrected in any way using ND data. $\Delta\kappa_{CP}^2$ is calculated for both neutrino mass hierarchies and separately

	FHC			RHC		
	Core	Peripheral	Total	Core	Peripheral	Total
Normal hierarchy	7.91	0.96	8.87	2.55	0.25	2.80
Inverted hierarchy	7.04	0.80	7.84	3.17	0.33	3.50

Table 6.1: $\Delta\kappa_{CP}^2$ values for the 2020 3-flavour analysis using FD only simulation (no extrapolation of ND data). The contributions of the core and peripheral samples to each of the totals is also given.

for FHC and RHC. This allows for easy identification of analysis changes which may affect one horn current more than the other.

Additionally, note that systematic uncertainties are not considered at this stage due to the large jump in resource requirements this would necessitate. Since the measurement of δ_{CP} is statistically limited, the first order consideration should be statistics. The effect of systematics on the metric are studied at a later stage (see Section 6.3.1), at which point the earlier stages of evaluation could be circled back to if necessary.

For the purpose of this thesis, methodology *A* (the ‘baseline’ analysis) will be the 2020 3-flavour analysis. The assumed exposure corresponds to the full 2020 neutrino and antineutrino FD dataset of 13.60×10^{20} and 12.50×10^{20} POT in FHC and RHC modes respectively. Table 6.1 shows the values of $\Delta\kappa_{CP}^2$ calculated for this baseline analysis, $\Delta\kappa_{CP_{2020}}^2$. All fractional improvements are calculated with respect to these numbers unless otherwise specified.

6.1.1 Searching for Gains

The driving factor governing sensitivity in the appearance analysis is the ability to separate signal and background. Introducing more variables to the FD spectra used in the fit for oscillations is one way that this could feasibly be achieved. Another way is by increasing the density of binning in a given variable. Additional information can also be passed to the fit by including events from less signal pure regions of phase space as separate samples. Practically, this is done by loosening the electron neutrino candidate selection cuts and using binning to separate the ‘new’ regions from existing ones. Each of these possibilities are explored separately and those which show the most promise are taken forward and

combined. The initial possibilities can be categorised into two focus groups: kinematic variables and particle identification classification scores. These groups are discussed next.

Kinematic variables:

As discussed in Section 4.4.2, at first order the 2020 analysis appearance FD spectra bin the electron neutrino candidates as a function of reconstructed neutrino energy between 1 and 4 GeV. The simplest way this scheme can be modified is by changing that energy dependence. Tables 6.2 and 6.3 show four such alterations, in FHC and RHC respectively. It can be seen that neither binning increasingly more finely (moving from result A to E) or widening the range of accepted energies to between 0 and 10 GeV, lead to relatively large increases in the value of $\Delta\kappa_{CP}^2$. Note that although more events are introduced by widening the energy range, negative changes in sensitivity can be seen in the two tables. These entries correspond to the cases where any gain in sensitivity due to this introduction is offset by energy binning which is coarser than the nominal configuration. In the core sample, therefore, the 2020 energy binning scheme of 0.5 GeV bins in reconstructed neutrino energy between 1 and 4 GeV is retained. However, introducing energy dependency to the peripheral sample results in a modest increase of approximately 1%, even with the coarsest of binnings. Unfortunately, since the peripheral sample by construction is made up of events which are not fully contained within the fiducial volume, the estimator for total neutrino energy is not expected to perform well. Peripheral events, however, must pass a very stringent electron neutrino PID cut. Namely, they are likely to have a well identified electron shower which carries a considerable fraction of the total event energy. With this in mind, the final rows of the two Tables show $\Delta\kappa_{CP}^2$ when the peripheral is binned as a function of *primary shower calorimetric energy*. As discussed in Section 4.5.1, ‘primary shower’ refers to the most energetic shower in the event and ‘calorimetric energy’ is a sum of calibrated hits, corrected for detector dead material. The simplicity of the variable and the specificity of the event type being quantified, make it a safer and more appropriate choice than using the energy of the whole event, where some might be missing. Approximately two thirds of the sensitivity from binning as function of neutrino energy is retained. Although this is a relatively small gain, given the simplicity of the change, the decision was made to implement three bins of primary shower calorimetric energy between 0 and 4.5 GeV in the peripheral sample. The lower and upper boundaries on the energy reflect the limits on the neutrino energy in the 2020 analysis. Figures 6.1a and 6.1b summarise the tabulated results graphically.




Figure 6.1: A summary of sensitivity improvements following changes to energy binning at the FD in FHC (**left**) and RHC (**right**). The normal and inverted neutrino mass hierarchy cases are shown by solid and dashed curves respectively. Each colour represents a proposed change to a current variable in, or, the introduction of a new variable to, the FD spectrum. Letters A - E correspond to how that variable is binned. A key can be found in Tables 6.2 and 6.3.

Binning as a function of energy plus other reconstructed energy variables such as electro-magnetic energy and reconstructed hadronic energy fraction was also trialled, to very little avail. Similarly, other kinematic variables at both the event and shower level such as total transverse momentum, transverse momentum fraction of the leading shower, total number of Michel electrons and primary shower direction, did not contribute any meaningful additional sensitivity.

Particle identification:

The investigation of how best to incorporate the PID variables into the analysis was done in three distinct steps. First, the effect of finer binning in the two PID variables already used in the appearance analysis, CVNe (CVN A) and cosmic rejection, was explored. Second, binning in terms of a second electron neutrino classification score not used in the main analysis, CVN B, was introduced. CVN B is directly analogous to CVN A but was trained on signal and background samples containing events satisfying slightly different loose kinematic cuts. Both of these changes result in mild to modest increases in $\Delta\kappa_{CP}^2$. Third, the lower bounds of the CVN A cut values were reduced to 0.5 and the cosmic rejection cut removed, thereby increasing the total electron neutrino appearance candidate

			Normal hierarchy		Inverted hierarchy	
	No. bins	Key	$\Delta\kappa_{CP}^2$	POT-equiv (%)	$\Delta\kappa_{CP}^2$	POT-equiv (%)
Neutrino energy (core)	12	A	8.89	0.23	7.88	0.51
	24	B	8.90	0.34	7.89	0.64
	48	C	8.91	0.45	7.89	0.64
	96	D	8.91	0.45	7.90	0.77
	192	E	8.92	0.56	7.91	0.89
Wider neutrino energy (core)	12	A	8.83	-0.40	7.76	-1.02
	24	B	8.89	0.23	7.86	0.26
	48	C	8.91	0.45	7.89	0.64
	96	D	8.92	0.56	7.90	0.77
	192	E	8.92	0.56	7.91	0.89
Neutrino energy (peri.)	2	A	8.96	1.01	7.94	1.28
	4	B	8.99	1.35	7.97	1.66
	8	C	9.00	1.47	7.99	1.91
	16	D	9.01	1.58	8.00	2.04
	32	E	9.01	1.58	8.00	2.04
Primary shower calorimetric energy (peri.)	2	A	8.94	0.79	7.92	1.02
	4	B	8.95	0.90	7.93	1.15
	8	C	8.96	1.01	7.94	1.28
	16	D	8.97	1.13	7.95	1.40
	32	E	8.97	1.13	7.95	1.40

Table 6.2: A summary of $\Delta\kappa_{CP}^2$ values for several analyses with alternative energy binning at the FD in FHC. The corresponding data equivalent sensitivity improvements (with respect to the 2020 analysis) are given as percentages in the column marked POT-equiv. This information is shown graphically in Figure 6.1a.

			Normal hierarchy		Inverted hierarchy	
	No. bins	Key	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)
Neutrino energy (core)	12	A	2.82	0.71	3.52	0.57
	24	B	2.83	1.07	3.53	0.86
	48	C	2.84	1.43	3.53	0.86
	96	D	2.84	1.43	3.53	0.86
	192	E	2.84	1.43	3.54	1.14
Wider neutrino energy (core)	12	A	2.75	-1.79	3.47	-0.86
	24	B	2.81	0.36	3.51	0.29
	48	C	2.84	1.43	3.53	0.86
	96	D	2.85	1.79	3.53	0.86
	192	E	2.85	1.79	3.53	0.86
Neutrino energy (peri.)	2	A	2.85	1.79	3.56	1.71
	4	B	2.86	2.14	3.57	2.00
	8	C	2.87	2.50	3.58	2.29
	16	D	2.88	2.86	3.58	2.29
	32	E	2.88	2.86	3.59	2.57
Primary shower calorimetric energy (peri.)	2	A	2.83	1.07	3.54	1.14
	4	B	2.84	1.43	3.55	1.43
	8	C	2.85	1.79	3.55	1.43
	16	D	2.85	1.79	3.56	1.71
	32	E	2.85	1.79	3.56	1.71

Table 6.3: A summary of $\Delta\kappa_{\text{CP}}^2$ values for several analyses with alternative energy binning at the FD in RHC. The corresponding data equivalent sensitivity improvements (with respect to the 2020 analysis) are given as percentages in the column marked POT-equiv. This information is shown graphically in Figure 6.1b.




Figure 6.2: A summary of sensitivity improvements following changes to PID binning at the FD in FHC (**left**) and RHC (**right**). The normal and inverted neutrino mass hierarchy cases are shown by solid and dashed curves respectively. Each colour represents a proposed change to a current variable in, or, the introduction of a new variable to, the FD spectrum. Letters A - F correspond to how that variable is binned. A key can be found in Tables 6.4 and 6.5.

sample size. The result of combining these looser cuts with binning in each of the PIDs is presented in Tables 6.4 and 6.5, and, Figures 6.2a and 6.2b. Similar partitioning of the peripheral sample was also attempted to no avail. Furthermore, note that to assess the importance of finer binning in the CVN variables, the fractional gain in sensitivity quoted in the first and second rows of the tables are calculated with respect to bin configuration ‘A’ in the first row. This configuration has the FD binning structure of the 2020 analysis but with relaxed cuts on cosmic rejection and CVN A. Any events with a CVN A score falling below the 2020 low threshold are put into an additional single selection bin with the standard reconstructed neutrino energy binning between 1 and 4 GeV. The extra events that enter the analysis due to the loosening of the cosmic rejection cut are distributed among all selection bins. Even the purity of the bin with the most stringent CVN requirement is affected.

It is clear from Figures 6.2a and 6.2b that there are gains to be had both by including the second CVN classification score (green) and by binning in, rather than simply cutting on, cosmic rejection score (red). With large numbers of bins, these changes result in sensitivity increases equivalent to analysing approximately 6.5% and 5% more FHC data respectively assuming the normal hierarchy. Unfortunately, introducing 96 bins in even

			Normal hierarchy		Inverted hierarchy	
	No. bins	Key	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)
CVN A (core)	3	A	8.50	0	7.44	0
	6	B	8.66	1.80	7.59	2.02
	12	C	8.71	2.47	7.65	2.82
	24	D	8.73	2.71	7.67	3.09
	48	E	8.73	2.71	7.68	3.23
	96	F	8.74	2.82	7.69	3.36
CVN B (core)	6	A	8.97	5.53	7.88	5.91
	12	B	9.00	5.88	7.91	6.32
	24	C	9.02	6.12	7.93	6.59
	48	D	9.04	6.35	7.95	6.85
	96	E	9.06	6.59	7.98	7.26
Cosmic rejection (core)	2	A	9.04	1.91	7.93	1.15
	4	B	9.19	3.48	8.06	2.81
	8	C	9.27	4.51	8.13	3.70
	16	D	9.30	4.62	8.15	3.95
	32	E	9.31	4.96	8.16	4.08
	64	F	9.32	5.07	8.18	4.34

Table 6.4: A summary of $\Delta\kappa_{\text{CP}}^2$ values for several analyses with alternative PID binning and cuts at the FD in FHC. The corresponding data equivalent sensitivity improvements are given as percentages in the column marked POT-equiv. These improvements are calculated with respect to the 2020 analysis in the case of cosmic rejection and to bin configuration A of CVN A for the others. See the main text for justification. This information is shown graphically in Figure 6.2a.

			Normal hierarchy		Inverted hierarchy	
	No. bins	Key	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)
CVN A (core)	3	A	2.72	0	3.41	0
	6	B	2.76	1.47	3.44	0.88
	12	C	2.77	1.84	3.46	1.47
	24	D	2.79	2.57	3.47	1.76
	48	E	2.79	2.57	3.47	1.76
	96	F	2.80	2.94	3.48	2.05
CVN B (core)	6	A	2.82	3.68	3.52	3.23
	12	B	2.82	3.68	3.53	3.52
	24	C	2.83	4.04	3.53	3.52
	48	D	2.84	4.41	3.54	3.81
	96	E	2.85	4.78	3.56	4.40
Cosmic rejection (core)	2	A	2.83	1.07	3.53	0.85
	4	B	2.86	2.14	3.57	2.00
	8	C	2.88	2.86	3.59	2.57
	16	D	2.89	3.21	3.60	2.86
	32	E	2.89	3.21	3.60	2.86
	64	F	2.89	3.21	3.61	3.14

Table 6.5: A summary of $\Delta\kappa_{CP}^2$ values for several analyses with alternative PID binning and cuts at the FD in RHC. The corresponding data equivalent sensitivity improvements are given as percentages in the column marked POT-equiv. These improvements are calculated with respect to the 2020 analysis in the case of cosmic rejection and to bin configuration A of CVN A for the others. See the main text for justification. This information is shown graphically in Figure 6.2b.

one extra variable is neither practical nor desirable. Bin combinatorics grow quickly with every new inclusion, leading to greater computer resource requirements and the need for larger MC samples. Furthermore, the appearance candidate sample in data contains of order 100 events in total. The inclusion of empty bins in FD spectra not only leads to an analysis which is more complex but contributes to longer fit times. Fit times grow approximately linearly with the number of bins. For these reasons it is sensible to impose an upper limit on the total number of analysis bins, while still drawing the best out of each of the most promising individual changes. Therefore, an FD core spectrum containing three bins in cosmic rejection and three bins in each CVN (162 total core bins including reconstructed neutrino energy) is taken forward for further optimisation. The optimisation procedure is discussed in the next section.

6.1.2 Optimisation of Bin Edges

In Section 6.1.1, all values of $\Delta\kappa_{CP}^2$ are calculated assuming each new variable is introduced in bins of equal size. As discussed in Section 6.1.1, for the core sample, the lower bounds on both CVNs were reduced to 0.5 and completely removed for cosmic rejection. As a result, the range of possible values that can be taken by either CVN is from 0.5 to 1. The cosmic rejection score can take values between 0.35 and 0.72 (see Figure 4.10). It is unlikely that the most optimal way to partition the sample is to make the selection bins equally spaced. A method has been developed to simultaneously locate the positions of all selection bin edges such that $\Delta\kappa_{CP}^2$ is maximised. Note the specification of *selection* bin edges. Bins of reconstructed neutrino energy are fixed at intervals of 0.5 GeV between 1 and 4 GeV.

For a set of variables A , each with N_a possible bin edge locations (not including the fixed upper and lower bounds) and k_a required bins, the number of possible binning schemes B is given by the following product of binomial coefficients

$$B = \prod_{a \in A} \text{Binom}(N_a, k_a - 1). \quad (6.7)$$

Even for the case of three variables, each with three floating bin edges across 50 possibly positions, B is greater than a billion. Within the bounds of standard CAFAna, making several FD predictions each with different binning, is difficult. Each prediction would require the filling of its own histogram. Instead, CAFAna is used to make a single spectrum which is binned very finely at the limit of resolution in each of the PID variables. This

single spectrum can generate four histograms, each one corresponding to an FD prediction at one of the test points used to calculate $\Delta\kappa_{CP}^2$, $(\delta_{CP} = \pi/2, \delta_{CP} = 3\pi/2) \otimes (\text{NH}, \text{IH})$. Evaluating a given binning scheme then equates to a ‘re-binning’ of these finely binned histograms in selection bin space. In practise, re-binning involves keeping track of a sum of all fine bins in each region of selection bin space bordered by the bin edges defining the binning scheme in consideration. This is a fast operation. The resulting pairs of re-binned histograms (corresponding to test points $\delta_{CP} = \pi/2$ and $\delta_{CP} = 3\pi/2$) can then be used to calculate values of $\Delta\kappa_{CP}^2$ for each selection bin scheme. The best scheme is the one which gives the greatest sensitivity, namely, the one which maximises $\Delta\kappa_{CP}^2$.

The distributions of CVN are reasonably flat at low values but show a lot of shape toward 1. For this reason, between 0.5 and 0.75, the fine binned predictions generated in CAFAna use eight bins. Between 0.75 and 1, possible bin edge positions are placed every 0.01 units. In the case of cosmic rejection, possible positions are placed every 0.01 units. This corresponds to roughly 150 million selection bin scenarios to be evaluated in each horn current¹. The selection bin scenario corresponding to the largest $\Delta\kappa_{CP}^2$ is denoted by the vertical blue line. The CVN and cosmic rejection bin edge positions for this scenario are given in the legend and summarised in Table 6.6. The resulting FD spectra are shown in Figures 6.3a and 6.3b. Since four variables are now being binned in (two CVNs, cosmic rejection and reconstructed neutrino energy), the spectra are presented as 1D functions of ‘Analysis Bins’. Regions of low, medium and high cosmic rejection score can be extracted by eye based on the contribution of cosmic events to each bin. Furthermore, bins appear to be in groups of six. These are the six reconstructed neutrino energy bins which retain the logical order within each group. Further optimisation by eye of this binning scheme to obtain the final version is discussed in the next section. A more intuitive, visual presentation of the analysis bins is given at this point.

A benefit of the algorithm discussed here is that it can be run in ‘bare’ C++. Once the four prediction histograms have been made using CAFAna and their information exported to a basic text file format, the code is implemented exclusively of even ROOT [99]. The operation performed to evaluate each binning scenario is a simple one and the total number of possible scenarios can easily be tackled by parallelisation. It has been demonstrated that around 100 million scenarios can be evaluated by 1000 nodes with 4 GB of memory in

¹The procedure was run under both the normal and inverted neutrino mass hierarchies. The optimal bin edge positions were found to be very similar for both.

	FHC	RHC
Cosmic rejection	0.35, 0.47, 0.52, 0.72	0.35, 0.49, 0.58, 0.72
CVN A	0.50, 0.89, 0.99, 1.00	0.50, 0.83, 0.97, 1.00
CVN B	0.50, 0.77, 0.94, 1.00	0.50, 0.93, 0.99, 1.00

Table 6.6: A summary of the selection bin edges which maximise the value of $\Delta\kappa_{CP}^2$.

2 to 3 hours. This excludes job queue time. The output file produced is also in a basic text file format with one line per bin edge combination evaluated. The algorithm is currently limited by its memory requirements. Each node running the code generates all possible bin edge permutations but only a subset (equal to the total number of permutations divided by the number of nodes being used) are ‘saved’ and taken forward for evaluation on that node. Therefore, there is no upper limit on the *total* number of permutations that can be evaluated, provided that the algorithm is run across a sufficient number of nodes. Refactoring the algorithm so that each permutation is evaluated on the fly may be possible in the future. The optimisation of cuts based on histogram input is a common task at NOvA, giving the algorithm the potential to be used more widely.

6.1.3 Optimisation by Eye

Although the FD spectra shown in Figures 6.3a and 6.3b reflect the most sensitive binning configuration, it is clear by eye that the total number of bins can be greatly reduced without impacting this sensitivity too greatly. There are areas of empty or nearly empty bins and specific regions of phase space which are completely dominated by cosmic background. Figures 6.4a and 6.4b can be used to assess which parts of the spectra are most important for sensitivity. They show the fractional contribution of each bin to the total $\Delta\kappa_{CP}^2$. In FHC, 99.9% of the total $\Delta\kappa_{CP}^2$ can be retained using just 67 bins (information in the plot legend), or, just under half of the current total. With this in mind and by evaluating a text file list of analysis bins ordered by their contribution to $\Delta\kappa_{CP}^2$, the total number of bins are reduced as follows:

- Complete removal of the lowest cosmic rejection bin (0.35 to 0.47) in FHC and (0.35 to 0.49) in RHC. This roughly corresponds to taking out analysis bins 0 to 50.
- Removal of heavily ‘off-diagonal’ selection bins, for example, (high CVN A, low CVN B) in FHC and (low CVN A, high CVN B) in RHC. This takes out the smaller regions of very little or no events.

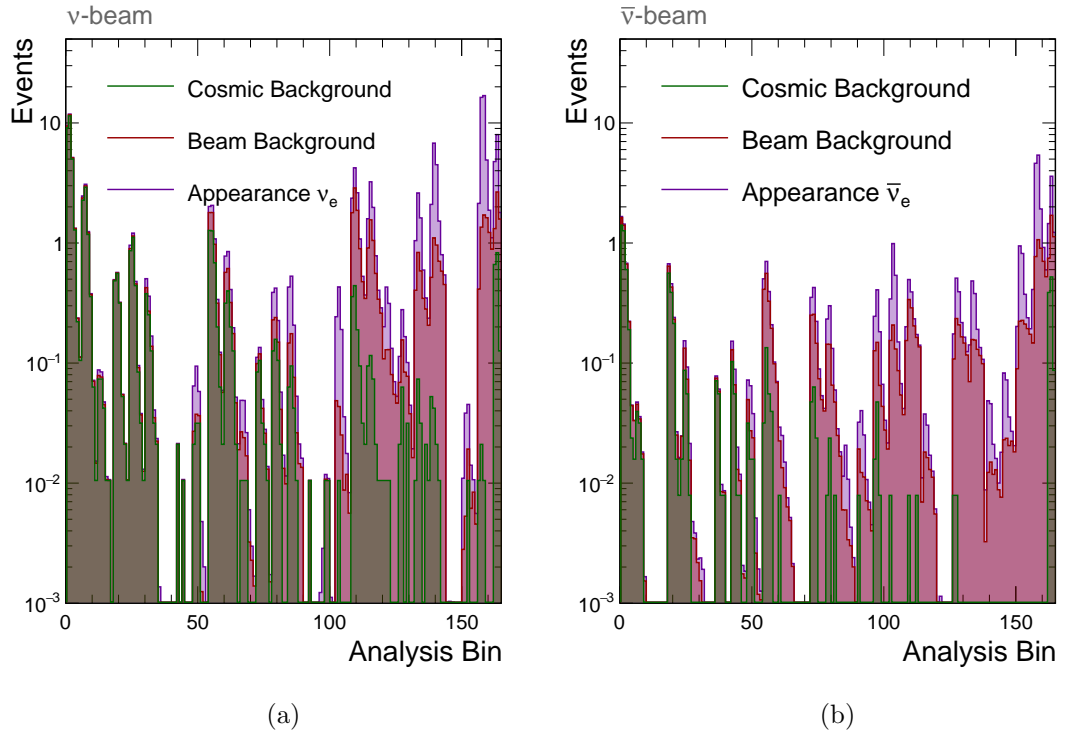



Figure 6.3: Prediction of events at the FD (core and peripheral) with selection bin edges that maximise the value of $\Delta\kappa_{CP}^2$ in FHC (**left**) and RHC (**right**). Events are binned as a function of four variables, shown here projected onto a single axis. The green, red and purple stacked histograms show the cosmic, beam background and signal components of the total prediction respectively.

Bin	PID description	Cosmic rejection	CVN A	CVN B
(a)	(low, low)	Low: ≥ 0.47 . High: ≥ 0.52	[0.50 - 0.89)	[0.50 - 0.77)
(b)	(low, medium) I	”	[0.50 - 0.89)	[0.77 - 0.94)
(c)	(low, medium) II	”	[0.89 - 0.99)	[0.50 - 0.77)
(d)	(medium, medium)	”	[0.89 - 0.99)	[0.77 - 0.94)
(e)	(low, high)	”	[0.50 - 0.89)	[0.94 - 1.00]
(f)	(medium, high)	”	[0.89 - 0.99)	[0.94 - 1.00]
(g)	(high, high)	”	[0.99 - 1.00]	[0.94 - 1.00]


Table 6.7: Selection bin boundaries for the fully optimised FD axes in FHC. Letters (a - g) are used to label all FD spectra included in the remainder of thesis. Figure 6.5 shows an example. These boundaries are also applicable to the ND with the exception of cosmic rejection. The ND sample is not binned according to cosmic rejection (see Section 6.2).

- Particularly in bins of lower cosmic rejection, the vast majority of sensitivity comes from the 1 to 2.5 GeV region. Bins above 2.5 GeV where this is true are removed.

The implementation of these changes results in FD spectra with 54 and 60 analysis bins in FHC and RHC respectively. Figures 6.5 and 6.6 present these spectra in their final form. Tables 6.7 and 6.8 define each of the selection bins labelled in the Figures (a - g, low and high cosmic rejection) in terms of their CVN and cosmic rejection boundaries for FHC and RHC respectively.



(a)



(b)

Figure 6.4: The contribution of each analysis bin to the total $\Delta\kappa_{CP}^2$ in FHC (**left**) and RHC (**right**). The normal and inverted hierarchy cases are shown in purple and pink respectively. The number of bins necessary to retain 95%, 99%, 99.5% and 99.9% of the total $\Delta\kappa_{CP}^2$ is given in the legend.




Figure 6.5: Distributions of reconstructed neutrino energy and reconstructed primary shower energy for the core and peripheral samples respectively at the FD in FHC without extrapolation. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.

Bin	PID description	Cosmic rejection	CVN A	CVN B
(a)	(low, low)	Low: ≥ 0.49 . High: ≥ 0.58	[0.50 - 0.83]	[0.50 - 0.93]
(b)	(low, medium)	”	[0.83 - 0.97]	[0.50 - 0.93]
(c)	(low, high)	”	[0.97 - 1.00]	[0.50 - 0.93]
(d)	(medium, medium)	”	[0.83 - 0.97]	[0.93 - 0.99]
(e)	(medium, high) I	”	[0.83 - 0.97]	[0.99 - 1.00]
(f)	(medium, high) II	”	[0.97 - 1.00]	[0.93 - 0.99]
(g)	(high, high)	”	[0.97 - 1.00]	[0.99 - 1.00]

Table 6.8: Selection bin boundaries for the fully optimised FD axes in RHC. Letters (a - g) are used to label all FD spectra included in the remainder of thesis. Figure 6.6 shows an example. These boundaries are also applicable to the ND with the exception of cosmic rejection. The ND sample is not binned according to cosmic rejection (see Section 6.2).




Figure 6.6: Distributions of reconstructed neutrino energy and reconstructed primary shower energy for the core and peripheral samples respectively at the FD in RHC without extrapolation. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.

	FHC			RHC		
	$\Delta\kappa_{\text{CP}_{2020}}^2$	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)	$\Delta\kappa_{\text{CP}_{2020}}^2$	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)
Normal hier.	8.87	9.36	5.5%	2.80	2.89	3.2%
Inverted hier.	7.84	8.25	5.2%	3.50	3.61	3.1%

Table 6.9: A comparison of the sensitivity achieved when using the thesis FD binning versus the 2020 FD binning in an FD only analysis with no systematic uncertainty. The corresponding data equivalent sensitivity improvements are given as percentages in the column marked POT-equiv.

A summary of the gain in sensitivity when using the FD ‘thesis’ binning derived in this section compared to the 2020 analysis axes is given in Table 6.9. The 2020 methodology would be required to analyse 5.5% more FHC and 3.2% more RHC data to achieve an equivalent sensitivity under the normal hierarchy. Note that these numbers are only an indication of the final sensitivity of the new analysis. They assume no systematic uncertainty and information from the ND has not been included. These two caveats are the subjects of the remaining sections in this chapter.

6.2 Studies at the ND

The ND plays a key role in reducing the effect of a number of systematic uncertainties on the FD predicted spectrum. Its energy spectra are divided up by PID in exactly the same way as the FD. As discussed in Section 4.5.1, a technique known as decomposition is used to assign any difference between data and simulation at the ND among the different beam components. In order to constrain the beam background in the new FD binning scheme, decomposition must be revisited.

Figures 6.7 and 6.8 show the energy spectra of the electron neutrino-like samples at the ND. These spectra show the results of the decomposition used in the final fit for oscillations, presented in Chapter 7. There is a full discussion of this decomposition in Section 7.1. This includes tabulated event counts for each selection bin and comparisons of data and simulation in several reconstructed variables, before and after decomposition. At this point, however, it is important to highlight particular features in this decomposition that warrants extra attention. First, note that due to the larger overburden of the ND compared to the FD, there is no need for a specially trained classifier to identify cosmic

interactions. Therefore, each selection bin (a) - (g) at the ND is used to correct *both* the low and high cosmic rejection bins of each analogous selection bin at the FD. Figures 6.9 through to 6.12 justify this choice. They show distributions of CVN A and CVN B in each selection bin at the FD, for just beam events. No extrapolation is used. It can be seen that the form of background is very similar in counterpart low and high cosmic rejection bins. It is appropriate, therefore, to correct each counterpart with the same ND information.

An issue that arises on close inspection of the ND data and MC is caused by the way that rock events are simulated. Take, for example, FHC selection bin (f) between 1 and 1.5 GeV shown in Figure 6.7. Overall, data is approximately 15% above simulation in this bin. This can be seen from the grey histogram in the ratio plot. The NC component of this bin in simulation (blue), however, is corrected *down*. As a consequence, the ν_μ -CC component (green) is corrected up by far more than 15% to compensate. The cause of this is found to be events with an interaction vertex outside the detector volume. This is discussed in detail in the next section.

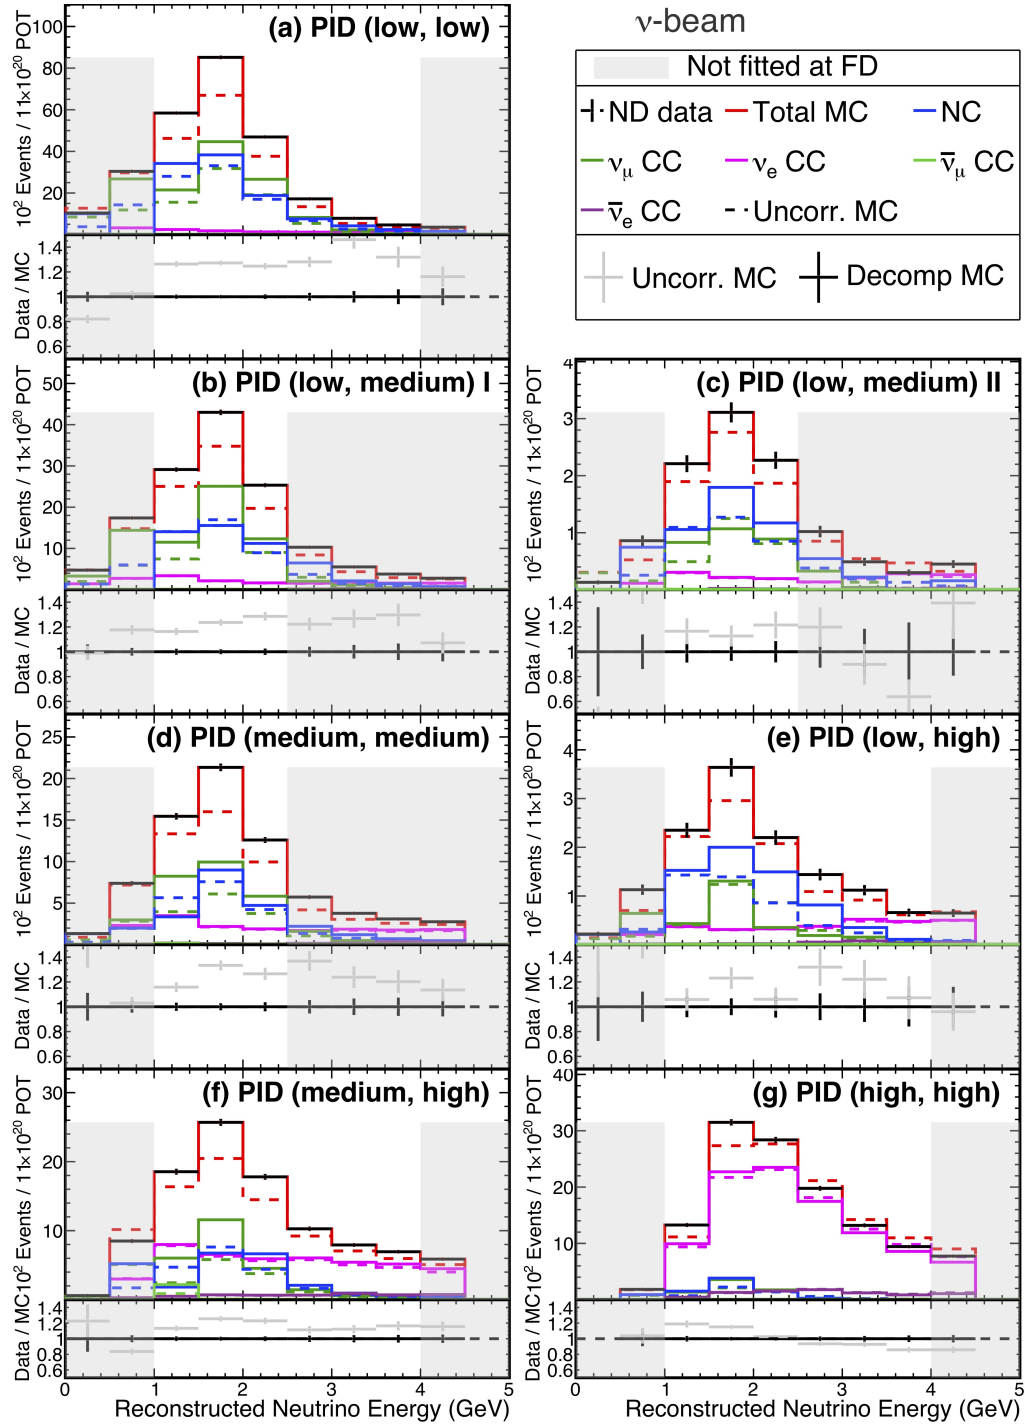


Figure 6.7: Distributions of reconstructed neutrino energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown. Any bin that is not extrapolated to the FD, that is to say, not included in the fit for oscillations, is lightly shaded.




Figure 6.8: Distributions of reconstructed neutrino energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown. Any bin that is not extrapolated to the FD, that is to say, not included in the fit for oscillations, is lightly shaded.




Figure 6.9: Distributions of CVN A for the core and peripheral samples respectively at the FD in FHC without extrapolation. Cosmic backgrounds are not shown. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.

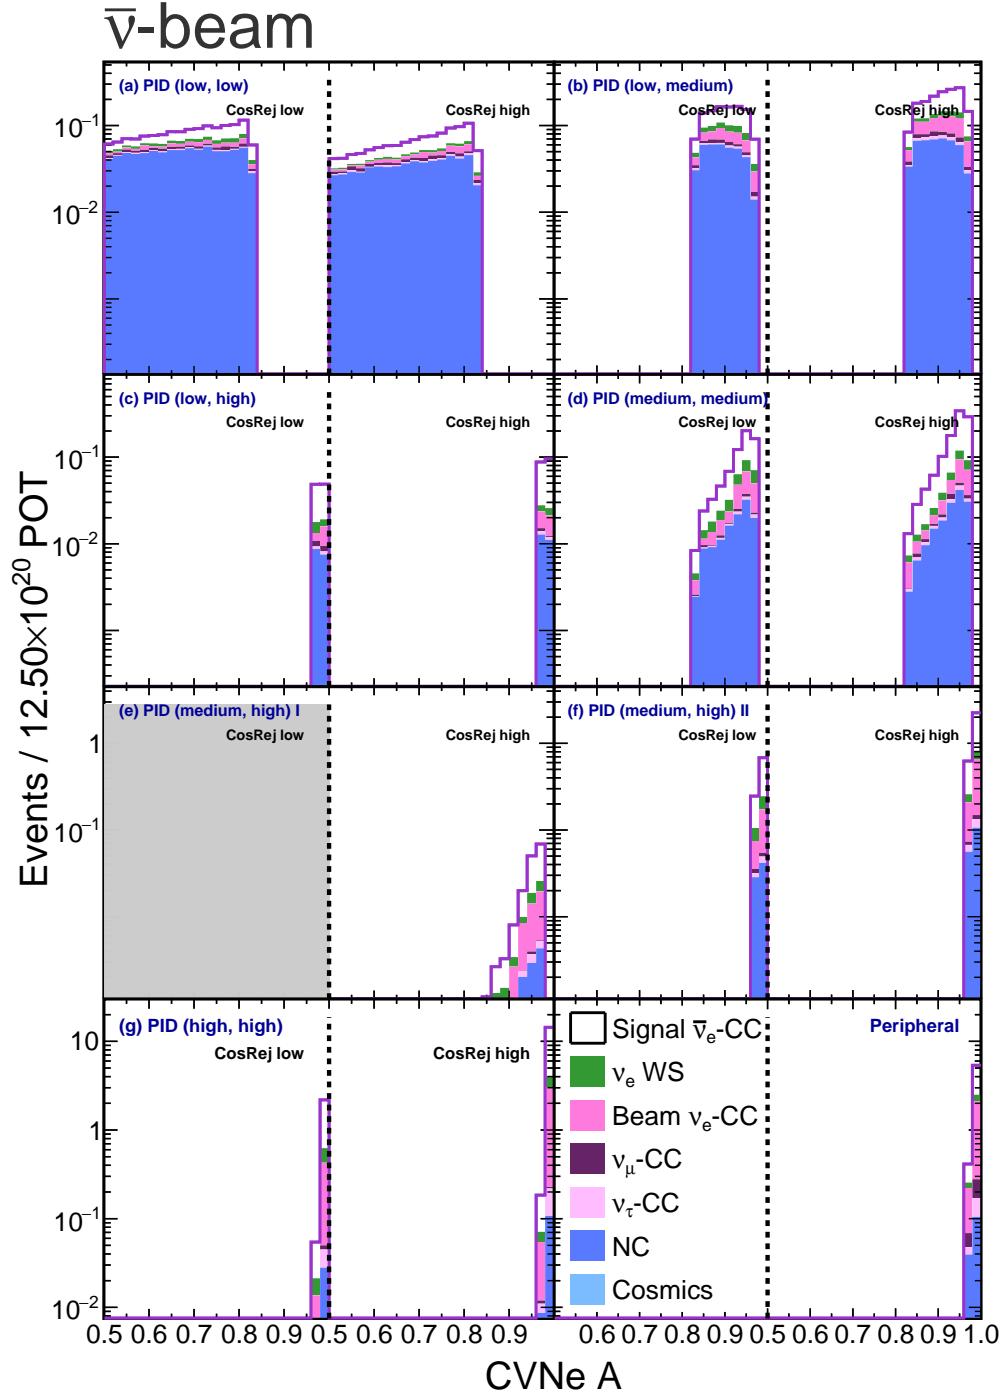


Figure 6.10: Distributions of CVN A for the core and peripheral samples respectively at the FD in RHC without extrapolation. Cosmic backgrounds are not shown. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.

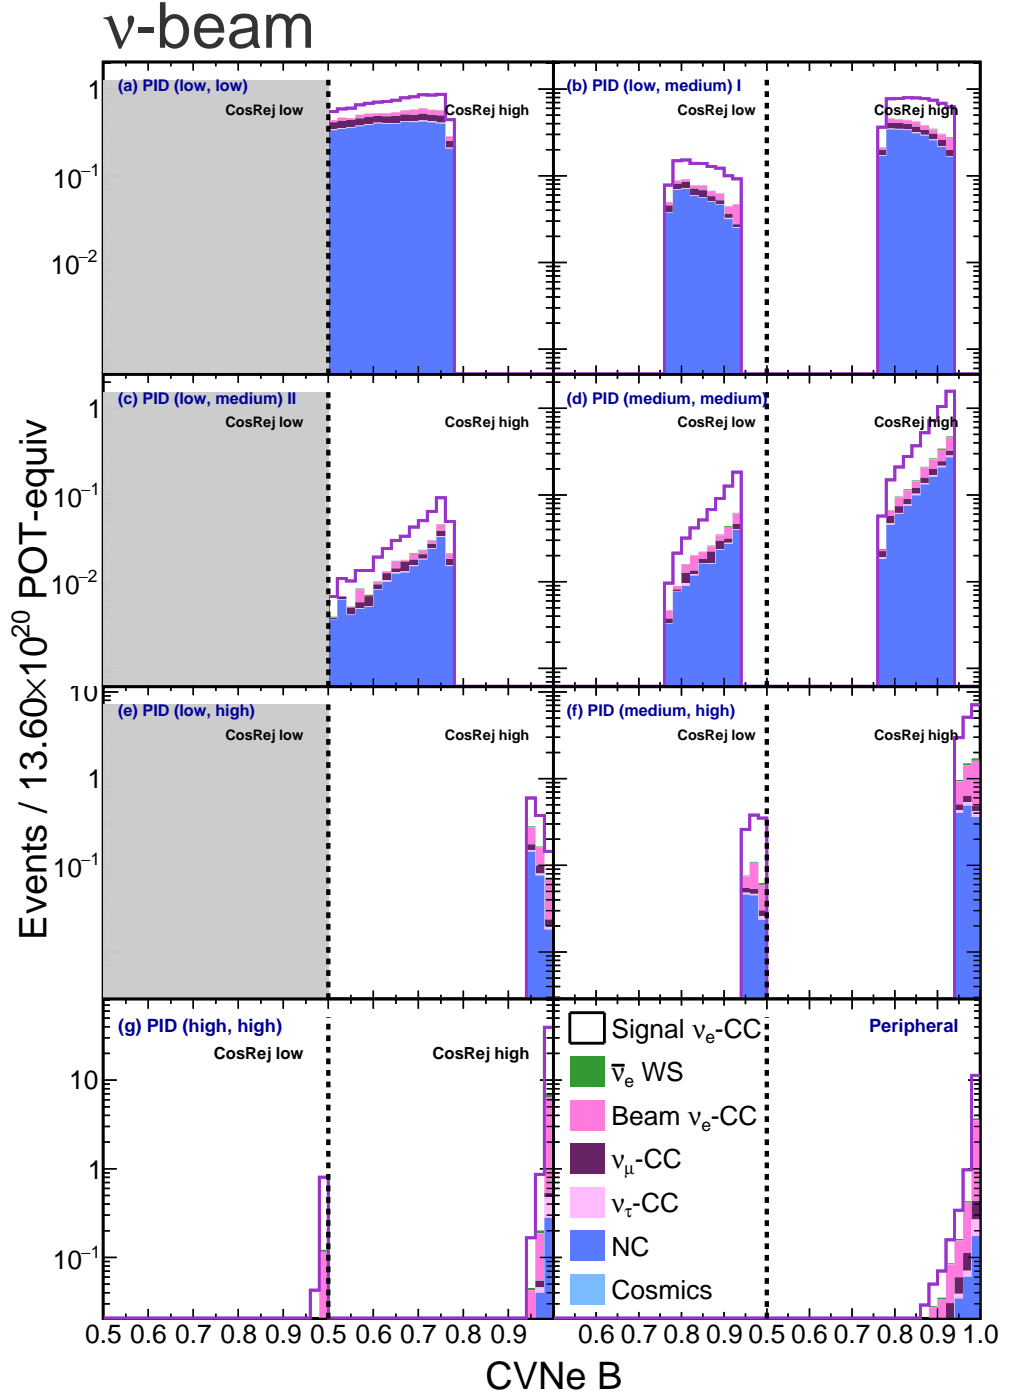


Figure 6.11: Distributions of CVN B for the core and peripheral samples respectively at the FD in FHC without extrapolation. Cosmic backgrounds are not shown. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.




Figure 6.12: Distributions of CVNe B for the core and peripheral samples respectively at the FD in RHC without extrapolation. Cosmic backgrounds are not shown. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.

6.2.1 Simulation of Rock Events

Anomalous behaviour in a small number of reconstructed energy bins following the decomposition of the electron neutrino-like sample at the ND is observed. A hypothesis based on simulated events with a true interaction vertex outside of the detector volume is investigated by ‘re-running’ the decomposition procedure specifically excluding these ‘rock events’. Figures 6.13 and 6.14 show the results of this decomposition. These plots should be interpreted carefully. The coloured histograms showing both uncorrected and corrected simulation include rock events. The decomposition weights, however, are calculated with rock events removed. Therefore, uncorrected simulation and its ratio to data are identical in these Figures to Figures 6.7 and 6.8. Conversely, corrected simulation (a sample containing rock events) has been corrected using weights assuming no rock events - namely, it is being ‘over-corrected’. This is the reason that total simulation now exceeds data everywhere. As a consequence, the extent to which this corrected simulation exceeds data (given by the solid black ratio), is purely a result of the influence of rock events. Any bins which are shaded grey are not used at the FD, so are of no concern. Considering bins which are extrapolated, it can be seen that the rock events are most influential at low energy.

Rock events introduce such behaviour as a result of the way that they are implemented into the ND simulation. They are generated separately and ‘overlaid’ onto the simulation of beam events interacting inside the detector. A large number of events are required, necessitating the re-use of the same rock interactions for different beam spills. On the whole, rock events are not selected themselves but are important due to their contribution to pile-up. However, when they are selected the same rock events can be selected multiple times, thereby distorting certain bins. Figure 6.15a shows the number of times each uniquely simulated rock event is selected, the so called ‘re-use rate’, in each FHC PID bin. In several cases, the same interaction is being selected greater than 100 times. Figure 6.15b shows that the large majority of selected rock events arise from muon neutrino CC events and that NC interactions account for approximately one third of the total. This matches the typical event type fractions, given the composition of the beam.

Introducing to the analysis samples in regions of phase space where statistics are limited amplifies the re-use effect. Areas of high statistics, regardless of their rock event purity, are less susceptible since the fractional contribution to the total events of the oc-

casional event being selected multiple times remains low. Figure 6.16a shows the fraction of each sample which are rock events. This is one reason that PID bin (a) remains largely unaffected. PID bin (g) on the other hand is largely unaffected since rock events are less likely to pass the stricter PID cuts. The story is similar in RHC. The decomposition procedure itself, however, is less susceptible to adverse behaviour from rock events since any total difference in data and simulation is shared equally between all beam components.

A number of ways to address the behaviour arising from the rock events have been considered. Removing the rock events completely or partially by, for example, limiting the re-use rate, are both possibilities. However, it was decided that a 100% uncertainty should be applied in the appearance analysis to all selected ND events with a true vertex outside of the detector. The systematic is implemented such that it accounts for two effects: uncertainty due to in the overall rate of accepted rock events and uncertainty in the fractional difference attributed to each beam component. The influence of this systematic on the appearance background prediction at the FD in FHC is shown in Figure 6.16b. As expected, the largest shifts in the nominal spectrum are observed in the low energy bins (bins centred at 9.5, 12.5, 15.5 and 39.5, denoted by arrows on the plot) which use extrapolation weights from the ND bins most heavily influenced by rock events. A map to translate from ‘Analysis Bin’ as shown in this plot to the presentational reconstructed neutrino energy binning used in Figures such as Figure 6.5 is given in Appendix B. The effect of the systematic on the ν_e -CC signal prediction is extremely small. Its overall effect on the sensitivity of the analysis to the oscillation parameters is addressed in the next section.

A final thought in this section relates to how rock events should be handled in the analysis in the future. Due to the larger surface area to volume ratio of the ND, events with a true interaction vertex outside of the detector are a much larger fraction of the selected sample at the ND compared to the FD. It is, therefore, undesirable to propagate the differences related to this effect via extrapolation. As can be seen from Figure 6.15b, the large majority of rock events are from muon neutrino CC interactions. To be selected, these events very likely contain a single photon that enters the detector unseen. The kinematics of the final state particles in the selected rock events may, therefore, be quite different to fiducial events. This suggests that it would be worth studying the possibility of subtracting off the rock events from the ND samples before performing the decomposition. To do this, much higher statistics samples with a lower rock event re-use rate would be

required.




Figure 6.13: Distributions of reconstructed neutrino energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Decomposition weights are calculated exclusively of rock events. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown. Any bin that is not extrapolated to the FD, that is to say, not included in the fit for oscillations, is lightly shaded.




Figure 6.14: Distributions of reconstructed neutrino energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Decomposition weights are calculated exclusively of rock events. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown. Any bin that is not extrapolated to the FD, that is to say, not included in the fit for oscillations, is lightly shaded.

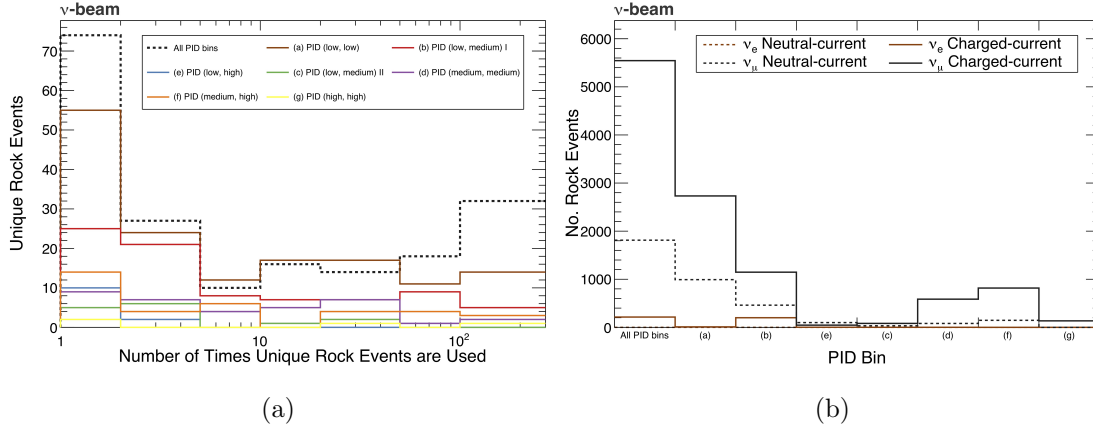


Figure 6.15: **Left:** The number of times that a uniquely simulated rock event is selected, the so called ‘re-use rate’, split by PID bin in FHC. **Right:** The number of rock events selected in each PID bin in FHC broken down by true interaction type including re-use.

6.3 Fake Data Sensitivities

Having established the decomposition procedure for the electron neutrino-like sample, the ND can now be included in sensitivity studies. Throughout the remainder of this thesis, any sensitivities (including those presented in Section 7.3) use *fake* FD data and *real* ND data². It is also possible to use fake ND data for sensitivity studies, however, this doubles the amount of computing resources required when making FD ‘prediction files’. Since FD predictions files made with real ND data are essential for the final oscillation fit, the decision was made not to produce prediction files using fake ND data.

6.3.1 The Effect of Systematic Uncertainty on Sensitivities

The level at which the simple sensitivity metric $\Delta\kappa_{CP}^2$ is affected by systematic uncertainty can be estimated by allowing the systematic parameters in the fit penalty terms to move away from zero. Specifically, a fit of the prediction at $\delta_{CP} = 3\pi/2$ to the fake data at $\delta_{CP} = \pi/2$ is performed, keeping the oscillation parameters fixed while minimising over all of the systematic penalty terms. $\Delta\kappa_{CP}^2$ can then be recalculated as in Equation 6.1 but with $E_i^{3\pi/2}$ replaced by $\bar{E}_i^{3\pi/2}$ - the predicted number of events in analysis bin i after minimising with respect to the systematic penalty terms. Figure 6.17 shows the results of such a fit. The difference between the predictions at $\delta_{CP} = 3\pi/2$ without (green) and with (purple) minimisation over all systematic parameters is most clear in the largest peak in FHC. There are two reasons why the difference between the two is relatively small.

²Recall that fake data at NOvA is constructed by taking the central value of a simulated FD spectrum, predicted at some set of oscillation parameters and scaled to a particular POT exposure.




Figure 6.16: **Left:** The fraction of the total selected events which are rock in each PID bin in FHC. A uniquely simulated rock event which is selected twice would contribute two events to both the numerator and denominator. **Right:** The effect of systematic uncertainty arising from the ND rock event rate on the appearance background prediction at the FD in FHC. The red histogram shows the nominal total predicted FD beam background and the $1\text{-}\sigma$ range of uncertainty is included as a purple band. The ratios of the $\pm 1\text{-}\sigma$ spectra to nominal are shown in the lower plot.

	FHC + RHC		
	$\Delta\kappa_{\text{CP}_{2020}}^2$	$\Delta\kappa_{\text{CP}}^2$	POT-equiv (%)
Normal hier.	10.53	11.16	6.0%
Inverted hier.	10.45	11.05	5.7%

Table 6.10: A comparison of the sensitivity achieved when using the thesis FD axes versus the 2020 FD axes including the considered systematic uncertainty. FHC and RHC spectra are evaluated simultaneously and extrapolation of ND data is included. The corresponding data equivalent sensitivity improvements are given as percentages in the column marked POT-equiv.

First, the Gaussian penalty applied to all systematic pulls provides an incentive for the systematic parameters not to wander too far from zero. Second, note that FHC and RHC are fitted simultaneously. Each sample is trying to pull in opposite directions. On the whole, it is favourable for the prediction to be decreased in FHC and increased in RHC, to move toward the fake data.

As with the statistics only case, the $\Delta\kappa_{\text{CP}}^2$ described here can be compared for the analysis using the new FD binning scheme and the 2020 scheme. These results are summarised in Table 6.10. FHC and RHC are now evaluated jointly and extrapolation of ND data is included. Using the new FD binning scheme equates to an effective exposure increase of approximately 6% compared to the 2020 analysis.

$\Delta\kappa_{\text{CP}}^2$ can also be calculated treating each systematic uncertainty separately. Comparing this number to the $\Delta\kappa_{\text{CP}}^2$ assuming no systematic uncertainty gives an indication as to which systematic parameter most affects the overall sensitivity. This can help to direct effort toward uncertainties that are the most influential. For a given systematic, Figure 6.18 shows the fractional difference between $\Delta\kappa_{\text{CP}}^2$ and $\Delta\kappa_{\text{CP}}^2$ assuming no systematic uncertainty. This can be interpreted as the data equivalent cost to the sensitivity due to that given systematic. Systematics associated with neutrino cross-section and detector energy scale are found to be the largest factors limiting the sensitivity.

6.3.2 $\Delta\chi^2$ Studies

The metric $\Delta\kappa_{\text{CP}}^2$ is a simple way to quickly compare the sensitivity of multiple analysis methodologies to the oscillation parameters, specifically δ_{CP} . A new FD binning scheme




Figure 6.17: A comparison of fake data at $\delta_{CP} = \pi/2$ (black) and predictions at $\delta_{CP} = 3\pi/2$ when all systematic parameters are (purple) and are not (green) minimised over. FHC and RHC are shown in the top and bottom windows respectively. Systematic parameters are minimised for both samples simultaneously.



Figure 6.18: The fractional difference in the sensitivity metric ($\Delta\kappa_{CP}^2$) assuming each systematic uncertainty individually versus no systematic uncertainty. Note that not all of the neutrino cross-section uncertainties listed on the horizontal axis are treated independently in the fit, many of them are combined as discussed in Section 4.6.2. A description of what each of the systematics are can be found in Section 4.6.

based on this metric has been developed in this Chapter and will be used to perform an oscillation analysis, the results of which are shown in Chapter 7. This section compares the performance of the new analysis to the 2020 analysis using ‘traditional’ measures of sensitivity such as $\Delta\chi^2$ as a function of $\sin^2\theta_{23}$ and δ_{CP} . All of the sensitivities presented use fake FD data at the 2020 3-flavour best fit³ (Section 2.5.5) and real ND data, as discussed at the top of Section 6.3.

Figures 6.19a and 6.19b show contours of constant $\Delta\chi^2$ (Equation 4.20) in the $\sin^2\theta_{23}$ versus δ_{CP} and Δm_{32}^2 versus $\sin^2\theta_{23}$ spaces respectively in the normal hierarchy. Figures 6.19c and 6.19d show the same thing but in the inverted hierarchy. Each point on the underlying surface is obtained by fixing the values of the oscillation parameters which are labelled on the axes in the fit to fake data. All other parameters, including those describing systematics are then profiled over. The global minimum χ^2 is then subtracted from the minimum χ^2 at that point in oscillation space. As described in Section 4.7, the borders of the 1-, 2- and 3- σ allowed regions are drawn using a Gaussian estimate. These contours are slightly tighter everywhere for the new analysis. The largest differences can be seen in $\sin^2\theta_{23}$ around $\delta_{CP} = 3\pi/2$ in the normal hierarchy and $\delta_{CP} = \pi/2$ in the inverted hierarchy. Furthermore, around maximal mixing δ_{CP} is clearly more constrained at 1- σ . As expected, there is not much improvement in the Δm_{32}^2 space since this is driven by the disappearance measurement.

Figures 6.20a to 6.20d show various sensitivities to different interesting observables in 1D. Note the vertical axis shows the *square-root* of $\Delta\chi^2$. The plots can, therefore, be interpreted as showing a Gaussian significance in σ . Each $\Delta\chi^2$ is calculated in the following way [138]:

- Figure 6.20a: $i \in [0, 2\pi)$, $\Delta\chi_i^2 = \min(\chi_{\delta=0}^2, \chi_{\delta=0}^2) - \chi_{\delta=i}^2$. The significance with which CP symmetry can be rejected as a function of true δ_{CP} . Δm_{32}^2 , $\sin^2\theta_{23}$, θ_{13} , the mass hierarchy and all systematic parameters are profiled over.
- Figure 6.20b: $i \in [0, 2\pi)$, $\Delta\chi_i^2 = \left(\chi_{\text{Wrong hierarchy}}^2 - \chi_{\text{Right hierarchy}}^2\right)_{\delta=i}$. The significance with which the incorrect mass hierarchy can be rejected as a function of true δ_{CP} . δ_{CP} , Δm_{32}^2 , $\sin^2\theta_{23}$, θ_{13} and all systematic parameters are profiled over in the other hierarchy.

³For the inverted hierarchy plots, the oscillation point corresponding to the local minimum in the region with negative Δm_{32}^2 is used.




Figure 6.19: Fake data sensitivities at the 2020 3-flavour best fit for the 2020 and new analyses. $\Delta\chi^2$ is calculated as different functions of the freely fitted oscillation parameters in the normal and inverted hierarchies. The borders of the 1-, 2- and 3- σ allowed Gaussian regions of the new analysis are shown in shades of dark blue (red) through to light blue (red) respectively for the normal (inverted) neutrino mass hierarchy. The corresponding borders of the allowed regions for the 2020 analysis are shown in lightening shades of grey.

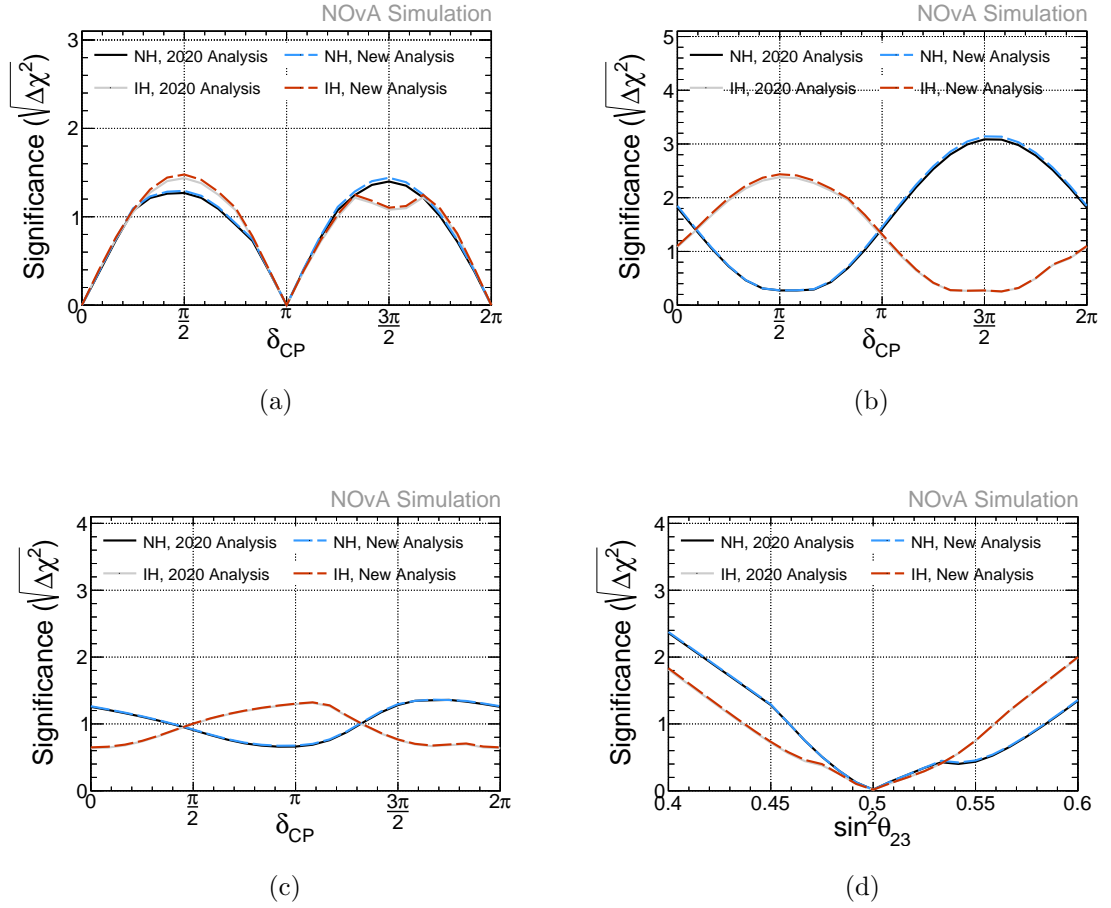


Figure 6.20: Sensitivity to: non-zero CP violation (**top left**), determine the neutrino mass hierarchy (**top right**) and reject the wrong octant of θ_{23} (**bottom**) as a function δ_{CP} (**left**) and $\sin^2 \theta_{23}$ (**right**). Results for the new and old analyses are shown by the dashed and solid curves respectively. Neutrino mass hierarchy information is given in the legend.

- Figure 6.20c: $i \in [0, 2\pi)$, $\Delta\chi_i^2 = \left(\chi_{\text{Wrong octant}}^2 - \chi_{\text{Right octant}}^2 \right)_{\delta=i}$. The significance with which the incorrect octant of θ_{23} can be rejected as a function of true δ_{CP} . δ_{CP} , Δm_{32}^2 , θ_{13} , the mass hierarchy and all systematic parameters are profiled over in the whole parameter space. $\sin^2 \theta_{23}$ is profiled within the limits of the respective wrong and right octants. This is so that the minima in each octant can be compared.
- Figure 6.20d: The significance with which the incorrect octant of θ_{23} can be rejected as a function of true $\sin^2 \theta_{23}$.

The new analysis shows rejection power greater than or equal to the old analysis in all of these metrics.

The Effect of Systematic Uncertainties on the Oscillation Parameters

A final consideration is how the sensitivity to the oscillation parameters is affected by systematic uncertainty. Figure 6.21 shows the estimated contribution of each category of systematic uncertainty to the total uncertainty, in each of the freely fitted oscillation parameters. A comparison to the statistical uncertainty and to the 2020 3-flavour analysis can also be seen. For a given oscillation parameter, a 1D function of $\Delta\chi^2$ (analogous to Figures 7.12, 7.13 and 7.14) is constructed assuming no systematic uncertainty. This is the so called, ‘statistics only’ fit. Next, assuming contributions from each category of systematic uncertainty individually, a set of analogous functions is constructed. These are the ‘systematic’ fits. Each systematic fit then has the statistics only fit subtracted from it in quadrature. The value of this difference at $\pm 1\text{-}\sigma$ is taken as the estimated systematic contribution.

With the exception of Δm_{32}^2 , it is clear that both the analysis presented in the (new) thesis and the (old) 2020 analysis are statistically limited in their measurements of the freely fitted oscillation parameters. However, there is a slight reduction in the statistical uncertainty of the most statistically limited parameter δ_{CP} in the new version versus the old. On the whole, both the shape (asymmetry of the uncertainty around zero) of the susceptibility and relative importance of each systematic group remains very similar between the two analyses. The total systematic contribution to the total uncertainty is also estimated to be almost identical, with the exception of Δm_{32}^2 . Here, there is a slight reduction in uncertainty in the new analysis compared to the old, coming mostly from reduced susceptibility to detector calibration and response. This may be due to the additional PID bins providing an extra constraint for the systematic parameters governing these two systematic groups.

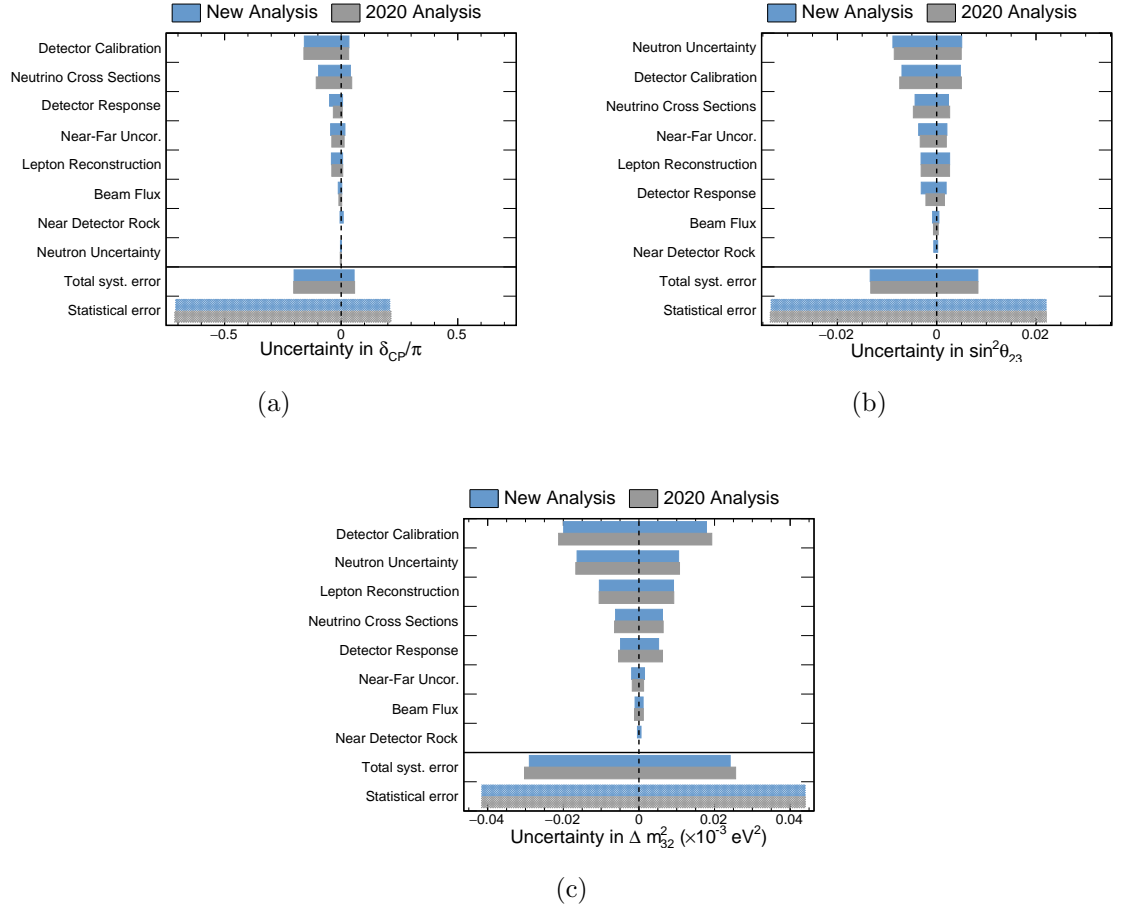


Figure 6.21: The simulated contribution of each category of systematic uncertainty to the total uncertainty on δ_{CP} (**top left**), $\sin^2 \theta_{23}$ (**top right**) and Δm_{32}^2 (**bottom**) at the 2020 analysis best fit. The new analysis (blue) can be compared to the 2020 3-flavour analysis (grey). An estimate of statistical uncertainty is shown by the hatched bars at the bottom of the Figure. An overview of the systematic groups can be found in Section 4.6.

Chapter 7

Results

This chapter presents the results of an improved analysis of NOvA’s 2020 3-flavour dataset [1] using the analysis structure outlined in Chapter 4 and the improvements outlined in Chapter 6. The dataset corresponds to a 14 ktonne equivalent exposure of 13.60×10^{20} and 12.50×10^{20} protons on target at the FD, in neutrino and antineutrino beam modes respectively. A measurement of the neutrino oscillation parameters δ_{CP} , $\sin^2 \theta_{23}$ and Δm_{32}^2 is performed, including the level at which the data disfavours one of the mass hierarchies and octants of θ_{23} .

Firstly, Section 7.1 discusses the results of the decomposition procedure for the electron neutrino-like sample at the ND, used to estimate the beam backgrounds in the appearance signal at the FD. Plots comparing data and simulation, before and after decomposition, and in several variables are also shown. Secondly, Section 7.2 presents the results of the 3-flavour model fit to the FD data. Candidate event counts, oscillation parameter best fit values, $1\text{-}\sigma$ allowed regions and several 1D and 2D contours in $\Delta\chi^2$ space, are all shown. There are also comparisons of FD data and simulation in several reconstructed variables. Finally, a close examination of the measurement’s systematic uncertainties is given in Section 7.3.

7.1 Studies at the ND

7.1.1 Decomposition

Tables 7.1 and 7.2, and, Figures 6.7 and 6.8 illustrate the final results of the electron neutrino-like sample decomposition at the ND, introduced in Section 6.2. The Figures show the distributions of reconstructed neutrino energy for events passing the electron

neutrino selection at the ND, before and after decomposition, in each selection bin (a - g). The tables give a detailed breakdown of the correction factors obtained for each interaction type in each selection bin. Overall, the nominal number of MC events are increased by 19.1% and 16.7% for FHC and RHC respectively. As explained in Section 4.5.1, in RHC the total discrepancy between data and simulation in each reconstructed energy bin is split equally between the beam components. In FHC, several other data driven techniques are used. Across all selection bins in FHC, the beam ν_e -CC component is increased by 1.2%. The NC and ν_μ -CC components are increased by 11.7% and 48.8% respectively. These corrections are similar to those obtained in the 2020 analysis (see Section 4.5.1) [102]. It is worth noting, however, that with the new analysis slightly fewer events are attributed to NC and slightly more attributed to ν_μ -CC. On closer inspection of Table 7.1 it can be seen that across individual selection bins, the corrections to the ν_μ -CC component differ by a reasonable amount. Given the different areas of phase space that each of the selection bins carve out, some variation is expected. Not all of the selection bins, for example, contain events within the same energy range. Bins (b), (c) and (d) do not have the ‘high energy’ tail of the others. Some of the variation may also be attributed to the decomposition procedure not performing optimally in certain reconstructed energy bins. The primary reason for this is the presence of poorly modelled rock events, described in detail in Section 6.2. The average corrections in the ‘worst affected’ selection bins, (d) and (f), are completely driven by the decomposition results in their lowest energy bins. The decomposition performs poorly here since the Michel spectrum, drawn up per reconstructed energy bin, is skewed by the presence of repeated rock events. The issues in these bins are covered by the rock systematic. Figure 6.16b shows large uncertainties on the corresponding FD bins (bins centred at 9.5, 12.5, 15.5 and 39.5).

Tables 7.3 and 7.4 show a breakdown of predicted events in simulation at the FD in FHC, with and without extrapolation at the 3-flavour 2020 best fit (Section 2.5.5). This illustrates to what order the decomposition affects the estimate of beam background and total prediction. Note that the $\sim 12\%$ change in NC at the ND translates to an almost identical change in NC at the FD. The $\sim 48\%$ change ν_μ -CC, however, translates to a $\sim 37\%$ change at the FD. This comes about due to the presence of oscillations causing the muon neutrinos to disappear. The corresponding FD spectrum with extrapolation is shown in Figure 7.3. Recall that each selection bin at the ND extrapolates to both the low and the high cosmic rejection equivalent selection bins at the FD. Tables 7.5 and 7.6,

	ND selection bin, FHC							
	(a)	(b)	(c)	(d)	(e)	(f)	(g)	All
Total data	22029	9746	759	4939	1141	8718	11553	58885
Total MC	17321	7953	652	3930	988	7355	11253	49452
Total corrected	22029	9746	759	4939	1141	8718	11553	58885
Total correction (%)	27.2	22.5	16.3	25.7	15.5	18.5	2.67	19.1
NC MC	8930	4000	322	1747	441	1948	517	17905
NC corrected	10567	4084	402	1707	630	1894	717	20002
NC correction (%)	18.3	2.10	25.1	-2.28	42.9	-2.77	38.8	11.7
ν_μ -CC+ $\bar{\nu}_\mu$ -CC MC	7426	3219	256	1397	296	1485	535	14616
ν_μ -CC+ $\bar{\nu}_\mu$ -CC corrected	10461	4917	280	2428	250	2695	713	21745
ν_μ -CC+ $\bar{\nu}_\mu$ -CC correction (%)	40.9	52.7	9.36	73.8	-15.8	81.5	33.0	48.8
ν_e -CC+ $\bar{\nu}_e$ -CC MC	964	734	74.5	786	250	3922	10200	16931
ν_e -CC+ $\bar{\nu}_e$ -CC corrected	1000	745	76.6	804	261	4128	10123	17138
ν_e -CC+ $\bar{\nu}_e$ -CC correction (%)	3.74	1.49	2.69	2.26	4.31	5.26	-0.75	1.22

Table 7.1: Decomposition results for the electron-like ND sample in FHC. All numbers are calculated using only the bins which are directly extrapolated to the FD and, therefore, influence the fit for oscillations. ‘Corrected’ refers to the particular quantity in simulation after decomposition weights have been applied. The percentage correction is calculated as the difference in the corrected and nominal MC, divided by the nominal MC. The corresponding reconstructed neutrino energy spectra at the ND are shown in Figure 7.1.

and, Figure 7.4 show the same information but in RHC. Note that the predicted number of selected, right sign, true electron neutrinos in this analysis and the 2020 analysis can be compared [104]. It can be seen that this analysis is 59.4% efficient in the core samples versus 56.0% efficient in the 2020 analysis at the 2020 best fit.

	ND selection bin, RHC							
	(a)	(b)	(c)	(d)	(e)	(f)	(g)	All
Total data	5297	3751	256	1567	96	2745	7377	21089
Total MC	4312	3249	186	1342	79	2201	6703	18073
Total corrected	5297	3751	256	1567	96	2745	7377	21089
Total correction (%)	22.8	15.5	37.4	16.7	21.9	24.715	10.0	16.7
NC MC	3015	1643	82	619	17	409	299	6085
NC corrected	3704	1915	114	719	21	529	336	7336
NC correction (%)	22.8	16.5	38.4	16.2	19.4	29.2	12.3	20.6
ν_μ -CC+ $\bar{\nu}_\mu$ -CC MC	924	675	37	219	6	178	184	2225
ν_μ -CC+ $\bar{\nu}_\mu$ -CC corrected	1152	748	48	252	8	227	208	2642
ν_μ -CC+ $\bar{\nu}_\mu$ -CC correction (%)	24.6	10.7	30.1	15.1	20.6	27.2	12.7	18.8
ν_e -CC+ $\bar{\nu}_e$ -CC MC	373	930	67	504	55	1613	6220	9763
ν_e -CC+ $\bar{\nu}_e$ -CC corrected	442	1088	94	596	67.7	1989	6833	11111
ν_e -CC+ $\bar{\nu}_e$ -CC correction (%)	18.4	17.0	40.3	18.2	22.9	23.3	9.86	13.8

Table 7.2: Decomposition results for the electron-like ND sample in RHC. All numbers are calculated using only the bins which are directly extrapolated to the FD and, therefore, influence the fit for oscillations. ‘Corrected’ refers to the particular quantity in simulation after decomposition weights have been applied. The percentage correction is calculated as the difference in the corrected and nominal MC, divided by the nominal MC. The corresponding reconstructed neutrino energy spectra at the ND are shown in Figure 7.2.




Figure 7.1: Distributions of reconstructed neutrino energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown. Any bin that is not extrapolated to the FD, that is to say, not included in the fit for oscillations, is lightly shaded.




Figure 7.2: Distributions of reconstructed neutrino energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown. Any bin that is not extrapolated to the FD, that is to say, not included in the fit for oscillations, is lightly shaded.

Event Type	Core	Peripheral	Total
Appearance ν_e	56.5	8.38	64.9
Wrong sign appearance	0.894	0.187	1.08
Beam $\nu_e + \bar{\nu}_e$	10.5	3.42	13.9
Total $\nu_\mu + \bar{\nu}_\mu$	3.11	0.267	3.38
Total $\nu_\tau + \bar{\nu}_\tau$	0.551	0.112	0.663
Neutral current	12.9	0.325	13.2
Total beam background	27.9	4.31	32.2
Cosmics	3.34	1.63	4.97
Total prediction	87.8	14.3	102

Table 7.3: The simulated breakdown of predicted events at the FD in FHC with extrapolation at the 2020 3-flavour best fit (see Section 2.5.5).

Event Type	Core	Peripheral	Total
Appearance ν_e	56.4	8.53	64.9
Wrong sign appearance	0.896	0.192	1.09
Beam $\nu_e + \bar{\nu}_e$	10.4	3.61	14.0
Total $\nu_\mu + \bar{\nu}_\mu$	2.21	0.256	2.46
Total $\nu_\tau + \bar{\nu}_\tau$	0.551	0.112	0.663
Neutral current	11.5	0.312	11.8
Total beam background	25.6	4.49	30.1
Cosmics	3.34	1.63	4.97
Total prediction	85.3	14.6	100

Table 7.4: The simulated breakdown of predicted events at the FD in FHC without extrapolation at the 2020 3-flavour best fit (see Section 2.5.5).




Figure 7.3: Distributions of reconstructed neutrino energy and reconstructed primary shower energy for the core and peripheral samples respectively at the FD in FHC with extrapolation. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.

Event Type	Core	Peripheral	Total
Appearance $\bar{\nu}_e$	17.0	2.89	19.8
Wrong sign appearance	1.92	0.404	2.32
Beam $\nu_e + \bar{\nu}_e$	5.52	2.10	7.61
Total $\nu_\mu + \bar{\nu}_\mu$	0.339	0.132	0.471
Total $\nu_\tau + \bar{\nu}_\tau$	0.300	0.074	0.374
Neutral current	3.65	0.152	3.80
Total beam background	11.7	2.86	14.6
Cosmics	0.498	0.973	1.47
Total prediction	29.2	6.72	35.9

Table 7.5: The simulated breakdown of predicted events at the FD in RHC with extrapolation at the 2020 3-flavour best fit .

Event Type	Core	Peripheral	Total
Appearance $\bar{\nu}_e$	17.6	3.03	20.6
Wrong sign appearance	1.98	0.419	2.39
Beam $\nu_e + \bar{\nu}_e$	4.86	1.99	6.85
Total $\nu_\mu + \bar{\nu}_\mu$	0.276	0.129	0.404
Total $\nu_\tau + \bar{\nu}_\tau$	0.300	0.074	0.374
Neutral current	2.97	0.144	3.12
Total beam background	10.4	2.76	13.1
Cosmics	0.498	0.973	1.47
Total prediction	28.5	6.75	35.2

Table 7.6: The simulated breakdown of predicted events at the FD in RHC without extrapolation at the 2020 3-flavour best fit (Section 2.5.5).




Figure 7.4: Distributions of reconstructed neutrino energy and reconstructed primary shower energy for the core and peripheral samples respectively at the FD in RHC with extrapolation. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. The purple histogram shows the total prediction, evaluated at the 2020 3-flavour best fit (Section 2.5.5). Different coloured stacked histograms show the various beam and cosmic background components of the total prediction. Appearance signal is shown in white.

7.1.2 Reconstructed Variables at the ND

Figures C.1 through C.18 (Appendix C) show distributions of various reconstructed quantities for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) and for both horn currents. In all plots, decomposition weights are applied as a function of reconstructed neutrino energy. The black ‘corrected’ ratio, therefore, is not expected to sit at exactly 1. Furthermore, events which fall outside of the reconstructed neutrino energy window applied at the FD are included in these spectra. As a result, the spectra are on the whole more susceptible to poorly modelled rock events. This can be seen, for example, in RHC selection bin (d) in transverse momentum. Despite this, all variables show improved agreement between data and simulation after decomposition.

7.2 Analysis of FD Data

7.2.1 Candidate Events & New Best Fit

The number of candidate ν_e -CC events in FHC data and candidate $\bar{\nu}_e$ -CC events in RHC data is found to be 102 and 34 respectively. The number of ν_μ -CC and $\bar{\nu}_\mu$ -CC candidates, of course, does not change from the analysis presented at the Neutrino 2020 conference (see Section 2.5.5). This leads to a best fit, when also minimising over all systematic parameters, in the normal neutrino mass hierarchy and lower octant of $\sin^2 \theta_{23}$. Specifically, the most favoured values of the oscillation parameters are found to be $\sin^2 \theta_{23} = 0.46$, $\delta_{CP} = 0.12\pi$ and $\Delta m_{32}^2 = 2.41 \times 10^{-3} \text{ eV}^2$. See Section 7.3 for a discussion of systematics. Throughout the rest of the document, this best fit will be referred to as the ‘thesis best fit’. Tables 7.7 and 7.8 show a breakdown of the predicted events in FD simulation at the thesis best fit in FHC and RHC respectively. The number of candidates observed in data are in general agreement with both of the predicted totals from simulation. Worthy of note is the peripheral sample in RHC, where 6.34 events are expected and only 2 are observed. However, the Poisson ‘up error’ on 2 events is approximately 2.2. This result, therefore, is still consistent at around $2\text{-}\sigma$.

Figures 7.5 and 7.6 show distributions of reconstructed neutrino energy and reconstructed primary shower energy for the core and peripheral samples respectively at the FD. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations.

Event Type	Core	Peripheral	Total
Appearance ν_e	51.7	7.87	59.6
Wrong sign appearance	0.822	0.194	1.02
Beam $\nu_e + \bar{\nu}_e$	10.7	3.75	14.5
Total $\nu_\mu + \bar{\nu}_\mu$	3.54	0.292	3.82
Total $\nu_\tau + \bar{\nu}_\tau$	0.538	0.113	0.651
Neutral current	14.3	0.387	14.7
Total beam background	29.9	4.74	34.6
Cosmics	3.34	1.63	4.97
Total prediction	84.9	14.2	99.2
Data	86	16	102

Table 7.7: The predicted number of events at the FD in FHC at the thesis best fit compared to data.

ν_e -CC and $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The $1\text{-}\sigma$ range of systematic uncertainty is represented by the purple band. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure 7.5: Distributions of reconstructed neutrino energy and reconstructed primary shower energy for the core and peripheral samples respectively at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The $1\text{-}\sigma$ range of systematic uncertainty is represented by the purple band. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

Event Type	Core	Peripheral	Total
Appearance $\bar{\nu}_e$	15.2	2.67	17.9
Wrong sign appearance	1.72	0.41	2.13
Beam $\nu_e + \bar{\nu}_e$	5.02	1.92	6.95
Total $\nu_\mu + \bar{\nu}_\mu$	0.352	0.132	0.484
Total $\nu_\tau + \bar{\nu}_\tau$	0.295	0.073	0.369
Neutral current	3.80	0.158	3.96
Total beam background	11.2	2.70	13.9
Cosmics	0.498	0.973	1.48
Total prediction	26.9	6.34	33.3
Data	32	2	34

Table 7.8: The predicted number of events at the FD in RHC at the thesis best fit compared to data.




Figure 7.6: Distributions of reconstructed neutrino energy and reconstructed primary shower energy for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The $1\text{-}\sigma$ range of systematic uncertainty is represented by the purple band. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

The so called ‘bi-event plot’, which shows the number of candidate appearance $\bar{\nu}_e$ -CC events versus ν_e -CC events, provides an intuitive way to compare old and new data and best fits. Comparing Figure 7.7 to its 2020 3-flavour counterpart Figure 2.16b, it can be seen that the data still sits in between the lower and upper octant ellipses but has shifted slightly closer to the lower one. The best fit moving from the upper to the lower octant reflects this change. Since this shift in the data is accompanied by only a small migration away from the central degenerate region, this implies the best fit value of δ_{CP} must also shift. It now sits in the top left, rather than bottom left, arc of the best fit ellipse.

7.2.2 Confidence Intervals

Figures 7.8 and 7.10 show contours of constant $\Delta\chi^2$ (Equation 4.20) in the $\sin^2\theta_{23}$ versus δ_{CP} and Δm_{32}^2 versus $\sin^2\theta_{23}$ spaces respectively in the normal hierarchy. Figures 7.9 and 7.11 show the same thing but in the inverted hierarchy. Each point on the underlying surface is obtained by fixing the values of the oscillation parameters which are labelled on the axes in the fit to data. All other parameters, including those describing systematics are then profiled over. The global minimum χ^2 is then subtracted from the minimum χ^2 at that point in oscillation space. As described in Section 4.7, the borders of the 1-, 2- and 3- σ allowed regions are drawn using a Gaussian estimate. A fair comparison of these contours with those from the 2020 analysis (Figure 2.16a and [1]) is not possible given that the ones from 2020 are Feldman-Cousins corrected¹ [139]. Broadly, however, it can be seen in Figure 7.8 that the range of $\sin^2\theta_{23}$ values allowed at 2- σ has shrunk in the region around $\delta_{CP} \sim 3\pi/2$. This is accompanied by an increase in the range of $\sin^2\theta_{23}$ values allowed at 1- σ around $\delta_{CP} \sim \pi/2$. Meanwhile in the inverse hierarchy, the 3- σ ‘islands’ observed in the 2020 analysis get much closer for high values of $\sin^2\theta_{23}$ around $\delta_{CP} \sim \pi/2$. Very little change is expected or observed in the Δm_{32}^2 versus $\sin^2\theta_{23}$. There does appear to be slightly more values of Δm_{32}^2 allowed at 1- σ around maximal mixing.

¹The Feldman-Cousins technique can be used to determine confidence limit boundaries for Poissonian processes. This is done by using pseudo-experiments to probe the range of possible log-likelihood ratios.




Figure 7.7: The total number of $\bar{\nu}_e$ -CC candidate events versus the total number of ν_e -CC candidate events in data and simulation. The black point shows FD data and the purple star is placed at the thesis best fit prediction. Ovals in the normal (blue) and inverted (red) neutrino mass hierarchies for the upper and low octants of $\sin^2 \theta_{23}$ are traced by varying δ_{CP} between 0 and 2π . The position of each oval is determined by fixing $\sin^2 \theta_{23}$ and Δm_{32}^2 at their ‘best fit values’ at each local minimum.




Figure 7.8: Allowed regions as a function of $\sin^2 \theta_{23}$ and δ_{CP} in the normal hierarchy. The borders of the 1-, 2- and 3- σ allowed Gaussian regions are shown in shades of dark blue through to light blue respectively. The thesis best fit is shown as a black star.




Figure 7.9: Allowed regions as a function of $\sin^2 \theta_{23}$ and δ_{CP} in the inverted hierarchy. The borders of the 1-, 2- and 3- σ allowed Gaussian regions are shown in shades of dark red through to light red respectively.




Figure 7.10: Allowed regions as a function of Δm_{32}^2 and $\sin^2 \theta_{23}$ in the normal hierarchy. The borders of the 1-, 2- and 3- σ allowed Gaussian regions are shown in shades of dark blue through to light blue respectively. The thesis best fit is shown as a black star.




Figure 7.11: Allowed regions as a function of Δm_{32}^2 and $\sin^2 \theta_{23}$ in the inverted hierarchy. The borders of the 1-, 2- and 3- σ allowed Gaussian regions are shown in shades of dark red through to light red respectively.

Figures 7.12, 7.13 and 7.14 show $\sqrt{\Delta\chi^2}$ as a function of δ_{CP} , $\sin^2\theta_{23}$ and Δm_{32}^2 respectively, for different choices of neutrino mass hierarchy and $\sin^2\theta_{23}$ octant. The plots are constructed in a way directly analogous to the 2D contours. Note the vertical axis shows the *square-root* of $\Delta\chi^2$. The plots can, therefore, be interpreted as showing the Gaussian significance with which a particular value of a given parameter is disfavoured. By looking at where particular curves pass $\sqrt{\Delta\chi^2} = 1$, $1\text{-}\sigma$ confidence intervals for each oscillation parameter can be extracted. These regions, summarised in Table 7.9, largely overlap with the $1\text{-}\sigma$ ranges obtained in the NOvA 2020 3-flavour analysis [140]. The point of maximal mixing ($\sin^2\theta_{23} = 0.5$) is allowed at $1\text{-}\sigma$.

It is apparent from any of the three 1D distributions that there is a set of oscillation parameters which are allowed at the $1\text{-}\sigma$ level, for all choices of hierarchy and octant. Figure 7.13 shows that for any fixed value of θ_{23} , there is always a choice of oscillation parameters (most likely driven by δ_{CP}) where the normal hierarchy is favoured over the inverted. Additionally, it can be seen that the preference for the lower versus the upper octant is mild. From Figure 7.12 it can be concluded that regions of phase space which would imply large asymmetry in the number of observed ν_e -CC versus $\bar{\nu}_e$ -CC (the ‘edges’ of the bi-event plot), are disfavoured. For example, $\delta_{CP} = \pi/2$ in the inverted hierarchy and $\delta_{CP} = 3\pi/2$ in the normal hierarchy are rejected at $> 3.8\sigma$ and $> 2.6\sigma$ respectively.




Figure 7.12: The Gaussian significance with which a particular value of δ_{CP} is disfavoured, for all choices of neutrino mass hierarchy and $\sin^2 \theta_{23}$ octant. The solid (dashed) blue (red) contour represents the upper (lower) octant in normal (inverted) hierarchy.




Figure 7.13: The Gaussian significance with which a particular value of $\sin^2\theta_{23}$ is disfavoured, for the normal (blue) and inverted (red) neutrino mass hierarchies.




Figure 7.14: The Gaussian significance with which a particular value of $|\Delta m^2_{32}|$ is disfavoured, for all choices of neutrino mass hierarchy and $\sin^2 \theta_{23}$ octant. The solid (dashed) blue (red) contour represents the upper (lower) octant in normal (inverted) hierarchy.

Best fit	
NH, $\sin^2 \theta_{23} = 0.46$, $\delta_{CP} = 0.12\pi$, $\Delta m_{32}^2 = 2.41 \times 10^{-3} \text{ eV}^2$	
Parameter	1- σ range
$\sin^2 \theta_{23}$ NH	$[0.43, 0.50] \cup [0.53, 0.59]$
$\sin^2 \theta_{23}$ IH	$[0.44, 0.49] \cup [0.53, 0.59]$
$\delta_{CP} (\pi)$ NHLO	$[0.00, 0.33] \cup [1.93, 2.00]$
$\delta_{CP} (\pi)$ NHUO	$[0.14, 0.94]$
$\delta_{CP} (\pi)$ IHLO	$[1.21, 1.53]$
$\delta_{CP} (\pi)$ IHUO	$[1.20, 1.75]$
$\Delta m_{32}^2 (\times 10^{-3} \text{ eV}^2)$ NHLO	$[2.35, 2.47]$
$\Delta m_{32}^2 (\times 10^{-3} \text{ eV}^2)$ NHUO	$[2.35, 2.47]$
$\Delta m_{32}^2 (\times 10^{-3} \text{ eV}^2)$ IHLO	$[-2.50, -2.42]$
$\Delta m_{32}^2 (\times 10^{-3} \text{ eV}^2)$ IHUO	$[-2.51, -2.41]$

Table 7.9: Summary of the 1- σ Gaussian allowed intervals for each freely fitted oscillation parameter. The global best fit is given in the second row.

7.2.3 Reconstructed Variables at the FD

Figures C.19 through C.44 (Appendix C) show the comparison of FD data against several reconstructed variables in simulation. With the exception of the RHC peripheral sample (already discussed in Section 7.2.1), there is good general agreement in all selection bin across all considered reconstructed variables. A single anomalous bin, however, in FHC selection bin (g) in the reconstructed vertex Z position spectra (Fig C.35) is identified. The probability of observing no events when 6 are predicted is approximately 1 in 400. It is found that there are 2 events just above the top bin edge. Additionally, none of the other selection bins show any anomalous behaviour in this particular range of vertex Z values. For these reasons and given the large number of plots shown in this section, it can be reasonably assumed to be a statistical fluctuation.

7.3 Systematic Uncertainties

Figure 7.15 shows the best fit value of each systematic parameter (pull) at the thesis best fit in units of σ . The pull can be seen as a measure of how much a prediction has to be shifted to improve the χ^2 between the data and the oscillation model due to given

systematic uncertainty [23]. The largest pulls arise from uncertainties associated with detector calibration and the ND light model. With the exception of the calibration shape parameter at $1.03\text{-}\sigma$, all pulls are less than $0.8\text{-}\sigma$.

Figures 7.16 through 7.18 show the estimated contribution of each broad category of systematic uncertainty to the total uncertainty, in each of the freely fitted oscillation parameters. A comparison to the statistical uncertainty and to the 2020 3-flavour analysis can also be seen. These plots are made using real ND data and fake FD data, produced at the thesis best fit. For a given oscillation parameter, a 1D function of $\Delta\chi^2$ (analogous to Figures 7.12, 7.13 and 7.14) is constructed assuming no systematic uncertainty. This is the so called, ‘statistics only’ fit. Next, assuming contributions from each category of systematic uncertainty individually, a set of analogous functions is constructed. These are the ‘systematic’ fits. Each systematic fit then has the statistics only fit subtracted from it in quadrature. The value of this difference at $\pm 1\text{-}\sigma$ is taken as the estimated systematic contribution. Figures 7.19 through 7.21 are constructed using the same methodology but each of the systematic parameters are presented individually rather than being grouped together in broad categories.

With the exception of Δm_{32}^2 , it is clear that both the analysis presented in the (new) thesis and the (old) 2020 analysis are statistically limited in their measurements of the freely fitted oscillation parameters. As expected, there is a slight reduction in the statistical uncertainty of the most statistically limited parameter δ_{CP} in the new analysis versus the old. This most likely a result of the extra events added to the analysis and the better separation of signal and background that the new analysis binning provides. It can also be seen that overall sensitivity to systematic uncertainty is very similar between the two versions. On the whole, each systematic category carries the same level of importance in both. In the case of δ_{CP} and Δm_{32}^2 in the new analysis, a small reduction in uncertainty due to neutrino cross-sections is observed. This most likely represents a mild underestimate of the true variance attributable to the ‘small’ cross-section uncertainties in the new analysis. Recall that the effect of these uncertainties is summed and treated as a single effect in the fit (see Section 4.6.2). In Chapter 8, a proposal to improve this treatment is discussed.




Figure 7.15: The best fit value of each systematic parameter (pull) at the thesis best fit in units of σ . An overview of the individual parameters listed on the x-axis can be found in Section 4.6.




Figure 7.16: The simulated contribution of each category of systematic uncertainty to the total uncertainty on δ_{CP} at the thesis best fit. The new analysis (blue) can be compared to the 2020 3-flavour analysis (grey). An estimate of statistical uncertainty is shown by the hatched bars at the bottom of the Figure. A breakdown of the systematic groups can be found in Section 4.6.




Figure 7.17: The simulated contribution of each category of systematic uncertainty to the total uncertainty on $\sin^2 \theta_{23}$ at the thesis best fit. The new analysis (blue) can be compared to the 2020 3-flavour analysis (grey). An estimate of statistical uncertainty is shown by the hatched bars at the bottom of the Figure. A breakdown of the systematic groups can be found in Section 4.6.




Figure 7.18: The simulated contribution of each category of systematic uncertainty to the total uncertainty on Δm_{32}^2 at the thesis best fit. The new analysis (blue) can be compared to the 2020 3-flavour analysis (grey). An estimate of statistical uncertainty is shown by the hatched bars at the bottom of the Figure. A breakdown of the systematic groups can be found in Section 4.6.




Figure 7.19: The simulated contribution of each systematic uncertainty to the total uncertainty on δ_{CP} at the thesis best fit. A breakdown (and shortname key) of the systematic parameters can be found in Section 4.6.



Figure 7.20: The simulated contribution of each systematic uncertainty to the total uncertainty on $\sin^2 \theta_{23}$ at the thesis best fit. A breakdown (and shortname key) of the systematic parameters can be found in Section 4.6.




Figure 7.21: The simulated contribution of each systematic uncertainty to the total uncertainty on Δm_{32}^2 at the thesis best fit. A breakdown (and shortname key) of the systematic parameters can be found in Section 4.6.

Chapter 8

Conclusion

This thesis has presented an improved analysis of NuMI data collected by the NOvA far detector between 2014 and 2020. The data corresponds to a 14 ktonne equivalent exposure of 13.60×10^{20} and 12.50×10^{20} POT, in neutrino and antineutrino beam modes respectively. Several methods for increasing NOvA's sensitivity to the 3-flavour oscillation parameters, particularly δ_{CP} and the mass hierarchy, were explored. It was found that the 2020 3-flavour analysis could be improved in three ways to achieve a gain in sensitivity equivalent to collecting 5 to 6% more data. First, in addition to binning the data as function of energy and a single particle identification score, two axes dependent on other particle identifiers were introduced. Second, more signal events are made available to the fit for oscillations by including regions of parameter space with a greater amount of beam and cosmic induced background. A bin edge optimisation procedure was developed to simultaneously establish the position of bin edges in every variable which maximised the sensitivity metric. Third, coarse binning in a basic energy variable targeting the most energetic shower in an event was introduced to the peripheral events sample. These three changes were combined and a fit for 3-flavour oscillations to the neutrino and antineutrino data was performed. A best fit at $\sin^2 \theta_{23} = 0.46$, $\delta_{CP} = 0.12\pi$ and $\Delta m_{32}^2 = 2.41 \times 10^{-3} \text{ eV}^2$ was obtained. Our data do not prefer a particular neutrino mass hierarchy or octant, but the points $\delta_{CP} = 3\pi/2$ in the inverted mass hierarchy and $\delta_{CP} = \pi/2$ in the normal mass hierarchy are excluded at 3.4σ and 2.4σ respectively. The $1\text{-}\sigma$ confidence ranges for the best fit combination of hierarchy and octant (normal hierarchy, lower octant) can be

summarised as follows:

$$\begin{aligned}\sin^2 \theta_{23} &\in [0.43, 0.50] \cup [0.53, 0.59] \\ \delta_{CP} (\pi) &\in [0.00, 0.33] \cup [1.93, 2.00] \\ \Delta m_{32}^2 (\times 10^{-3} \text{ eV}^2) &\in [2.35, 2.47].\end{aligned}$$

Despite favouring the lower octant over the upper octant, these $1\text{-}\sigma$ ranges overlap significantly with those of the NOvA 2020 3-flavour analysis [140]. Notably, however, maximal mixing ($\sin^2 \theta_{23} = 0.5$) is permitted in the upgraded analysis at $1\text{-}\sigma$.

If I were to develop the analysis further, I would return to the results of the bin edge optimisation procedure and choose the combination of edges to take forward for further analysis more carefully. Instead of my approach in this thesis which was to simply pick the combination which maximised a sensitivity metric, I would also take into account how ‘off-diagonal’ the two sets of bin edges corresponding to electron CVN score were with respect to each other. Areas of low statistics at the ND arose from selecting bin edges corresponding to narrow regions of phase space. This in turn necessitated the introduction of a systematic to cover against uncertainty in the true rate of “rock” events originating outside of the ND. Secondly, given additional time and resources, instead of combining the effect of all cross-section systematics deemed small in the 2020 3-flavour analysis into a single systematic uncertainty, I would establish a new list of the most important cross-section uncertainties given the updated selection. I would then use a principle component analysis to obtain a set of uncorrelated systematic parameters to include in the fit, covering some high percentage of the total variance. This would increase the robustness of the new analysis method.

In addition to sensitivity improvements, a data-driven study was undertaken for this thesis to quantify an important and potential systematic effect arising in the ND. Due to the high intensity of the beam at the location of the ND close to the source, multiple neutrino interactions occur in the detector and in the surrounding rock during every spill. This pile-up activity is not present in the FD and must be well modelled in ND simulation to avoid biasing the measurements. Single MC neutrino interactions were overlaid into ND spills in both data and simulation and re-reconstructed. Selection efficiencies in data and MC for these single $\nu_\mu\text{-CC}$ and $\nu_e\text{-CC}$ events (in both horn currents) were determined and studied as a function of spill intensity and true neutrino energy. Differences were found to be small, showing that pile-up, noise and other context effects are sufficiently well modelled

in simulation. A further study of neutrino energy resolution reached the same conclusions.

Looking toward the future, NOvA plans to roughly double or triple its current FD dataset in both horn currents by 2025, allowing the experiment to reach $3\text{-}\sigma$ hierarchy sensitivity for 30 - 50% of δ_{CP} values [141]. As more data is gathered, the 3-flavour analysis becomes more affected by systematic, rather than statistical, uncertainty. As can be seen from Figure 7.18, this is particularly true of the variable Δm_{32}^2 , which is driven by the disappearance measurement and is sensitive to the detector energy scale. Significant effort is being made to reduce the level of uncertainty associated with this systematic. Analysis of data from a dedicated test beam facility at Fermilab is currently on-going [142]. It is unlikely, however, that NOvA will truly become systematically limited in its measurement of δ_{CP} or θ_{23} . Recent tension between the best fit values of δ_{CP} measured by T2K [48] and by NOvA [1] has brought renewed emphasis as to whether 3-flavour oscillations fully describe the data or whether additional, new physics could be involved. The value of δ_{CP} measured in this thesis maintains the tension. A joint fit of NOvA and T2K data is currently under way that will quantify the tension. Future long-baseline experiments, with more powerful beams and larger detectors, will be able to make far more sensitive measurements. DUNE, for example, has the ability to exclude the CP conserved case at more than $5\text{-}\sigma$ for 50% of δ_{CP} 's range after 10 years of running [51]. These future experiments will start taking data towards 2030. In the meantime, the upgraded NOvA analysis with more data will continue to provide world leading insight into the physics of neutrinos.

Acronyms

1D One Dimensional.

2D Two Dimensional.

3D Three Dimensional.

4D Four Dimensional.

ADC Analogue to Digital Converter.

APD Avalanche Photodiode.

ASIC Application Specific Integrated Circuit.

BDT Boosted Decision Tree.

BEN Beam Electron Neutrino.

CAF Common Analysis Format.

CAFAna Common Analysis Format Analysis.

CC Charged Current.

CCQE Charged Current Quasi-elastic.

CMC Comprehensive Model Configuration.

CNN Convolutional Neural Network.

CP Charge-Parity.

CRY Cosmic Ray Shower Library.

CVN Convolutional Visual Network.

CVNe electron neutrino CVN.

DCM Data Concentrator Module.

DCS Dual Correlated Sampling.

DIS Deep Inelastic Scattering.

DONUT Direct Observation of the Nu Tau.

DQ Data Quality.

DUNE Deep Underground Neutrino Experiment.

FD Far Detector.

FEB Front End Board.

Fermilab Fermi National Accelerator Laboratory.

FHC Forward Horn Current.

FPGA Field Programmable Gate Array.

FSI Final State Interactions.

GPS Global Positioning System.

KamLAND the Kamioka Liquid-scintillator Antineutrino Detector.

LINAC Linear Accelerator.

MC Monte Carlo.

MEC Meson Exchange Current.

MI Main Injector.

MID Michel electron Identifier.

MINOS Main Injector Neutrino Oscillation Search.

MSW Mikheyev-Smirnov-Wolfenstein.

MTDU Master Timing Distribution Unit.

NC Neutral Current.

ND Near Detector.

NDOS Near Detector on the Surface.

Neutrino 2020 The XXIX International Conference on Neutrino Physics and Astrophysics.

NOvA NuMI Off-Axis Neutrino Experiment.

NuMI Neutrinos from the Main Injector.

OMSD One Mass Scale Dominance.

PCA Principle Component Analysis.

PDG Particle Data Group.

PE photo-electrons.

PID Particle Identification.

PMNS Pontecorvo-Maki-Nakagawa-Sakata.

PMTs Photo-multiplier Tubes.

POT Protons on Target.

PPFX Package to Predict the Flux.

PPO 2,5-Diphenyloxazole.

PVC Polyvinyl Chloride.

QE Quasi-elastic.

QED Quantum Electrodynamics.

ReMId Reconstructed Muon Identifier.

RENO Reactor Experiment for Neutrino Oscillation.

RES Baryon Resonance Production.

RHC Reverse Horn Current.

RMS Root Mean Square.

RPA Random Phase Approximation.

SM Standard Model of Particle Physics.

SNO Sudbury Neutrino Observatory.

T2K Tokai to Kamioka.

TDR Technical Design Report.

Bibliography

- [1] A. Himmel *et al.*, “New oscillation results from the nova experiment.” The XXIX International Conference on Neutrino Physics and Astrophysics. indico.fnal.gov/event/43209, 2020.
- [2] J. Chadwick, “The intensity distribution in the magnetic spectrum of beta particles from radium (b + c),” *Verh. Phys. Gesell.*, vol. 16, pp. 383–391, 1914.
- [3] W. Pauli, “Open letter to the gauverein meeting in tubingen.”
- [4] E. Fermi, “An attempt of a theory of beta radiation. 1.,” *Z. Phys.*, vol. 88, pp. 161–177, 1934.
- [5] F. Reines and C. Cowan, “Detection of the free neutrino,” *Phys. Rev.*, vol. 92, pp. 830–831, 1953.
- [6] C. Cowan, F. Reines, F. Harrison, H. Kruse, and A. McGuire, “Detection of the free neutrino: A confirmation,” *Science*, vol. 124, pp. 103–104, 1956.
- [7] G. Danby, J. Gaillard, K. A. Goulianos, L. Lederman, N. B. Mistry, M. Schwartz, and J. Steinberger, “Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos,” *Phys. Rev. Lett.*, vol. 9, pp. 36–44, 1962.
- [8] K. Kodama *et al.*, “Observation of tau neutrino interactions,” *Phys. Lett. B*, vol. 504, pp. 218–224, 2001.
- [9] K. Kodama *et al.*, “Final tau-neutrino results from the donut experiment,” *Phys. Rev. D*, vol. 78, p. 052002, 2008.
- [10] B. Pontecorvo, “Mesonium and anti-mesonium,” *Sov. Phys. JETP*, vol. 6, p. 429, 1957.
- [11] S. S. Z. Maki, M. Nakagawa, “Remarks on the unified model of elementary particles,” *Prog. Theor. Phys.*, vol. 28, pp. 870–880, 1962.

- [12] B. Pontecorvo, “Neutrino experiments and the problem of conservation of leptonic charge,” *Sov. Phys. JETP*, vol. 26, pp. 984–988, 1968.
- [13] J. Davis, Raymond, D. S. Harmer, and K. C. Hoffman, “Search for neutrinos from the sun,” *Phys. Rev. Lett.*, vol. 20, pp. 1205–1209, 1968.
- [14] J. N. Bahcall, N. A. Bahcall, and G. Shaviv, “Present status of the theoretical predictions for the cl-36 solar neutrino experiment,” *Phys. Rev. Lett.*, vol. 20, pp. 1209–1212, 1968.
- [15] K. Hirata *et al.*, “Observation of b-8 solar neutrinos in the kamiokande-ii detector,” *Phys. Rev. Lett.*, vol. 63, p. 16, 1989.
- [16] A. Abazov *et al.*, “Search for neutrinos from sun using the reaction ga-71 (electron-neutrino e-) ge-71,” *Phys. Rev. Lett.*, vol. 67, pp. 3332–3335, 1991.
- [17] P. Anselmann *et al.*, “Solar neutrinos observed by gallex at gran sasso,” *Phys. Lett. B*, vol. 285, pp. 376–389, 1992.
- [18] K. Hirata *et al.*, “Observation of a small atmospheric muon-neutrino / electron-neutrino ratio in kamiokande,” *Phys. Lett. B*, vol. 280, pp. 146–152, 1992.
- [19] Y. Fukuda *et al.*, “Evidence for oscillation of atmospheric neutrinos,” *Phys. Rev. Lett.*, vol. 81, pp. 1562–1567, 1998.
- [20] Q. Ahmad *et al.*, “Direct evidence for neutrino flavor transformation from neutral current interactions in the sudbury neutrino observatory,” *Phys. Rev. Lett.*, vol. 89, p. 011301, 2002.
- [21] T. N. Prize, “The nobel prize in physics 2015.” Link. Accessed: 2020-Nov-02.
- [22] M. Thomson, *Modern particle physics*. New York: Cambridge University Press, 2013. pages 350.
- [23] D. P. Mendez Mendez, *Comparison of oscillation parameters measured from ν_μ and $\bar{\nu}_\mu$ disappearance in the NOvA experiment*. PhD thesis, U. of Sussex, 10 2019.
- [24] C. Andreopoulos *et al.*, “The genie neutrino monte carlo generator,” *Nucl. Instrum. Meth. A*, vol. 614, pp. 87–104, 2010.
- [25] P. Lasorak, *A search for neutrino-induced single photons and measurement of oscillation analysis systematic errors with electron and anti-electron neutrino selections*,

- using the off-axis near detector of the Tokai to Kamioka experiment.* PhD thesis, Queen Mary, U. of London, 2018.
- [26] M. Tanabashi *et al.*, “Review of particle physics,” *Phys. Rev. D*, vol. 98, p. 030001, Aug 2018. doi:10.1103/PhysRevD.98.030001. pages 26, 41, 47, 253.
- [27] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 4 2007.
- [28] P. Hernandez, “Neutrino physics,” in *8th CERN–Latin-American School of High-Energy Physics*, pp. 85–142, 2016.
- [29] S. Mikheyev and A. Smirnov, “Resonance amplification of oscillations in matter and spectroscopy of solar neutrinos,” *Sov. J. Nucl. Phys.*, vol. 42, pp. 913–917, 1985.
- [30] L. Wolfenstein, “Neutrino oscillations in matter,” *Phys. Rev. D*, vol. 17, pp. 2369–2374, 1978.
- [31] E. D. Niner, *Observation of Electron Neutrino Appearance in the NuMI Beam with the NOvA Experiment*. PhD thesis, Indiana U., 2015.
- [32] H. Nunokawa, S. J. Parke, and J. W. Valle, “Cp violation and neutrino oscillations,” *Prog. Part. Nucl. Phys.*, vol. 60, pp. 338–402, 2008.
- [33] E. Catano-Mur, “Dpf2017 - physics reach of electron neutrino measurements,” *Internal Document NOvA DocDB 21764*, August 2017.
- [34] K. Abe *et al.*, “Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam,” *Phys. Rev. Lett.*, vol. 107, p. 041801, 2011.
- [35] B. Abi *et al.*, “Deep underground neutrino experiment (dune), far detector technical design report, volume iv: Far detector single-phase technology,” *JINST*, vol. 15, no. 08, p. T08010, 2020.
- [36] I. Esteban, M. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, “The fate of hints: updated global analysis of three-flavor neutrino oscillations,” *JHEP*, vol. 09, p. 178, 2020.
- [37] A. Gando *et al.*, “Reactor on-off antineutrino measurement with kamland,” *Phys. Rev. D*, vol. 88, no. 3, p. 033001, 2013.

- [38] Y. Nakajima *et al.*, “Recent results and future prospects from super-kamiokande.” The XXIX International Conference on Neutrino Physics and Astrophysics. indico.fnal.gov/event/43209, 2020.
- [39] F. An *et al.*, “New measurement of antineutrino oscillation with the full detector configuration at daya bay,” *Phys. Rev. Lett.*, vol. 115, no. 11, p. 111802, 2015.
- [40] Y. Abe *et al.*, “Indication of reactor $\bar{\nu}_e$ disappearance in the double chooz experiment,” *Phys. Rev. Lett.*, vol. 108, p. 131801, 2012.
- [41] H. de Kerret *et al.*, “Double chooz θ_{13} measurement via total neutron capture detection,” *Nature Phys.*, vol. 16, no. 5, pp. 558–564, 2020.
- [42] F. An *et al.*, “Observation of electron-antineutrino disappearance at daya bay,” *Phys. Rev. Lett.*, vol. 108, p. 171803, 2012.
- [43] P. Adamson *et al.*, “Measurement of the neutrino mass splitting and flavor mixing by minos,” *Phys. Rev. Lett.*, vol. 106, p. 181801, 2011.
- [44] J. Ahn *et al.*, “Observation of reactor electron antineutrino disappearance in the reno experiment,” *Phys. Rev. Lett.*, vol. 108, p. 191802, 2012.
- [45] D. Adey *et al.*, “Measurement of the electron antineutrino oscillation with 1958 days of operation at daya bay,” *Phys. Rev. Lett.*, vol. 121, no. 24, p. 241805, 2018.
- [46] M. Aartsen *et al.*, “Measurement of atmospheric tau neutrino appearance with icecube deepcore,” *Phys. Rev. D*, vol. 99, no. 3, p. 032007, 2019.
- [47] P. Adamson *et al.*, “Precision constraints for three-flavor neutrino oscillations from the full minos+ and minos dataset,” *Phys. Rev. Lett.*, vol. 125, no. 13, p. 131802, 2020.
- [48] P. Dunn *et al.*, “Latest neutrino oscillation results from T2K.” The XXIX International Conference on Neutrino Physics and Astrophysics. indico.fnal.gov/event/43209, 2020.
- [49] K. Abe *et al.*, “Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations,” *Nature*, vol. 580, no. 7803, pp. 339–344, 2020. [Erratum: *Nature* 583, E16 (2020)].
- [50] M. Ross-Lonergan, “Atmospheric parameters: Friends.” Link. Accessed: 2020-Nov-10.

- [51] B. Abi *et al.*, “Long-baseline neutrino oscillation physics potential of the dune experiment,” *Eur. Phys. J. C*, vol. 80, no. 10, p. 978, 2020.
- [52] A. Back, “Blessing package for all-time pot exposure for ana2020,” *Internal Document NOvA DocDB 46337*, June 2020.
- [53] D. Ayres *et al.*, “The NOvA technical design report,” *FERMILAB-DESIGN-2007-01*, October 2007. doi:10.2172/935497. pages 2-2, 2-3, 5-8, 5-12, 10-7, 11-4,.
- [54] P. Adamson *et al.*, “The NuMI neutrino beam,” *Nucl. Instrum. Meth. A*, vol. 806, pp. 279–306, 2016. doi:10.1016/j.nima.2015.08.063.
- [55] Fermilab Creative Services, “FCS still photos.” <https://mod.fnal.gov>. Accessed: 2020-Sep-02.
- [56] M. D. Tutto, “Blessing package - beam simulation plots,” *Internal Document NOvA DocDB 13524*, June 2016.
- [57] M. D. Tutto, “NOvA near and far detector neutrino energy spectra for several off-axis angles,” *Internal Document NOvA DocDB 13455*, June 2015.
- [58] P. Adamson *et al.*, “Measurement of neutrino and antineutrino oscillations using beam and atmospheric data in MINOS,” *Phys. Rev. Lett.*, vol. 110, p. 251801, Jun 2013. doi:10.1103/PhysRevLett.110.251801.
- [59] S. Agostinelli *et al.*, “Geant4—a simulation toolkit,” *Nucl. Instrum. Meth. A*, vol. 506, pp. 250–303, 2003.
- [60] Fermilab, “G4numi.” Link. Accessed: 2020-Sep-16.
- [61] C. Alt *et al.*, “Inclusive production of charged pions in p+c collisions at 158-GeV/c beam momentum,” *Eur. Phys. J. C*, vol. 49, pp. 897–917, 2007.
- [62] L. Aliaga *et al.*, “Neutrino flux predictions for the numi beam,” *Phys. Rev. D*, vol. 94, no. 9, p. 092005, 2016. [Addendum: Phys.Rev.D 95, 039903 (2017)].
- [63] L. Aliaga, “2017-2018 beam plots,” *Internal Document NOvA DocDB 20843*, May 2018.
- [64] J. Nieves *et al.*, “Inclusive quasi-elastic neutrino reactions,” *Phys. Rev. C*, vol. 70, p. 055503, 2004. [Erratum: Phys.Rev.C 72, 019902 (2005)].

- [65] S. Meyer *et al.*, “Deuterium target data for precision neutrino-nucleus cross sections,” *Phys. Rev. D*, vol. 93, no. 11, p. 113015, 2016.
- [66] J. Nieves *et al.*, “Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV,” *Phys. Rev. D*, vol. 88, no. 11, p. 113007, 2013.
- [67] C. Berger and L. Sehgal, “Lepton mass effects in single pion production by neutrinos,” *Phys. Rev. D*, vol. 76, p. 113004, 2007.
- [68] A. Bodek and U. Yang, “Modeling deep inelastic cross-sections in the few GeV region,” *Nucl. Phys. B Proc. Suppl.*, vol. 112, pp. 70–76, 2002.
- [69] C. Hagmann *et al.*, “Cosmic-ray shower library.” <https://nuclear.llnl.gov/simulation/main.html>. Accessed: 2020-Sep-17.
- [70] M. Martinez-Casales, “Cross section tuning 2020 blessing package,” *Internal Document NOvA DocDB 45359*, June 2020.
- [71] C. Backhouse, “The CAFAna framework,” *Internal Document NOvA DocDB 9222*, November 2014. .
- [72] D. Pershey, J. Huang, and M. Judah, “Tdslicer technote,” *Internal Document NOvA DocDB 27689*, February 2019.
- [73] A. Rodriguez and A. Laio, “Clustering by fast search and find of density peaks,” *Science*, vol. 344, no. 6191, pp. 1492–1496, 2014.
- [74] R. C. Prim, “Shortest connection networks and some generalizations,” *The Bell System Technical Journal*, vol. 36, no. 6, pp. 1389–1401, 1957.
- [75] Dcoetzee, “Cosmic-ray shower library.” <https://commons.wikimedia.org/w/index.php?curid=29652883>. Accessed: 2021-Sep-21.
- [76] F. Psihas, “Event displays for nue selected events,” *Internal Document NOvA DocDB 15647*, August 2016.
- [77] P. V. Hough, “Method and means for recognizing complex patterns,” 12 1962.
- [78] L. A. Fernandes and M. M. Oliveira, “Real-time line detection through an improved hough transform voting scheme,” *Pattern Recognition*, vol. 41, no. 1, pp. 299 – 314, 2008.

- [79] M. Ohlsson, “Extensions and explorations of the elastic arms algorithm,” *Comput. Phys. Commun.*, vol. 77, pp. 19–32, 1993.
- [80] R. Krishnapuram and J. M. Keller, “A possibilistic approach to clustering,” *IEEE Transactions on Fuzzy Systems*, vol. 1, no. 2, pp. 98–110, 1993.
- [81] M. Baird, J. Bian, M. Messier, E. Niner, D. Rocco, and K. Sachdev, “Event reconstruction techniques in nova,” *J. Phys. Conf. Ser.*, vol. 664, no. 7, p. 072035, 2015.
- [82] N. Raddatz, “Kalmantrack technical note,” *Internal Document NOvA DocDB 13545*, June 2015.
- [83] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Journal of Basic Engineering*, vol. 82, pp. 35–45, 03 1960.
- [84] B. Rebel, “A window tracking algorithm for cosmic ray muons,” *Internal Document NOvA DocDB 15977*, August 2016.
- [85] C. Backhouse, A. Radovic, P. Singh, and M. Campbell, “The attenuation and threshold calibration of the NOvA detectors,” *Internal Document NOvA DocDB 13579*, September 2017.
- [86] L. Vinton, “Far detector stopping muon energy scale calibration,” *Internal Document NOvA DocDB 13475*, November 2017.
- [87] L. Vinton, “Calorimetric energy scale calibration of the NOvA detectors,” *Internal Document NOvA DocDB 13579*, July 2015.
- [88] C. Backhouse, A. Radovic, P. Singh, and M. Campbell, “Calibration tech note 2020 analysis,” *Internal Document NOvA DocDB 43820*, March 2020.
- [89] S. Yu, “Ana2020 nue energy estimator technote,” *Internal Document NOvA DocDB 43814*, May 2020.
- [90] F. Psihas, “Electron neutrino energy estimator,” *Internal Document NOvA DocDB 14733*, January 2016.
- [91] W. Mu, “Blessed package: Numu energy spline plots 2020,” *Internal Document NOvA DocDB 46283*, June 2020.

- [92] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, “A convolutional neural network neutrino event classifier,” *Journal of Instrumentation*, vol. 11, pp. P09001–P09001, sep 2016.
- [93] F. Psihas and A. Radovic, “Assorted cvn plots for blessing,” *Internal Document NOvA DocDB 15639*, November 2017.
- [94] F. Psihas, *Measurement of Long Baseline Neutrino Oscillations and Improvements from Deep Learning*. PhD thesis, Indiana U., 2018.
- [95] E. Catano-Mur, *Constraints on neutrino oscillation parameters with the NOvA experiment*. PhD thesis, Iowa State U. (main), 2018.
- [96] J. Porter, “ReMId retraining 2019 technote,” *Internal Document NOvA DocDB 42277*, December 2019.
- [97] T. Warburton, “Numu cosmic rejection bdt technote for prod 5,” *Internal Document NOvA DocDB 42358*, December 2019.
- [98] R. P. Gandrajula and L. Kolupaeva, “Cosmic rejection BDT for ν_e appearance analysis,” *Internal Document NOvA DocDB 42473*, December 2019.
- [99] R. Team, “Root data analysis framework.” <https://root.cern>. Accessed: 2021-Jan-05.
- [100] F. James and M. Roos, “Minuit: A system for function minimization and analysis of the parameter errors and correlations,” *Comput. Phys. Commun.*, vol. 10, pp. 343–367, 1975.
- [101] A. Back, R. Bowles, M. Elkins, B. Ramson, and S. Sanchez, “Data quality 2020 technote,” *Internal Document NOvA DocDB 44226*, March 2020.
- [102] M. Baird, L. Suter, and J. Wolcott, “Executive summary for 2020 3-flavor analysis review,” *Internal Document NOvA DocDB 44422*, April 2020.
- [103] A. Sutton, “3Flavor 2020 cut flow chart,” *Internal Document NOvA DocDB 44349*, August 2020.
- [104] T. Warburton, “3 flavour cut flow tables,” *Internal Document NOvA DocDB 45880*, August 2020.
- [105] L. Vinton, *Measurement of Muon Neutrino Disappearance with the NOvA Experiment*. PhD thesis, U. of Sussex, 2018.

- [106] A. Sutton, “Ana2020 numu reconstructed energy blessing package,” *Internal Document NOvA DocDB 45809*, July 2020.
- [107] T. Warburton, “Ana 20 numu quantile plots,” *Internal Document NOvA DocDB 47242*, May 2021.
- [108] A. M. Hall, “Ana2020 far detector mc breakdown blessing package,” *Internal Document NOvA DocDB 45892*, July 2020.
- [109] J. Lozier, “Modularextrap technical note,” *Internal Document NOvA DocDB 12563*, December 2014.
- [110] A. Mislivec, “Tech note: Near-to-far extrapolation in transverse momentum for the nova 2020 3-flavor analysis,” *Internal Document NOvA DocDB 44401*, May 2020.
- [111] S. Yu, “Blessing package: Extrapolation cartoon for ana2018,” *Internal Document NOvA DocDB 29612*, June 2018.
- [112] A. Mislivec, “Pt extrapolation blessed plots for the 2020 analysis,” *Internal Document NOvA DocDB 45980*, September 2020.
- [113] D. Pershey, “Technote for the extrapolation for the 2017 nue appearance analysis,” *Internal Document NOvA DocDB 23108*, September 2017.
- [114] S. Yu, “Bendecomp ana2018 technote,” *Internal Document NOvA DocDB 27833*, April 2018.
- [115] I. Anghel, Z. Djurcic, and M. Sanchez, “A technote on development of beam nue estimation method in nova oscillation analysis,” *Internal Document NOvA DocDB 15442*, June 2016.
- [116] M. Elkins, “Ben+Michel decomposition blessing package 2020,” *Internal Document NOvA DocDB 45865*, July 2020.
- [117] D. Pershey, “Technote for MichelDecomp2017,” *Internal Document NOvA DocDB 22523*, September 2017.
- [118] C. Backhouse and A. Himmel, “Nue 2017 executive summary,” *Internal Document NOvA DocDB 22475*, October 2017.
- [119] D. Pershey, “MEFinder technote,” *Internal Document NOvA DocDB 14789*, February 2016.

- [120] L. Kolupaeva and M. Elkins, “Cosmic predictions for 2020 three-flavor analysis,” *Internal Document NOvA DocDB 44316*, May 2020.
- [121] L. Kolupaeva, “Cosmic systematic,” *Internal Document NOvA DocDB 47305*, September 2020.
- [122] N. Nayak, “Technote : PCA flux systematics for the 2020 oscillation analysis,” *Internal Document NOvA DocDB 44049*, March 2020.
- [123] K. Bays, M. Dolce, S. MacDonald, M. Casales, A. Mislivec, G. Pawloski, and J. Walcott, “Nova 2020 cross-section tune tech note,” *Internal Document NOvA DocDB 43962*, September 2020.
- [124] P. Rodrigues, C. Wilkinson, and K. McFarland, “Constraining the genie model of neutrino-induced single pion production using reanalyzed bubble chamber data,” *Eur. Phys. J. C*, vol. 76, no. 8, p. 474, 2016.
- [125] K. Abe *et al.*, “Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector,” *Phys. Rev. D*, vol. 96, no. 9, p. 092006, 2017. [Erratum: *Phys.Rev.D* 98, 019902 (2018)].
- [126] J. Hewes and O. Samoylov, “Nova production 5 detector simulation changes and systematics,” *Internal Document NOvA DocDB 43935*, April 2020.
- [127] A. M. Hall, “Light level down / FD ageing comparison,” *Internal Document NOvA DocDB 41070*, October 2019.
- [128] T. Alion, “Absolute calibration and shape systematics for 3a,” *Internal Document NOvA DocDB 20318*, June 2017.
- [129] Z. Vallari, “Technote: Nue acceptance syst for ana2020,” *Internal Document NOvA DocDB 43917*, April 2020.
- [130] M. Strait, “Muon energy scale systematic,” *Internal Document NOvA DocDB 20816*, November 2020.
- [131] D. E. Groom, N. V. Mokhov, and S. I. Striganov, “Muon stopping power and range tables 10-MeV to 100-TeV,” *Atom. Data Nucl. Data Tabl.*, vol. 78, pp. 183–356, 2001.

- [132] N. T. Force, “NOvA neutron systematic for 3-flavor analysis,” *Internal Document NOvA DocDB 43811*, February 2020.
- [133] M. Baird, “Update on gg-mismatch problem and other normalization issues for the 3f group,” *Internal Document NOvA DocDB 44238*, March 2020.
- [134] S. Baker and R. D. Cousins, “Clarification of the use of χ^2 and likelihood functions in fits to histograms,” *Nucl. Instrum. Meth.*, vol. 221, pp. 437–442, 1984.
- [135] A. Radovic, “Oscillation parameters for first NOvA analyses,” *Internal Document NOvA DocDB 13640*, July 2015.
- [136] C. J. Backhouse and A. Himmel, “Nue 2018 executive summary,” *Internal Document NOvA DocDB 26699*, May 2018.
- [137] T. Alion, “Data-mc differences in normalization,” *Internal Document NOvA DocDB 27821*, May 2018.
- [138] L. Kolupaeva and T. Warburton, “2020 sensitivity and fitting technote,” *Internal Document NOvA DocDB 44043*, May 2020.
- [139] G. J. Feldman and R. D. Cousins, “A unified approach to the classical statistical analysis of small signals,” *Phys. Rev. D*, vol. 57, pp. 3873–3889, 1998.
- [140] S. Y. Calvez, “Blessing package : Joint fit ana2020 results,” *Internal Document NOvA DocDB 46331*, September 2020.
- [141] L. Kolupaeva, “Future 3 flavour sensitivities - 2020 analysis,” *Internal Document NOvA DocDB 46343*, December 2020.
- [142] M. Wallbank, “The NOvA test beam program,” in *Proceedings of 40th International Conference on High Energy physics — PoS(ICHEP2020)*, vol. 390, p. 188, 2021.

Appendix A

Neutrino Physics Supplementary Material

A.1 Explicit Calculation of 3-flavour Transition Probability

It is shown in Chapter 2 that for a neutrino produced in definite flavour state λ , the amplitude corresponding to detecting that neutrino in flavour state λ' after propagating a distance L in vacuum, is given by

$$\mathcal{A}(\nu_\lambda \rightarrow \nu_{\lambda'}) = \sum_{j=1}^N U_{\lambda j}^* U_{\lambda' j} e^{-iP_j \cdot X}. \quad (\text{A.1})$$

Lemma: The corresponding transition probability can be expressed

$$\begin{aligned} \mathcal{P}(\nu_\lambda \rightarrow \nu_{\lambda'}) &= \delta_{\lambda\lambda'} - 4 \sum_{j < k} \text{Re} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin^2 \left(\frac{\phi_{jk}}{2} \right) \\ &\quad - 2 \sum_{j < k} \text{Im} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin(\phi_{jk}) \quad , \quad \phi_{jk} := \frac{\Delta m_{jk}^2 L}{2E}. \end{aligned} \quad (\text{A.2})$$

The following identities and definitions will be used:

$$|z_1 + z_2 + z_3|^2 = z_1^2 + z_2^2 + z_3^2 + 2\text{Re} [z_1 z_2^* + z_1 z_3^* + z_2 z_3^*] \quad , \quad z_l \in \mathbb{C} \quad (\text{A.3})$$

$$e^{i\theta} = \cos(\theta) + i \sin(\theta) \quad (\text{A.4})$$

$$\text{Re}[A] = \text{Re}[A^*] \quad (\text{A.5})$$

$$\text{Re}[AB] = \text{Re}[A] \text{Re}[B] - \text{Im}[A] \text{Im}[B] \quad (\text{A.6})$$

$$U_{PMNS}^{-1} = U_{PMNS}^\dagger \quad (\text{A.7})$$

$$\cos(2\theta) = 1 - 2 \sin^2 \left(\frac{\theta}{2} \right). \quad (\text{A.8})$$

$$\begin{aligned}
\mathcal{P}(\nu_\lambda \rightarrow \nu_{\lambda'}) &= |\mathcal{A}(\nu_\lambda \rightarrow \nu_{\lambda'})|^2 \\
&= \left| \sum_{j=1}^N U_{\lambda j}^* U_{\lambda' j} e^{-i P_j \cdot X} \right|^2 \\
&= \left| U_{\lambda 1}^* U_{\lambda' 1} e^{-i P_1 \cdot X} + U_{\lambda 2}^* U_{\lambda' 2} e^{-i P_2 \cdot X} + U_{\lambda 3}^* U_{\lambda' 3} e^{-i P_3 \cdot X} \right|^2 \\
&= (U_{\lambda 1}^* U_{\lambda' 1} e^{-i P_1 \cdot X}) (U_{\lambda 1}^* U_{\lambda' 1} e^{-i P_1 \cdot X})^* + (U_{\lambda 2}^* U_{\lambda' 2} e^{-i P_2 \cdot X}) (U_{\lambda 2}^* U_{\lambda' 2} e^{-i P_2 \cdot X})^* \\
&\quad + (U_{\lambda 3}^* U_{\lambda' 3} e^{-i P_3 \cdot X}) (U_{\lambda 3}^* U_{\lambda' 3} e^{-i P_3 \cdot X})^* \\
&\quad + 2\text{Re} \left[(U_{\lambda 1}^* U_{\lambda' 1} e^{-i P_1 \cdot X}) (U_{\lambda 2}^* U_{\lambda' 2} e^{-i P_2 \cdot X})^* \right. \\
&\quad + (U_{\lambda 1}^* U_{\lambda' 1} e^{-i P_1 \cdot X}) (U_{\lambda 3}^* U_{\lambda' 3} e^{-i P_3 \cdot X})^* \\
&\quad \left. + (U_{\lambda 2}^* U_{\lambda' 2} e^{-i P_2 \cdot X}) (U_{\lambda 3}^* U_{\lambda' 3} e^{-i P_3 \cdot X})^* \right] \\
&= |U_{\lambda 1}|^2 |U_{\lambda' 1}|^2 + |U_{\lambda 2}|^2 |U_{\lambda' 2}|^2 + |U_{\lambda 3}|^2 |U_{\lambda' 3}|^2 \\
&\quad + 2\text{Re} \left[U_{\lambda 1}^* U_{\lambda 2} U_{\lambda' 1} U_{\lambda' 2}^* e^{i\phi_{21}} \right. \\
&\quad + U_{\lambda 1}^* U_{\lambda 3} U_{\lambda' 1} U_{\lambda' 3}^* e^{i\phi_{31}} \\
&\quad \left. + U_{\lambda 2}^* U_{\lambda 3} U_{\lambda' 2} U_{\lambda' 3}^* e^{i\phi_{32}} \right] \\
&= |U_{\lambda 1} U_{\lambda' 1}^*|^2 + |U_{\lambda 2} U_{\lambda' 2}^*|^2 + |U_{\lambda 3} U_{\lambda' 3}^*|^2 \\
&\quad + 2\text{Re} \left[U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2} e^{-i\phi_{21}} \right. \\
&\quad + U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3} e^{-i\phi_{31}} \\
&\quad \left. + U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3} e^{-i\phi_{32}} \right] \\
&= |U_{\lambda 1} U_{\lambda' 1}^* + U_{\lambda 2} U_{\lambda' 2}^* + U_{\lambda 3} U_{\lambda' 3}^*|^2 \\
&\quad - 2\text{Re} [U_{\lambda 1} U_{\lambda' 1}^* U_{\lambda 2}^* U_{\lambda 2} + U_{\lambda 1} U_{\lambda' 1}^* U_{\lambda 3}^* U_{\lambda 3} + U_{\lambda 2} U_{\lambda' 2}^* U_{\lambda 3}^* U_{\lambda 3}] \\
&\quad + 2\text{Re} [U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2} (\cos(\phi_{12}) + i \sin(\phi_{12})) \\
&\quad + U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3} (\cos(\phi_{13}) + i \sin(\phi_{13})) \\
&\quad + U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3} (\cos(\phi_{23}) + i \sin(\phi_{23}))] \\
&= \delta_{\lambda\lambda'} - 2\text{Re} [U_{\lambda 1} U_{\lambda' 1}^* U_{\lambda 2}^* U_{\lambda 2} + U_{\lambda 1} U_{\lambda' 1}^* U_{\lambda 3}^* U_{\lambda 3} + U_{\lambda 2} U_{\lambda' 2}^* U_{\lambda 3}^* U_{\lambda 3}] \\
&\quad + 2\text{Re} \left[U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2} \left(1 - 2 \sin^2 \left(\frac{\phi_{12}}{2} \right) + i \sin(\phi_{12}) \right) \right. \\
&\quad + U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3} \left(1 - 2 \sin^2 \left(\frac{\phi_{13}}{2} \right) + i \sin(\phi_{13}) \right) \\
&\quad \left. + U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3} \left(1 - 2 \sin^2 \left(\frac{\phi_{23}}{2} \right) + i \sin(\phi_{23}) \right) \right]
\end{aligned}$$

cont.

$$\begin{aligned}
\mathcal{P}(\nu_\lambda \rightarrow \nu_{\lambda'}) &= \delta_{\lambda\lambda'} - 4\text{Re} \left[U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2} \sin^2 \left(\frac{\phi_{12}}{2} \right) \right. \\
&\quad + U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3} \sin^2 \left(\frac{\phi_{13}}{2} \right) \\
&\quad \left. + U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3} \sin^2 \left(\frac{\phi_{23}}{2} \right) \right] \\
&\quad + 2(\text{Re} [U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2} i \sin(\phi_{12})] \\
&\quad + \text{Re} [U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3} i \sin(\phi_{13})] \\
&\quad + \text{Re} [U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3} i \sin(\phi_{23})]) \\
&= \delta_{\lambda\lambda'} - 4 \sum_{j < k} \text{Re} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin^2 \left(\frac{\phi_{jk}}{2} \right) \\
&\quad + 2(\text{Re} [U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2}] \text{Re} [i \sin \phi_{12}] - \text{Im} [U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2}] \text{Im} [i \sin \phi_{12}] \\
&\quad + \text{Re} [U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3}] \text{Re} [i \sin \phi_{13}] - \text{Im} [U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3}] \text{Im} [i \sin \phi_{13}] \\
&\quad + \text{Re} [U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3}] \text{Re} [i \sin \phi_{23}] - \text{Im} [U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3}] \text{Im} [i \sin \phi_{23}]) \\
&= \delta_{\lambda\lambda'} - 4 \sum_{j < k} \text{Re} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin^2 \left(\frac{\phi_{jk}}{2} \right) \\
&\quad - 2(\text{Im} [U_{\lambda 1} U_{\lambda 2}^* U_{\lambda' 1}^* U_{\lambda' 2}] \text{Im} [i \sin \phi_{12}] \\
&\quad + \text{Im} [U_{\lambda 1} U_{\lambda 3}^* U_{\lambda' 1}^* U_{\lambda' 3}] \text{Im} [i \sin \phi_{13}] \\
&\quad + \text{Im} [U_{\lambda 2} U_{\lambda 3}^* U_{\lambda' 2}^* U_{\lambda' 3}] \text{Im} [i \sin \phi_{23}]) \\
&= \delta_{\lambda\lambda'} - 4 \sum_{j < k} \text{Re} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin^2 \left(\frac{\phi_{jk}}{2} \right) \\
&\quad - 2 \sum_{j < k} \text{Im} [U_{\lambda j} U_{\lambda' j}^* U_{\lambda k}^* U_{\lambda' k}] \sin(\phi_{jk}),
\end{aligned}$$

where by Equation 2.19

$$\phi_{jk} := \frac{\Delta m_{jk}^2 L}{2E}.$$

■

Appendix B

Analysis Bin Mapping

Table B.1 is a map to translate from ‘Analysis Bin’ as shown in Figure 6.16b to the presentational reconstructed neutrino energy binning used in Figures such as Figure 6.5.

1D Bin Number	Presentational Reconstructed Neutrino Energy Binning
1 to 3	PID (low, medium) I, CosRej low, 1 to 2.5 GeV
4 to 6	PID (low, medium) I, CosRej high, 1 to 2.5 GeV
7 to 9	PID (low, medium) II, CosRej high, 1 to 2.5 GeV
10 to 12	PID (medium, medium), CosRej low, 1 to 2.5 GeV
13 to 15	PID (medium, medium), CosRej high, 1 to 2.5 GeV
16 to 21	PID (medium, high), CosRej low, 1 to 4 GeV
22 to 27	PID (high, high), CosRej low, 1 to 4 GeV
28 to 33	PID (low, low), CosRej high, 1 to 4 GeV
34 to 39	PID (low, high), CosRej high, 1 to 4 GeV
40 to 45	PID (medium, high), CosRej high, 1 to 4 GeV
46 to 51	PID (high, high), CosRej high, 1 to 4 GeV
52 to 57	Unused
58 to 60	Peripheral

Table B.1: A map to translate from ‘Analysis Bin’ in FHC as shown in Figure 6.16b to the presentational reconstructed neutrino energy binning used extensively from Chapter 6 onwards.

Appendix C

Supplementary Data to Simulation Comparison Plots

This appendix contains a suite of figures comparing data and simulation at the ND (Section § C.1) and FD (Section § C.2). They are described and discussed fully in Chapter 7.

C.1 ND

Figures begin on the next page.

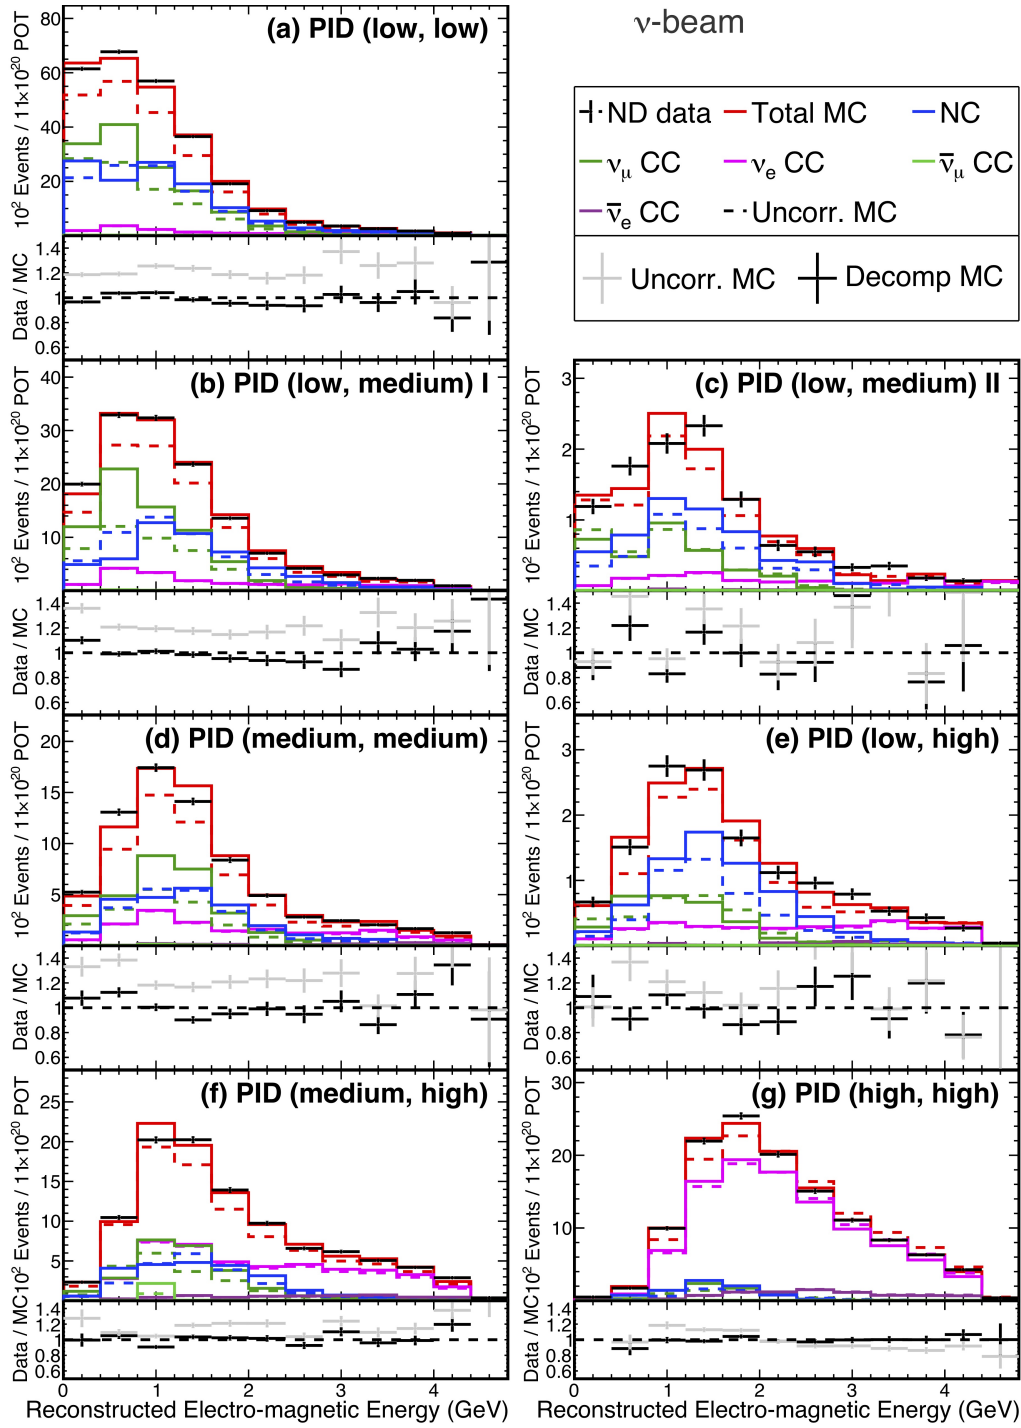


Figure C.1: Distributions of reconstructed electro-magnetic energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.




Figure C.2: Distributions of reconstructed electro-magnetic energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.




Figure C.3: Distributions of reconstructed hadronic energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

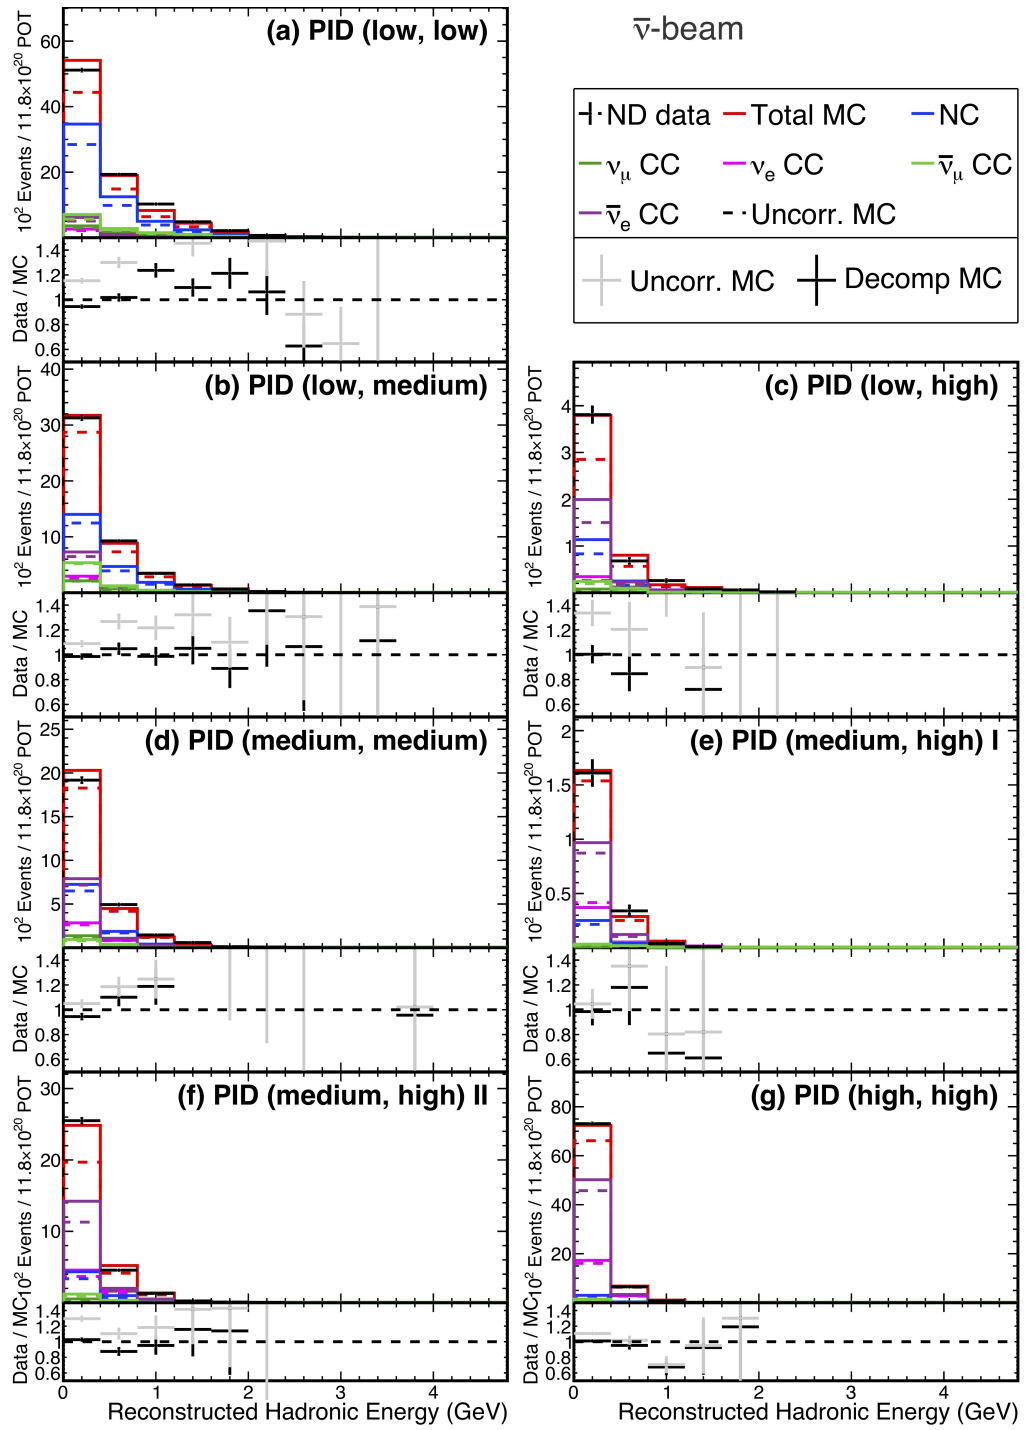


Figure C.4: Distributions of reconstructed hadronic energy for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

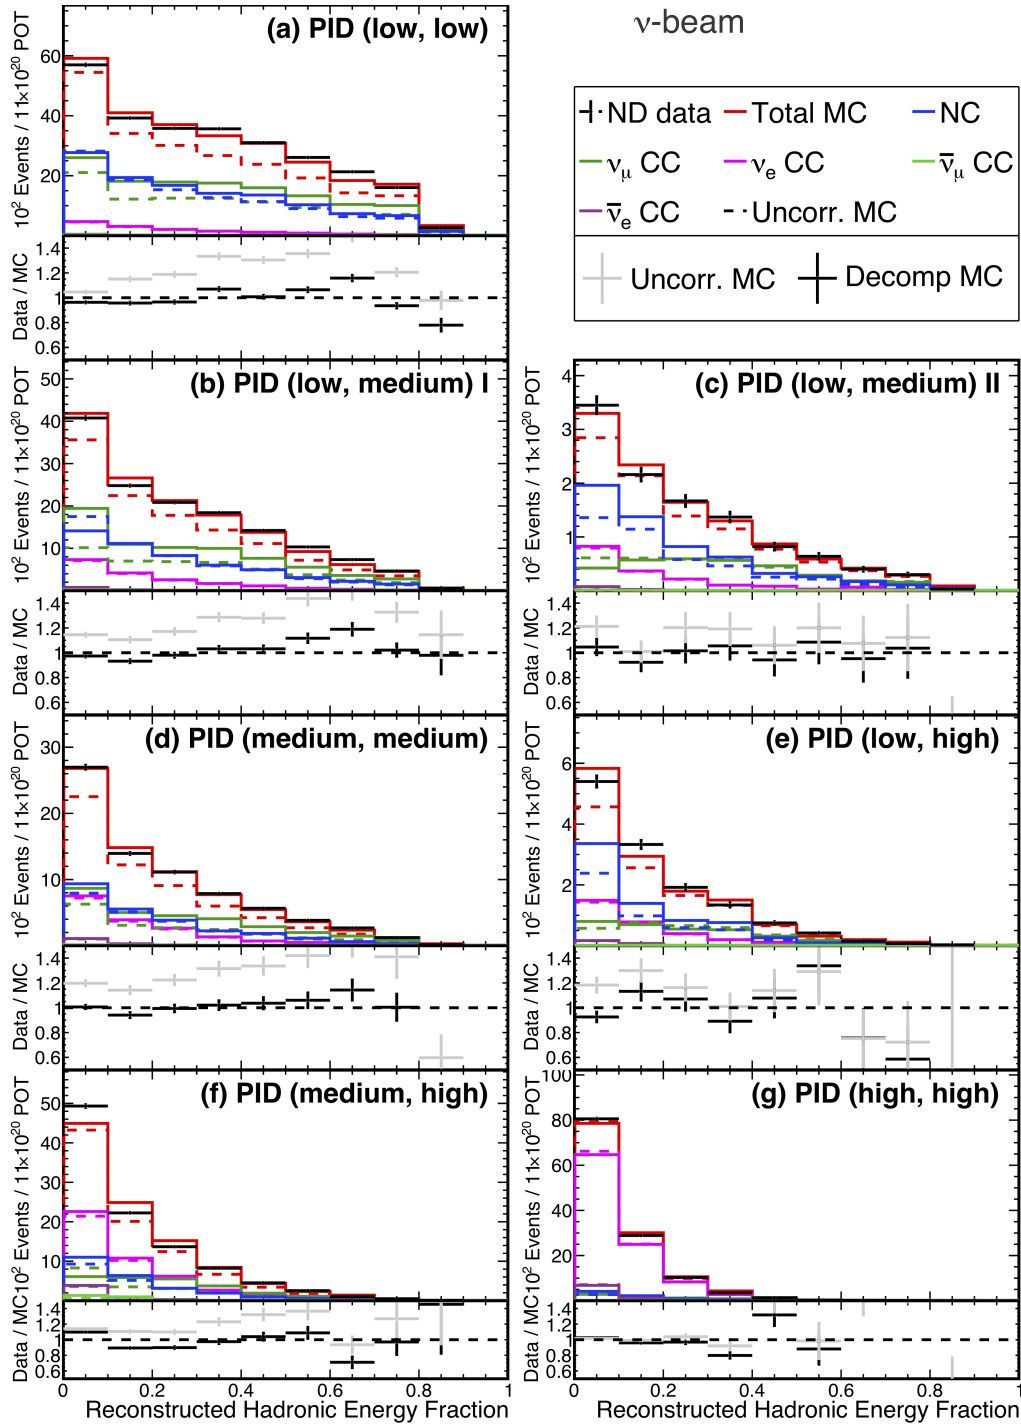


Figure C.5: Distributions of reconstructed hadronic energy fraction for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

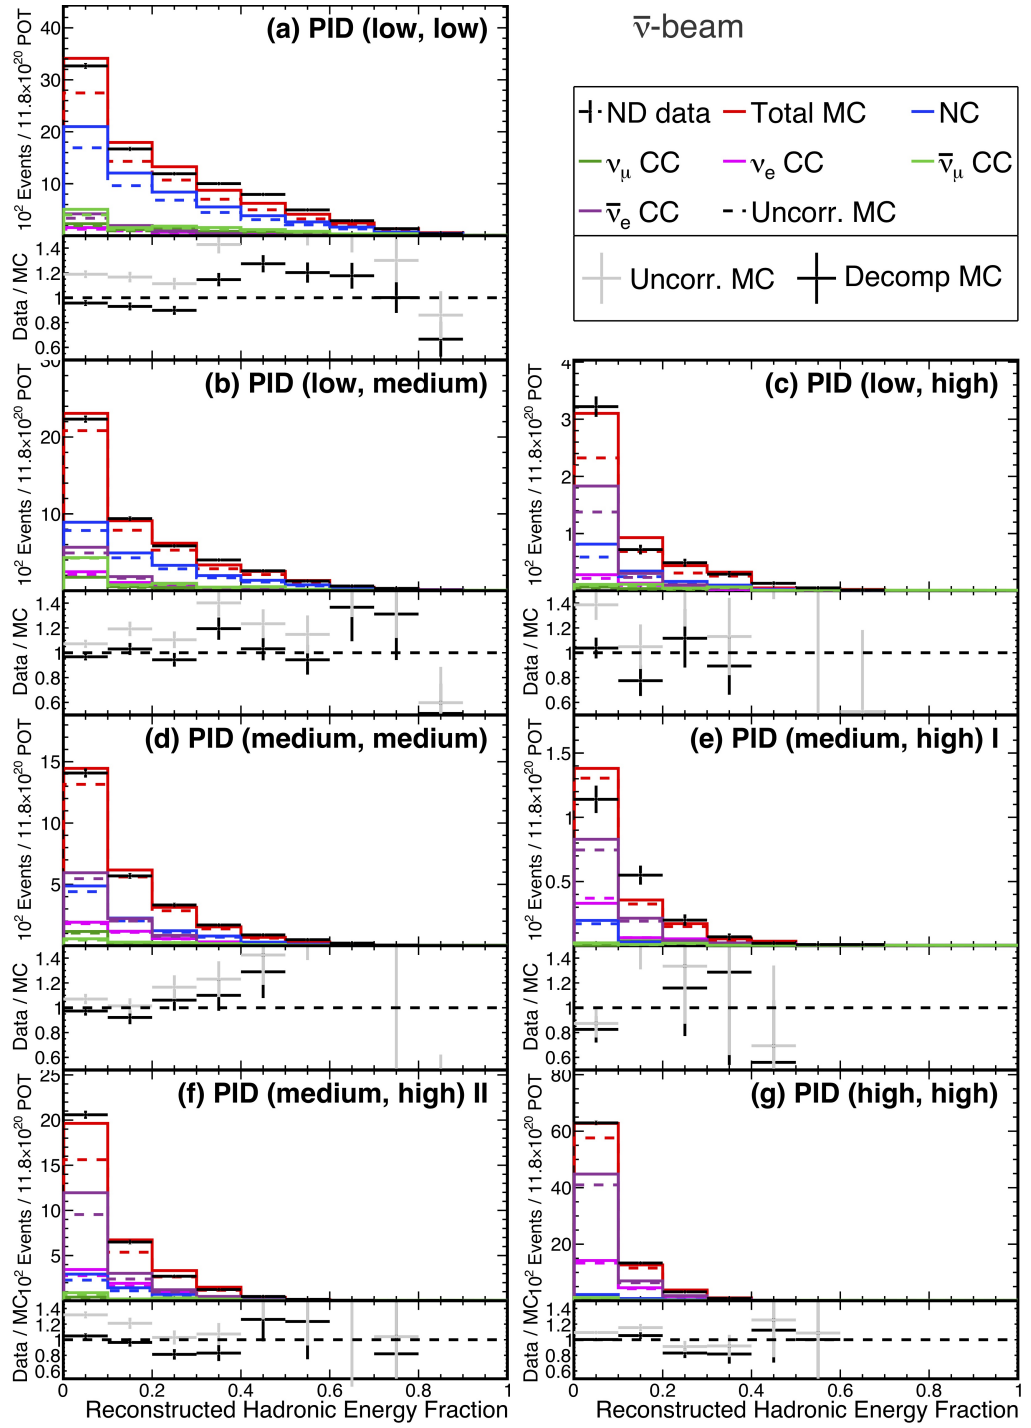


Figure C.6: Distributions of reconstructed hadronic energy fraction for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.




Figure C.7: Distributions of electron neutrino CVN A score for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

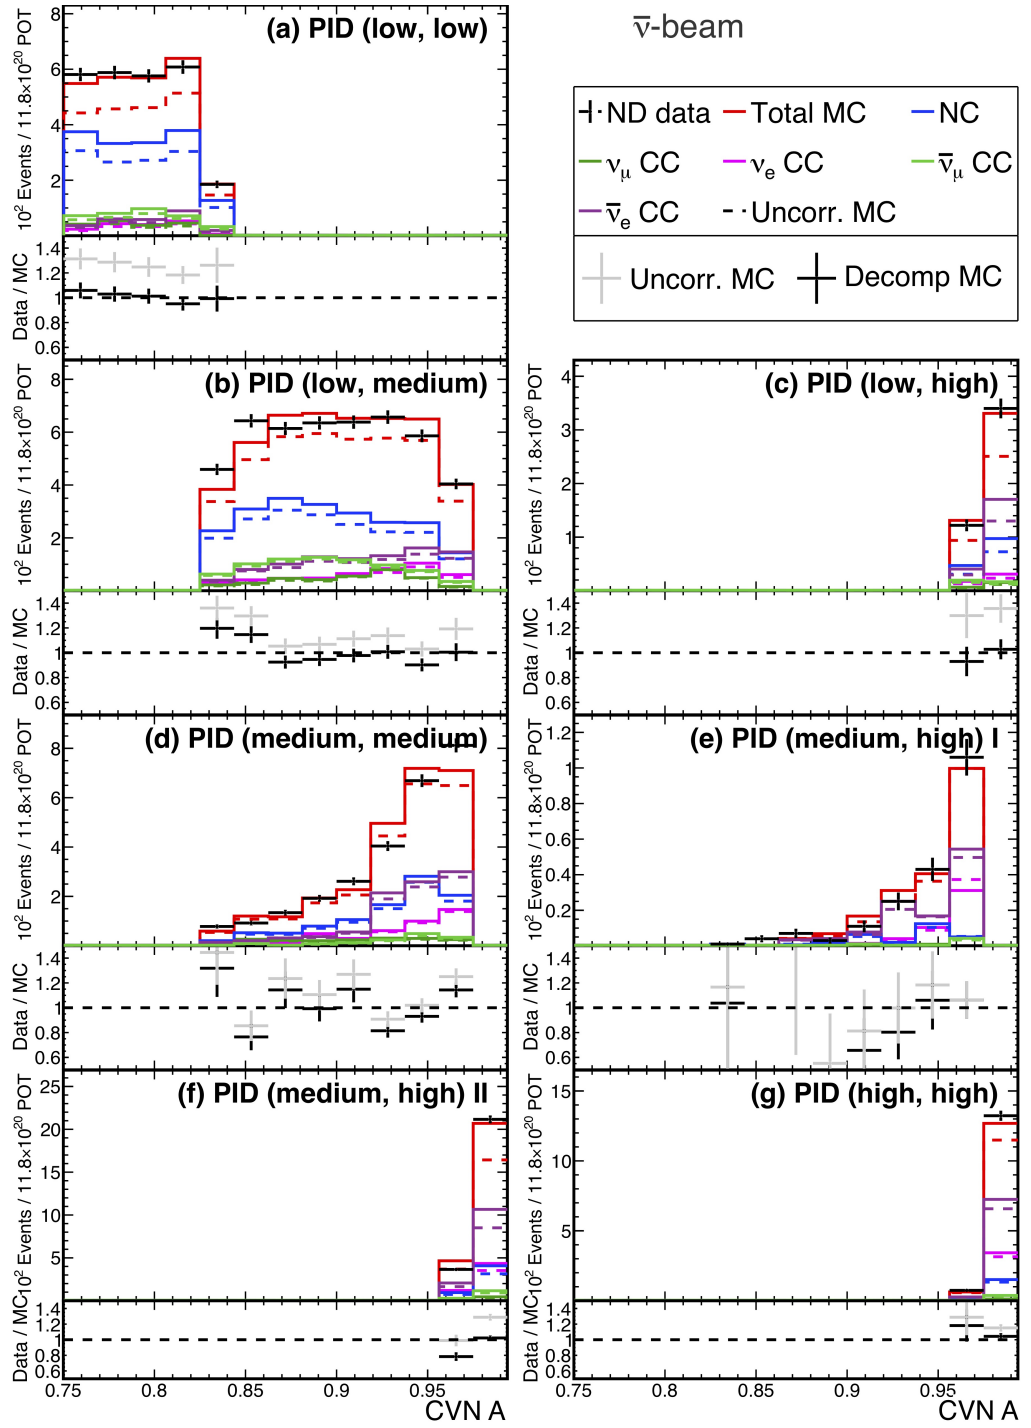


Figure C.8: Distributions of electron neutrino CVN A score for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.




Figure C.9: Distributions of reconstructed vertex X position for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

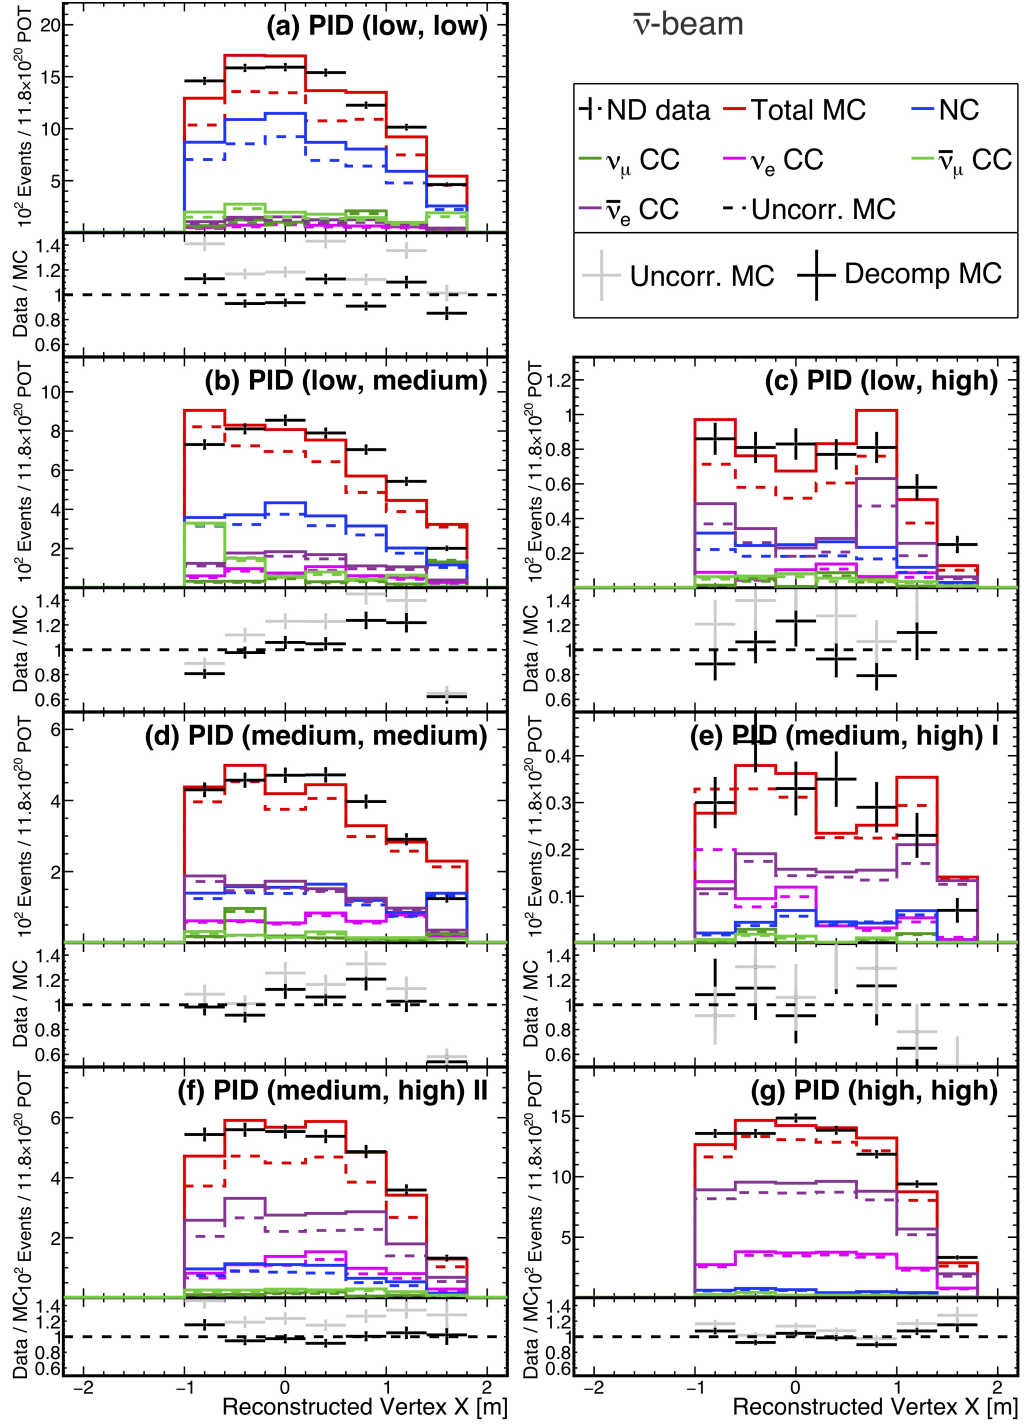


Figure C.10: Distributions of reconstructed vertex X position for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.




Figure C.11: Distributions of reconstructed vertex Y position for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

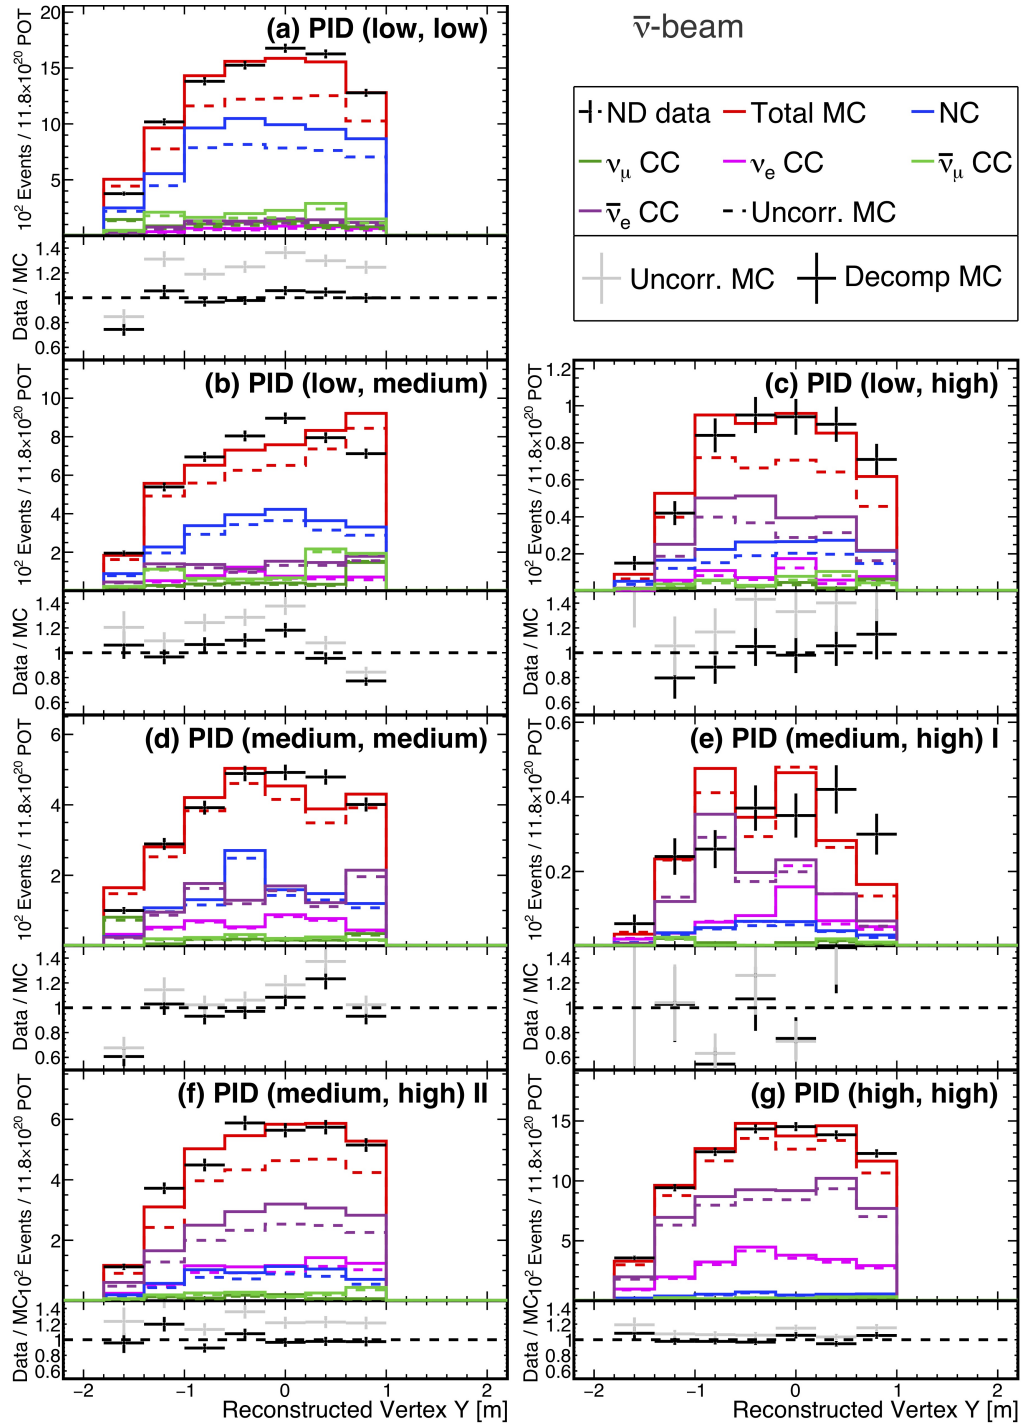


Figure C.12: Distributions of reconstructed vertex Y position for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

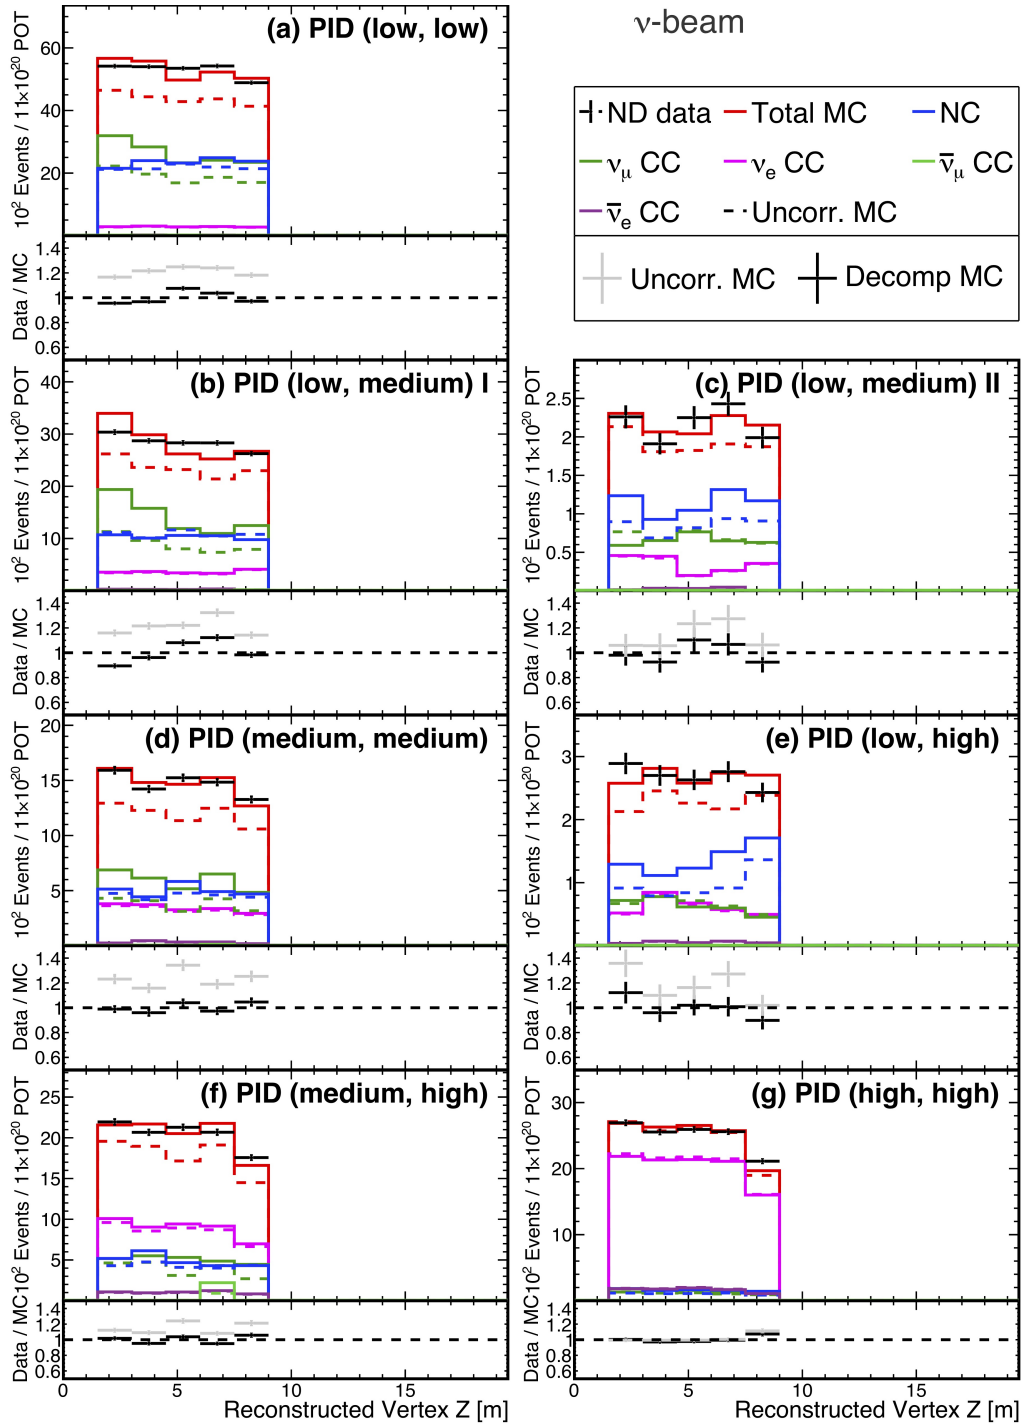


Figure C.13: Distributions of reconstructed vertex Z position for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

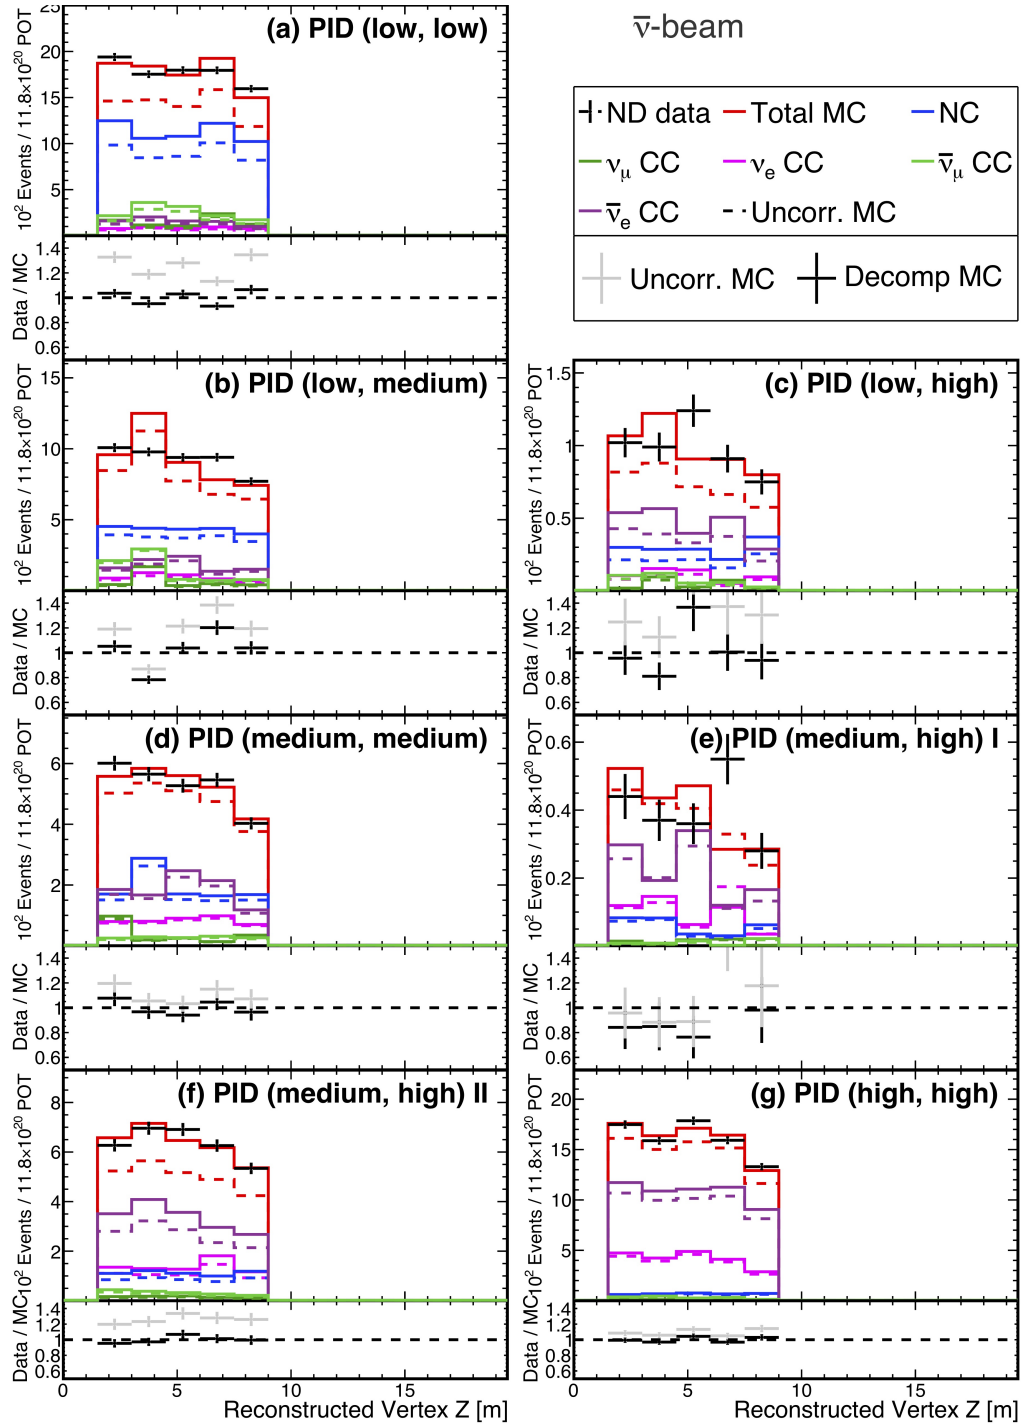


Figure C.14: Distributions of reconstructed vertex Z position for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

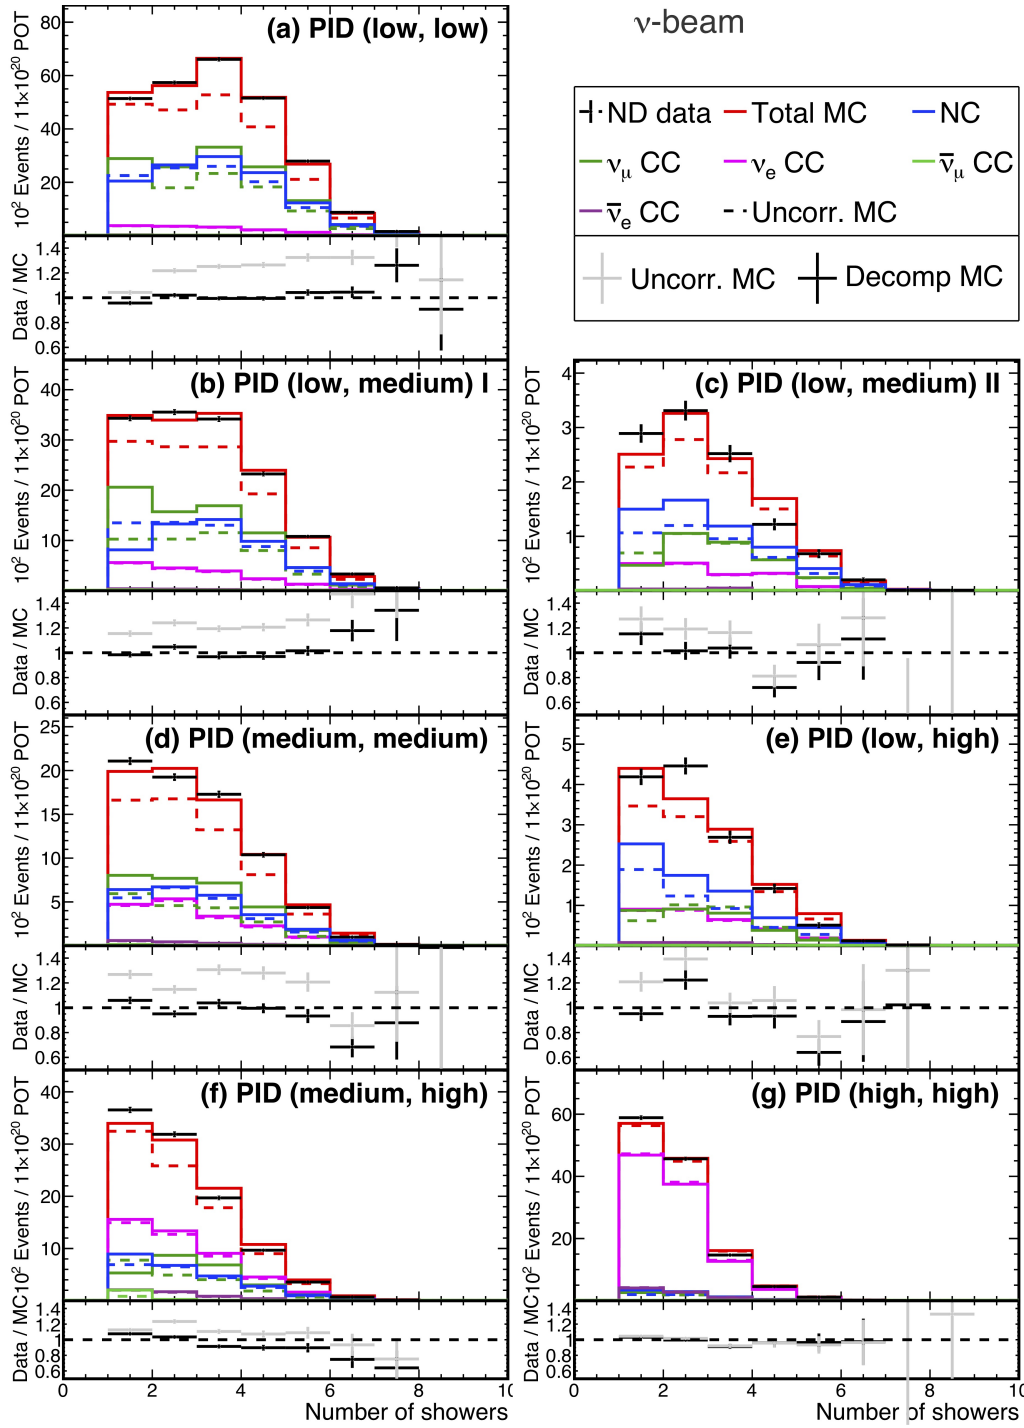


Figure C.15: Distributions of the number of reconstructed showers in a slice for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

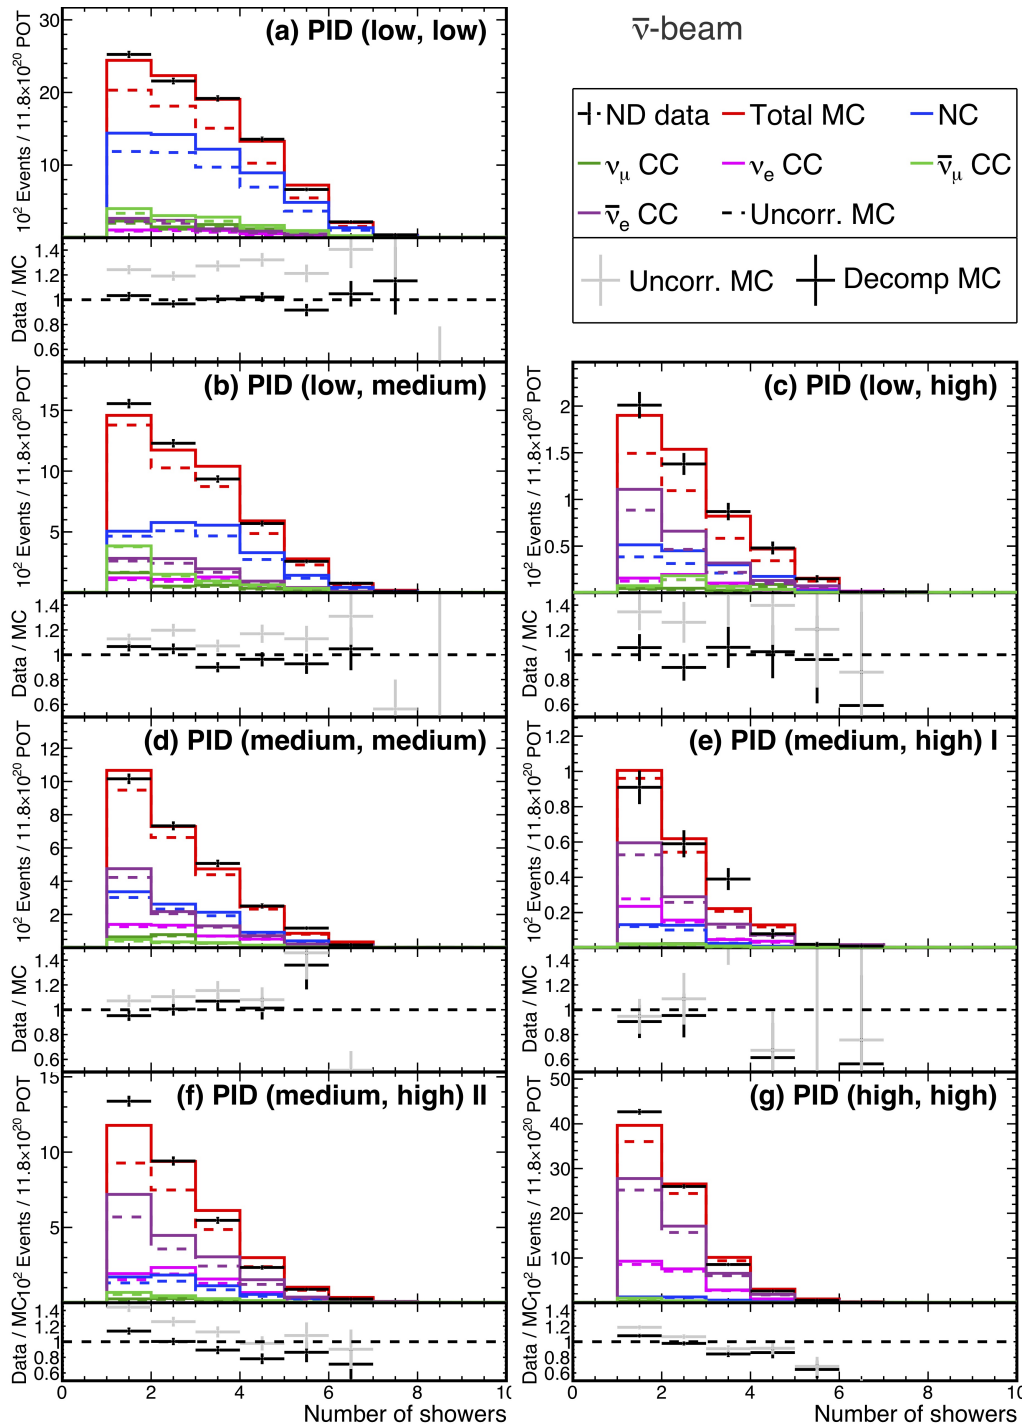


Figure C.16: Distributions of the number of reconstructed showers in a slice for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

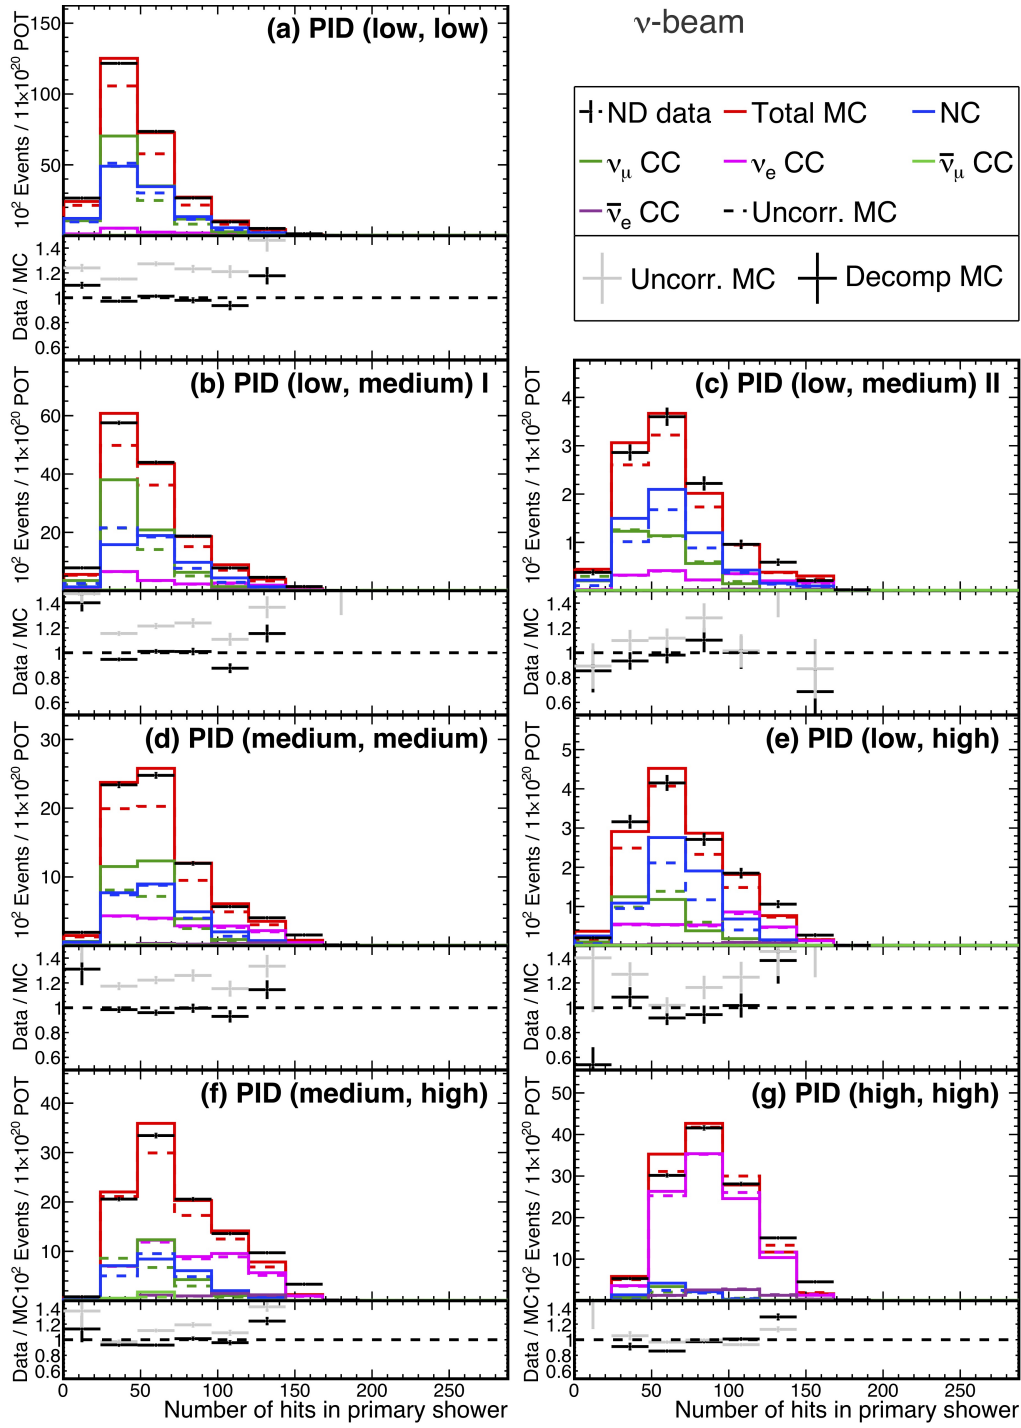


Figure C.17: Distributions of the number of hits in the primary shower for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for FHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

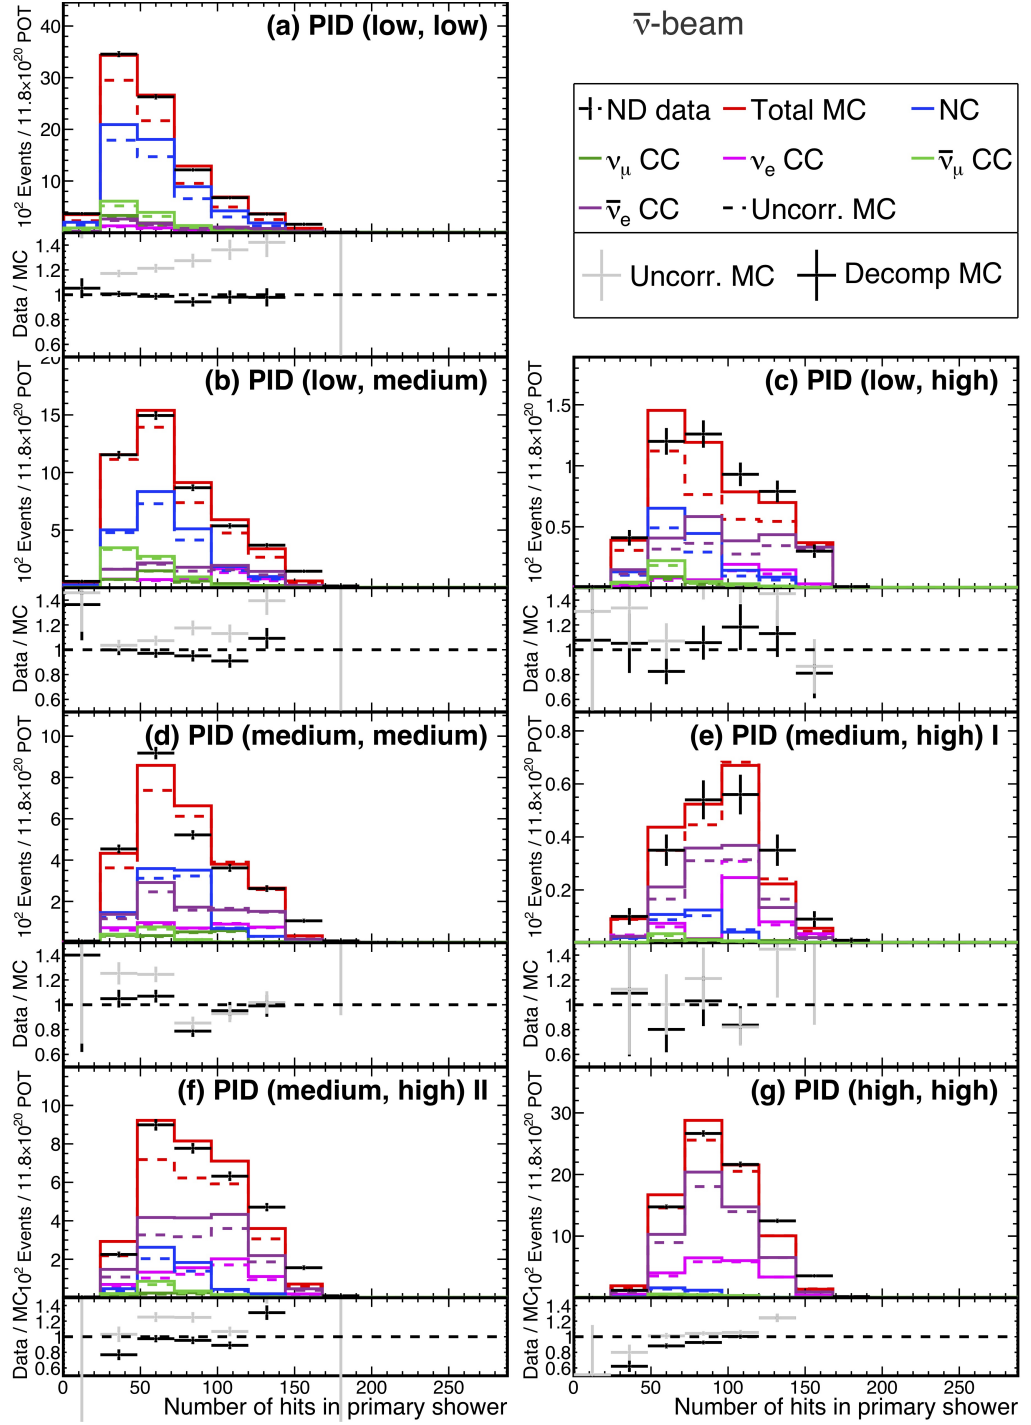


Figure C.18: Distributions of the number of hits in the primary shower for events passing the electron neutrino selection at the ND, before and after decomposition, in each selection bin (a - g) for RHC beam. Data is shown in black. Corrected and uncorrected simulated beam components are shown in the remaining solid and dashed histograms respectively. The ratios of the data to corrected (black) and uncorrected (grey) total simulated events are also shown.

C.2 FD

Figures begin on the next page.

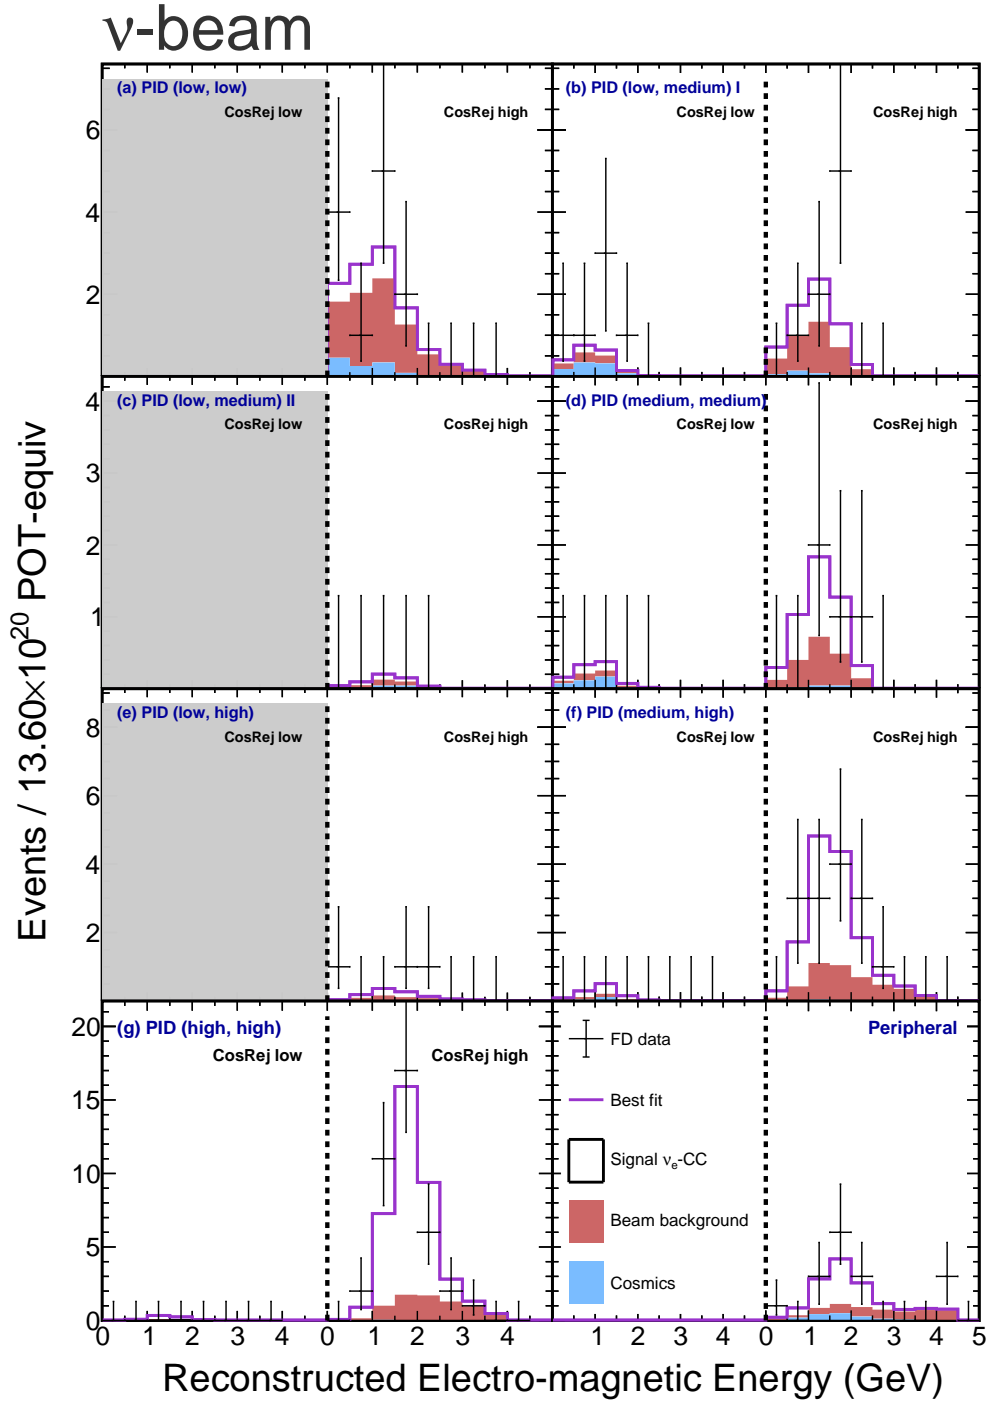


Figure C.19: Distributions of reconstructed electro-magnetic energy for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

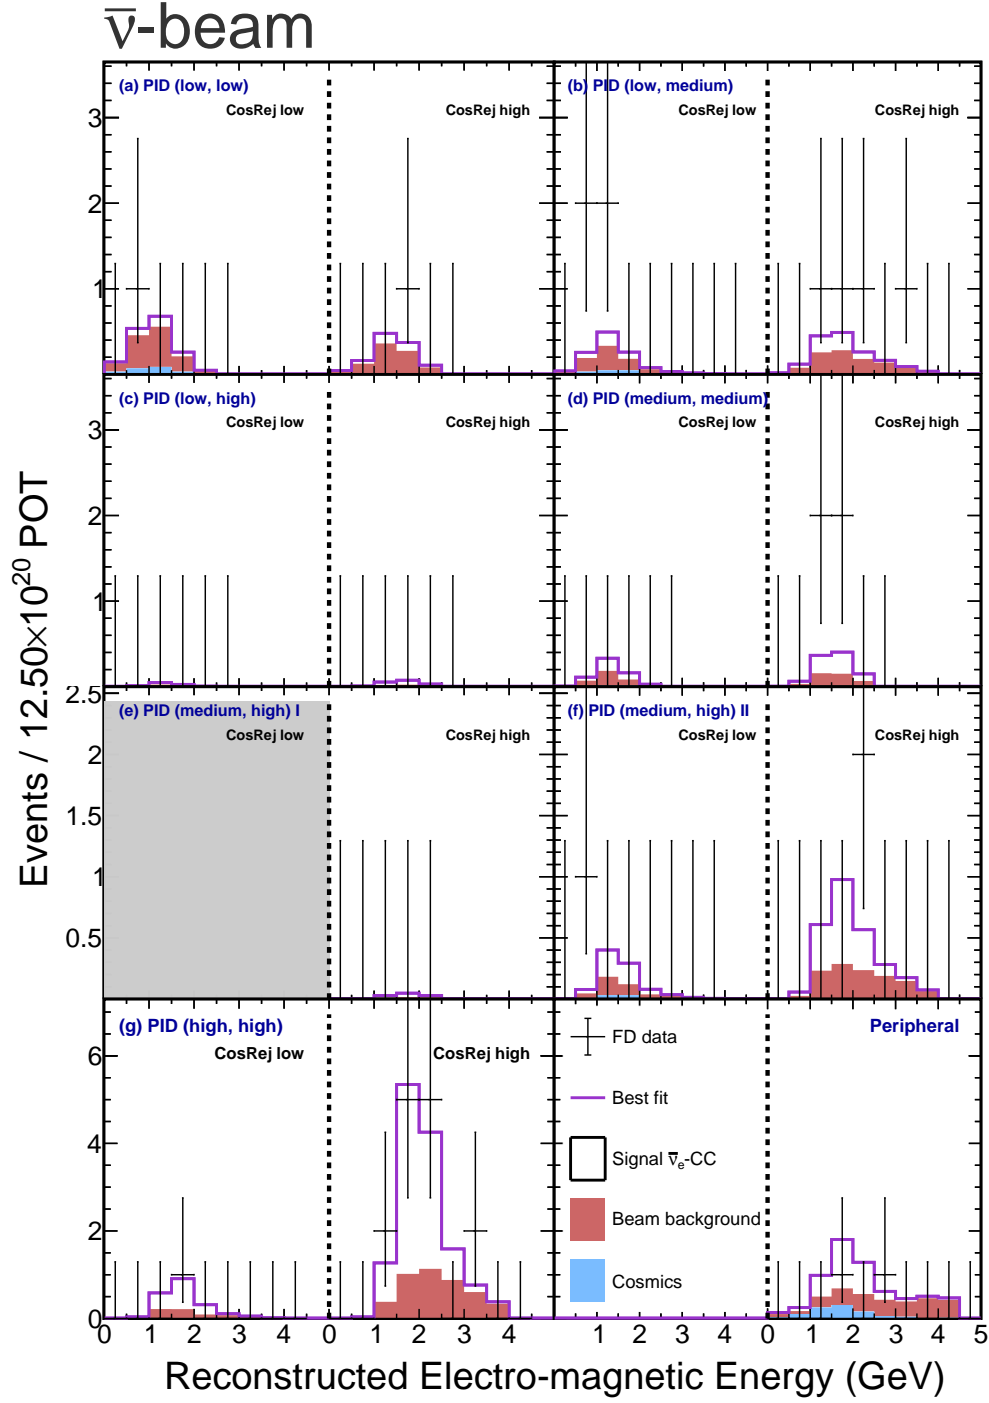


Figure C.20: Distributions of reconstructed electro-magnetic energy for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.21: Distributions of reconstructed hadronic energy for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

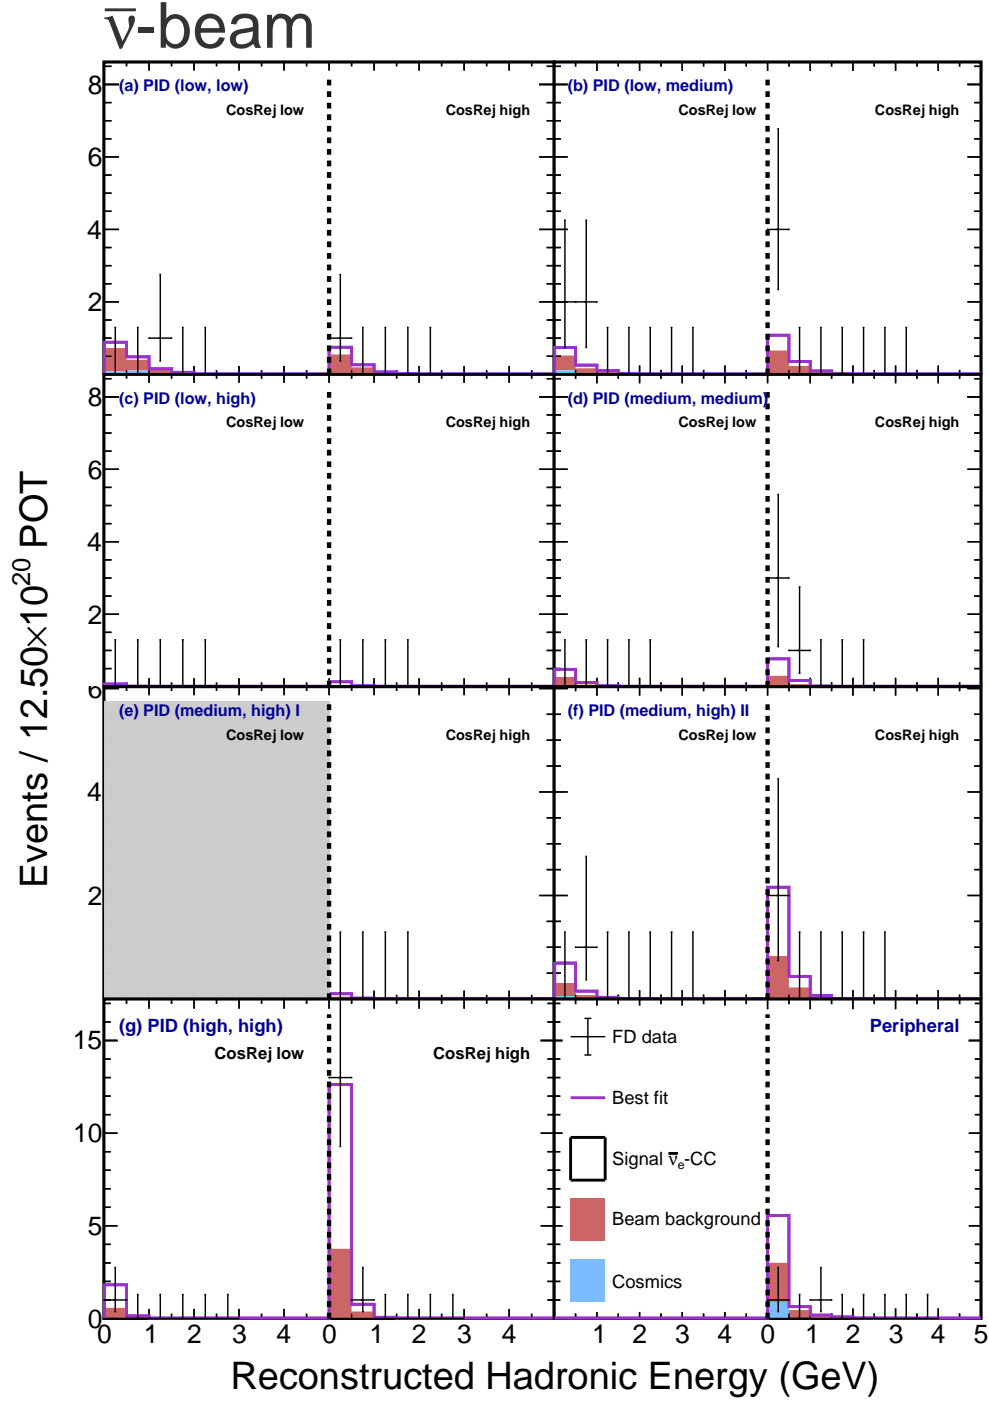


Figure C.22: Distributions of reconstructed hadronic energy for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.23: Distributions of reconstructed hadronic energy fraction for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

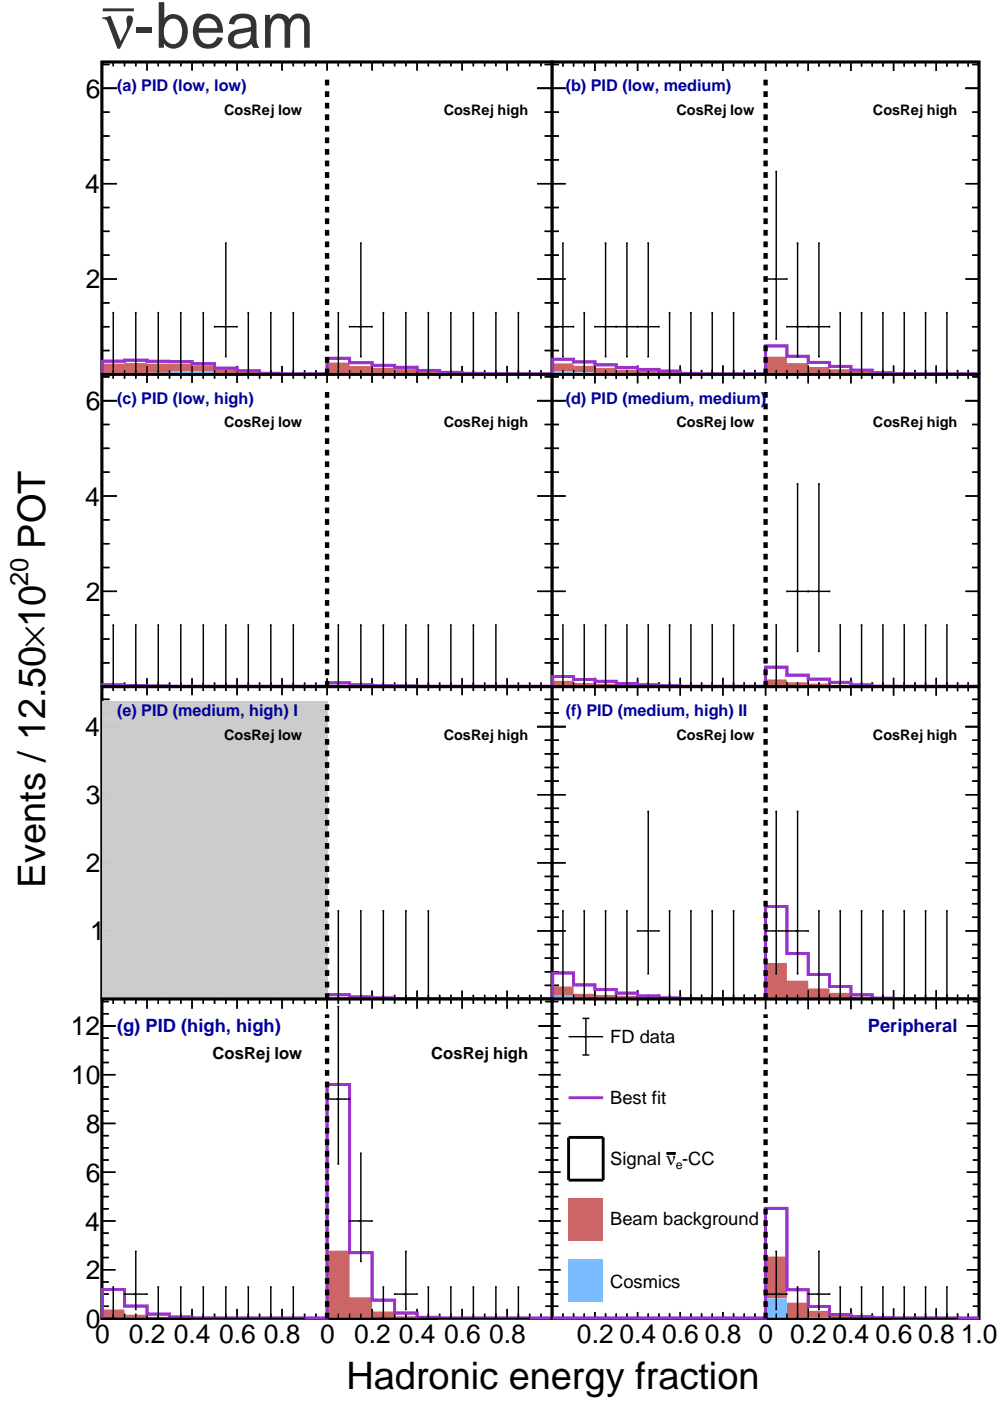


Figure C.24: Distributions of reconstructed hadronic energy fraction for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

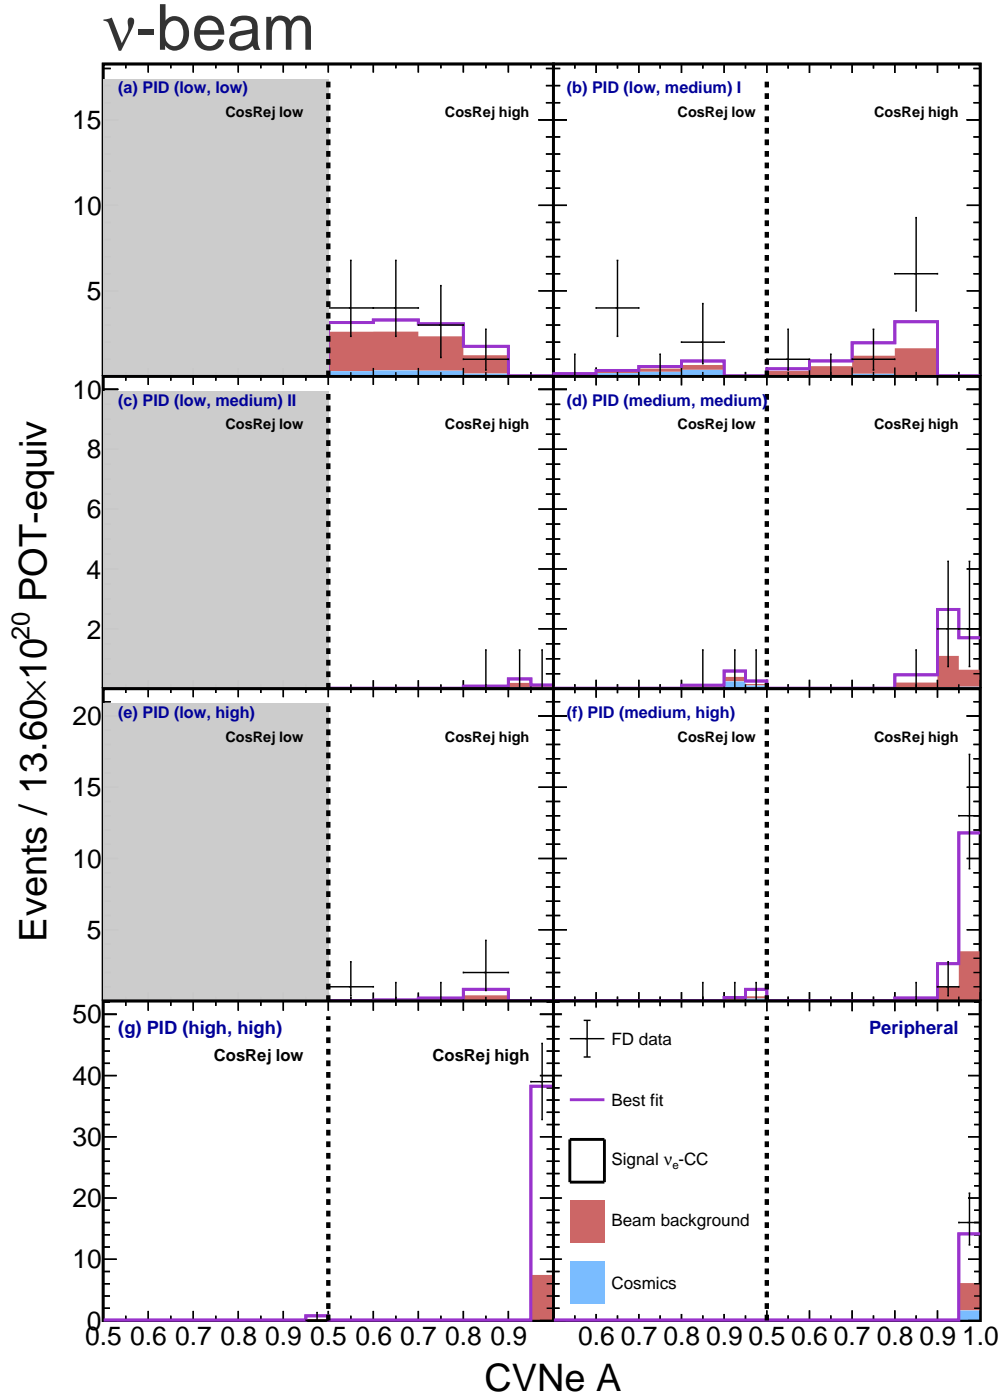


Figure C.25: Distributions of CVN A score for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.26: Distributions of CVN A score for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

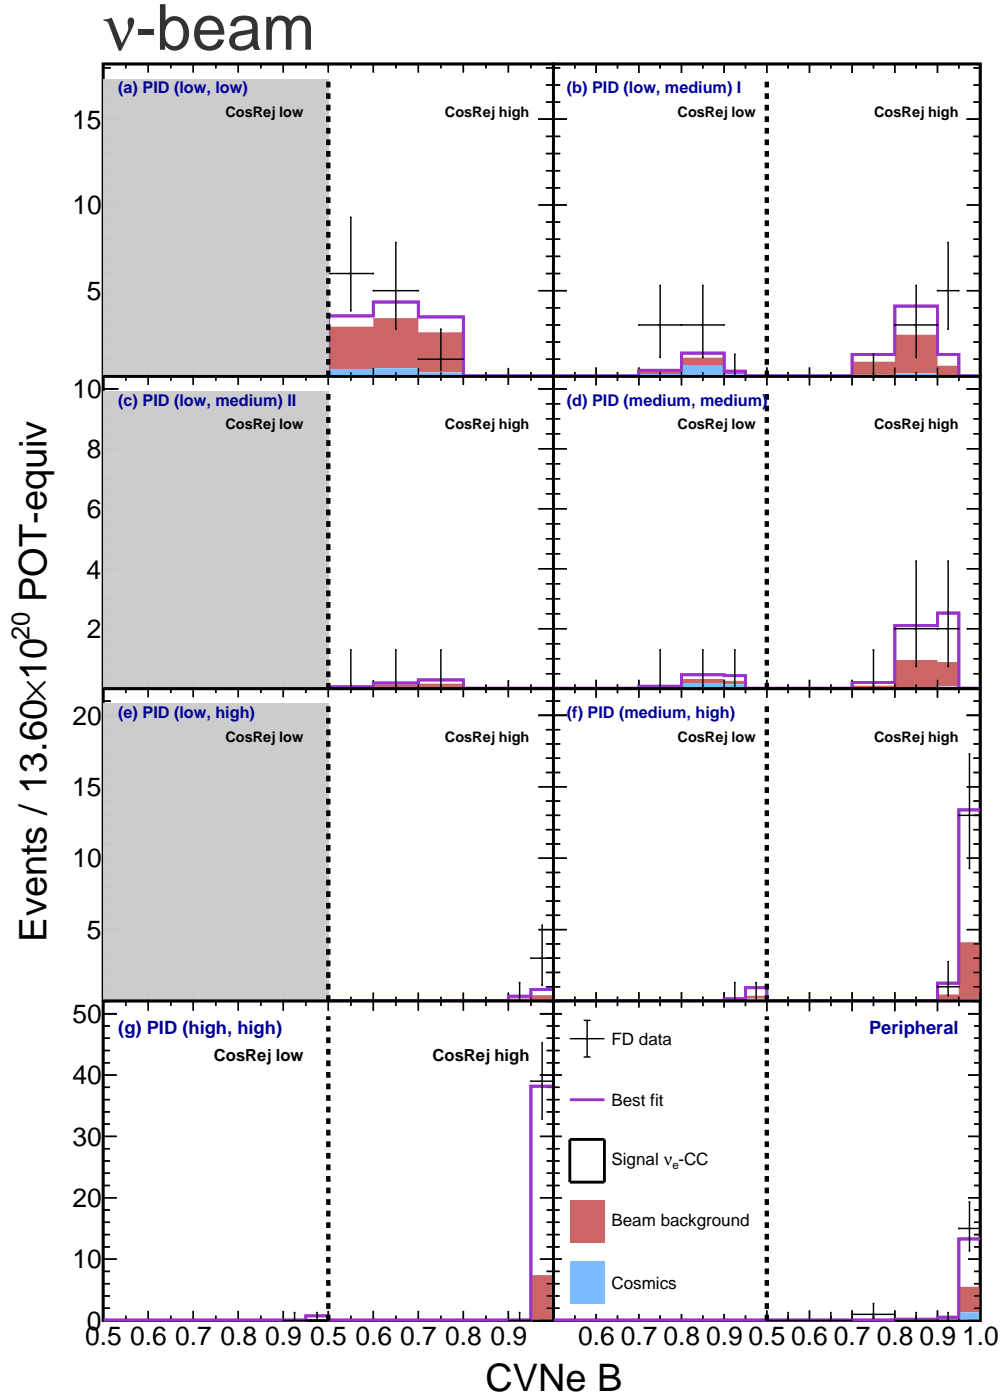


Figure C.27: Distributions of CVN B score for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.28: Distributions of CVN B score for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

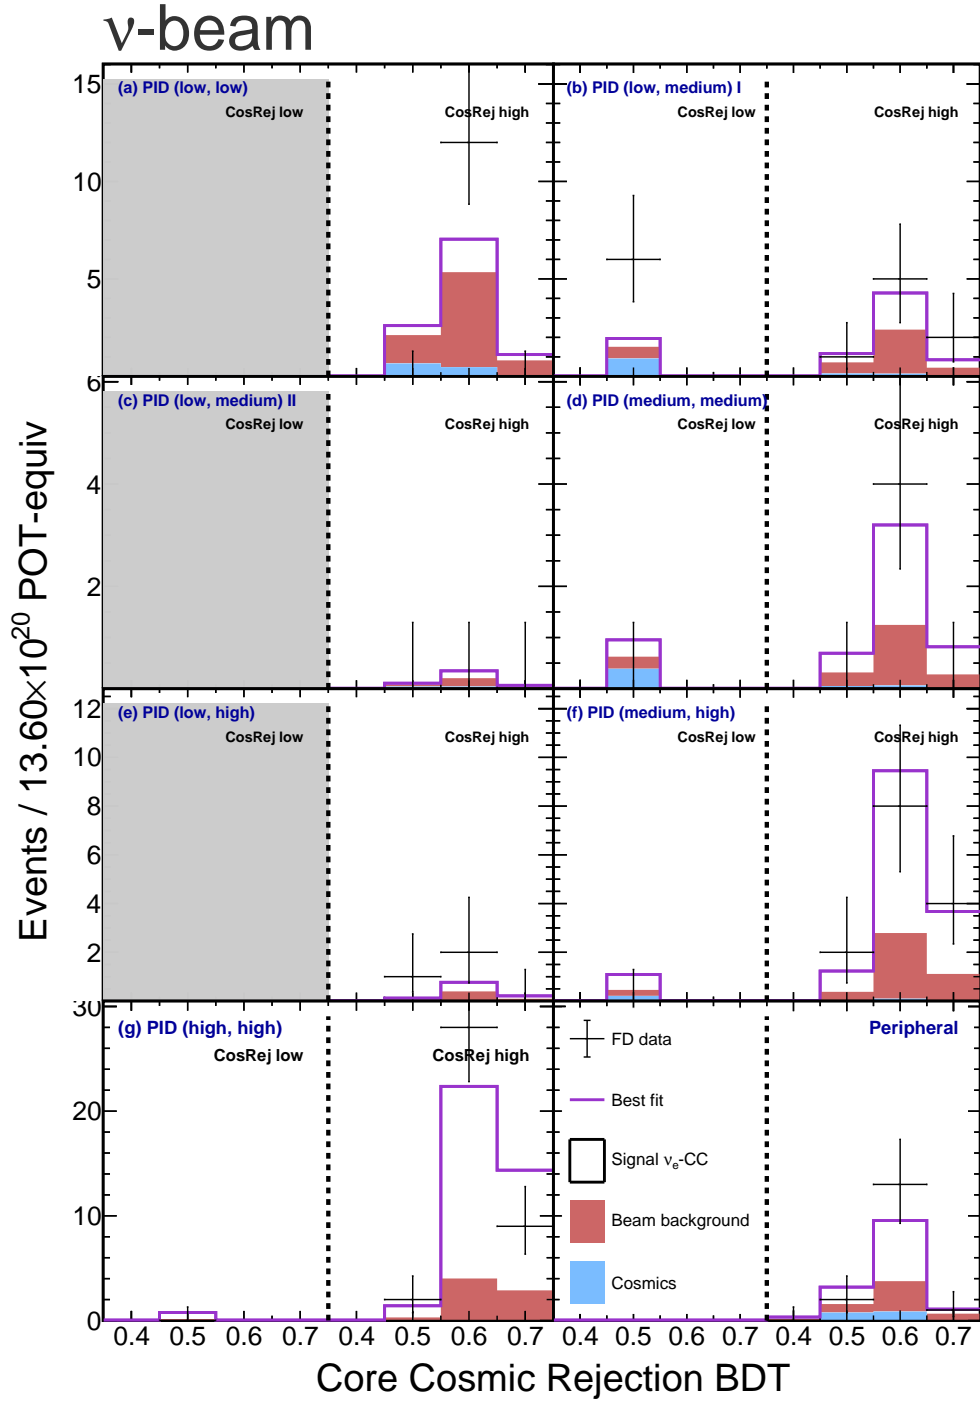


Figure C.29: Distributions of the cosmic rejection BDT score for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

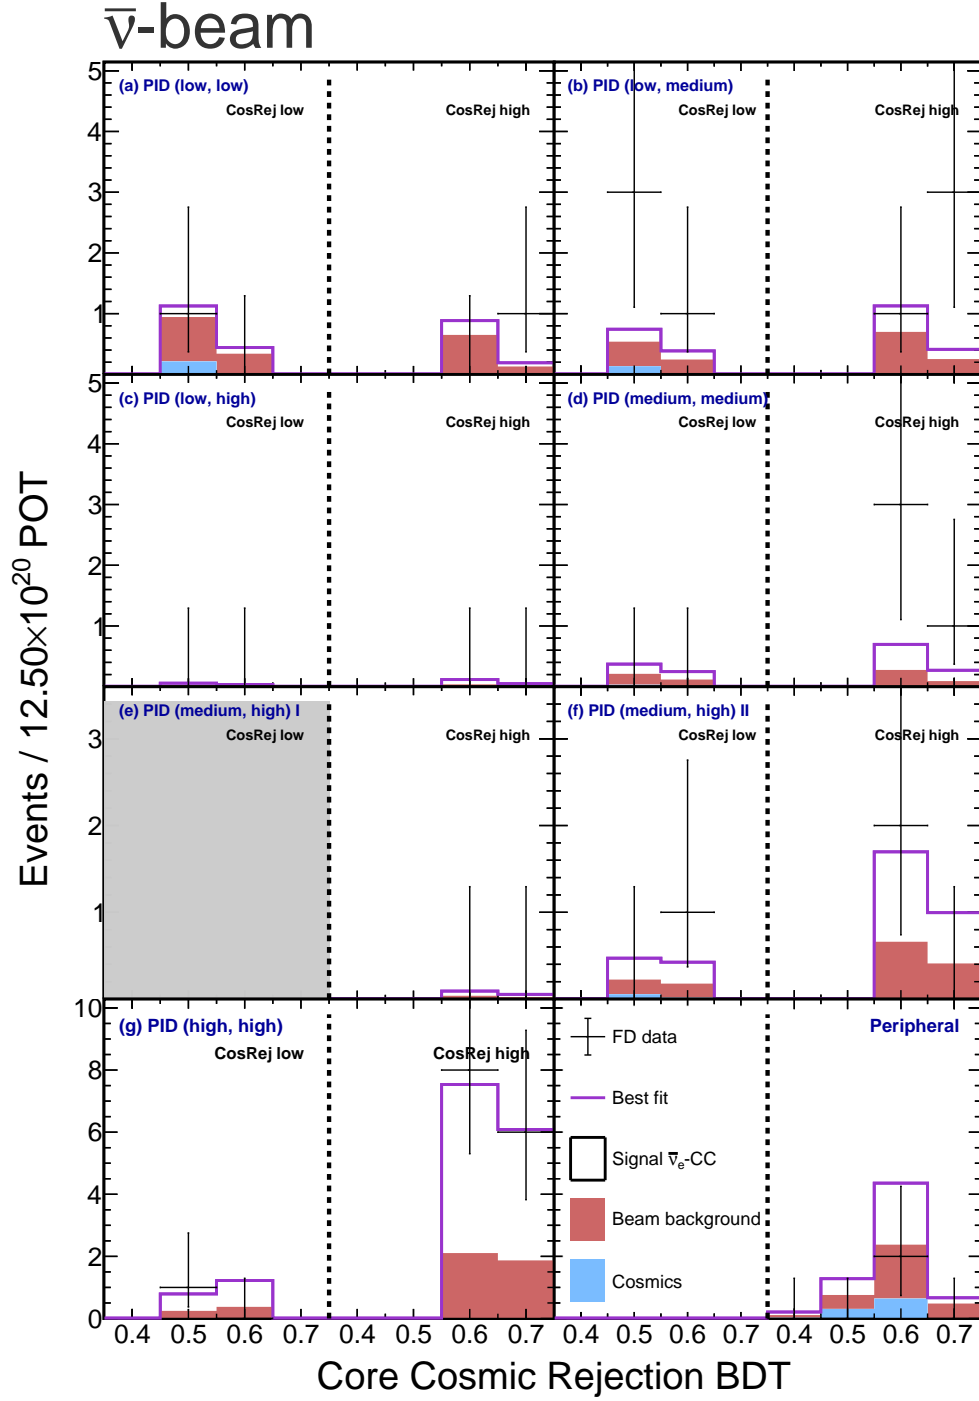


Figure C.30: Distributions of the cosmic rejection BDT score for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

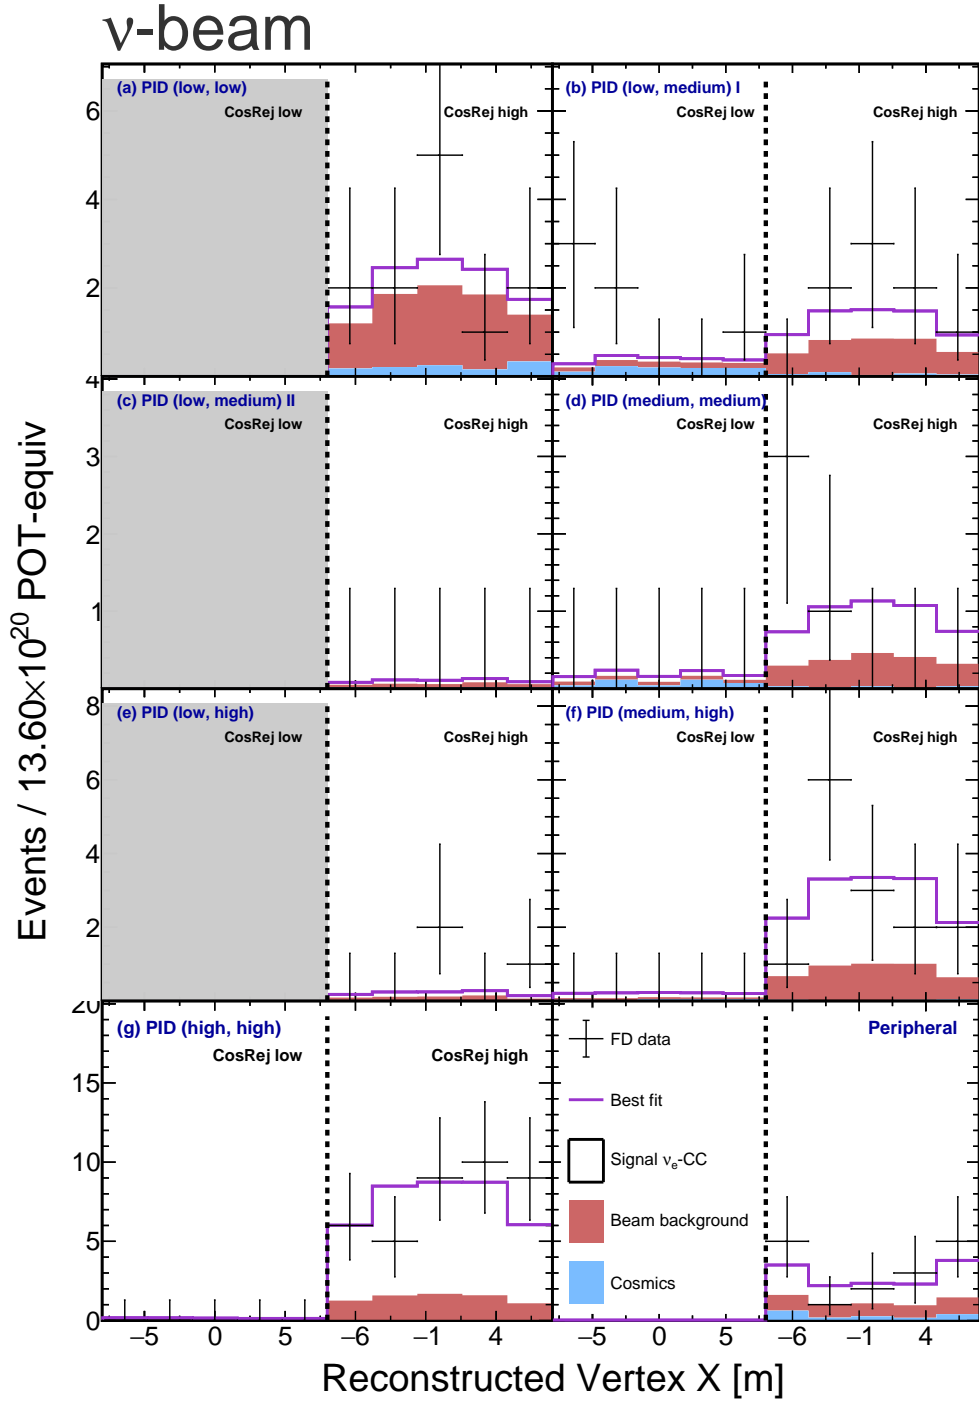


Figure C.31: Distributions of the reconstructed vertex X position for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.32: Distributions of the reconstructed vertex X position for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.33: Distributions of the reconstructed vertex Y position for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

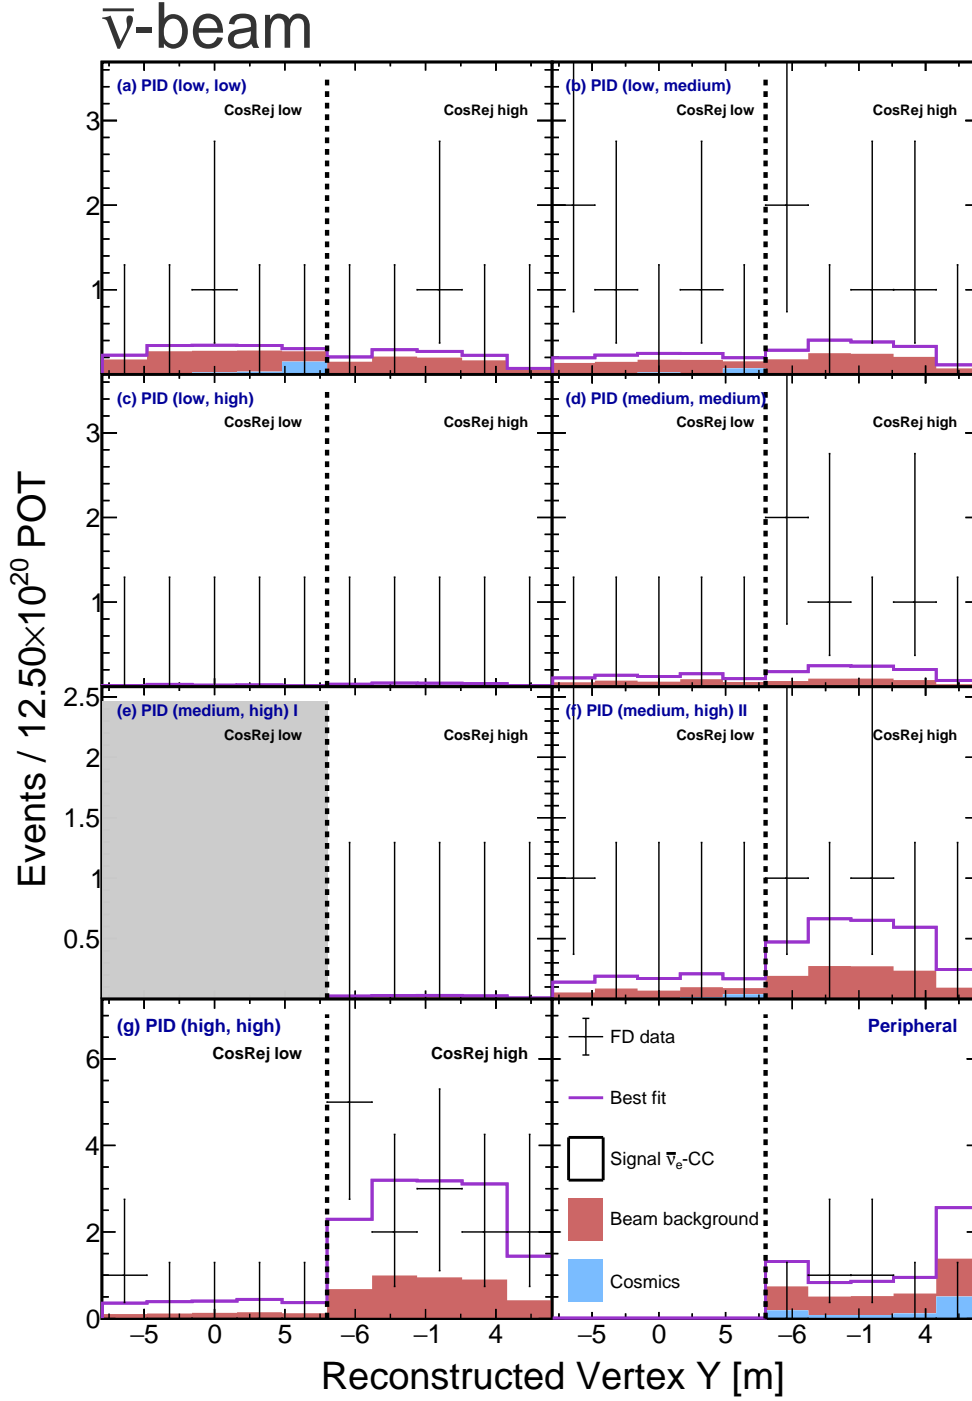


Figure C.34: Distributions of the reconstructed vertex Y position for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.35: Distributions of the reconstructed vertex Z position for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.36: Distributions of the reconstructed vertex Z position for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

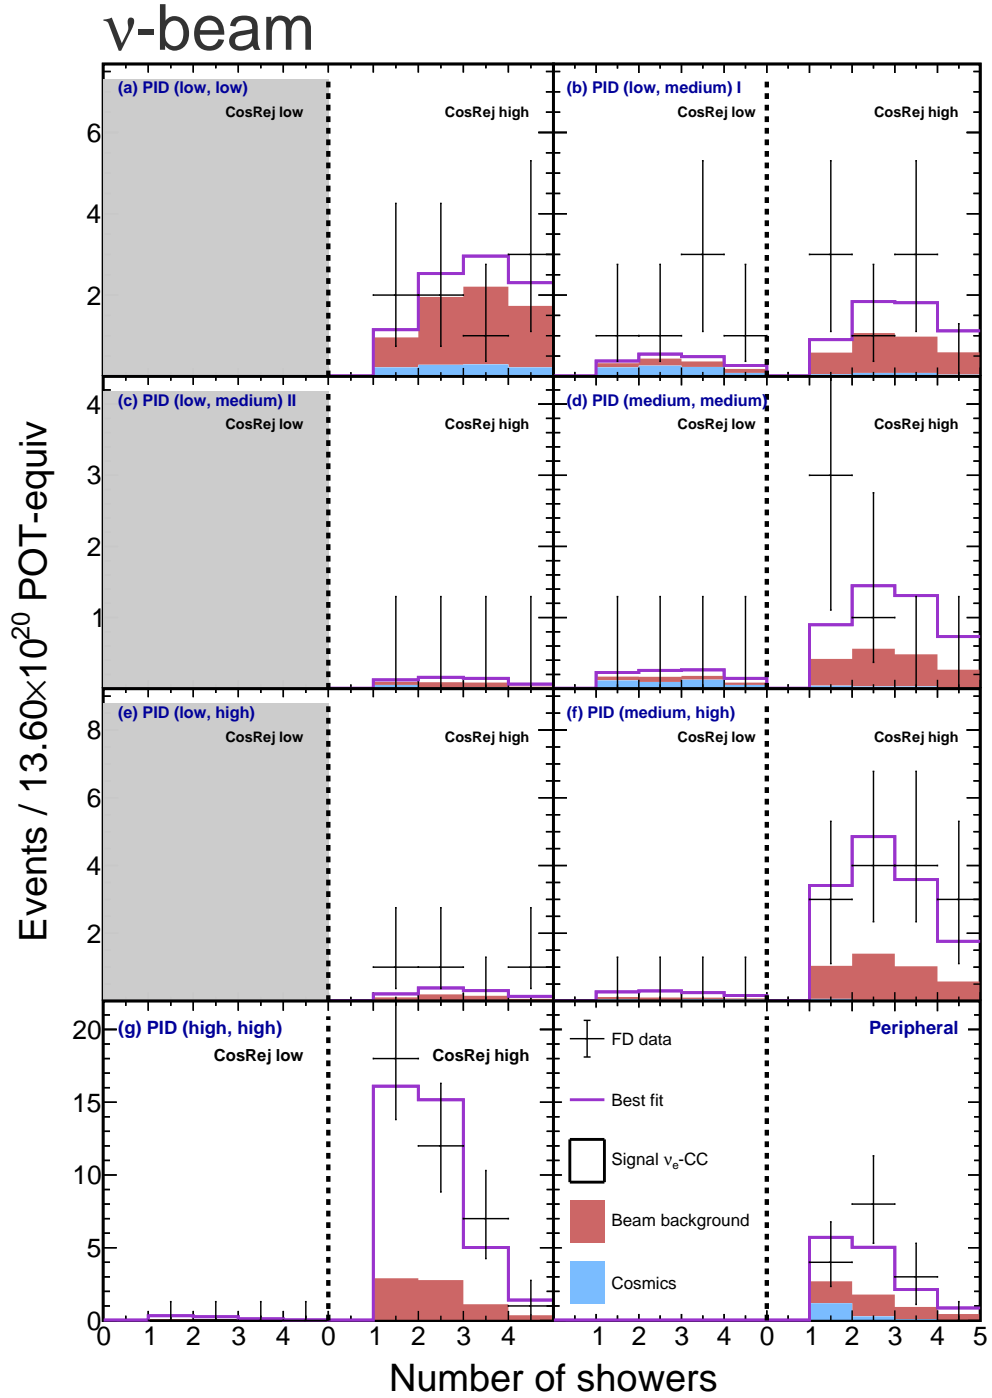


Figure C.37: Distributions of the number of reconstructed showers for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.38: Distributions of the number of reconstructed showers for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.39: Distributions of the reconstructed calorimetric energy of the primary shower for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

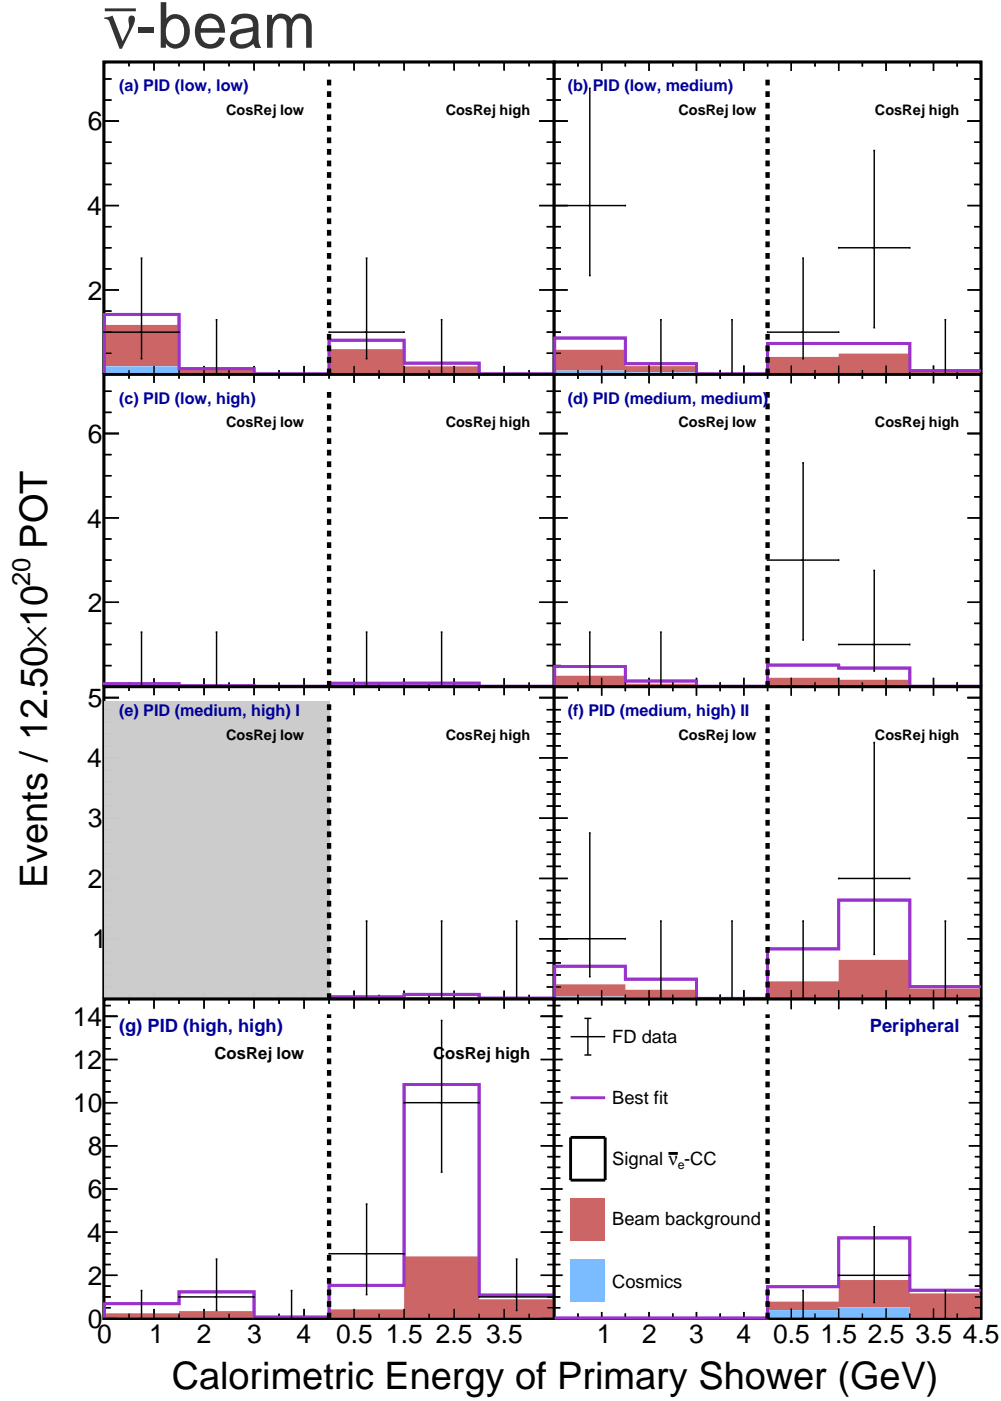


Figure C.40: Distributions of the reconstructed calorimetric energy of the primary shower for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

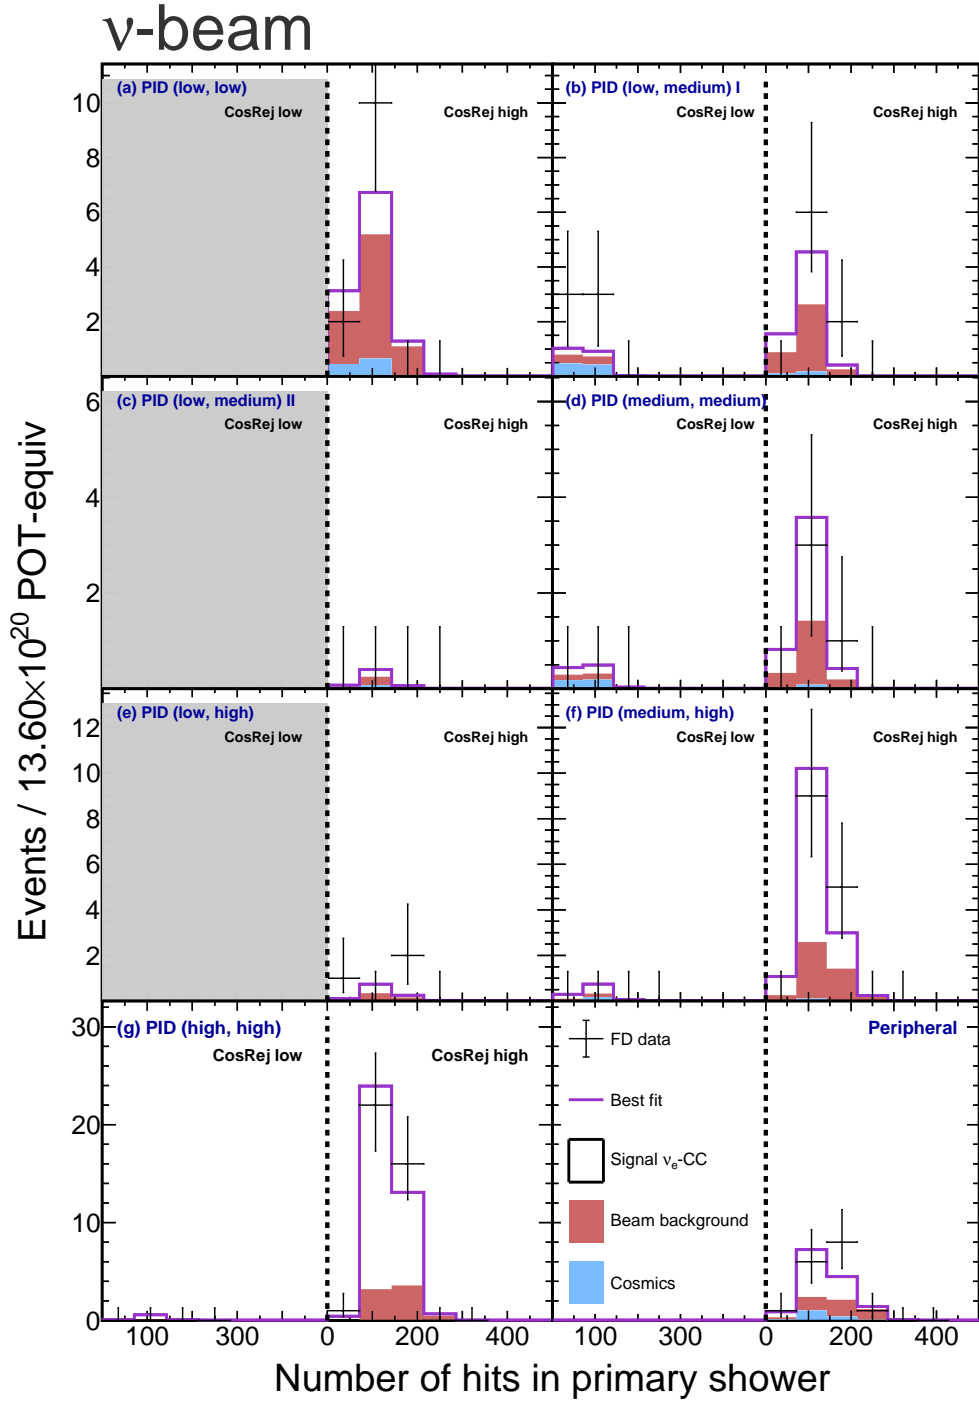


Figure C.41: Distributions of the number of hits in the primary shower for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.42: Distributions of the number of hits in the primary shower for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.

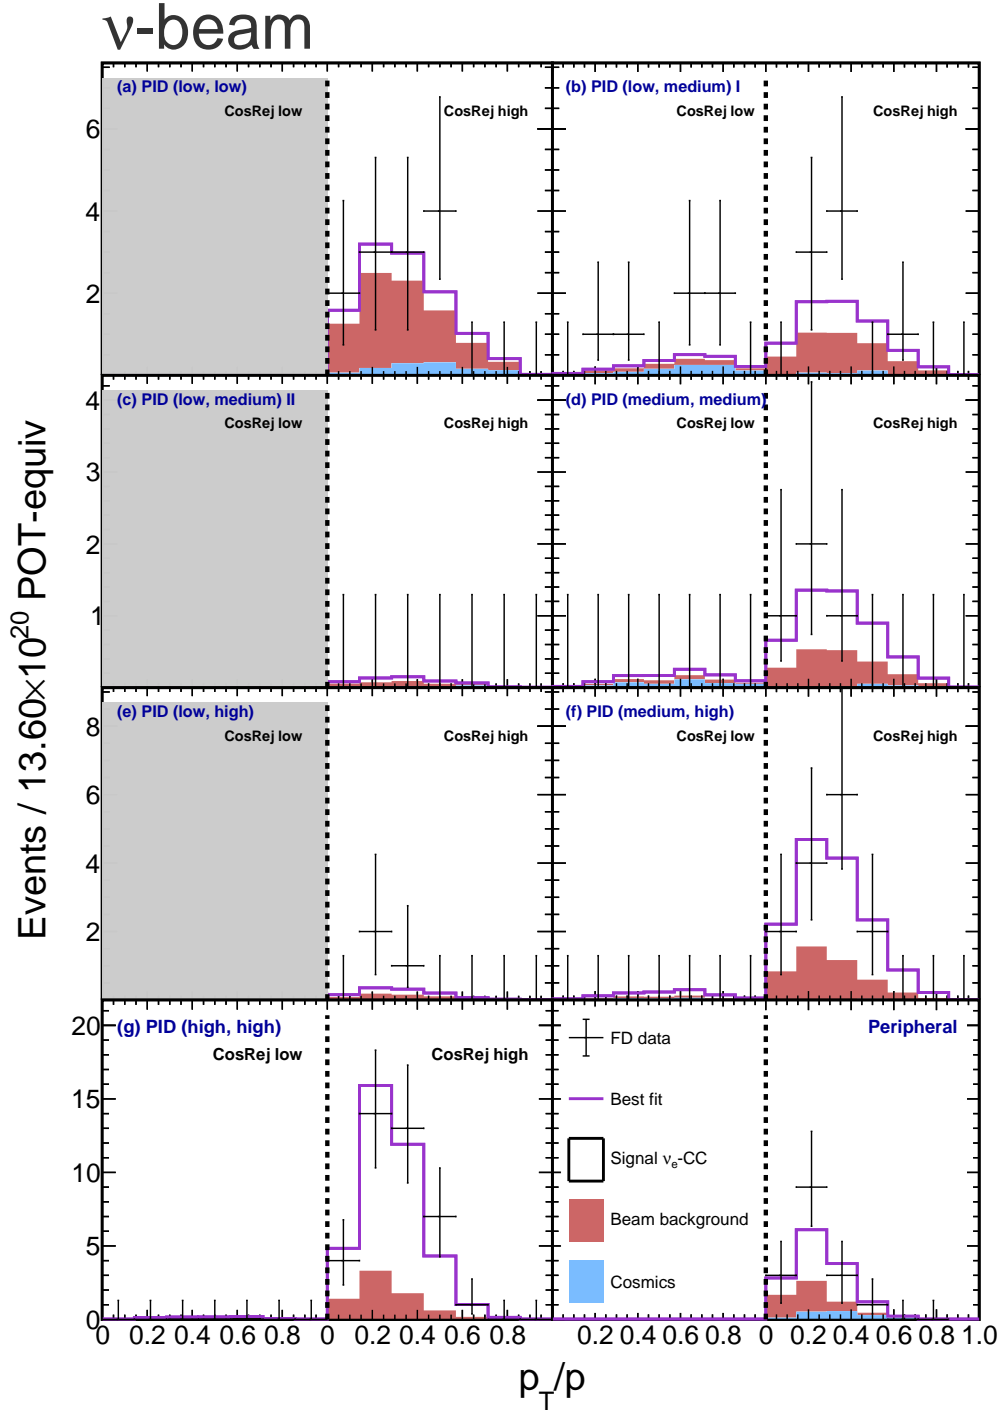


Figure C.43: Distributions of the reconstructed transverse momentum fraction of the primary shower for the core and peripheral samples at the FD in FHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. ν_e -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.




Figure C.44: Distributions of the reconstructed transverse momentum fraction of the primary shower for the core and peripheral samples respectively at the FD in RHC. The core sample is split into selection bins (a - g) and additionally by low and high cosmic rejection score. Lightly shaded regions are not included in the fit for oscillations. $\bar{\nu}_e$ -CC appearance candidates in data are shown in black. The purple histogram shows the total prediction, evaluated at the thesis best fit. The white, red and blue stacked histograms show the signal, beam background and cosmic background components respectively of the total prediction.