

Measurement of the muon neutrino charged-current single charged pion production cross-section in the NOvA Near Detector

Cathal Sweeney

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of
University College London.

Department of Physics and Astronomy
University College London

September 18, 2023

Declaration

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

Cathal Sweeney

Abstract

NOvA is a two-detector long-baseline neutrino oscillation experiment. It consists of a Near and Far Detector, which sit 1 km and 810 km, respectively, from the source of the NuMI neutrino beam. Both detectors sit 14 mrad off-axis, resulting in a beam of neutrinos with energy narrowly peaked around 2 GeV.

This thesis describes a cross-section analysis done at the Near Detector. A cross-section measurement was made for ν_μ charged-current interactions with a single charged pion in the final state. This measurement, the first of its kind in the NOvA Near Detector, is flux integrated and reported in bins of pion angle. The measurement was made using 13.8×10^{20} protons on target of neutrino beam, collected between 2014 and 2020. This measurement is sensitive to the transition region between the quasi-elastic and deep inelastic scattering regimes of neutrino-nucleus interactions.

Impact Statement

This thesis describes research undertaken as part of the NOvA collaboration. The work described in this thesis will aid both NOvA and future neutrino oscillation experiments in making more precise and accurate neutrino oscillation measurements. The primary motivation for this is to better understand how neutrinos behave. This is of interest for pure science reasons. Neutrinos are the most abundant matter particle in the universe. Studying them may tell us why the universe is composed of matter rather than anti-matter. Answering questions like this is culturally enriching for humankind.

Though not the primary motivation, many advancements in particle physics have benefited broader society. Previous examples include medical imaging techniques, developments in grid computing, and the World Wide Web.

Acknowledgements

Firstly, I would like to thank my parents, Anne and Harry, for supporting me and encouraging me in whatever I do. I would not be the person I am today without you. Thank you to my brothers Colm, Kevin, and Eoin for always keeping me entertained (and for keeping me humble). Thank you to Anabel for your endless support.

Sincere thanks to my supervisor Ryan Nichol for his insightful comments on my work and invaluable career guidance. This thesis would be far worse without his attention to detail. Thanks also to the UCL NuMI group and the other UCL HEP PhD students for being around to talk about work (or to avoid doing work).

I would like to thank the many conveners who guided my work during my PhD. Thanks to the Reconstruction and Deep Learning group conveners and the Detector Systematics group conveners. Special thanks go to the conveners of the Near Detector group who have helped me enormously: Jon Paley, Linda Cremonesi, Matt Muether, Leo Aliaga, Matt Judah, Matt Wetstein, and Greg Pawloski.

Finally, I would like to thank the many NOvA collaborators I have had the pleasure of working with. They have always been incredibly friendly, and I consider many of them friends.

Contributions

This thesis contains both original work completed during the PhD programme, and necessary background information on work done by others. The purpose of this section is to clearly outline which work is original.

Chapters 1-3 do not contain original work. These chapters give the reader the necessary background in neutrino physics and the NOvA experiment to understand the work in the subsequent chapters.

Chapter 4 consists entirely of original work done to improve reconstruction in the NOvA detectors. Other collaborators gave feedback on the work but did not contribute. All parts of Chapter 5 pertain to original work, except for Section 5.1 which describes the Geant4Reweight software package. Though I made some contributions to the Geant4Reweight project, most of the development was done by others. Chapter 6, which describes a cross-section analysis, consists entirely of original work. Chapter 7, which shows the results of this analysis, is also original work.

During the period of study, a qualification task was undertaken, which is not described elsewhere in this thesis. The qualification task was to work as release manager for one year. Responsibilities included ensuring that nightly builds of the NOvA software were successfully distributed, and creating tagged releases of the NOvA software. In addition, I led several upgrade tasks, including migrating the NOvA software to Scientific Linux 7.

Contents

Acronyms	26
1 The neutrino	28
1.1 The Standard Model	28
1.2 History	29
1.2.1 Genesis	30
1.2.2 Observing the neutrino	31
1.2.3 Oscillations	33
1.3 Oscillation physics	35
1.3.1 Oscillations in a vacuum	35
1.3.2 Oscillations in matter	38
1.4 Neutrino interactions	40
1.4.1 Interaction modes	41
1.4.2 Nuclear effects	45
1.5 Why measure cross-sections?	49
1.5.1 Cross-section definition	49
1.5.2 Inclusive measurements	50
1.5.3 Non-inclusive measurements	51
1.6 Summary	53
2 The NOvA experiment	54
2.1 The NuMI beam	54
2.1.1 Fermilab accelerator complex	55

2.1.2	Target and horns	55
2.1.3	Decay pipe and absorber	57
2.1.4	Off-axis design	57
2.1.5	Beam characteristics	58
2.2	The NOvA detectors	59
2.2.1	Scintillator	60
2.2.2	PVC cells	61
2.2.3	Wavelength-shifting fibre	61
2.2.4	Avalanche photodiode	61
2.2.5	Detector geometry	62
2.2.6	Near Detector	63
2.2.7	Far Detector	63
2.3	Data acquisition system	64
2.3.1	Triggering	65
2.4	Summary	65
3	NOvA offline software	66
3.1	Simulation	66
3.1.1	Flux	67
3.1.2	Event generation	67
3.1.3	Detector simulation	68
3.1.4	GENIE tuning	69
3.2	Reconstruction	70
3.2.1	Hit clustering	72
3.2.2	Vertexing	74
3.2.3	Prong formation	75
3.2.4	Particle identification	79
3.3	Calibration	80
3.3.1	Relative calibration	80
3.3.2	Absolute calibration	82
3.4	Summary	82

4	OPALProng: An improved prong view-matching algorithm	84
4.1	Motivation	84
4.2	Methods	85
4.2.1	Evaluating prong match quality	85
4.2.2	Finding a global solution	87
4.2.3	Determining number of hits to overlap	88
4.2.4	Sharing hits	89
4.3	Evaluation	92
4.3.1	Results	92
4.3.2	Event Displays	93
4.4	Summary	93
5	A systematic uncertainty on hadron scattering	104
5.1	Geant4Reweight	104
5.2	Implementation in NOvA	105
5.2.1	Generating weights	106
5.2.2	Fitting	107
5.3	Usage	108
5.3.1	Multiverse systematic	108
5.3.2	Varying parameters manually	108
5.4	Summary	108
6	Measurement of the ν_μ CC $1\pi^\pm$ cross-section	114
6.1	Background and motivation	114
6.1.1	Similar measurements	115
6.1.2	Analysis deliverable	116
6.1.3	How to measure a cross-section	117
6.2	Analysis signal definition	118
6.2.1	Pion kinematic thresholds	119
6.2.2	Muon kinematic phase space	121
6.3	Event selection	122

6.3.1	Background categories	124
6.3.2	Number of prongs	124
6.3.3	Fiducial volume	126
6.3.4	Containment	127
6.3.5	Particle identification	128
6.4	Resolution and binning	130
6.4.1	Pion angle	131
6.4.2	Muon kinematics	131
6.5	Systematic uncertainties	133
6.5.1	Cross-section modelling (GENIE)	133
6.5.2	Hadron production uncertainty (PPFX)	133
6.5.3	Beam transport	134
6.5.4	Hadron reinteraction (Geant4)	134
6.5.5	Calibration up/down	135
6.5.6	Calibration shape	135
6.5.7	Light level up/down	135
6.5.8	Cherenkov up/down	135
6.5.9	Detector ageing	135
6.6	Data-driven purity estimation	135
6.6.1	Template fit	135
6.6.2	Global template fit	138
6.6.3	Implementation	139
6.6.4	Fake data studies	142
6.7	Unfolding	148
6.8	Target counting	152
6.9	Flux estimation	154
6.10	Validation of analysis framework	155
6.10.1	In-out test	156
6.10.2	Poisson fluctuated fake data	156
6.10.3	Untuned GENIE fake data	160

6.11 Summary	160
7 Results and conclusions	161
7.1 Template fit	161
7.2 Differential cross-section measurement	164
7.3 Conclusions	165
Bibliography	167

List of Figures

1.1	The 17 elementary particles in the Standard Model. Figure from [1]. . . .	29
1.2	Continuous β -decay spectrum observed by Ellis and Wooster [2]. A single peak at the right-most edge of the graph is expected for a two-body decay. . . .	31
1.3	Design of the Reines and Cowan experiment, used for the first-ever detection of neutrinos. Figure from [3].	32
1.4	Average cross-section of the process $e^+e^- \rightarrow$ hadrons, as measured at LEP. Also overlaid is the theoretical prediction for two, three, and four generations of light neutrino. Figure from [4].	33
1.5	Zenith angle distributions observed in Super-Kamiokande for: (a) sub-GeV e-like, (b) sub-GeV μ -like, (c) multi-GeV e-like and (d) multi-GeV μ -like events. The value $\cos \Theta = 1$ means downward-going particles. The histograms with shaded error bars show the MC prediction with their statistical errors for the no neutrino oscillation case. The dotted histograms show the Monte Carlo prediction for $\nu_\mu \leftrightarrow \nu_\tau$ oscillation with a certain choice of oscillation parameters. Figure from [5].	34
1.6	Coherent forward scattering can occur for ν_e in matter via a Charged-current (left) and Neutral-current (right) interaction.	38
1.7	Representation of the two possible neutrino mass orderings. On the left, we have the normal ordering ($\Delta m_{atm}^2 > 0$). Shown on the right is the inverted ordering ($\Delta m_{atm}^2 < 0$). The colours represent the contribution of ν_e (red), ν_μ (blue), and ν_τ (green) in a given a mass state. Here the labels Δm_{atm}^2 and Δm_{sol}^2 refer to Δm_{32}^2 and Δm_{21}^2 , respectively. Figure from [6].	39

-
- 1.8 Charged-current (left) and Neutral-current (right) neutrino-nucleus interactions on a down quark. 40
- 1.9 Neutrino-nucleus interaction cross-section divided by neutrino energy. Shown is the total cross-section as well as the cross-section broken down by modes - quasi-elastic (Quasi-elastic (QE)), resonant (Resonant (RES)), and deep inelastic scattering (Deep inelastic scattering (DIS)). This figure is over ten years old at the time of writing, so the latest model predictions may differ somewhat. Also, some rarer interaction modes are not shown. Nonetheless, this figure gives a sense of the relevant energy scales. Figure from [7]. 42
- 1.10 Diagram of a QE neutrino interaction. For the antineutrino mode, we instead have $p \rightarrow n$ 42
- 1.11 Diagram of resonant single pion production in a Charged-current interaction on a proton. 43
- 1.12 Diagram of a meson exchange current event, which knocks out two particles (nucleons) and leaves behind two holes. The contact current (left) and pion-in-flight current (right) are shown. 43
- 1.13 Flux-unfolded MiniBooNE ν_μ CCQE cross-section per neutron as a function of neutrino energy, along with results from the LSND and NOMAD experiments. Also shown are predictions from the NUANCE [8] simulation for an RFG model with two different parameter variations and for scattering from free nucleons with the world-average M_A value. Figure adapted from [9]. A portion of the original figure, which is not relevant to the current discussion, has been cropped out. 44
- 1.14 Diagram of coherent pion production. 45
- 1.15 Diagram of a Charged-current (CC) Deep Inelastic Scattering interaction . 45

- 1.16 Nuclear potential well for protons and neutrons. Also marked are the neutron and proton Fermi energies E_F^n and E_F^p , respectively. E_B is the nuclear binding energy. Note the proton potential well is different from the neutron potential well, as protons are subject to electromagnetic forces. Figure from [10]. 46
- 1.17 Comparison of the Fermi momentum for global and local Fermi gas in the case of carbon. Carbon is not chosen for any particular reason, rather this figure serves to generically highlight the difference between global and local Fermi gases. Figure from [10]. 47
- 1.18 Ratio of GENIE simulations with and without RPA corrections included. The blue curve was made using a weighting just in Q^2 , whereas the black curve was made using weights in energy and three-momentum transfer space. Figure from [11]. 48
- 1.19 Example of FSI manifesting through pion absorption. 48
- 1.20 Impact on the results if a different generator is used to compute the true and fitted rates in the analysis. The shaded areas show the confidence regions at 1, 2, and 3σ that would be obtained in the $\theta_{23} - \Delta m_{31}^2$ plane if the true and fitted rates are generated using the same set of migration matrices (obtained from GiBUU, with oxygen as the target nucleus). The coloured lines show the same confidence regions if the true rates are generated using matrices produced with GiBUU, but the fitted rates are computed using matrices produced with GENIE. Both sets of matrices are generated using oxygen as the target nucleus. The red dot indicates the true input value, while the black triangle shows the location of the best fit point. The value of the χ^2 at the best fit is also shown, together with the number of degrees of freedom. Figure adapted from [12]. A portion of the original figure, which is not relevant to the current discussion, has been cropped out. . . . 52
- 2.1 An illustration of the various components of the NuMI beamline. Figure from [13]. 55

2.2	A diagram of the Fermilab accelerator program, showing the various beams provided to the end-user experiments. Figure from [14].	56
2.3	A diagram of the magnetic focusing horns, with possible hadron trajectories drawn. Figure from [13].	57
2.4	Neutrino energy as a function of parent pion energy for various off-axis angles	58
2.5	NuMI flux at 14 mrad off-axis for (left) Forward Horn Current (FHC) and (right) Reverse Horn Current (RHC).	58
2.6	Drawing of the Near and Far Detectors next to each other for size comparison. Figure from [15].	59
2.7	Schematic of the interior of NOvA cell, with the path of an ionising particle through the scintillator overlaid. Scintillation light, marked in purple, is collected by the wavelength-shifting fibres in green. Figure from [15]. . .	60
2.8	Schematic of a module consisting of two 16 cell extrusions glued together. The module is sealed on one end with an end plate and on the other with a manifold cover. Figure from [16].	62
2.9	Schematic of a NOvA detector showing alternating horizontal and vertical planes. The vertical planes give a view of the event as viewed from above, and the horizontal planes give a view of the event as viewed from the side. Figure from [17].	63
2.10	A set of four consecutive ADC readings marked with black circles. A fit is performed to these readings, marked in red.	64
3.1	Simulated neutrino interaction at the Far Detector with cosmic ray data overlaid. The simulated particles have truth vectors drawn. Figure from [18].	68
3.2	Weights applied to Meson exchange current (MEC) events during tuning, as a function of energy and momentum transfer. Figure from [19].	69
3.3	Distribution of MEC events in simulation (left) before tuning and (right) after tuning. Figures from [19].	70
3.4	Event displays with true particle labels overlaid. We see the characteristic energy depositions of different particles in the NOvA detectors.	71

- 3.5 (Left) A scatter plot of data to be clustered. Numbers correspond to the density ordering. (Right) A plot of density vs isolation. We see that hits 1 and 10 have a high density and isolation, and hence are identified as centroids. Figure from [20]. 72
- 3.6 (Left) Collection of hits with straight lines overlaid indicating the output of the Hough transform. (Right) Hough parameter space plot with two hotspots visible, showing the location and orientation of the prominent lines in the event. Figure from [21]. 74
- 3.7 A simulated ν_e CC quasielastic interaction in the Far Detector with completed 3D prong reconstruction from the fuzzy k-means algorithm. The prong outlined in red in each view is the reconstructed electron, and outlined in green is the reconstructed proton. The corresponding energy profile histograms used to compute the suitable 3D prong matches is shown in Figure 3.8. Figure and caption from [22]. 77
- 3.8 Cumulative energy profile as a function of path length along a prong for prospective 3D match candidates shown in Figure 3.7. The red curves are for prongs in the XZ (vertical planes) view, and the blue is for the YZ (horizontal planes) view. The upper-left and lower-right panels show the preferred matches with similar energy profiles that result in the green and red tracks, respectively. The off-diagonal elements illustrate the difference in energy profile shape for the wrong combinations. Figure and caption from [22]. 78
- 3.9 Schematic of a tri-cell hit, with the cell width d marked. A particle trajectory is overlaid, with the path length through the cell l marked. . . . 81
- 3.10 PE/cm vs W for a single cell at the Near Detector. A fitted function of the form in Equation 3.17 is shown in red. An additional fit, shown in blue, is done to account for the “roll-off” near the edges. 82

- 4.1 **Left:** An event display where two particle tracks are visible in the x view, but the tracks overlap in the y view. **Centre:** Example reconstruction under the standard view-matching algorithm. One prong is formed (orange), but we are forced to orphan one 2D prong. **Right:** Reconstruction possible with OPALProng. Both particles are reconstructed. Hits in the y view are shared between the orange and blue prongs. 85
- 4.2 **Left:** Cumulative energy fraction of a pair of x and y prongs. The Kuiper score is low, so this is a good match **Right:** Cumulative energy fraction of a different pair of x and y prongs. The Kuiper score is high, so this is a bad match. 86
- 4.3 The four possible solutions when matching two x prongs with one y prong. Each table is one solution; boxes in green represent matches made 87
- 4.4 An event display of an interaction which will be reconstructed with overlapping prongs. It is necessary to determine how many hits in the y view will be included when making a 3D prong with the short 2D prong marked X_1 89
- 4.5 Toy example showing a hit in the y view (circled in red), which must be shared between two 3D prongs α and β . The upper prong in the x view deposits 50 MeV of calorimetric energy in both planes on either side of the hit to be weighted. The lower prong deposits 100 MeV in both of these planes. With reference to Equation 4.4 and Equation 4.6 we see that this hit will be shared with weights $w_{n,\alpha} = 0.33$ and $w_{n,\beta} = 0.67$ 90
- 4.6 Prong purity in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms. 94
- 4.7 Prong completeness in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms. 95
- 4.8 Prong calorimetric energy in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms. 96
- 4.9 Prong length in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms. 97

4.10	Cosine of reconstructed prong angle w.r.t. the beam in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms.	98
4.11	Number of hits in prong in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms.	99
4.12	The difference between calorimetric energy calculated under the preexisting and OPAL reconstructions plotted vs true energy for a sample of muon prongs.	100
4.13	Event display showing 3D prongs made using regular reconstruction (top) and OPAL reconstruction (bottom). In the y view of the OPAL-reconstructed event an extra 3D prong can be seen, marked in blue. This prong is not visible in the x view as it is drawn behind the prong marked in pink.	101
4.14	Event display showing 3D prongs made using regular reconstruction (top) and OPAL reconstruction (bottom). In the x view of the OPAL-reconstructed event an extra 3D prong can be seen, marked in blue. This prong is not visible in the y view as it is drawn behind the prong marked in pink.	102
4.15	Event display showing 3D prongs made using regular reconstruction (top) and OPAL reconstruction (bottom). In the x view of the OPAL-reconstructed event an extra 3D prong can be seen, marked in blue. This prong is not visible in the y view as it is drawn behind the prong marked in pink.	103
5.1	Inclusive and exclusive cross-sections for π^+ inelastic scattering in carbon. Pre- and post-fit predictions are shown, with fit data overlaid. No data was found for double charge exchange or pion production at the time of writing.	109
5.2	Inclusive and exclusive cross-sections for π^- inelastic scattering in carbon. Pre- and post-fit predictions are shown with fit data overlaid. At the time of writing, no data was found for double charge exchange, pion production, or pion absorption.	110

- 5.3 Inclusive cross-section for neutron inelastic scattering in carbon. Shown in blue is the pre-fit cross-section, and the post-fit cross-section with the error band is shown in red. Fit data is overlaid. 111
- 5.4 Results of fit for π^+ parameters. The best fit values and uncertainties are shown on the left, and the fit covariance matrix is shown on the right. . . . 111
- 5.5 Results of fit for π^- parameters. The best fit values and uncertainties are shown on the left, and the fit covariance matrix is shown on the right. . . . 112
- 5.6 Result of fit for the neutron parameter. The best fit parameter value and uncertainty is shown. 112
- 5.7 Simulated events with a multiverse error band drawn. The only cut applied here is that the true neutrino interaction vertex is inside the active detector limits 112
- 5.8 Simulated prong length histograms with shifts applied to proton and piplus parameters separately. 113
- 6.1 Diagram of resonant single pion production in a CC interaction on a proton. 114
- 6.2 Charged-current neutrino interaction cross section as a function of energy. Figure from [7]. 115
- 6.3 MINERvA measurement of CC $1\pi^\pm d\sigma/d\theta_\pi$ compared to the GENIE, ACS, NEUT, and NuWro models. The inner (outer) error bars correspond to the statistical (total) uncertainties. Figure adapted from [23]. A portion of the original figure, which is not relevant to the current discussion, has been cropped out. 116
- 6.4 (Left) Plot of (reconstructed - true) pion energy vs true pion energy. Here calorimetric energy in the pion prong is used to estimate pion energy. (Right) Bias of pion energy estimation. 117
- 6.5 The number of signal events in simulation, broken down by interaction mode. (Left) A ν_μ CC inclusive signal definition, containing 1.1×10^7 events, of which 39 % are resonant interactions. (Right) A ν_μ CC $1\pi^\pm$ signal definition, containing 1.1×10^6 events, of which 46 % are resonant interactions. 118

- 6.6 The number of pions in the final state in a sample of $\nu_{\mu}\text{CC}$ quasi-elastic interactions. Shown in black is the sample of all $\nu_{\mu}\text{CC}$ quasi-elastic interactions in the fiducial volume outlined in Section 6.3.3. For the histogram in red we further require that the muon and pions be within the kinematic thresholds outlined in Sections 6.2.2 and 6.2.1. 119
- 6.7 Event display of a simulated signal event. 120
- 6.8 The fraction of pions which form a prong as a function of pion kinetic energy (left) and angle (right). The signal thresholds are shown with a dashed line. To make these plots, a simple $\nu_{\mu}\text{CC}1\pi^{\pm}$ signal definition is used, with no kinematic thresholds. Fiducial, containment, and quality selection cuts are applied. 120
- 6.9 Fractional uncertainty on the cross-section in bins of muon kinematics. The muon phase space cut is indicated with a solid black border. The full signal and selection cuts are used to make this plot, excluding the muon phase space cut. 121
- 6.10 The number of hits in true pion prongs which are below (red) and above (blue) the pion energy threshold in the signal definition. To make these plots, a simple $\nu_{\mu}\text{CC}1\pi^{\pm}$ signal definition is used, with no kinematic thresholds. Fiducial, containment, and quality selection cuts are applied. . 125
- 6.11 Fractional uncertainty on the integrated cross-section as a function of vertex x , y , and z components. To make this plot, a simple $\nu_{\mu}\text{CC}1\pi^{\pm}$ signal definition is used, with no kinematic thresholds and with true interaction vertex inside the detector. The selection cuts applied are an unoptimised muon identification cut, reconstructed vertex inside the detector, containment, and quality cuts. 126
- 6.12 Muon energy resolution as a function of distance from the detector edge. The containment limits are marked with a red dashed line. These plots were made with fiducial, quality, and unoptimised muon ID cuts applied. When optimising, e.g. prong x containment limits, unoptimised containment cuts are applied in the y and z coordinates. 127

6.13	Fractional uncertainty on the integrated cross-section as a function of cut values. For muon and pion ID, we determine a minimum CVN score. For pion rejection, we determine a maximum CVN score. The cut values are marked with a dashed line. To make these plots, preselection cuts are applied. Here, the signal is $\nu_{\mu}CC1\pi^{\pm}$, with simple muon kinematic thresholds.	129
6.14	Resolution (left) and bias (right) of reconstructed pion angle, as a function of true pion angle.	131
6.15	Resolution (left) and bias (right) of $\cos \theta_{\mu}$, as a function of true $\cos \theta_{\mu}$. . .	131
6.16	Resolution (left) and bias (right) of reconstructed muon kinetic energy, as a function of true muon kinetic energy.	132
6.17	Example templates with fake data overlaid. (Left) pre-fit templates. (Right) post-fit templates	136
6.18	Concatenated fake data and pre-fit Monte Carlo as used in the global template fit. The x-labels “Bin 1” etc., are analysis bin labels (in this case, they are pion angle bins). i.e. Bin 1: $(0 < \theta_{\pi} < 6)$, Bin 2: $(6 < \theta_{\pi} < 12)$, etc. The fake data was made from Monte Carlo simulation (MC) by artificially increasing (decreasing) the amount of signal (background). Hence there is an excess (deficit) of events in the signal (background) rich regions.	138
6.19	Selected background event (pion is below kinetic energy threshold in signal definition). We see that the prong identified by prong CVN as the pion prong is, in fact, a proton.	139
6.20	Training sample input variables to BDT.	140
6.21	Output of BDT evaluated on the testing sample.	141
6.22	Total covariance (left) and correlation matrix (right).	142
6.23	Each systematic component of the covariance matrix.	143
6.24	Each systematic component of the correlation matrix.	143
6.25	From a single fake data universe, pre-fit (left) and post-fit (right) templates with fake data overlaid. Shown here are the first five analysis bins.	145

6.26	From a single fake data universe, pre-fit (left) and post-fit (right) templates with fake data overlaid. Shown here are the last five analysis bins.	146
6.27	Pre- and post-fit purity estimates from a single fake data universe. The calculated χ^2 comparing true and pre-fit purity is 39.8, whereas, for true and post-fit purity, the χ^2 is 3.7, with 10 degrees of freedom.	147
6.28	Comparison of pre-fit (left) and post-fit (right) purity estimates on an ensemble of 200 fake data universes, with scale factors of 1.2 and 0.8 applied to signal and background templates, respectively. (Top) Mean and RMS of data-MC purity ratio from each universe. (Bottom) Mean and RMS of fractional uncertainty on purity from each universe.	147
6.29	Comparison of pre-fit (left) and post-fit (right) purity estimates on an ensemble of 200 fake data universes, each made from a different systematically shifted GENIE universe. (Top) Mean and RMS of data-MC purity ratio from each universe. (Bottom) Mean and RMS of fractional uncertainty on purity from each universe.	148
6.30	Pion angle smearing matrix. Note that the overflow is included, which appears as a large bin in the top right of the figure.	149
6.31	Average MSE as a function of the number of unfolding iterations.	150
6.32	An ensemble of unfolded universes after zero iterations of unfolding (left) and two iterations of unfolding (right).	150
6.33	Average MSE as a function of the number of unfolding iterations. Here fake data is made by applying random shifts to the GENIE cross-section modelling parameters, but the matrix used for unfolding is made from the nominal simulation.	151
6.34	Closure test of unfolding.	152
6.35	Predicted ν_μ flux with associated systematic error bars.	154
6.36	Fractional uncertainty on flux, broken down by source of systematic uncertainty	155
6.37	Cross-section calculated using Asimov fake data and full uncertainty calculation. Overlaid is the predicted true cross-section from the simulation.	156

6.38	Cross-section calculation using fake data with Poisson fluctuations applied. Error bars include all statistical and systematic uncertainties. Overlaid is the predicted true cross-section from the simulation.	157
6.39	Fractional uncertainty on the cross-section measurement, broken down by sources of uncertainty.	158
6.40	Total covariance (left) and correlation (right) matrices for the fake data cross-section measurement.	158
6.41	Covariance matrix for each source of uncertainty.	159
6.42	Cross-section calculation using fake data made from untuned GENIE simulation. Error bars include all statistical and systematic uncertainties. Overlaid is the predicted true cross-section from tuned and untuned simulation.	160
7.1	Pre-fit (left) and post-fit (right) templates with data overlaid. Error bars represent statistical uncertainty. Shown here are the first five analysis bins.	162
7.2	Pre-fit (left) and post-fit (right) templates with data overlaid. Error bars represent statistical uncertainty. Shown here are the last five analysis bins.	163
7.3	Pre and post-fit estimates of purity. Error bars include both statistical and systematic contributions.	164
7.4	Differential cross-section measurement with the NOvA GENIE prediction overlaid. Error bars represent uncertainty from unfolding, statistical, and systematic sources.	165
7.5	Total covariance (left) and correlation (right) matrices for the cross-section measurement.	165

List of Tables

2.1	Composition of the NOvA scintillator.	60
4.1	Kuiper score table for an example with two prongs in the x view and one prong in the y view. For every combination of prongs in one view paired with prongs in the other view, a Kuiper score is calculated. Some (nonsensical) combinations are never considered and hence are greyed out. The NULL row and column are included to allow prongs to be orphaned.	87
4.2	With reference to Figures 4.6-4.11, ratio of the number of prongs formed under the OPAL and regular view-matching algorithms	92
5.1	Table of exclusive inelastic scattering channels for pions in Geant4Reweight. Here X represents any number of additional non-pion particles present in the final state	106
5.2	χ^2 per degree of freedom determined from fits to scattering data.	107
6.1	Existing ν_μ CC $1\pi^\pm$ cross-section measurements.	116
6.2	Cutflow table showing purity, efficiency, and relative efficiency at each stage in the selection.	123
6.3	The number of events from simulation at each step in the selection, broken down by signal and multiple background samples. Also shown is the fraction of the total number of events in each sample at each step in the selection.	123
6.4	All systematic uncertainties used in this analysis. Also listed is whether each is a reweight or file-based uncertainty.	133
6.5	Composition of the fiducial volume, broken down by material.	153

6.6	Composition of the fiducial volume, broken down by element.	153
-----	---	-----

Acronyms

ADC Analogue-to-digital convertor

APD Avalanche photodiode

CC Charged-current

CVN Convolutional Visual Network

DAQ Data acquisition system

DCM Data concentrator module

DCS Dual correlated sampling

DDT Data-driven trigger

DIS Deep inelastic scattering

FD Far Detector

FEB Front end board

FHC Forward Horn Current

FPGA Field-programmable gate array

FSI Final state interactions

MC Monte Carlo simulation

MEC Meson exchange current

NC Neutral-current

ND Near Detector

NOvA NuMI Off-axis ν_e Appearance experiment

NuMI Neutrinos at the Main Injector

PE Photoelectrons

POT Protons on target

PPFX Package to Predict the Flux

PVC Polyvinyl chloride

QE Quasi-elastic

RES Resonant

RHC Reverse Horn Current

RPA Random phase approximation

Chapter 1

The neutrino

The neutrino is the most abundant matter particle in the universe. The earth is constantly bombarded with neutrinos from the sun. An outstretched human hand would expect about 10 trillion solar neutrinos to pass through it every second. However, they interact so rarely that only approximately one of these solar neutrinos interact inside the human body per day¹. This small interaction probability means they remain among the most challenging standard model particles to study.

In this chapter, I will give a brief overview of the Standard Model, followed by an outline of the neutrino's theoretical and experimental history. I will also describe neutrino oscillation physics and neutrino-nucleus interactions.

1.1 The Standard Model

The Standard Model is one of the most successful theories in all of science. It is a quantum field theory which describes particles and their interactions, and has symmetries given by the gauge group $SU(3) \times SU(2) \times U(1)$. The Standard Model contains 12 matter particles (fermions), 4 force carriers (gauge bosons), and a Higgs boson, as shown in Figure 1.1.

Fermions have half-integer spin. They can be further split into quarks and leptons. Quarks have electromagnetic charge and weak isospin. As a result, they feel both the electromagnetic and weak forces. Quarks also have colour charge, so they are subject to the strong force. In nature, quarks are always found in colourless combinations. Quarks

¹Very roughly, $(10^{11} \nu/\text{cm}^2/\text{s solar flux}) \times (10^{-44} \text{cm}^2 \text{ cross-section}) \times (10^{28} \text{ nucleons in the body}) \times (10^5 \text{ s in a day}) = 1 \text{ interaction per day}$

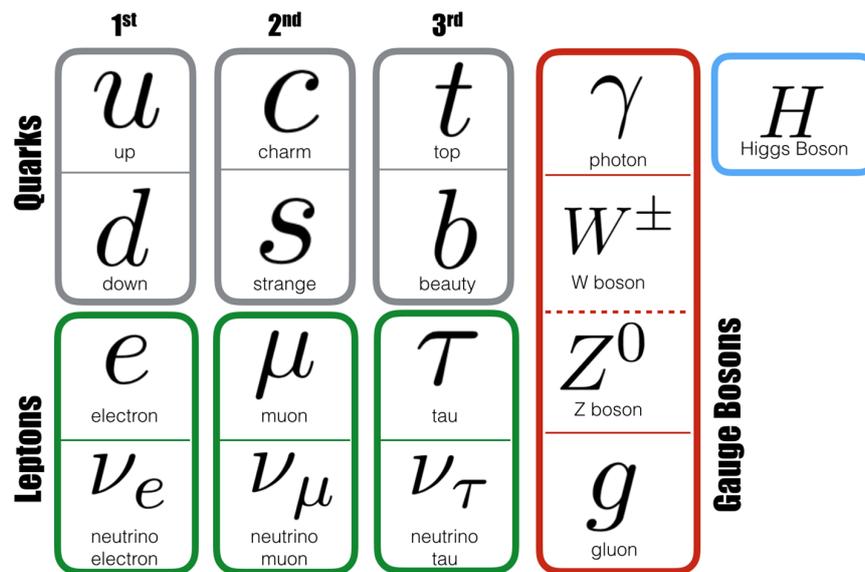


Figure 1.1: The 17 elementary particles in the Standard Model. Figure from [1].

are divided into three generations - up and down, charm and strange, and top and bottom. Atomic matter is mainly comprised of up and down quarks. Leptons have weak isospin but no colour charge. Three leptons have an electromagnetic charge, the electron, muon, and tau. The remaining three leptons are the neutrinos, which come in electron, muon, and tau flavours, and are electrically neutral.

The gauge bosons are the force carriers of the Standard Model. The photon carries the electromagnetic force, the W and Z bosons carry the weak force, and the gluon carries the strong nuclear force. All gauge bosons have spin 1.

The Higgs boson is the only scalar boson (spin 0) in the Standard Model. It is an excitation of the Higgs field, which gives mass to the fundamental particles.

Neutrinos were predicted in the Standard Model to be massless. Experimental evidence of neutrino oscillation shows that the neutrino mass must be non-zero. Neutrino oscillation is the most significant discrepancy from Standard Model prediction observed to date. Hence neutrino physics is an area of intense research today.

1.2 History

This section will outline the theoretical and experimental history of the neutrino - from first its proposal to the early days of neutrino detection and later the study of its properties.

1.2.1 Genesis

In the early 20th century, β -decay was believed to be a two-body decay, i.e.

$$N \rightarrow N' + e^- \quad (1.1)$$

As such, it was expected that the emitted electrons would be mono-energetic. We can see this if we consider the four-momenta of the initial state $P = (M, \vec{0})$, and the final state particles p_1 and p_2 , where $P = p_1 + p_2$. Let us call p_2 the four-momentum of the electron. We have

$$\begin{aligned} p_1^2 &= (P - p_2)^2 \\ m_1^2 &= M^2 - 2ME_2 + m_2^2 \\ E_2 &= \frac{M^2 + m_2^2 - m_1^2}{2M} \end{aligned}$$

Assuming fixed masses, the outgoing electron energy E_2 is fixed. In 1914 James Chadwick observed a continuous energy spectrum from β -decay [24]. Given the assumptions of the time, this implied a violation of the law of conservation of energy². Figure 1.2 shows an example of this continuous energy spectrum.

At a 1930 meeting in Gauverrein, a letter written by Wolfgang Pauli, who could not make it in person to the meeting, was read to the attendees [25]. In this letter, he proposed a “desperate remedy” to the β -decay problem. He hypothesised that an electrically neutral, spin 1/2 particle existed inside the nucleus, which he named the neutron. This new particle could therefore carry away the “missing energy” observed in β -decay experiments.

Just a few years later, in 1934, Enrico Fermi published his theory of four-fermion β -decay [26], in which he posited that the new mystery particle, which he named the neutrino³, was created during the decay, rather than residing inside the nucleus. This gave a theoretical framework to explain the energy spectrum of β -decay. However, Bethe and Peierls were quick to point out that the neutrino would have a penetrating power of 10^{16}

²The decay also indicated non-conservation of angular momentum, but vanishing energy is already troubling enough.

³The name neutron had since been taken by the nucleon we know and love today.

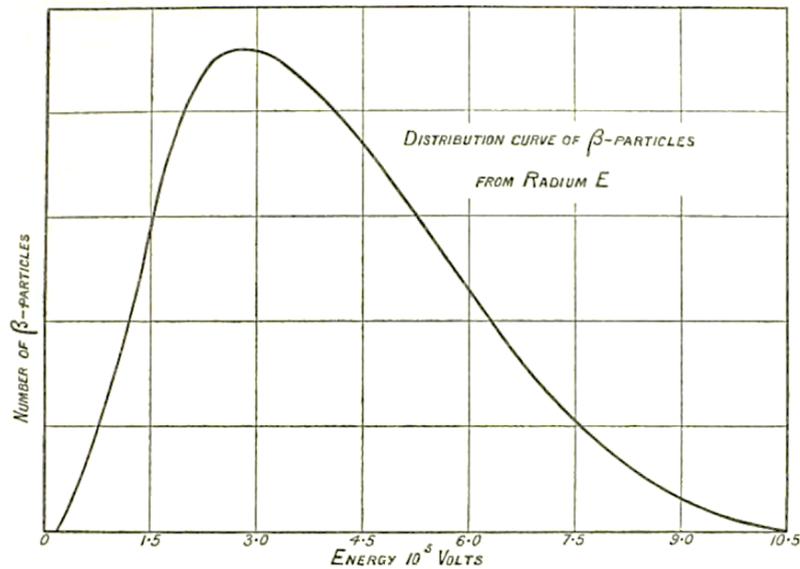


Figure 1.2: Continuous β -decay spectrum observed by Ellis and Wooster [2]. A single peak at the right-most edge of the graph is expected for a two-body decay.

km in solid matter [27]. Hence they concluded that there was “no practically possible way of observing the neutrino”.

1.2.2 Observing the neutrino

Luis Alvarez was among the first to propose that neutrinos could be detected⁴ [29]. The proposal was to bombard a tank containing chlorine with an intense neutrino beam (perhaps from a nuclear reactor). Chlorine atoms could be transmuted by converting neutrons to protons



After irradiation, the ${}^{37}\text{Ar}$ would be separated, and one could count the number of argon atoms and therefore infer the number of neutrino interactions.

Fred Reines and Clyde Cowan were the first to conclusively observe the neutrino in 1956 [30]. Their experiment, located underground near the Savannah River Plant reactor, relied on observing the coincidence of photon emission from prompt positron annihilation

⁴Alvarez’s work was based on an earlier lecture by Pontecorvo which the US government quickly classified [28].

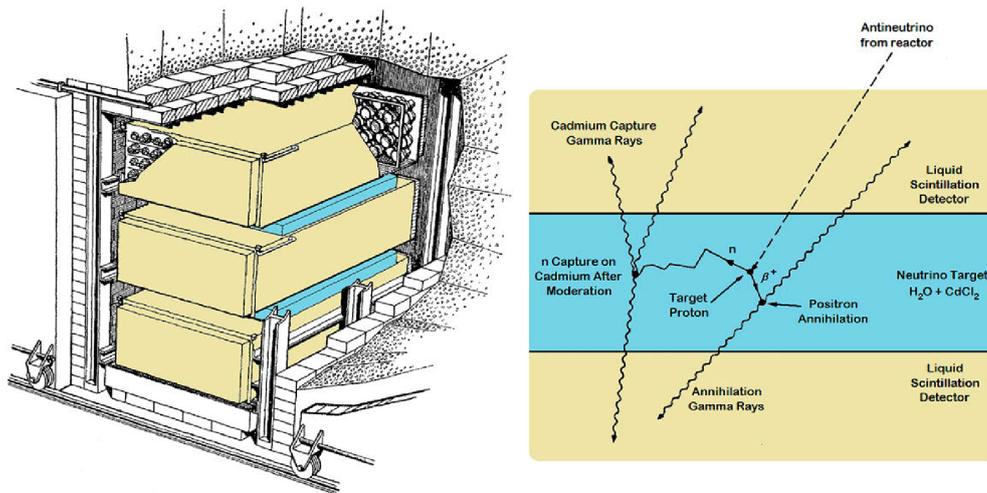


Figure 1.3: Design of the Reines and Cowan experiment, used for the first-ever detection of neutrinos. Figure from [3].

followed by delayed neutron capture in cadmium-doped water, produced via the reaction



The experiment consisted of a multi-layered “sandwich” of water detectors doped with cadmium (to capture neutrons) between liquid scintillator tanks viewed by photomultiplier tubes, as shown in Figure 1.3. The entire detector was encased in lead to reduce backgrounds and employed veto detectors outside the shielding. For this work, Reines was awarded the 1995 Nobel Prize in Physics ⁵.

A 1962 experiment by Danby *et al.* at the AGS in Brookhaven was the first to show that there was more than one type of neutrino [31]. This experiment also pioneered the accelerator neutrino beam method still used today (see Section 2.1). Using a high-energy proton beam incident on a target, pions and kaons were produced, which decayed into muons and neutrinos. These particles impinged on a 13.5 m iron shield, which absorbed the muons, leaving a beam of neutrinos. It was observed that when these beam neutrinos later interacted, they only ever produced muons in the detector. They concluded that the neutrinos they were producing were muon-flavoured.

Following the 1975 discovery of the tau lepton [32], it was anticipated that there would be a tau-flavoured neutrino. The DONUT experiment at Fermilab first observed this

⁵Cowan had died in 1974.

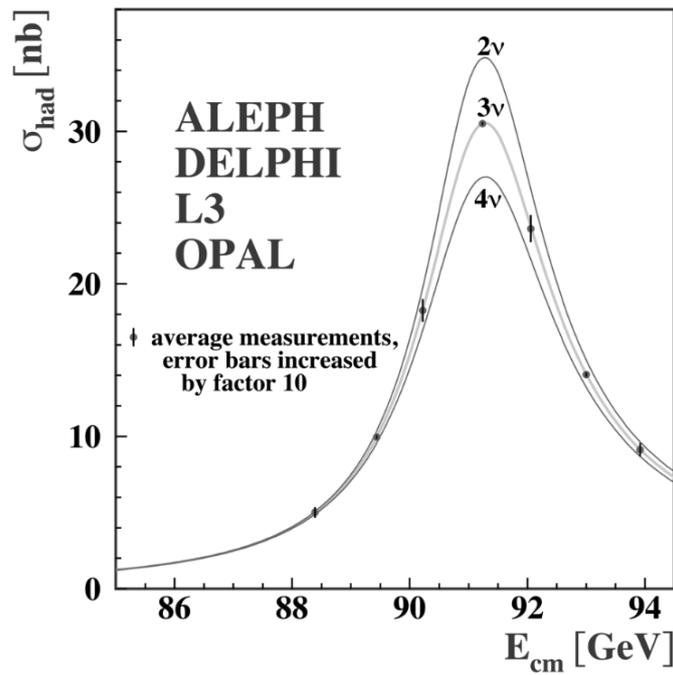


Figure 1.4: Average cross-section of the process $e^+e^- \rightarrow \text{hadrons}$, as measured at LEP. Also overlaid is the theoretical prediction for two, three, and four generations of light neutrino. Figure from [4].

third neutrino flavour in 2001 [33]. Measurements of the $e^+e^- \rightarrow \text{hadrons}$ cross-section by the OPAL, ALEPH, L3, and DELPHI experiments at LEP put a combined limit on the number of light neutrinos ($m_\nu < m_Z/2$) at 2.9840 ± 0.0082 [4] (see Figure 1.4).

1.2.3 Oscillations

Ray Davis designed an experiment in the Homestake mine in Lead, South Dakota to measure the flux of solar neutrinos (neutrinos produced in the nuclear reactions occurring in the sun) [34]. The experiment was based on the Pontecorvo/Alvarez design described in Section 1.2.2, where the chlorine in a tank of C_2Cl_4 was transmuted to argon via neutrino capture (see Equation 1.2). John Bahcall had previously performed theoretical calculations which predicted a significantly larger flux than was observed [35], an anomaly which persisted for the decades-long run of Davis' experiment. Pontecorvo noted that this anomaly could be explained by neutrino flavour oscillation⁶ [36], an idea he had developed along with Maki, Nakagawa, and Sakata in 1967. The suggestion was that the number

⁶Pontecorvo and Gribov proposed a two flavour oscillation framework, as neither the τ lepton nor ν_τ had yet been discovered.

of ν_e emitted from the sun matched Bahcall's prediction, but during the 1.5×10^8 km journey from the sun to the earth many of these ν_e oscillate into ν_μ and ν_τ which were undetectable with the experimental methods used⁷. In the decades following Davis' initial observation, other experiments also observed a solar neutrino deficit, including SAGE [37], KamiokaNDE-II [38], and GALLEX [39]. Davis and Masatoshi Koshiya were awarded the 2002 Nobel Prize in Physics for their work on measuring solar neutrinos.

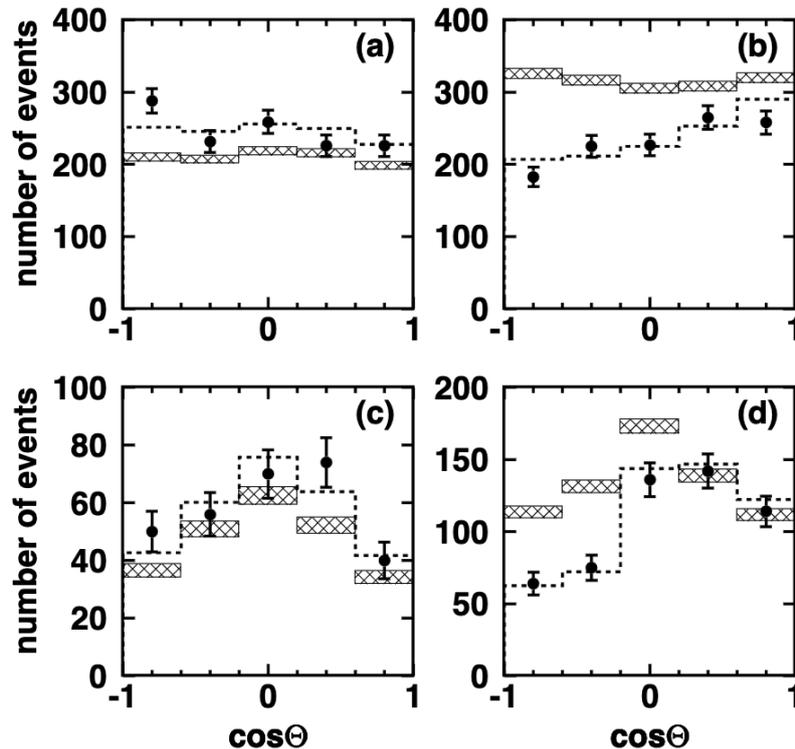


Figure 1.5: Zenith angle distributions observed in Super-Kamiokande for: (a) sub-GeV e-like, (b) sub-GeV μ -like, (c) multi-GeV e-like and (d) multi-GeV μ -like events. The value $\cos \Theta = 1$ means downward-going particles. The histograms with shaded error bars show the MC prediction with their statistical errors for the no neutrino oscillation case. The dotted histograms show the Monte Carlo prediction for $\nu_\mu \leftrightarrow \nu_\tau$ oscillation with a certain choice of oscillation parameters. Figure from [5].

In 1998 Super-Kamiokande observed a deficit in atmospheric neutrinos⁸ (neutrinos produced when cosmic rays interact in the upper atmosphere), which increased with zenith angle [5]. The deficit was only observed for ν_μ . They observed very little deficit of downward-going ν_μ (coming from straight above) and a significant deficit of upward-going

⁷The energy of the solar neutrino flux is too low for μ^\pm or τ^\pm to be created in a Charged-current interaction in the detector.

⁸Other experiments such as Soudan 2 had also observed an atmospheric neutrino deficit [40].

ν_μ (coming up through the earth), as shown in Figure 1.5. This demonstrated that the ν_μ disappearance depended on the distance travelled from the point of production. This observation could be explained by neutrino oscillation.

In 2002 the solar neutrino problem was solved by SNO, which measured both the ν_e and total solar neutrino flux [41]. The ν_e flux agreed with previous measurements, and the total neutrino flux matched the theoretical prediction. This conclusively showed that solar ν_e were disappearing, i.e. neutrinos oscillate. For their work in discovering neutrino oscillations Takaaki Kajita and Arthur McDonald were awarded the 2015 Nobel Prize in Physics.

1.3 Oscillation physics

One of the most intense areas of neutrino research today is the study of oscillations. A better understanding of neutrino oscillations may lead to an explanation of the observed matter-antimatter asymmetry in the universe or point to new physics. In this section, I will give an outline of neutrino oscillation physics.

1.3.1 Oscillations in a vacuum

The three neutrino generations can be described by either of two valid bases - a flavour basis (states of fixed lepton flavour) or a mass basis (states of fixed mass). These bases are orthonormal and related to each other via a rotation matrix. Neutrinos are created and observed as flavour states but propagate through space as mass states. The mass states are eigenstates of the free Hamiltonian \hat{H} . This can be expressed in the Dirac bra-ket notation as

$$\hat{H} |\nu_i\rangle = E_i |\nu_i\rangle \quad (1.4)$$

We can express a flavour state as a superposition of mass states

$$|\nu_\alpha\rangle = \sum U_{\alpha i}^* |\nu_i\rangle \quad (1.5)$$

where ν_α is one of the flavour states (ν_e, ν_μ, ν_τ), ν_i is one of the mass states (ν_1, ν_2, ν_3), and U is a rotation matrix between bases. Consider a neutrino created with a fixed flavour

α at time t_0 . Time evolution is given by the operator $e^{-i\hat{H}(t-t_0)}$. Hence the state at some time t is given by

$$|\nu_\alpha(t)\rangle = e^{-i\hat{H}(t-t_0)} |\nu_\alpha(t_0)\rangle = \sum U_{\alpha i}^* e^{-iE_i(t-t_0)} |\nu_i\rangle \quad (1.6)$$

Remembering the orthogonality relation $\langle \nu_i | \nu_j \rangle = \delta_{ij}$, we can therefore express the probability that a neutrino produced with flavour α will be observed with flavour β at some time t

$$P(\nu_\alpha \rightarrow \nu_\beta)(t) = |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 = \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-iE_i(t-t_0)} \right|^2 \quad (1.7)$$

Neutrinos are relativistic ($E \gg m$); hence we make the following approximation using a Taylor series expansion

$$E_i = \sqrt{p^2 + m_k^2} \simeq |p| + \frac{m_k^2}{2|p|} \quad (1.8)$$

We can also assume that $L \simeq (t - t_0)$ in natural units, where L is the distance travelled. We therefore have

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_{i,j} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{ij}^2 L}{2|p|}} \quad (1.9)$$

where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$. Using the unitarity of U , and making the approximation $|p| \simeq E$ we may write

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) = & \delta_{ab} - 4 \sum_{i>j} \text{Re} \left[U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \right] \sin^2 \left(\Delta m_{ij}^2 \frac{L}{4E} \right) \\ & + 2 \sum_{i>j} \text{Im} \left[U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \right] \sin \left(\Delta m_{ij}^2 \frac{L}{2E} \right) \end{aligned} \quad (1.10)$$

If we had derived this equation for antineutrinos instead of neutrinos, we would have arrived at the same expression except for the exchange $U \rightarrow U^*$. Hence we can see that the

first two terms on the RHS of Equation 1.10 are CP conserving⁹, as they are the same for neutrinos and antineutrinos. The final term on the RHS of Equation 1.10 is CP violating, as it changes sign when we exchange neutrinos for antineutrinos.

So we see that neutrinos created with a particular flavour can be expressed as a superposition of mass states. The neutrino propagates via the mass basis, with each mass state travelling at a different speed due to its unique mass. This causes each mass state to acquire a different phase during propagation. Suppose the neutrino is later detected (interacts). In that case, the superposition of mass states will be different from when it was created, and hence there is some probability for the neutrino to be detected in a different flavour state from the one it was created in. This phenomenon is known as neutrino oscillation.

The matrix U , which rotates us from the mass basis to the flavour basis, is known as the PMNS matrix, named for Pontecorvo, Maki, Nakagawa, and Sakata. It is typically expressed as follows

$$\begin{aligned}
 U_{PMNS} &= \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{bmatrix} \\
 &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (1.11)
 \end{aligned}$$

where $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$. The PMNS matrix is characterized by three mixing angles ($\theta_{12}, \theta_{13}, \theta_{23}$) and one Dirac CP-violating phase δ_{CP} . Equation 1.11 is factorised such that each matrix term describes a distinct sector of neutrinos physics. That is to say - each neutrino oscillation experiment tends to be most sensitive to the parameters in just one of the three matrices. From left to right these matrix factors represent the atmospheric, reactor, and solar sectors, respectively. These names refer to the different

⁹CP is the combination of a charge conjugation and parity transformation. If a process is CP conserving, it is unchanged under the effect of these two transformations. This implies that the process is the same for matter and antimatter.

sources of neutrinos used to study oscillations.

1.3.2 Oscillations in matter

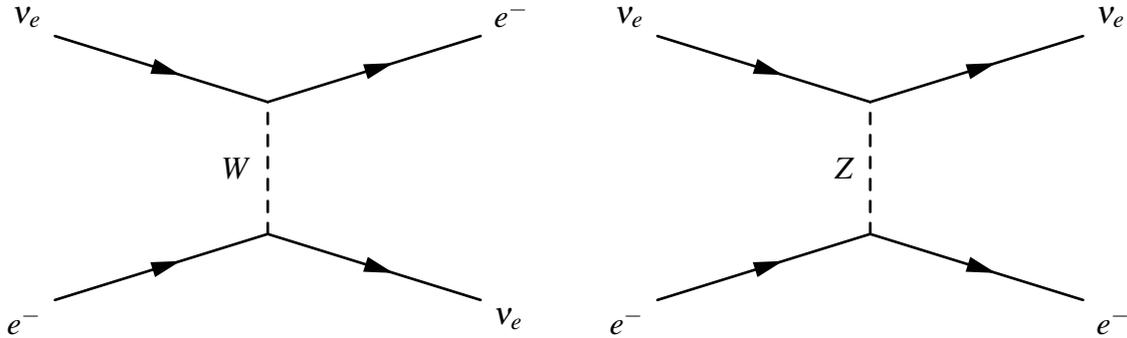


Figure 1.6: Coherent forward scattering can occur for ν_e in matter via a Charged-current (left) and Neutral-current (right) interaction.

The oscillation probabilities derived thus far assume propagation through a vacuum. When travelling through matter, neutrinos can scatter elastically off electrons and nucleons, preserving their momentum and spin. This process, known as coherent forward scattering, alters the neutrino's effective mass. This is known as the MSW effect, after Mikheyev, Smirnov, and Wolfenstein, who were the first to calculate it [42] [43]. Furthermore, because matter is rich in electrons, ν_e can scatter coherently via Charged-current and Neutral-current interactions, as shown in Figure 1.6 (see Section 1.4 for further explanation of these terms). Other flavours of neutrino can only scatter coherently via Neutral-current interactions. This can be understood with reference to Figure 1.6. If the incoming ν_e is swapped with another flavour of neutrino, for example ν_μ , diagram on the right of Figure 1.6 still represents elastic scattering, as the same particles are in the initial and final state. The same is not true of the diagram on the left, hence only ν_e are able to scatter elastically via Charged-current interactions in an electron-rich environment. The MSW effect also causes ν_e and $\bar{\nu}_e$ to oscillate differently, due to the high density of electrons as opposed to positrons in matter. This can cause differences in neutrino and antineutrino oscillation for expected reasons, which must be disambiguated from CP violation coming from δ_{CP} .

Matter effects make oscillation sensitive to the mass ordering of neutrino states, that is to say - which neutrino mass states are the heaviest/lightest. It is unknown whether ν_3 is the heaviest or lightest mass state. This is known as the mass hierarchy problem, illustrated

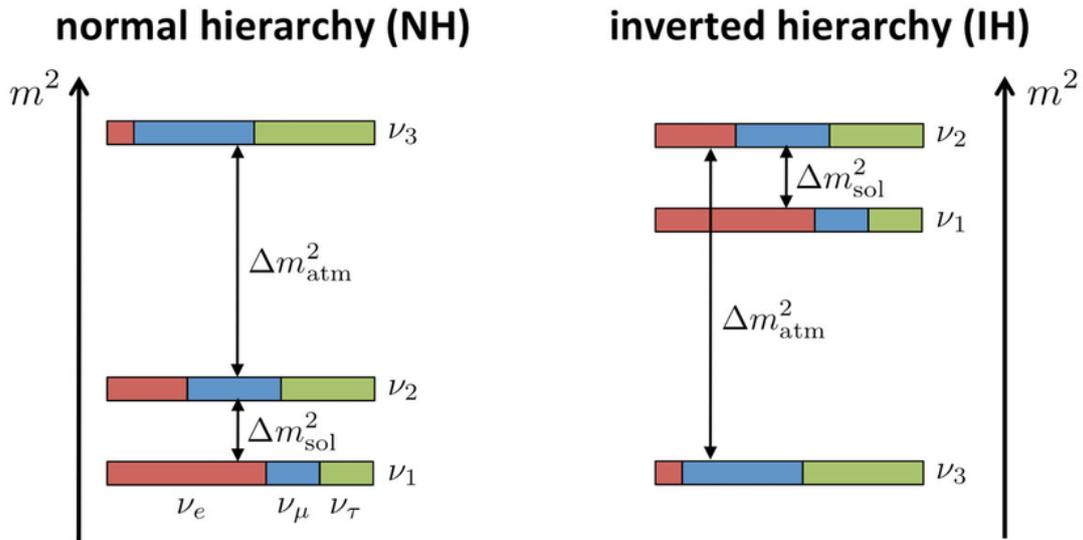


Figure 1.7: Representation of the two possible neutrino mass orderings. On the left, we have the normal ordering ($\Delta m_{\text{atm}}^2 > 0$). Shown on the right is the inverted ordering ($\Delta m_{\text{atm}}^2 < 0$). The colours represent the contribution of ν_e (red), ν_μ (blue), and ν_τ (green) in a given mass state. Here the labels Δm_{atm}^2 and Δm_{sol}^2 refer to Δm_{32}^2 and Δm_{21}^2 , respectively. Figure from [6].

in Figure 1.7.

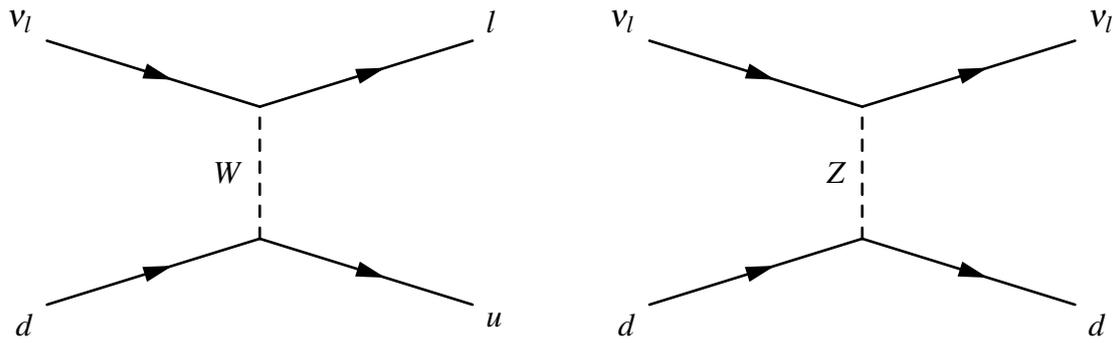


Figure 1.8: Charged-current (left) and Neutral-current (right) neutrino-nucleus interactions on a down quark.

1.4 Neutrino interactions

Neutrinos carry neither electric nor colour charge and hence do not feel the electromagnetic and strong forces. Therefore, when we observe neutrino interactions on Earth, it is due to their weak interactions with matter. The weak force is mediated by two gauge bosons, the W and Z bosons. As a result, there are two types of neutrino interactions we can observe, Charged-current (CC) and Neutral-current (NC), which are shown in Figure 1.8.

Charged-current interactions involve changing electrical charge by one unit and can induce flavour changes in the quark sector. There is an outgoing charged lepton of the same flavour as the incoming neutrino, which is easily detected in modern detectors. Note that for a CC interaction to occur, the energy of the incoming neutrino must be greater than the mass of the corresponding charged lepton plus the change in nuclear mass. This leads to strong suppression of ν_τ CC interactions from most artificial neutrino sources. Many of the experiments discussed in Sections 1.2.2 and 1.2.3 relied on CC interactions. Neutrino flavour identification is possible in the CC channel by determining the flavour of the outgoing charged lepton.

NC interactions were first discovered by the Gargamelle experiment at CERN [44] when they observed hadronic activity caused by the interaction of a neutral particle inside their detector, with no observed charged lepton in the final state. NC interactions involve no change of quark flavour or electric charge. They are only detectable through the momentum they impart to the hadronic system. Neutrino energy reconstruction is very difficult for NC events, as the outgoing neutrino carries a considerable fraction of energy away invisibly.

No flavour identification is possible using NC events. However, it is possible to observe NC events from all three neutrino flavours in artificial (and indeed solar) neutrino beams. The SNO experiment determined the total solar neutrino flux using NC events, as discussed in Section 1.2.3.

Most neutrino experiments rely on observing neutrino interactions with nuclei in detectors. It is helpful to describe neutrino-nucleus interactions in terms of interaction modes, which correspond to different energy regimes and final state topologies. Recall the de Broglie relation $\lambda = h/p$, where λ is the wavelength of a quantum particle, h is Planck's constant, and p is the momentum of the particle. It follows that different energy (momentum) scales correspond to different length scales. At low neutrino energies (large λ), the neutrino “sees” the nucleus as a whole or a collection of nucleons. At high neutrino energy (small λ), the neutrino can resolve (interact with) individual quarks within the nucleons.

1.4.1 Interaction modes

Much of this section's information on different models was sourced from the GENIE physics and user manual [45].

1.4.1.1 Quasi-elastic scattering

Quasi-elastic scattering (QE) is the dominant CC interaction mode at sub-GeV neutrino energy, as seen in Figure 1.9. QE events are interactions with a neutrino and a nucleon in the initial state and an outgoing lepton and nucleon in the final state, as shown in Figure 1.10. If an NC interaction occurs via this mode it is truly elastic. However, for a CC interaction, there is a change in nucleon mass and isospin ($n \rightarrow p$) as well as lepton mass, so we describe the process as *quasi*-elastic. QE interaction models include Llewellyn-Smith [46], Smith-Moniz [47], and Valencia models[†] [48].

1.4.1.2 Resonant interactions

When a sufficiently energetic probe (e.g. γ, e^-, ν) interacts with a nucleus, the nucleus can be excited into a resonant state. There are many such resonances, which can decay into one or multiple pions or kaons. However, of most interest is the lowest energy resonance

[†]Model used in NOvA simulation.

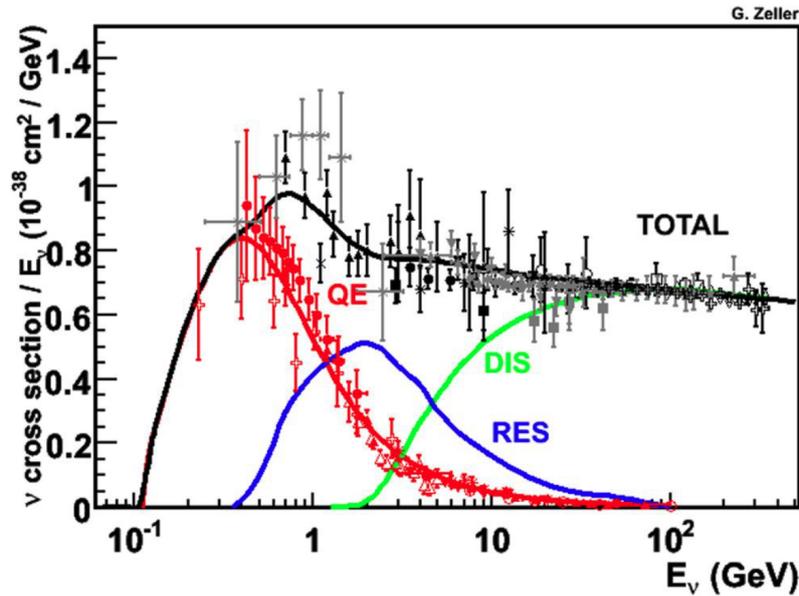


Figure 1.9: Neutrino-nucleus interaction cross-section divided by neutrino energy. Shown is the total cross-section as well as the cross-section broken down by modes - quasi-elastic (QE), resonant (RES), and deep inelastic scattering (DIS). This figure is over ten years old at the time of writing, so the latest model predictions may differ somewhat. Also, some rarer interaction modes are not shown. Nonetheless, this figure gives a sense of the relevant energy scales. Figure from [7].

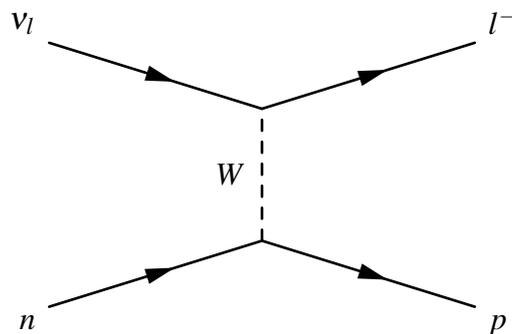


Figure 1.10: Diagram of a QE neutrino interaction. For the antineutrino mode, we instead have $p \rightarrow n$.

$\Delta(1232)$ ¹⁰, which typically decays into a single pion, as seen for example in Figure 1.11. Resonant single pion production can occur in ν_l CC interactions via the following channels:

$$\nu_l + p \xrightarrow{\Delta^{++}} l^- + p + \pi^+ \quad (1.12)$$

¹⁰The name $\Delta(1232)$ collectively describes the Δ^- , Δ^0 , Δ^+ , and Δ^{++} baryon resonances.

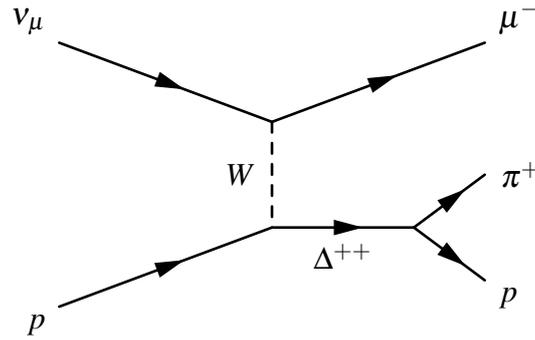


Figure 1.11: Diagram of resonant single pion production in a Charged-current interaction on a proton.

$$\nu_l + n \xrightarrow{\Delta^+} l^- + p + \pi^0 \tag{1.13}$$

$$\nu_l + n \xrightarrow{\Delta^+} l^- + n + \pi^+ \tag{1.14}$$

Commonly used models of resonant neutrino-nucleus interactions include the Rein-Sehgal [49] and Berger-Sehgal[†] [50] models.

1.4.1.3 Meson exchange current

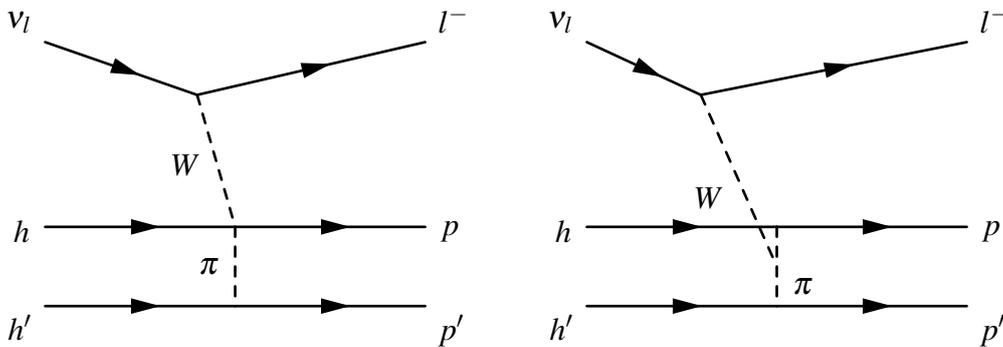


Figure 1.12: Diagram of a meson exchange current event, which knocks out two particles (nucleons) and leaves behind two holes. The contact current (left) and pion-in-flight current (right) are shown.

If a neutrino interacts with a correlated pair of nucleons, it can knock both nucleons out of the nucleus leaving behind two “holes” where the nucleons used to be, as shown in Figure 1.12. These are known as “two particle two hole” (2p2h) events. The nucleons are correlated by exchanging a meson (usually a pion), leading to the alternate name meson

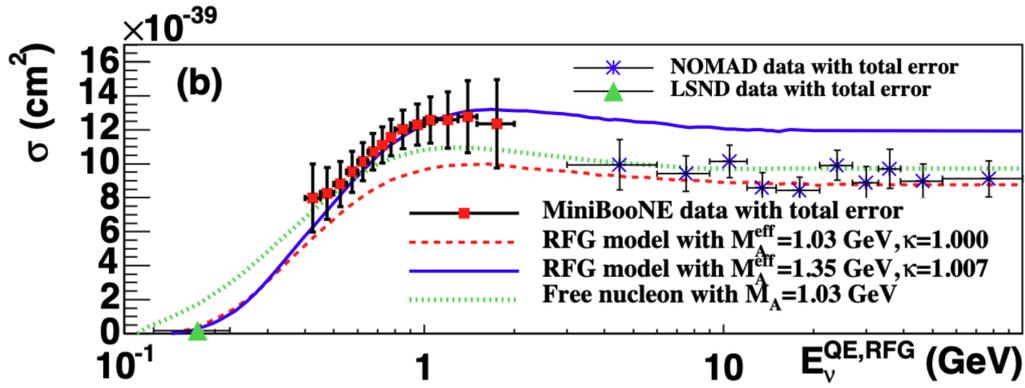


Figure 1.13: Flux-unfolded MiniBooNE ν_μ CCQE cross-section per neutron as a function of neutrino energy, along with results from the LSND and NOMAD experiments. Also shown are predictions from the NUANCE [8] simulation for an RFG model with two different parameter variations and for scattering from free nucleons with the world-average M_A value. Figure adapted from [9]. A portion of the original figure, which is not relevant to the current discussion, has been cropped out.

exchange current (MEC).

The history of MEC in neutrino physics is interesting. The MiniBooNE experiment observed a significant excess of QE-like events [9], which they could not explain by varying parameters of their QE model, as shown in Figure 1.13. It turns out that an additional non-QE mode was contributing to their sample. This led to the re-discovery of MEC in neutrino physics, even though electron scattering researchers had long been aware of this process [51].

Common models of MEC interactions include Valencia[†] [52], SuSa [53], and Empirical GENIE [54] models. Note that the kinematics and rate of MEC interactions in neutrino-nucleus interactions are not strongly constrained by measurements at the time of writing. As a result, NOvA tunes MEC interactions to adjust the shape and normalisation of the simulation. The tuning is done by reweighting MEC events in energy and three-momentum transfer space (see Section 3.1.4 for more details).

1.4.1.4 Coherent pion production

Sometimes a neutrino can interact with an entire nucleus coherently, leaving the nucleus intact and producing a pion in the final state as shown in Figure 1.14. These interactions are characterised by a very forward-going pion and small four-momentum transfer to the nucleus (i.e. small value of the Mandelstam variable t). The Rein-Sehgal [49] and

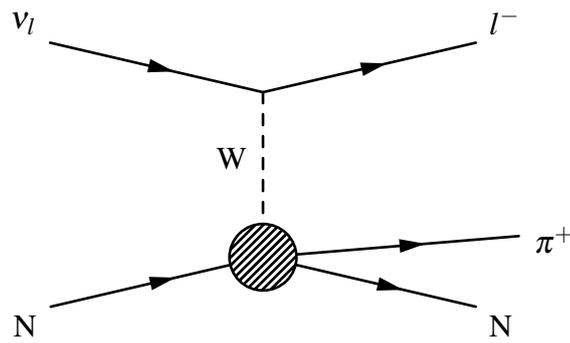


Figure 1.14: Diagram of coherent pion production.

Berger-Sehgal[†] [50] models are often used to simulate coherent interactions.

1.4.1.5 Deep inelastic scattering

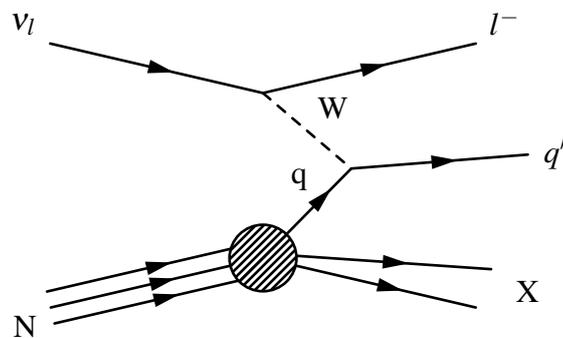


Figure 1.15: Diagram of a CC Deep Inelastic Scattering interaction

In deep inelastic scattering (DIS) interactions, the neutrino has sufficiently high energy to interact with individual quarks in a nucleon, as shown in Figure 1.15. The interaction breaks apart the nucleus, producing many hadrons in the final state. At high neutrino energies (and therefore high momentum transfer to the nucleus), DIS is the dominant interaction mode. The Bodek-Yang model[†] [55] is often used to simulate DIS interactions.

1.4.2 Nuclear effects

Modern neutrino experiments do not typically involve interactions on free nucleons (i.e. hydrogen detectors). Heavier elements such as carbon, oxygen, iron, and argon are found in modern neutrino detectors. Ab-initio calculations of interaction cross-sections are impossible due to the poorly understood nuclear effects introduced by interactions in heavy nuclei. In this section, I will outline some of these nuclear effects. Other works partially inspired this section [56] [57].

1.4.2.1 Fermi motion

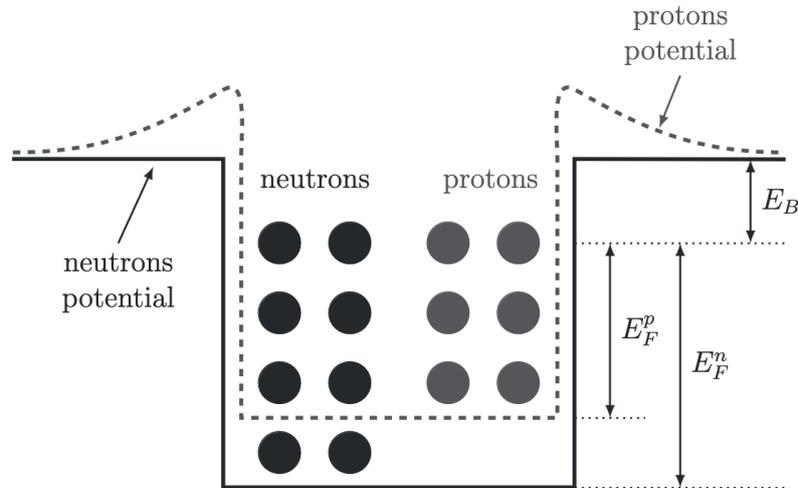


Figure 1.16: Nuclear potential well for protons and neutrons. Also marked are the neutron and proton Fermi energies E_F^n and E_F^p , respectively. E_B is the nuclear binding energy. Note the proton potential well is different from the neutron potential well, as protons are subject to electromagnetic forces. Figure from [10].

Protons and neutrons are fermions (they have half-integer spin) and so must obey the Pauli exclusion principle. As a result, the energy of nucleons in a nucleus is determined by Fermi-Dirac statistics. This means nucleons fill up energy levels sequentially, starting with the lowest up to the Fermi energy, as shown in Figure 1.16. So even at absolute zero, nucleons are not at rest; they have Fermi motion due to their finite energy. Neutrino interactions receive a boost in the lab frame from the Fermi motion of the struck nucleon. This boost depends on the speed and direction of the struck nucleon and will differ for every interaction.

The motion of nucleons is modelled by Fermi gas models. In a Fermi gas, nucleons are free to move within the nuclear volume under a binding potential. Due to the Pauli exclusion principle, an interaction can not occur if the struck nucleon is promoted to an energy level that is already filled. If a nucleon receives enough energy to promote it above the Fermi and binding energy, it will be ejected from the nucleus. Otherwise, the interaction is suppressed. As a result, the kinematics of neutrino interactions are impacted. This phenomenon is called Pauli blocking.

Some models assume that there is a constant nuclear potential that all nucleons feel throughout the nucleus and that the potential drops to zero at the boundary of the nucleus.

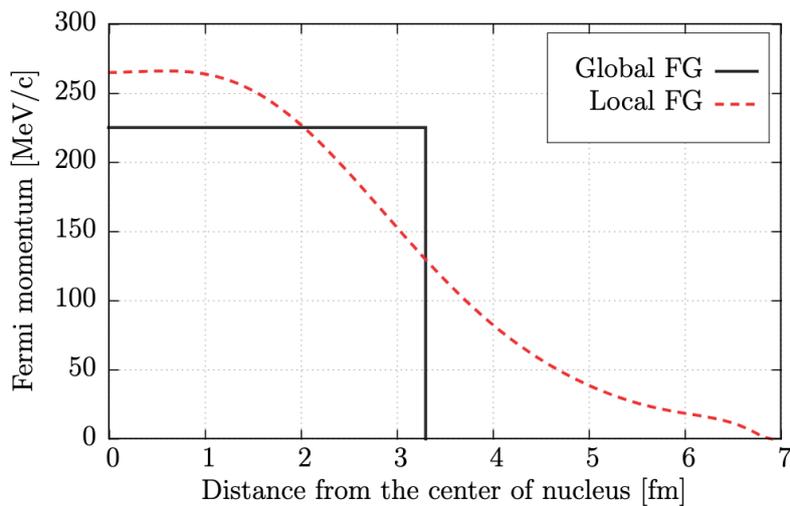


Figure 1.17: Comparison of the Fermi momentum for global and local Fermi gas in the case of carbon. Carbon is not chosen for any particular reason, rather this figure serves to generically highlight the difference between global and local Fermi gases. Figure from [10].

This is known as a Global Fermi Gas model. Other models take a more nuanced approach, assuming the nuclear potential is strongest at the centre and weakens as we travel radially outward. These are known as Local Fermi Gas models. We can see a comparison of these in Figure 1.17.

1.4.2.2 Random phase approximation

The random phase approximation (RPA) goes beyond the treatment of nucleons as non-interacting particles in a mean-field as described in Section 1.4.2.1. RPA is introduced to account for nucleon-nucleon interactions. The RPA framework assumes that collective excitations of the nucleus can be described by a linear combination of one-particle-one-hole (1p-1h) excitations. The collective excitation states are treated as orthogonal, and the phases are assumed to be random.

RPA can cause a modification of the electroweak coupling strength for a many-nucleon system. The complex nuclear correlations considered by RPA introduce a residual inter-nucleon interaction which causes a suppression of interactions at low Q^2 , as shown in Figure 1.18. There is a slight enhancement at higher Q^2 . RPA has almost no effect at high Q^2 . At higher energy scales, nucleons can be treated as approximately free, and the effect of the nuclear medium is less important.

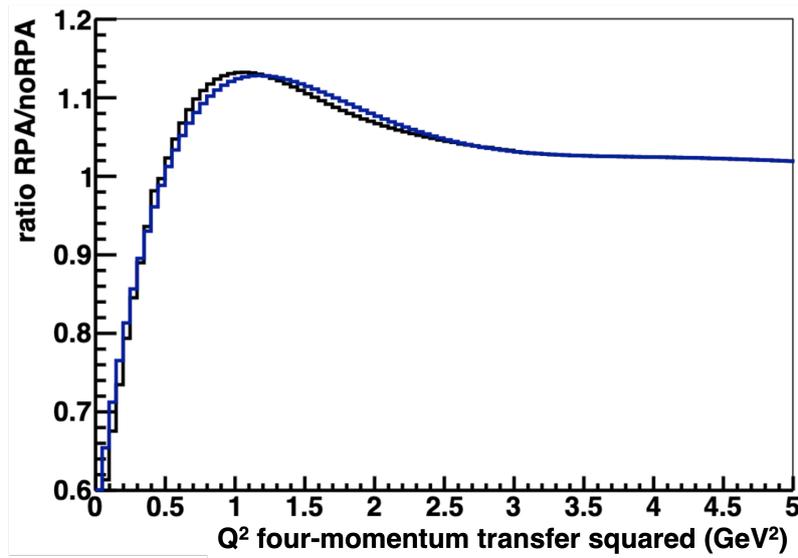


Figure 1.18: Ratio of GENIE simulations with and without RPA corrections included. The blue curve was made using a weighting just in Q^2 , whereas the black curve was made using weights in energy and three-momentum transfer space. Figure from [11].

1.4.2.3 Final state interactions

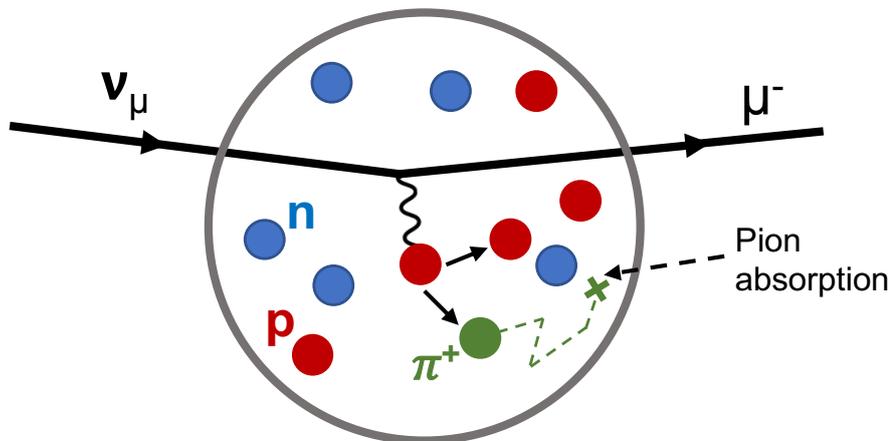


Figure 1.19: Example of FSI manifesting through pion absorption.

When a neutrino interacts with a nucleus, the struck nucleon may be ejected from the nucleus. Additional final state particles such as photons and hadrons (e.g. pions) may be produced. These particles must traverse and exit the nucleus before being detected. Interactions between the particles and the nuclear medium are known as final state interactions (FSI), and they can change the experimental observables. FSI can

1. Alter kinematics (energy and angle)

2. Change particle type e.g. $\pi^+ \rightarrow \pi^0$ (charge exchange)
3. Cause final state particles to be absorbed inside the nucleus
4. Cause the emission of additional particles, not coming from the primary neutrino interaction

FSI makes the precise determination of interaction modes very difficult. Consider the example shown in Figure 1.19, where a pion is produced (perhaps through a resonant interaction) but is absorbed before it exits the nucleus. In an experiment, we might only observe the outgoing muon and assume that this is a quasi-elastic interaction, even though it is not. So we see that there is confusion between the QE and RES interaction modes. As a result, experimentalists try to define interactions in terms of final state observables, e.g. “one proton in the final state”, “zero pions in the final state”, etc.

1.5 Why measure cross-sections?

The main physics goal of neutrino oscillation experiments, including NOvA, is to measure neutrino oscillation. This raises the question - why do we measure neutrino-nucleus interaction cross-sections? It turns out that precise knowledge of cross-sections is necessary to extract the oscillation parameters, as will be explained in this section. A publication by Benhar *et. al* [58] was a helpful reference in writing this section.

1.5.1 Cross-section definition

The cross-section σ is a measure of the probability of a particle-particle interaction. The term originates in reference to classical scattering of hard spheres. For example, consider firing small projectiles at a target. Some projectiles will miss the target, and others will collide. If the target were bigger, we would expect a larger fraction of the projectiles to collide with it. We see that the probability of interacting depends on the cross-sectional area of the target.

This analogy is retained for interactions at a distance, such as those seen in the Standard Model. The cross-section still has units of area, but it no longer represents the size of the target. Rather it represents the probability of interaction. For a beam of particles

incident on a collection of targets, the cross-section is defined as

$$\sigma = \frac{\text{Number of scatters}}{\text{Flux of beam particles} \times \text{Number of target particles}} \quad (1.15)$$

1.5.2 Inclusive measurements

When the signal definition of a cross-section measurement does not specify the final state particles, the measurement is described as “inclusive”. For example, a ν_μ CC inclusive measurement tells us how often we expect a ν_μ to interact in the detector, producing a muon and zero or more other particles.

The inclusive interaction cross-section is a key ingredient in oscillation measurements. In oscillation experiments, we create a beam of flavour α and measure the number of interactions of flavour β during some exposure¹¹. Exposure is measured in protons on target (POT) - the number of protons fired at a target to create the neutrino beam used. The number of POT is proportional to the number of neutrinos produced. The number of observed interactions of flavour β depends on several factors

$$N_{\nu_\beta} \propto \Phi_\alpha(E) \times \sigma_\beta(E) \times \varepsilon_\beta(E) \times P(\nu_\alpha \rightarrow \nu_\beta) \quad (1.16)$$

where Φ_α is the flux, σ_β is the interaction cross-section, ε_β is the detection efficiency, and $P(\nu_\alpha \rightarrow \nu_\beta)$ is the oscillation probability. Long-baseline oscillation experiments have a Near Detector (ND) and Far Detector (FD), where nominally, the ND measures the unoscillated beam and the FD measures the oscillated beam. Oscillation experiments benefit from the partial cancellation of systematic uncertainties when taking a ratio of event rates at the ND and FD. We assume that there is no oscillation at the ND

$$P^{ND}(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} \quad (1.17)$$

The ND is usually a larger angular target than the FD. As a result, the flux Φ is different at the ND and FD¹². The FD typically will be much larger than the ND, leading to differing acceptances and detection efficiencies ε .

¹¹It is possible to have $\alpha = \beta$, for example in the case of ν_μ survival.

¹²There is also a $1/L^2$ attenuation of the flux, but that is not particularly important for this discussion.

Even still, it might seem that the cross-section term would cancel when, e.g. trying to calculate the ν_μ survival probability using $N_{\nu_\mu}^{FD}/N_{\nu_\mu}^{ND}$. However, the number of interactions is always calculated as an integral over some bin width in neutrino energy, so the cross-section terms are now inextricable from the other terms, and cancellation is no longer possible

$$\frac{N_{\nu_\mu}^{FD}(E_1 < E < E_2)}{N_{\nu_\mu}^{ND}(E_1 < E < E_2)} = \frac{\int_{E_1}^{E_2} dE \Phi_\mu^{FD}(E) \sigma_\mu(E) \epsilon_\mu^{FD}(E) P(\nu_\mu \rightarrow \nu_\mu)}{\int_{E_1}^{E_2} dE \Phi_\mu^{ND}(E) \sigma_\mu(E) \epsilon_\mu^{ND}(E)} \quad (1.18)$$

Therefore it is necessary to know σ_μ (σ_e) to calculate the muon survival (electron appearance) probability. As we will see in Section 1.5.3, however, the dominant impact of cross-sections on oscillation measurements is in migration matrices.

1.5.3 Non-inclusive measurements

Often neutrino experiments report measurements of cross-sections where the final state topology is restricted, e.g. “zero pions in the final state” or “only one proton in the final state”. These non-inclusive measurements¹³ are a probe of the interaction modes described in Section 1.4.1. For example, a measurement of interactions with many hadrons in the final state is a probe of DIS. Understanding the breakdown of the inclusive cross-section by interaction modes is critical for accurate energy estimation. To see this, we will focus on a specific example.

Neutrino energy reconstruction is most straightforward for QE events. As there are only two particles in the final state (muon and nucleon), the kinematics are sufficiently constrained such that there is a simple equation for the incoming neutrino energy [58]

$$E_\nu^{QE} = \frac{2(M_n - E_b)E_l - (E_b^2 - 2M_n E_b + \Delta M^2)}{2(M_n - E_b - E_l + p_l \cos \theta_l)} \quad (1.19)$$

where M_n is the free neutron mass, E_b is the nuclear binding energy, $\Delta M^2 = M_n^2 - M_p^2$, and E_l , p_l , θ_l are the outgoing lepton’s energy, momentum, and angle respectively. Energy reconstruction is more difficult for other modes, and here we rely more heavily on the

¹³There are a variety of subgroupings in this category, e.g. exclusive, semi-inclusive. We will call them all collectively “non-inclusive”.

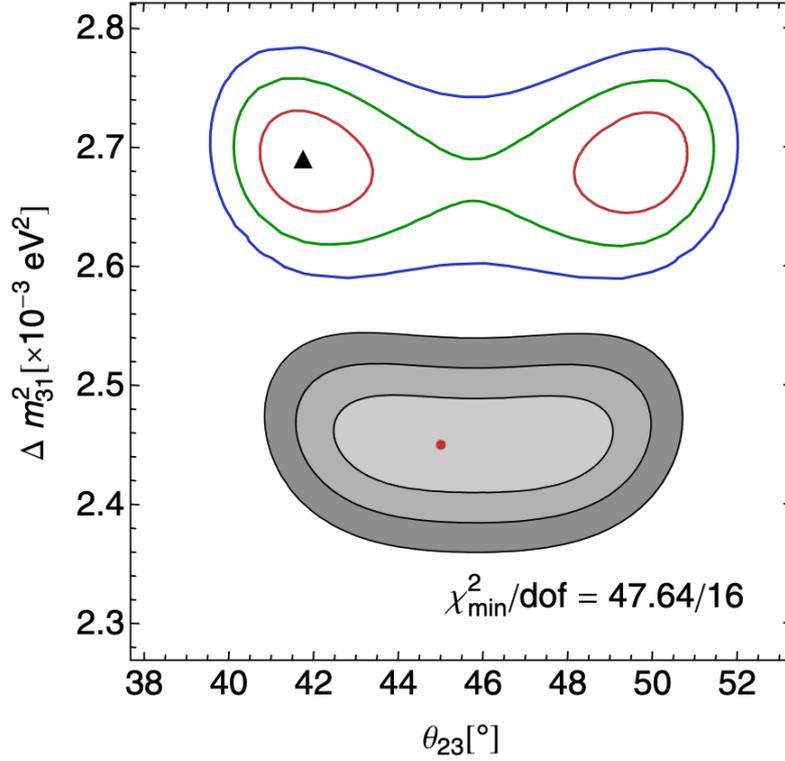


Figure 1.20: Impact on the results if a different generator is used to compute the true and fitted rates in the analysis. The shaded areas show the confidence regions at 1, 2, and 3 σ that would be obtained in the $\theta_{23} - \Delta m_{31}^2$ plane if the true and fitted rates are generated using the same set of migration matrices (obtained from GiBUU, with oxygen as the target nucleus). The coloured lines show the same confidence regions if the true rates are generated using matrices produced with GiBUU, but the fitted rates are computed using matrices produced with GENIE. Both sets of matrices are generated using oxygen as the target nucleus. The red dot indicates the true input value, while the black triangle shows the location of the best fit point. The value of the χ^2 at the best fit is also shown, together with the number of degrees of freedom. Figure adapted from [12]. A portion of the original figure, which is not relevant to the current discussion, has been cropped out.

migration matrices obtained from simulation. Imagine some QE-like selection. There will undoubtedly be some non-QE backgrounds, and so the predicted number of events from simulation in some bin i of reconstructed energy is given by the following equation [12]

$$N_i^{QE-like} = \sum_j M_{ij}^{QE} N_j^{QE} + \sum_{non-QE} \sum_j M_{ij}^{non-QE} N_j^{non-QE} \quad (1.20)$$

where M_{ij} is a migration matrix describing the probability of an event with true energy in bin j being reconstructed in bin i . Thus the predicted energy spectrum depends on the number of QE, RES, etc. events (which is a function of the cross-sections) and also the

migration matrices. If the model used during simulation is significantly different from the true underlying physics, then the translation from reconstructed energy to true energy outlined in Equation 1.20 will be imperfect. This poor estimation of true neutrino energy caused by imperfect modelling leads to poor estimation of neutrino oscillation parameters. This is highlighted in Figure 1.20, where we see that the oscillation results obtained depend strongly on which models are used during simulation. The discrepancy between the true and extracted values in Figure 1.20 is mostly due to model dependence, as, to make this figure, no smearing due to finite detector resolution is done for signal¹⁴ events. Refer to the paper for more details [12].

Measurement of non-inclusive cross-sections gives feedback to the model-building and event generator community, which leads to better simulation and, therefore, more accurate and precise measurements of the oscillation parameters.

1.6 Summary

This chapter gives an overview of the history of the neutrino, from its proposal by Pauli to the cutting-edge research of today. Neutrino research has changed our understanding of the universe, yet many open questions remain. Answering these questions requires state-of-the-art detectors and intense neutrino beams. The next chapter will explore these topics.

¹⁴Here signal is defined as ν_μ CC interactions. A small background of NC events are also considered with a gaussian energy smearing of width 85 MeV.

Chapter 2

The NOvA experiment

This chapter will give details of the detector design and neutrino source for the NOvA experiment.

The Neutrinos at the Main Injector (NuMI) beam is the most intense artificial neutrino beam in the world, achieving a record power of 900 kW [59]. The NuMI Off-axis ν_e Appearance experiment (NOvA) is a long-baseline neutrino oscillation experiment. It consists of a Near and Far Detector, which are functionally identical tracking calorimeters. See Section 2.2 for further details on the detector technology. NOvA is designed to look for the disappearance of muon (anti)neutrinos and the appearance of electron (anti)neutrinos in the NuMI beam. The Near Detector sees a mostly unoscillated muon (anti)neutrino beam. The Far Detector sees an oscillated beam of muon (anti)neutrinos and electron (anti)neutrinos¹. The Near and Far Detectors are ~ 1 km and 810 km from the neutrino source, respectively.

2.1 The NuMI beam

This section was written with reference to a NuMI design and operations paper [13]. The NuMI beam has historically been used for several neutrino experiments, though NOvA is the primary user at the time of writing. The neutrino beam is created by impinging 120 GeV protons from the Main Injector on a graphite target, as shown in Figure 2.1. The protons come in bunches known as spills. Each spill lasts about $10 \mu\text{s}$, contains 4.9×10^{13}

¹There are also many tau (anti)neutrinos in the oscillated beam. However, the neutrinos are below the CC interaction threshold, so NOvA can not identify them

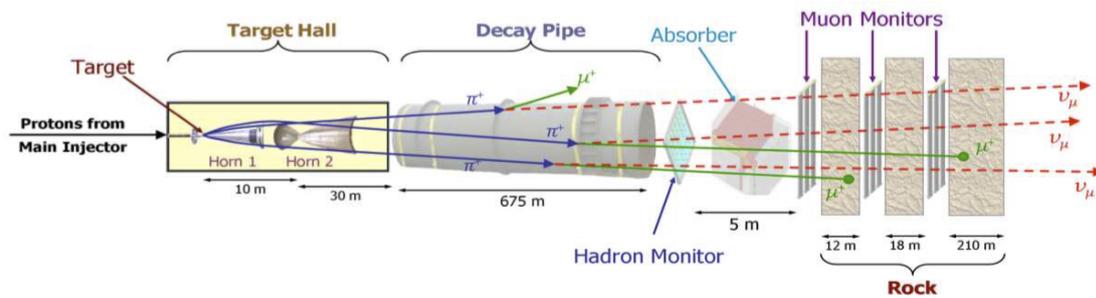


Figure 2.1: An illustration of the various components of the NuMI beamline. Figure from [13].

protons, and occurs approximately every 1.1 - 1.3 s.

This section will explain the various stages used to create the beam and briefly characterise the beam energy and composition.

2.1.1 Fermilab accelerator complex

Many of the details in this section are sourced from a technical report [60]. The Fermilab accelerator complex, shown in Figure 2.2, provides beam to many experiments, including NOvA. The beam is first created by accelerating H^- ions to 750 keV before being sent to the Linear accelerator (Linac), where they are further accelerated to 400 MeV. A carbon foil then transforms the H^- ions to H^+ ions (protons). The protons pass to the rapid-cycling-synchrotron Booster, where the protons are accelerated to 8 GeV. From here, some protons are sent to create the low-energy Booster Neutrino Beam, while others are sent to the Recycler. The Recycler is big enough to accommodate six batches of 84 Booster bunches each; however, using a procedure known as “slip-stacking” six more batches can be injected for a total of 12 batches. Some of this stacked proton beam is sent to the Fermilab Test Beam Facility at this stage. The remainder of the beam is then sent for final acceleration at the Main Injector, where the proton energy reaches 120 GeV. These high-energy protons are then used to create muons and neutrinos for Fermilab’s high-energy experiments.

2.1.2 Target and horns

The NuMI target is comprised of 47 graphite fins stacked together for a total target length of 95.38 cm. When the 120 GeV protons from the Main Injector impinge on the target, a spray of hadrons is produced - primarily pions, kaons, and protons. These hadrons are

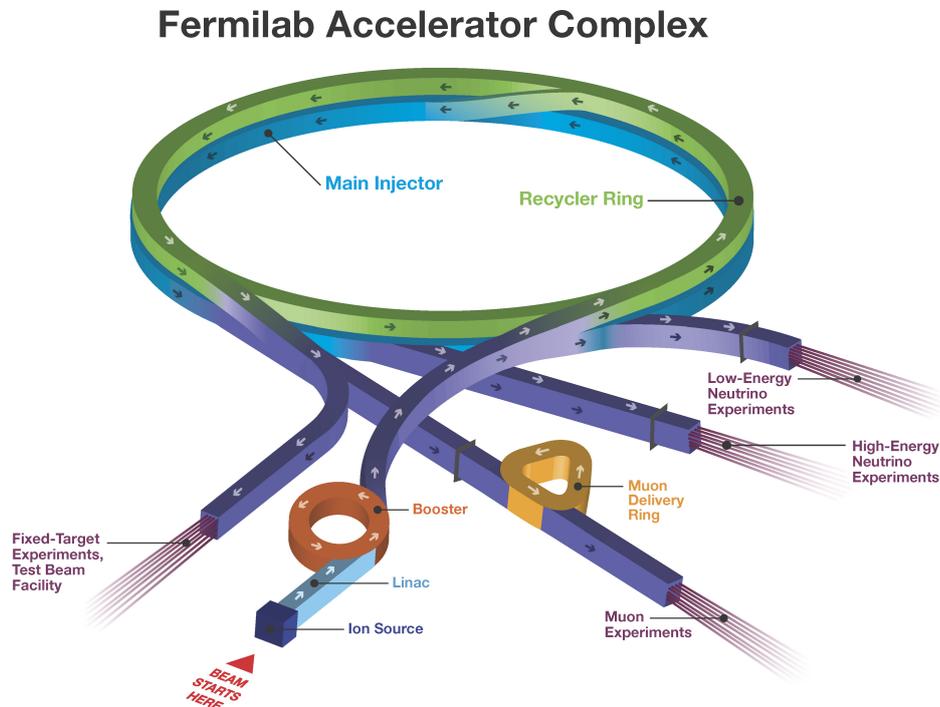


Figure 2.2: A diagram of the Fermilab accelerator program, showing the various beams provided to the end-user experiments. Figure from [14].

focused by the magnetic horns, as shown in Figure 2.3. The horns are made of an outer and inner conductor, which produce strong magnetic fields when pulsed with a current. The horns act as a hadron lens, where the lens's focal length is proportional to the hadron's momentum. A second horn is placed 10 m downstream of the first horn to correct for over/under-focusing. Including a second horn improves the efficiency of the focusing system by $\sim 50\%$.

When operating in FHC mode, the horns focus π^+ and K^+ and defocus (deflect) π^- and K^- ². By reversing the polarity, the horns can operate in RHC mode, whereby negative hadrons are focused, and positive hadrons are defocused.

²There is still some contamination of π^- and K^- travelling directly along the beam axis, which are undeflected.

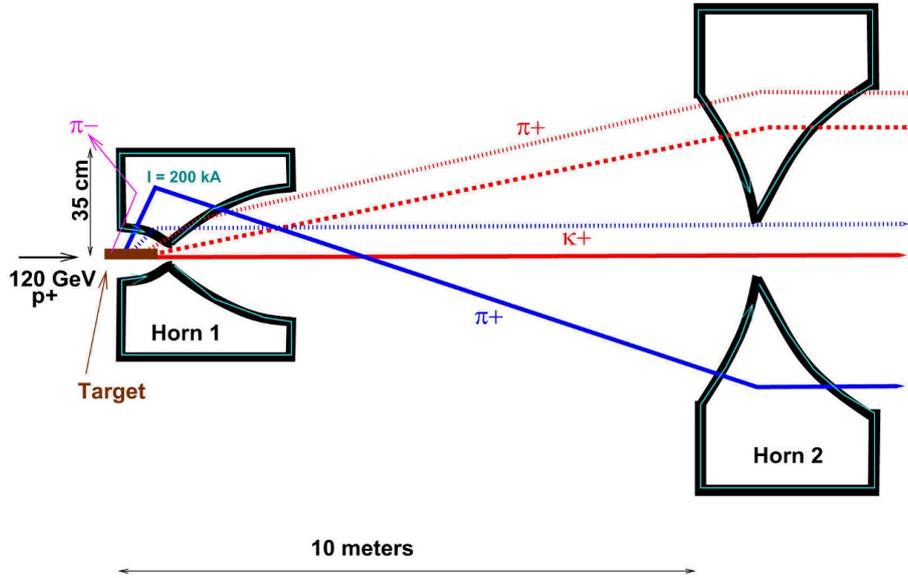


Figure 2.3: A diagram of the magnetic focusing horns, with possible hadron trajectories drawn. Figure from [13].

2.1.3 Decay pipe and absorber

The focused hadron beam travels through a decay pipe. The decay pipe is filled with helium, which provides structural integrity for the pipe while having a long interaction length, to allow the hadrons time to decay. The hadrons decay primarily as follows

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

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

and a CP flipped version producing $\bar{\nu}_\mu$ in RHC mode. Following the decay pipe, the muons and remaining hadrons are absorbed in either the absorber or 240 m of dolomite rock. The absorber is a structure made of aluminium, concrete, and steel, which measures 5.5 m wide \times 5.6 m tall \times 8.5 m long. It is designed to absorb any remaining hadrons in the beam.

2.1.4 Off-axis design

For NOVA's fixed baseline of 810 km, the oscillation probability is maximal for neutrinos of energy ~ 1.6 GeV. To this end, both detectors are placed 14 mrad off the beam axis. Due to the kinematics of relativistic two-body decay, it is possible to express the energy of

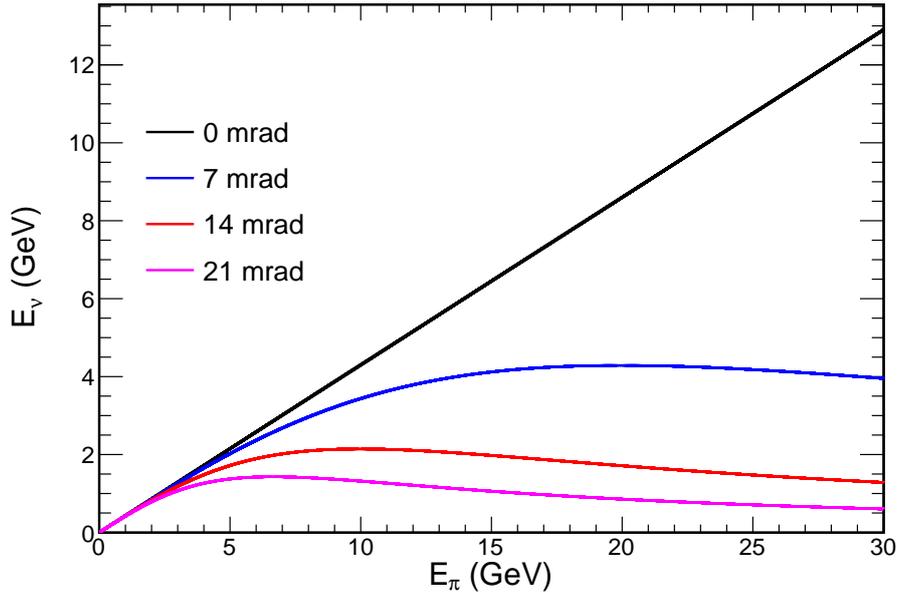


Figure 2.4: Neutrino energy as a function of parent pion energy for various off-axis angles

the neutrino as a function of the emission angle θ with respect to the pion momentum

$$E_\nu \simeq \left(1 - \frac{m_\mu^2}{m_\pi^2}\right) \frac{E_\pi}{1 + \gamma^2 \theta^2} \quad (2.3)$$

where $\gamma = E_\pi/m_\pi$ is the Lorentz factor. Figure 2.4 shows E_ν as a function of E_π for various off-axis angles. By placing the detectors off-axis, we see that the neutrino energy is more closely peaked near the energy of maximum oscillation.

2.1.5 Beam characteristics

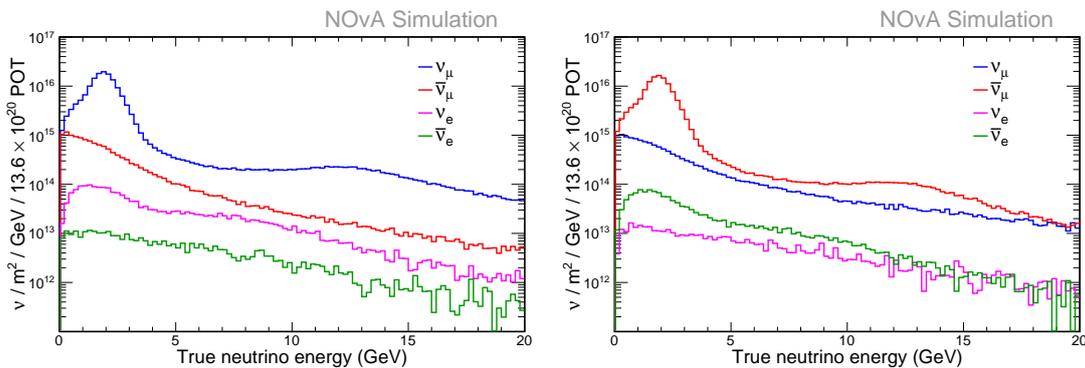


Figure 2.5: NuMI flux at 14 mrad off-axis for (left) FHC and (right) RHC.

The NuMI off-axis beam has a peak at ~ 2 GeV coming from pion decay and a

smaller peak at ~ 13 GeV from kaon decay, as shown in Figure 2.5. The FHC (RHC) beam has a $\bar{\nu}_\mu$ (ν_μ) “wrong-sign” contamination, due to imperfect defocussing of wrong-sign hadrons. There is also ν_e contamination from the following decay chain

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (2.4)$$

$$\mu^+ \rightarrow \bar{\nu}_\mu + e^+ + \nu_e \quad (2.5)$$

as well as the CP flipped and kaon decay versions of the above. There is also ν_e contamination from, for example

$$K_L^0 \rightarrow \pi^- + e^+ + \nu_e \quad (2.6)$$

$$K^+ \rightarrow \pi^0 + e^+ + \nu_e \quad (2.7)$$

2.2 The NOvA detectors

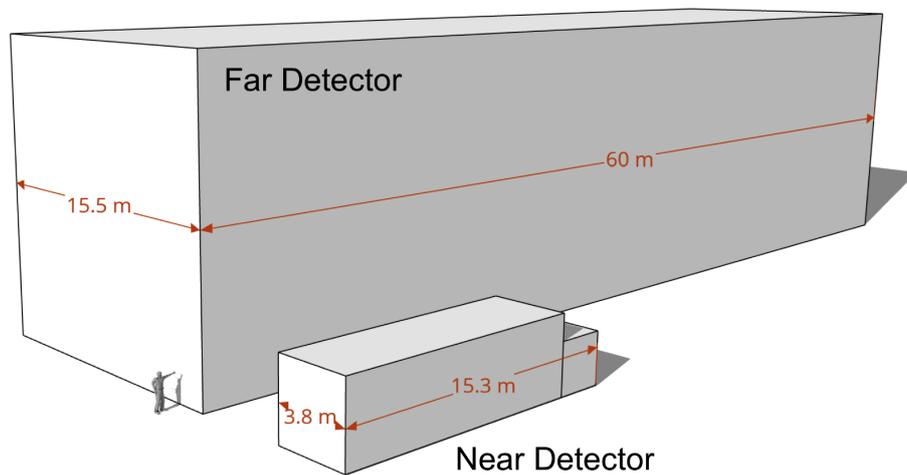


Figure 2.6: Drawing of the Near and Far Detectors next to each other for size comparison. Figure from [15].

The Near and Far Detectors are functionally identical tracking calorimeters. Here “functionally identical” means that the same detector materials are used for both detectors. However, the FD is approximately four times bigger than the ND in each dimension (~ 64

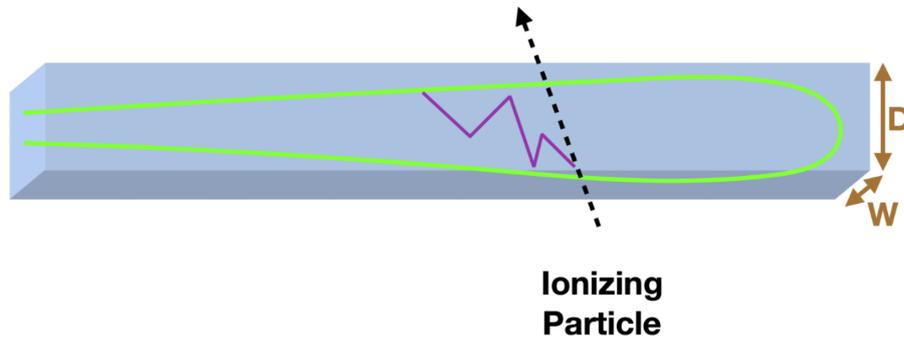


Figure 2.7: Schematic of the interior of NOvA cell, with the path of an ionising particle through the scintillator overlaid. Scintillation light, marked in purple, is collected by the wavelength-shifting fibres in green. Figure from [15].

times the ND volume). The detector design consists of long PVC cells filled with liquid scintillator. Figure 2.7 shows a schematic of the interior of a NOvA cell. This section will give a brief overview of detector elements and was written with reference to the NOvA Technical Design Report [16].

2.2.1 Scintillator

Component	Purpose	Mass fraction
Mineral oil	solvent	95.8%
Pseudocumene	scintillant	4.1%
PPO	waveshifter	0.091%
bis-MSB	waveshifter	0.0013%
Stadis-425	antistatic agent	0.0003%
tocopherol (Vit.E)	antioxidant	0.0003%

Table 2.1: Composition of the NOvA scintillator.

The NOvA detectors are approximately 69% liquid scintillator by mass. See Table 2.1 for the composition of the liquid scintillator. The mineral oil acts as a solvent for scintillant and waveshifters. NOvA uses 1,2,4-Trimethylbenzene (also known as pseudocumene) as the scintillant. Pseudocumene produces scintillation light of wavelength 360-390 nm when a charged particle passes. This light is shifted to 400-450 by other doping elements in the mineral oil, leading to better absorption in the fibres. Charged particles travelling faster than the speed of light in the medium also emit some Cherenkov light, which increases the light yield by a few percent.

2.2.2 PVC cells

The long rectangular cell boundaries are made from rigid polyvinyl chloride (PVC), loaded with titanium dioxide to increase reflectivity. Increased reflectivity increases the chances for scintillation light to be collected by the fibres. The PVC walls are 2-4.5 mm thick. The interior dimensions of the cell are 5.9 cm along the beam direction and 3.8 cm transverse to the beam direction. The cells are 15.5 cm long at the FD and 3.6 cm long at the ND. The cell is filled with liquid scintillator, and a wavelength-shifting fibre is looped through with both ends going to a single avalanche photodiode (APD) pixel detector, as shown in Figure 2.7.

2.2.3 Wavelength-shifting fibre

The fibres are 0.7 mm in diameter. The fibre core consists of polystyrene doped with a 300 ppm concentration of R27 dye. The 400-450 nm scintillation light is absorbed by the dye and re-emitted in the 490-550 nm range. Photons at this shifted wavelength are less attenuated in the fibres during the journey to the photodiode readout. The fibre core is clad in two layers of acrylic with lower refractive indices than the polystyrene core. This facilitates total internal reflection in the fibres, which improves transmission.

2.2.4 Avalanche photodiode

The scintillation light collected by the fibres is detected by an APD. The APD consists of a reverse-biased p-n junction. Incident photons create an electron-hole pair which causes a current to flow. The APDs, which are manufactured by Hamamatsu, have a quantum efficiency of 85% for photons in the range 520-550 nm. The number of photoelectrons produced by an incident light is magnified by a factor of 150 by applying a voltage of 425 V. To reduce thermal noise the APDs are cooled to -15 °C using a thermoelectric cooler. The cooling, however, increases the risk of condensation forming on the APDs. To mitigate this, a stream of dry air is blown on them. APDs are arranged in 32 pixel arrays, one to read out each of 32 cells.

2.2.5 Detector geometry

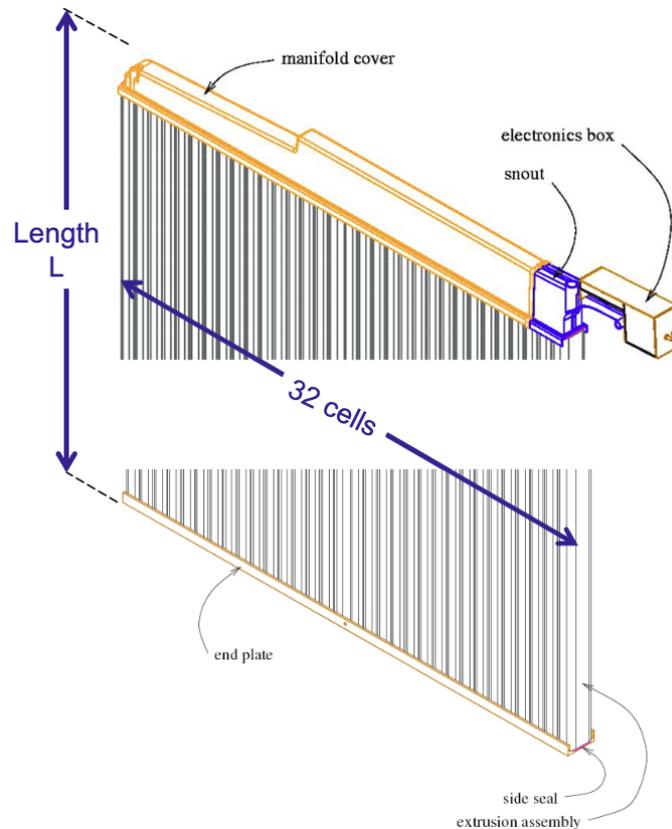


Figure 2.8: Schematic of a module consisting of two 16 cell extrusions glued together. The module is sealed on one end with an end plate and on the other with a manifold cover. Figure from [16].

PVC cells are created in 16 cell extrusions. Two extrusions are glued together to form a module, as shown in Figure 2.8. The end plate seals one end of the module. The other end is sealed by a manifold cover, which directs 64 fibre ends to the APD (two fibre ends per cell). A plane at the Far Detector is made by glueing together 12 modules. At the Near Detector, a plane is made by glueing 3 modules together. Planes are stacked along the beam direction in an alternating horizontal and vertical configuration, as shown in Figure 2.9. The y - and z -position of a charged particle can be determined when it passes through a horizontal plane. Similarly, the x - and z -position can be determined when passing through a vertical plane. By combining this information, a full 3D reconstruction of the interaction can be done.

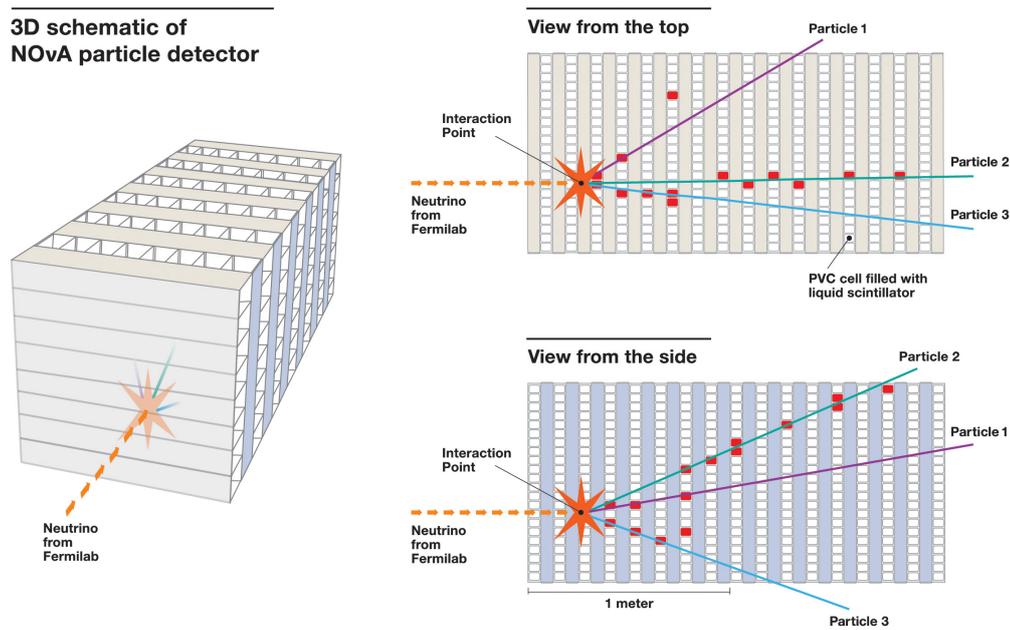


Figure 2.9: Schematic of a NOvA detector showing alternating horizontal and vertical planes. The vertical planes give a view of the event as viewed from above, and the horizontal planes give a view of the event as viewed from the side. Figure from [17].

2.2.6 Near Detector

The 290 ton Near Detector sits 105 m below the surface and 1.05 km from the NuMI target. The overburden³ significantly reduces cosmogenic backgrounds in the detector. The detector sits 14 mrad off-axis in a high-intensity beam. There are several neutrino interactions per spill, though most are uncontained. In order to aid containment, the ND has a muon catcher at the end. The muon catcher consists of layers of steel between 10 scintillator planes. Its purpose is to range out muons in a high-density material before they exit the detector. The muon catcher is shorter than the rest of the detector, as shown in Figure 2.6.

2.2.7 Far Detector

The 14 kton Far Detector sits on the surface, 810 km from the NuMI target. Just like the ND, it is placed 14 mrad off-axis. The FD sits in a purpose-built facility in Ash River, northern Minnesota. There is a 122 cm concrete enclosure, with an additional 15 cm barite overburden. The FD sees a high rate of background interactions from cosmic rays.

³Overburden is the name for any shielding above a detector. In this case it is 105 m of rock.

Backgrounds are mitigated by using the beam spill timing. The beam is much less intense at the FD than the ND; hence a bigger detector is required.

2.3 Data acquisition system

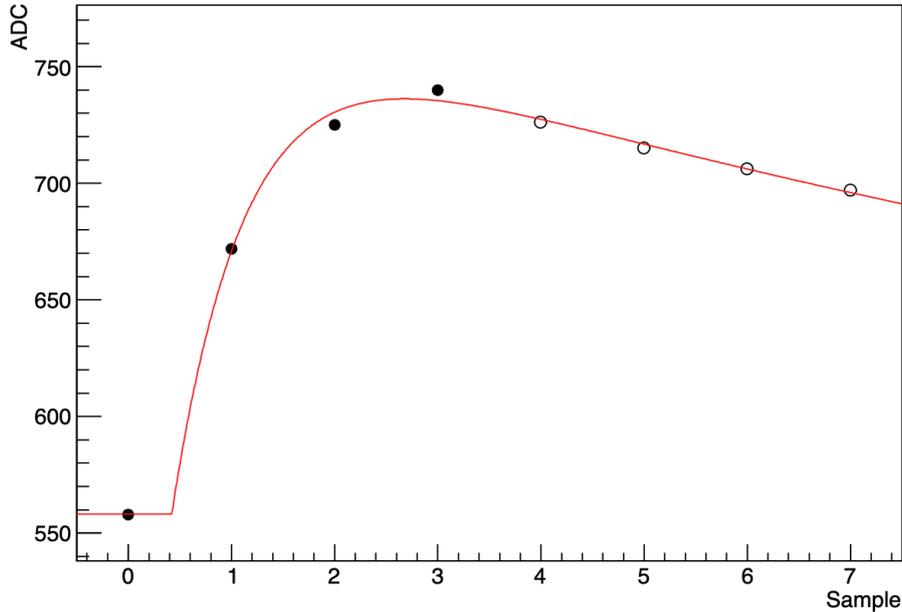


Figure 2.10: A set of four consecutive ADC readings marked with black circles. A fit is performed to these readings, marked in red.

The data acquisition system (DAQ) saves interesting data from the NOvA detectors. The signals from 32 APDs are read out by a single front end board (FEB). Each FEB performs shaping of the signal. It also digitises the APD signal using an analogue-to-digital convertor (ADC).

After digitisation, the signals are sent to a field-programmable gate array (FPGA) to look for hits above threshold. NOvA uses dual correlated sampling (DCS) to determine whether a signal should be kept. The ADC outputs a series of readings α_i in real time. If $\alpha_i - \alpha_{i-3}$ is above some preset threshold, then α_i and the three preceding samples are recorded, as shown in Figure 2.10. A fit to these samples allows precise timing. At the FD the timing resolution is ~ 10 ns, and at the ND it is ~ 7 ns. Each cell has its own DCS threshold, which is implemented using the FPGA.

The data from up to 64 FEBs is sent to a data concentrator module (DCM). The DCM collates data into 5 ms “millislices”, which are sent to a buffer farm. Data sit on the buffer

for up to 20 minutes until we either receive a trigger and save it or, alternatively, discard it. A circular buffer is used, with a first-in-first-out policy. If the buffer is full when new data arrives, the oldest data in the buffer is overwritten.

2.3.1 Triggering

Data is saved only if a trigger is received. There are three types of triggers used in NOvA. Only one trigger is required to record data. The first type of trigger is a beam spill trigger, used to record beam neutrino interactions. A beam spill time stamp is sent from the accelerator and compared to the data timestamps in the buffer farm. Even though a spill lasts only $10\ \mu\text{s}$, if the data is within a $500\ \mu\text{s}$ window around the beam spill, the data is saved. This provides robustness against timing drift and allows for the use of out-of-spill backgrounds to constrain in-spill backgrounds. For beam triggers, the neutrino time-of-flight is taken into account ($\sim 3\ \text{ms}$ at the FD).

The second type of trigger writes out $500\ \mu\text{s}$ of data at a regular 10 Hz frequency. The purpose of this trigger is to save cosmic ray interactions for calibration.

The third type of trigger is the data-driven trigger (DDT), which decides to save collections of hits with specific interesting properties. The purpose of the DDT is to save data for exotic/new physics analyses. These include searches for magnetic monopoles, dark matter, and supernova neutrinos.

2.4 Summary

This chapter describes the beam and detector technologies which enable scientists in the NOvA collaboration to study the properties of the neutrino. The accelerator neutrino beam method, first employed by Danby *et al.* [31], has revolutionised the field of neutrino physics. Advances in accelerator technologies have led to more intense neutrino beams, albeit still operating under the same principle. Detector technologies have continually advanced, leading to the collection of ever more precise neutrino interaction data. However, making full use of this data requires an advanced suite of offline software. Chapter 3 will describe the offline software used by NOvA.

Chapter 3

NOvA offline software

Modern particle physics experiments rely not only on advanced detector hardware but also on advanced software. Simulation and reconstruction are two software-based processes which are vital to interpreting data collected by a detector. This chapter will outline the various Monte Carlo simulations used by NOvA. Then a description of the reconstruction methods will be given, which are applied to both data and simulation. Finally, a brief description of the calibration procedure will be given.

3.1 Simulation

MC uses randomness to simulate complex phenomena. In the context of particle physics, this allows us to generate predictions for interactions with many degrees of freedom by relying on repeated random sampling of probability distributions. Simulation provides a way of testing analyses using truth information. For example, the validity of an analysis framework can be tested by performing an in-out test, where we use analysis machinery to try recover the same physics parameters as used in the simulation. MC can be used to develop a blind analysis, whereby data is not looked at until the analysis is very mature to avoid unconscious bias. MC also allows us to quantify systematic uncertainties by changing some simulation parameters within errors and re-simulating, or reweighting events to reproduce a change in simulation parameters. This section will outline the stages of the NOvA simulation chain.

3.1.1 Flux

The NuMI flux is simulated using G4NuMI, a custom simulation based on Geant4 [61]. G4NuMI simulates protons hitting the graphite target, followed by hadron production. It also simulates hadron focusing in the magnetic horns and hadron decay in the decay pipe. The output of G4NuMI is a prediction of the NuMI neutrino flux and some information on the neutrino parent. Parent information is used for later reweighting by the Package to Predict the Flux (PPFX) [62]. PPFX is a software package developed by the MINERvA collaboration. It aims to reweight the NuMI flux prediction to incorporate constraints on hadron production obtained from external experiments. PPFX provides a central value weight which adjusts the flux prediction. In addition, PPFX is used to obtain a systematic error band on the flux prediction due to uncertainties in hadron production.

3.1.2 Event generation

Neutrino interactions in NOvA are simulated using the GENIE event generator [54]. The simulation used in this thesis was generated using GENIE v3.0.6.

GENIE takes the simulated flux and geometry as input. GENIE uses its internal models (see Section 1.4.1) to create neutrino interactions on nuclei in the detector and surroundings. FSI is simulated using the hN-intranuke model¹. The output of GENIE is the kinematics of all particles exiting the nucleus, as well as a truth label for the interaction, e.g. DIS.

Interactions are simulated in the surrounding rock, as well as the detector. Simulation is time-consuming, so rock events are re-used multiple times in simulation files to reproduce the expected rate of rock interactions. At the Far Detector, measured cosmic ray data is overlaid on simulated interactions, as shown in Figure 3.1. This not only provides realistic backgrounds in simulation but also saves computing time. The overlaid cosmic ray data is taken from the 10 Hz trigger described in Section 2.3.1.

GENIE provides physics parameters which can be changed to alter the interaction models within some uncertainty. These parameters can introduce both shape and normalization changes. NOvA uses the NOvARwgt framework [63] to reweight events using the

¹hN-intranuke is an internal model in GENIE. It simulates the interaction of hadrons (h) with nucleons (N).

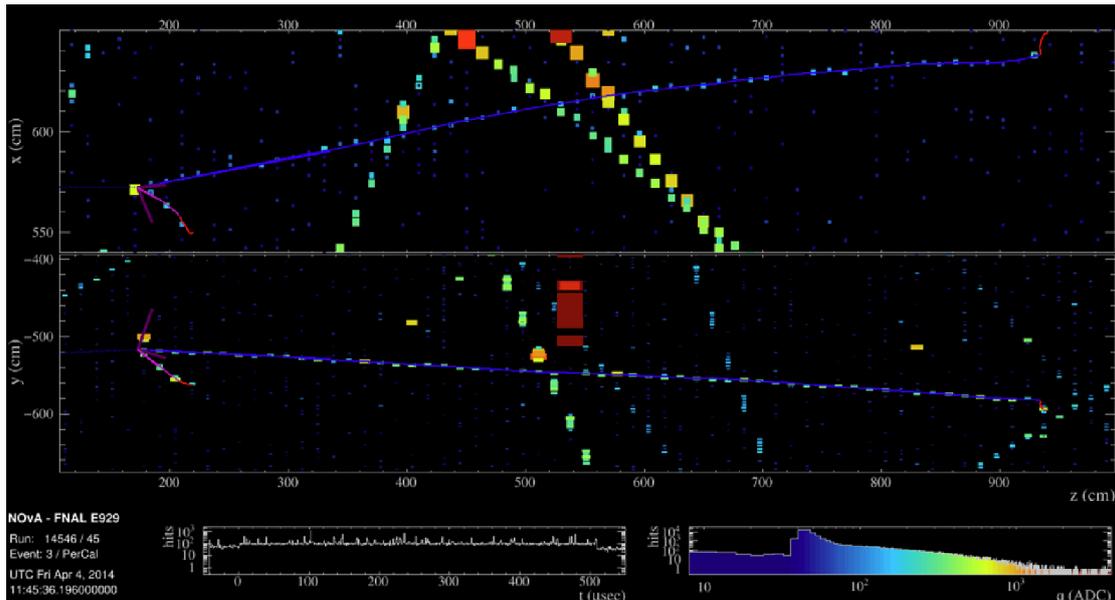


Figure 3.1: Simulated neutrino interaction at the Far Detector with cosmic ray data overlaid. The simulated particles have truth vectors drawn. Figure from [18].

GENIE parameters to estimate a cross-section modelling systematic uncertainty.

3.1.3 Detector simulation

The final state particles simulated by GENIE are handed to Geant4. Geant4 tracks the particles through the detector and surroundings. A description of the dimensions and atomic composition of the detector and detector hall components is provided to Geant4. Geant4 simulates particle propagation, production of scintillation light, as well as hard and soft scatters. The output of Geant4 is the path and energy deposition of particles traversing the detector.

Geant4 uses different models depending on energy scales. Users can choose which models (known as “physics lists”) they want to run. NOvA uses the QGSP_BERT_HP physics list:

- **QGSP:** The Quark-Gluon String Precompound model simulates the development of high-energy hadronic showers.
- **BERT:** The Bertini cascade is used to model intranuclear cascades of particles with energy less than ~ 9.5 GeV.
- **HP:** High precision set of cross-sections for neutrons with energy less than 20 MeV.

NOvA then uses a simulation-derived template to calculate the photon collection rate as a function of distance along the fibre. The simulation assumes that half of the collected photons go in each direction along the fibre. The simulated light is attenuated as it travels along the fibre, according to results obtained from a test stand.

The final outputs are simulated detector readouts in the same format as data but with the addition of truth information.

3.1.4 GENIE tuning

NOvA applies a custom tune to GENIE simulation to incorporate constraints from internal and external data [64]. The tune consists of two components - the FSI tune and the MEC tune.

NOvA tunes the hN FSI model to match external pion scattering data. The four parameters adjusted are the pion mean free path, and the fraction of pions which undergo absorption, charge exchange, and quasi-elastic scattering. Weights are generated using a boosted decision tree to reproduce the effect of resimulating with altered FSI physics.

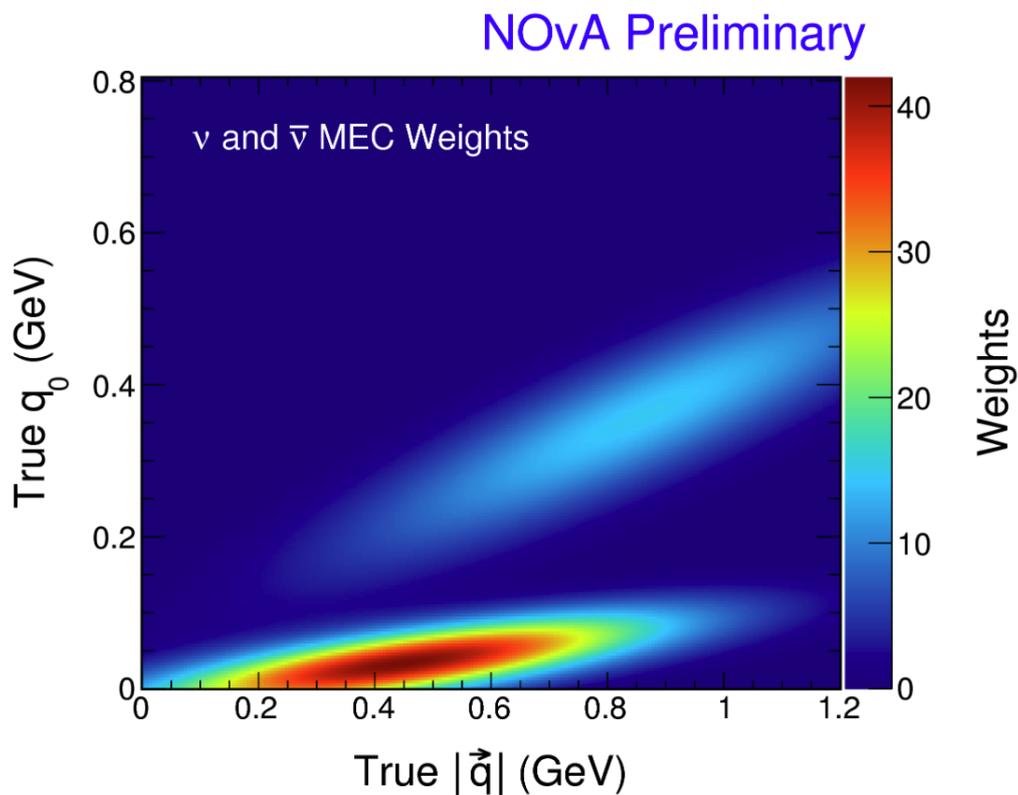


Figure 3.2: Weights applied to MEC events during tuning, as a function of energy and momentum transfer. Figure from [19].

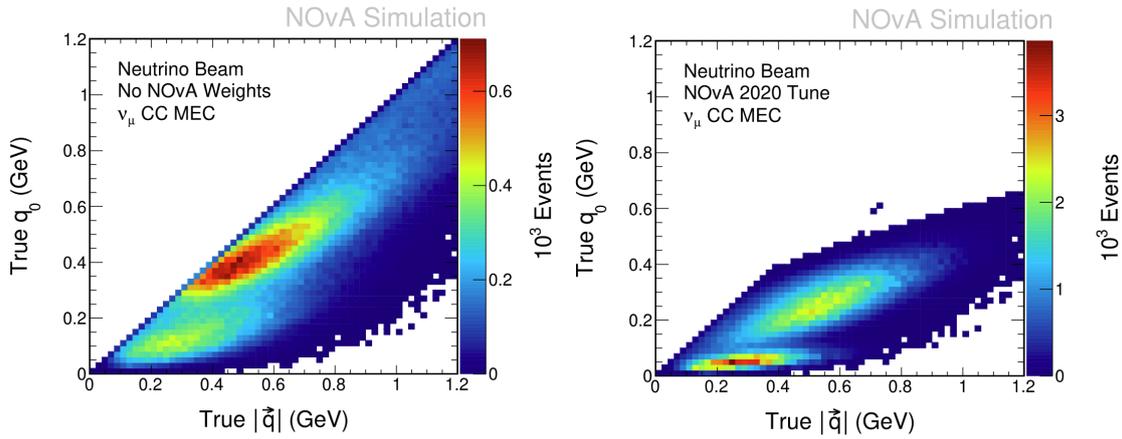


Figure 3.3: Distribution of MEC events in simulation (left) before tuning and (right) after tuning. Figures from [19].

The MEC tune is used to account for the remaining discrepancy with ND data. Here we adjust the normalisation and shape of MEC interactions in simulation using weights. The reweighting is done in energy and three-momentum transfer space ($q_0/|\vec{q}|$). The functional form of the weights is given by the sum of two 2D Gaussian distributions. A fit is performed to find the parameters which make the tuned MC look most like the ND data. For each Gaussian, the fit parameters are the mean and standard deviation in both q_0 and $|\vec{q}|$, the normalisation, and the correlation between q_0 and $|\vec{q}|$. This yields a total of 12 fit parameters. Figure 3.2 shows the weighting function used. Figure 3.3 shows the distribution of MEC events pre- and post-tuning.

3.2 Reconstruction

Reconstruction is the process of translating low-level information (e.g., cell hits) into high-level information (e.g., particle energy, angle, tracks). Reconstruction is done offline, and so NOvA can constantly develop new and improved reconstruction techniques. Reconstruction is critical to all NOvA physics analyses, from oscillation measurements to cross-section measurements to searches for exotic particles. Some reconstruction methods are tuned specifically for certain types of analyses or, in fact, even certain particle species.

Particles leave characteristic energy depositions in the NOvA, as shown in Figure 3.4. Muons tend to produce long, relatively straight tracks. Protons have a high dE/dx and produce short tracks in the NOvA detectors. Charged pions have a similar dE/dx as

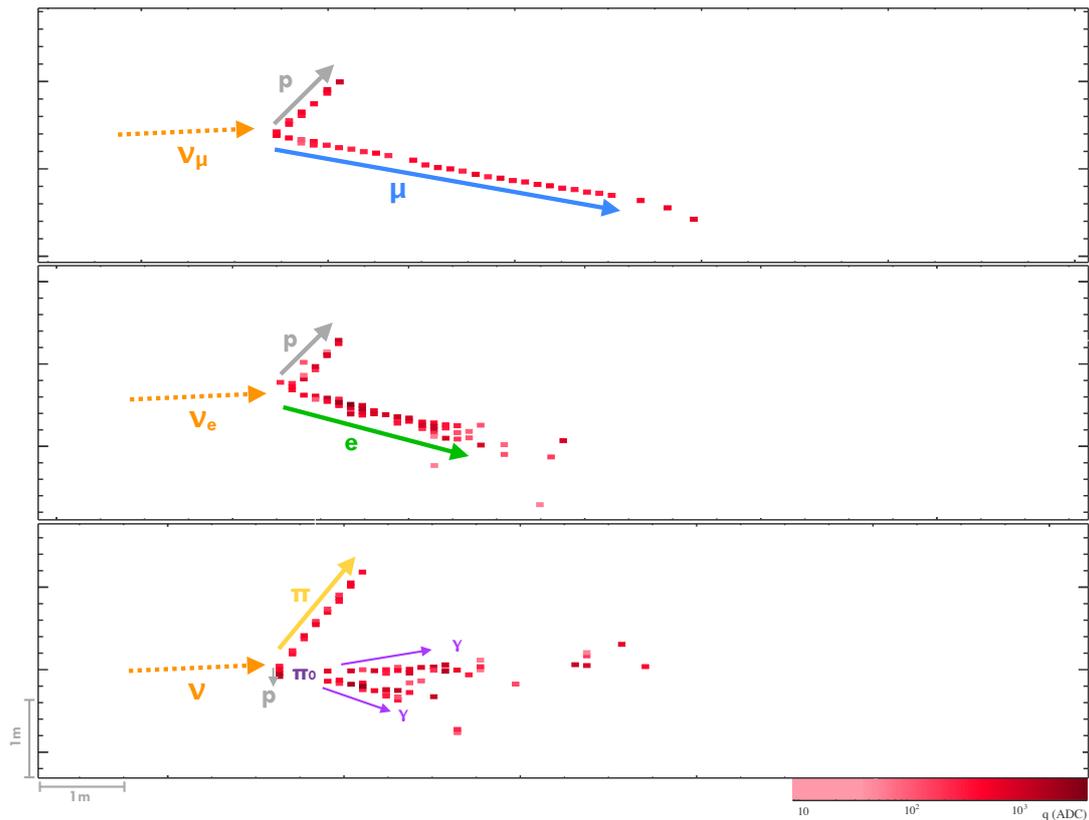


Figure 3.4: Event displays with true particle labels overlaid. We see the characteristic energy depositions of different particles in the NOvA detectors.

muons but are prone to hard scatters off nuclei and hence rarely travel as far as muons. As electrons pass through the detector, they emit photons via bremsstrahlung. These photons travel ~ 40 cm before converting to an e^+e^- pair, which both radiate further photons. As a result, electrons produce showers in the detector. Neutral pions travel some distance before decaying into two photons. Hence they are identified by a pair of electromagnetic showers displaced from the interaction vertex.

Reconstruction algorithms are applied identically to data and MC. This enables the use of MC in place of data for analysis studies. In addition, reconstruction algorithms can be tuned and validated using truth information from MC. This section will describe aspects of the NOvA reconstruction chain, starting with the most low-level and working towards high-level reconstruction.

3.2.1 Hit clustering

The first step in reconstruction is clustering together spatially and temporally correlated hits. In NOvA this is known as “slicing”, where a slice is a group of hits that belong together (e.g. all the hits from a neutrino interaction). The algorithm used is called Time Density Slicer (TDSlicer) [65]. There are three steps in the algorithm:

1. Find centroids in the detector x and y views independently.
2. Build clusters around the centroids by grouping nearby hits.
3. Merge clusters from the x and y views to form a 3D cluster.

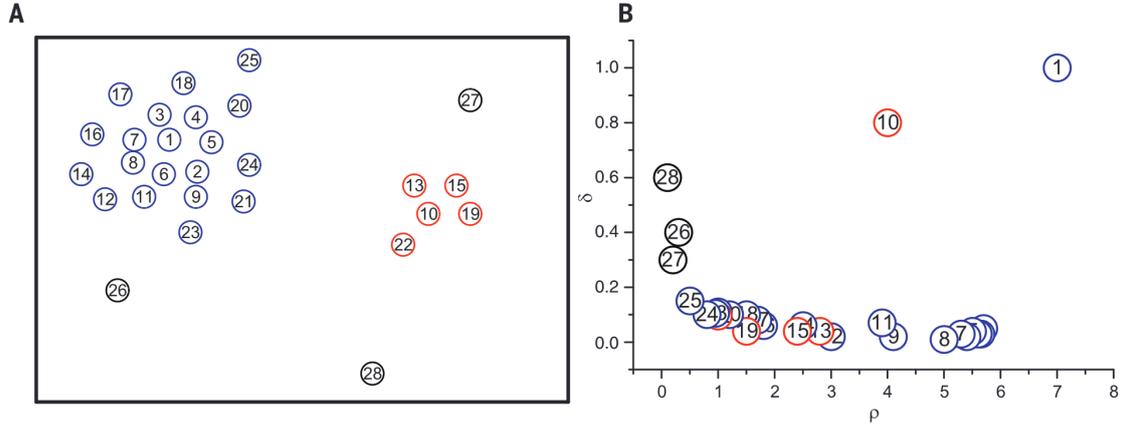


Figure 3.5: (Left) A scatter plot of data to be clustered. Numbers correspond to the density ordering. (Right) A plot of density vs isolation. We see that hits 1 and 10 have a high density and isolation, and hence are identified as centroids. Figure from [20].

Centroids are found in xzt and yzt space independently, using a method described by Rodriguez and Laio [20]. For each hit, the density of surrounding hits is given by

$$\rho_i = \sum_j \exp(-d_{ij}^2) \quad (3.1)$$

where d_{ij} gives the Euclidean distance between a pair of hits i and j in xzt (yzt) space, defined as follows

$$d_{ij} = \frac{||\Delta t| - r/c|}{\tau} \quad (3.2)$$

Here r is the spatial separation, Δt is the time difference, and τ is a configurable scale parameter which is set to 16 (80) ns at the ND (FD). The scale parameter is different for the two detectors because the clock tick for sampling from the DAQ is faster at the ND². In addition to a density score, each hit is assigned an isolation score. The isolation score δ_i for some hit i is the distance to the closest hit with a higher density than hit i

$$\delta_i = \min_{j:\rho_j > \rho_i} (d_{ij}) \quad (3.3)$$

Cluster centroids are found by looking for hits with high density and isolation. See Figure 3.5 for an example.

Once centroids have been identified, any hit i is added to a cluster centred on hit j if $d_{ij} < 10$. The value 10 was chosen by trying different values and studying how resilient the slicer was to pile-up when the beam intensity was increased. Clusters are further built up using Prim's algorithm [66] by successively adding the hit outside the cluster closest to some hit inside the cluster.

After forming clusters in the x and y views, a simple view-merging step is run. Clusters from the x and y detector views are merged by comparing average zt values for all pairs of clusters from opposite views. The output of TDSlicer is a 3D slice. Any slices not successfully matched with a slice in the other view are treated as noise.

Completeness and purity are two standard metrics in reconstruction. In the context of slicing, they are defined as follows

$$\text{completeness} = \frac{\text{Energy deposition from neutrino interaction contained in slice}}{\text{Total energy deposition from neutrino interaction}} \quad (3.4)$$

$$\text{purity} = \frac{\text{Energy deposition from neutrino interaction contained in slice}}{\text{Total energy deposition in slice}} \quad (3.5)$$

For $\nu_\mu\text{CC}$ GENIE interactions passing containment at the FD, both the average

²In the documentation the authors say that the value of the scale parameter is linked to the timing resolution, but the methodology of how the scale parameter values were chosen is not discussed.

completeness and purity are greater than 0.96.

3.2.2 Vertexing

A central assumption when reconstructing neutrino interactions is that the interaction products (leptons and hadrons) radiate outward from the point of neutrino-nucleus interaction. The purpose of vertexing is to estimate the point of interaction by looking at the paths of particles through the detector. Vertexing is done using the elastic arms algorithm [67] [68]. Elastic arms is seeded with the output of a Hough transform [69].

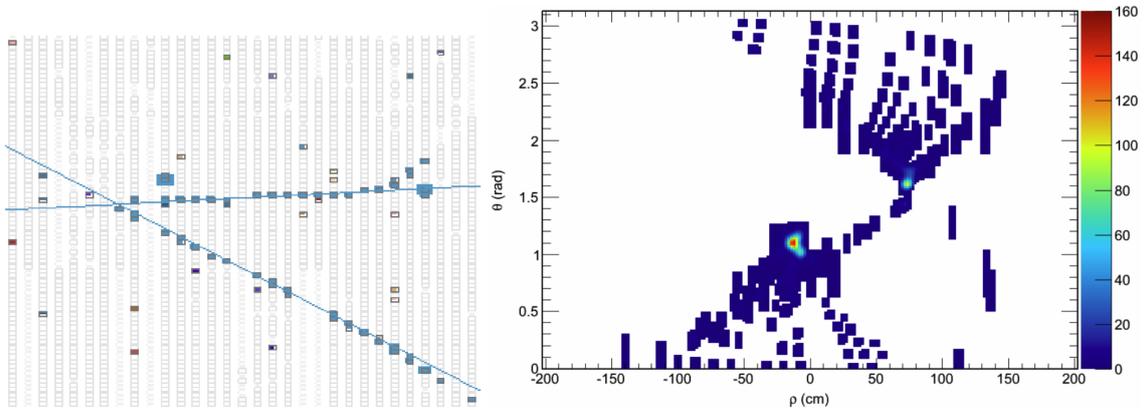


Figure 3.6: (Left) Collection of hits with straight lines overlaid indicating the output of the Hough transform. (Right) Hough parameter space plot with two hotspots visible, showing the location and orientation of the prominent lines in the event. Figure from [21].

The Hough transform is done in the x and y views separately. Lines are drawn between all pairs of hits in the slice and plotted in ρ and θ , where ρ is the perpendicular distance from the line to the origin, and θ is the angle between the line and the z -axis (parallel to the beam). Straight-line features are identified as hotspots in the parameter space plot, as shown in Figure 3.6.

The global interaction vertex is then found using the elastic arms algorithm. We assume that the event is well-described by a single vertex at (x_0, y_0, z_0) with N arms radiating out, each parameterised by a polar and azimuthal angle θ_i and ϕ_i . Any point a distance s along arm i is given by

$$x_i(s) = x_0 + s \sin \theta_i \cos \phi_i \quad (3.6)$$

$$y_i(s) = y_0 + s \sin \theta_i \sin \phi_i \quad (3.7)$$

$$z_i(s) = z_0 + s \cos \theta_i \quad (3.8)$$

The algorithm works to find the parameters $(x_0, y_0, z_0, \vec{\theta}, \vec{\phi})$ that best describe the event. The best fit parameters are found by minimising an “energy function”. This energy function requires proper seeding, which is provided by the output of the Hough transform. The Hough transform sets the number of arms and an initial guess for the vertex position and arm angles.

For ν_μ CC interactions, elastic arms yields a vertex estimate with a resolution of ~ 4 cm in x and y, and ~ 8 cm in z. For comparison, the cell width and depth are 3.8 cm and 5.9 cm respectively.

3.2.3 Prong formation

Particles in the NOvA detectors are described by reconstructed objects called prongs. A prong is a collection of hits emanating from the reconstructed vertex in one direction. The goal is for a prong to contain all the hits a single particle creates as it traverses the detector. Prong formation is done using a modified version of the fuzzy k-means algorithm [70] [22]. Here “fuzzy” describes the membership of hits to a prong. Hits can belong to multiple prongs; hence they have fuzzy membership. In this context, the k in k-means is the number of prongs formed.

We first form 2D prongs in the x and y views separately. The central assumption of prong formation is that hits emanate radially from the vertex. Each hit i in the slice is converted into an angle θ_i with respect to the detector z-axis by drawing a straight line between the hit and the vertex. The prong angles, which are yet to be determined, are given by ω_j . An angle of zero is parallel to the z-axis (approximately parallel to the beam axis). The upper half of the plane is described by angles from zero to π , and the lower half is described by angles from 0 to $-\pi$. Each hit is assigned an angular uncertainty σ based on

the distance s from the hit to the vertex

$$\sigma_i = \frac{1.745}{s} + 0.0204 + 0.000173s \quad (3.9)$$

The coefficients were determined through studies of simulated 1-2 GeV muons. The aim of prong formation is to minimise the angular distance from n hits to k prong centres. The angular distance between hit i and prong j is given by

$$d_{ij} = \left(\frac{\theta_i - \omega_j}{\sigma_i} \right)^2 \quad (3.10)$$

The membership of hit i to prong j is given by

$$\mu_{ij} = \exp \left(- \frac{m d_{ij} \sqrt{k}}{\beta} \right) \quad (3.11)$$

where k is the number of prongs and m is a fuzziness factor set to 2 in NOvA simulation. For larger m , the prong associations become fuzzier. The parameter β is a normalisation factor representing the expected spread of hits around a prong centre (we use $\beta = 4$).

Next, the prong centres are updated as follows

$$\omega'_i = \omega_i + \frac{\sum_{j=1}^n \frac{\mu_{ij}^m}{\sigma_j^2} (\theta_j - \omega_i)}{\sum_{j=1}^n \frac{\mu_{ij}^m}{\sigma_j^2}} \quad (3.12)$$

This iterative process is repeated until the difference between the old and new prong centres is less than 1×10^{-7} radians. This step is followed by a clean-up step to determine if any prongs should be merged together.

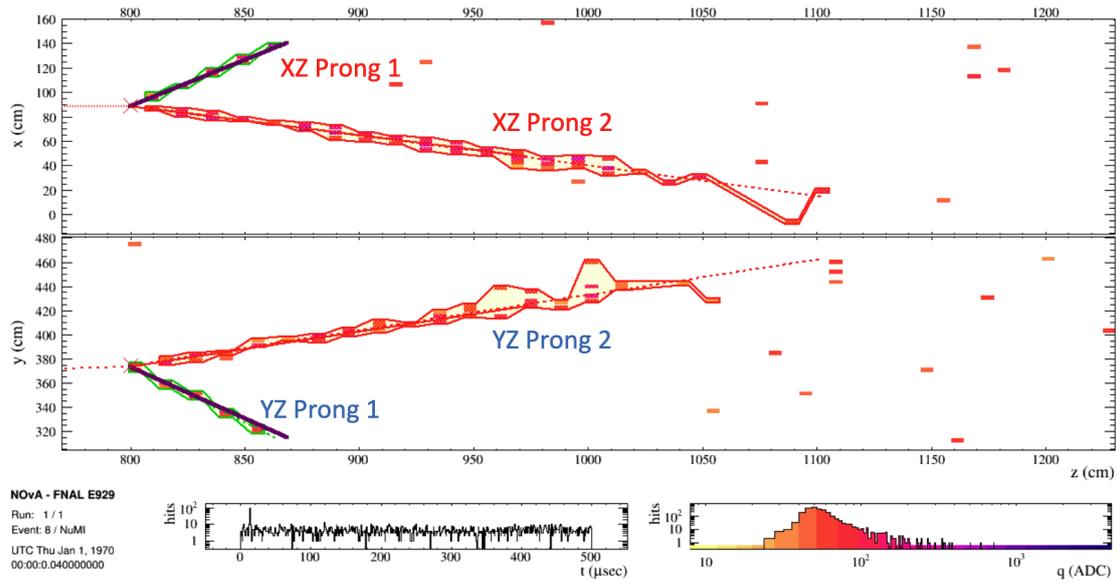


Figure 3.7: A simulated ν_e CC quasielastic interaction in the Far Detector with completed 3D prong reconstruction from the fuzzy k-means algorithm. The prong outlined in red in each view is the reconstructed electron, and outlined in green is the reconstructed proton. The corresponding energy profile histograms used to compute the suitable 3D prong matches is shown in Figure 3.8. Figure and caption from [22].

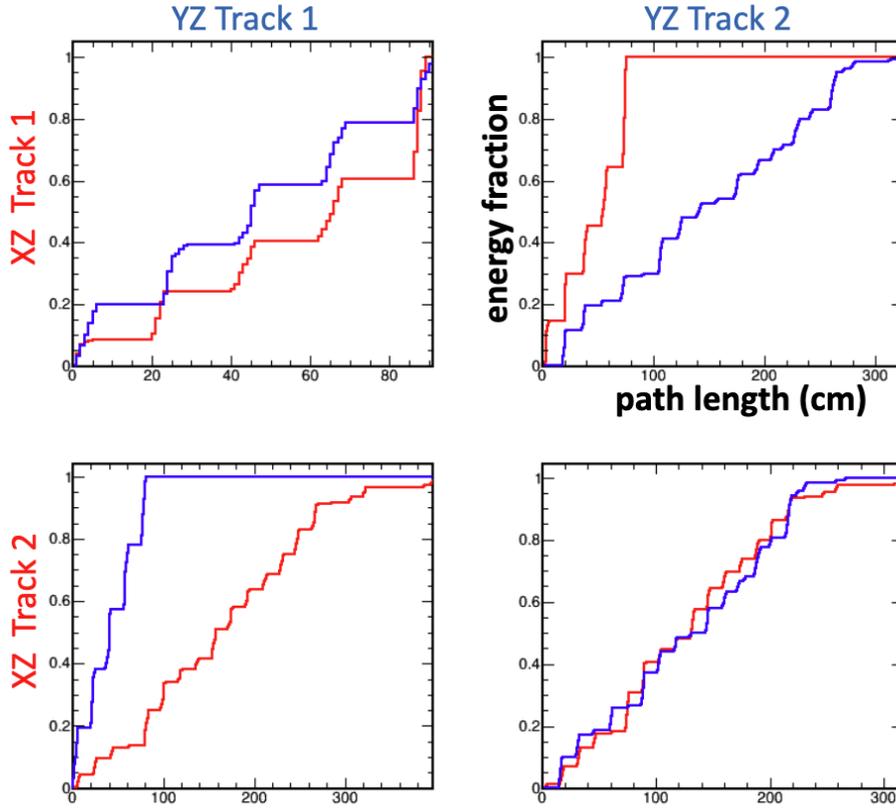


Figure 3.8: Cumulative energy profile as a function of path length along a prong for prospective 3D match candidates shown in Figure 3.7. The red curves are for prongs in the XZ (vertical planes) view, and the blue is for the YZ (horizontal planes) view. The upper-left and lower-right panels show the preferred matches with similar energy profiles that result in the green and red tracks, respectively. The off-diagonal elements illustrate the difference in energy profile shape for the wrong combinations. Figure and caption from [22].

Once 2D prongs have been formed in the x and y views, the next step is to combine prongs from each view to form 3D prongs. This procedure is known as view-matching. For a prong X_i from the x view and a prong Y_j from the y view, we can perform a Kuiper test [71], defined as

$$K_{ij} = D_{ij}^+ + D_{ij}^- \quad (3.13)$$

$$D^+ = \max [E_{X_i}(s) - E_{Y_j}(s)] \quad (3.14)$$

$$D^- = \max [E_{Y_j}(s) - E_{X_i}(s)] \quad (3.15)$$

where $E_{X_i}(s)$ and $E_{Y_j}(s)$ are the fractional energy depositions in the respective detector views as a function of the distance s along the prong. See Figures 3.7 and 3.8 for an event

display and fractional energy depositions. Here a lower value of K_{ij} is better. The pair of prongs which yield the lowest score are merged into a single 3D prong and removed from further consideration. Then the next best pair of prongs are merged. This process is repeated until all 2D prongs have been merged into 3D prongs. If there is an unequal number of prongs formed in the x and y views, then some 2D prongs will be orphaned³ at the end of this process.

Reconstructed prongs for simulated 1 GeV muons have an average completeness and purity greater than 0.9.

3.2.4 Particle identification

NOvA employs machine learning methods in its oscillation measurements. A convolution neural network is used to identify event topologies based on pixel maps, which are collections of all hits in an event arranged as an image [72]. Within NOvA, the name Convolutional Visual Network (CVN) is given to trained models of convolutional neural networks.

The versions of CVN used in oscillation measurements are trained on entire simulated GENIE neutrino interactions. As a result, the network outputs may be biased towards the underlying MC model. This makes the network unsuitable for use in measurements which are explicitly trying to constrain these models, for example, cross-section measurements. For this reason, a single particle trained version of CVN was created [73]. This network performs particle identification on a prong-by-prong basis. No context information from the rest of the event is used to compute network outputs. To avoid generator bias, the network was trained on simulated single particles passing through the detector. The particles were simulated with uniform kinematics to avoid assumptions about the link between kinematics and particle type. The network was trained on electrons, muons, protons, photons, and charged pions⁴. The network output is five scores, one for each particle species the network was trained on. For each prong, the sum of all five scores equals 1.0. The higher the score is for a given particle category, the more confident CVN is that the prong in question belongs to that category.

³An orphaned prong is not matched with a prong in the other view and hence remains as a 2D prong.

⁴There was a 50/50 split between π^+ and π^-

3.3 Calibration

The NOvA detectors must be calibrated such that a detector reading can be consistently translated into an amount of energy in GeV. The calibration procedure is split into relative and absolute parts. The relative calibration corrects ADC readings to ensure that a fixed energy deposition looks the same in every part of the detector. The absolute calibration translates detector readings into GeV.

3.3.1 Relative calibration

For a fixed energy deposition within a given cell, the number of photoelectrons (PE) produced at the APD decreases as the distance from the readout increases. This is due to the attenuation of the signal as it travels along the fibre. The relative calibration aims to correct for this effect so that a consistent reading is obtained no matter where energy is deposited in the cell. This relative calibration is done on a cell-by-cell basis to account for natural variations in electronics, scintillator, and PVC reflectivity.

Energy deposition of cosmic ray muons is well-described by the Bethe-Bloch equation, and so they can be used as standard candles. Only through-going muons are used for the relative calibration to avoid the sharp rise in energy deposition at the end of the muon track. We call the distance along the length of the cell W , where $W = 0$ is the centre of the cell, and large W are closer to the readout. The purpose of the relative calibration is to describe the attenuation of the signal as a function of W .

Calibration is done using so-called tri-cell hits. A tri-cell hit is a hit where the two adjacent cells in the same plane also have a hit, as shown in Figure 3.9. Tri-cell hits are used because it simplifies the calculation of the particle path length through the cell. In a tri-cell hit, we know the particle must have travelled through two opposite cell walls. Hence, given the angle of the particle track, we can determine the path length through the cell. This allows us to convert PE to PE/cm⁵. With reference to Figure 3.9, the path length l is given by

$$l = \frac{d}{\cos \theta} \tag{3.16}$$

⁵The energy deposition in a given cell increases approximately linearly with path length through the cell, and we need to account for this.

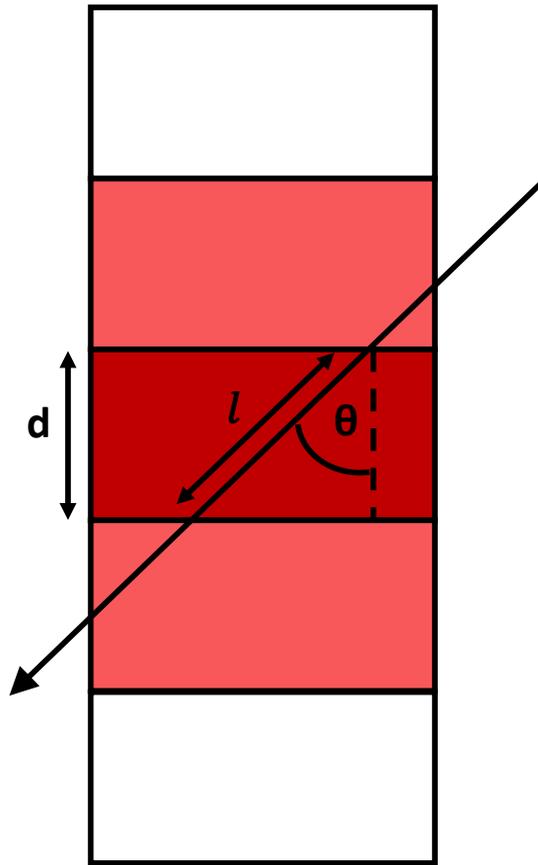


Figure 3.9: Schematic of a tri-cell hit, with the cell width d marked. A particle trajectory is overlaid, with the path length through the cell l marked.

For each cell, a plot of PE/cm vs W is made, as shown in Figure 3.10. We see the number of photoelectrons per centimetre increases with W . A fit of the form

$$y = C + A \left(e^{\frac{W}{X}} + e^{-\frac{L+W}{X}} \right) \quad (3.17)$$

is done, where y is PE/cm, L is the cell length, and A , C , and X are fit parameters. The second exponential term accounts for the path of photons which travel down the fibre away from the readout before looping back to the readout. This fit is done for values of W more than 50 cm from the end of a cell. The remaining range of W is known as the “roll-off” region, and locally weighted scatterplot smoothing (LOWESS) [74] is used to capture the behaviour in this region. The results of both of these fits can be seen in Figure 3.10.

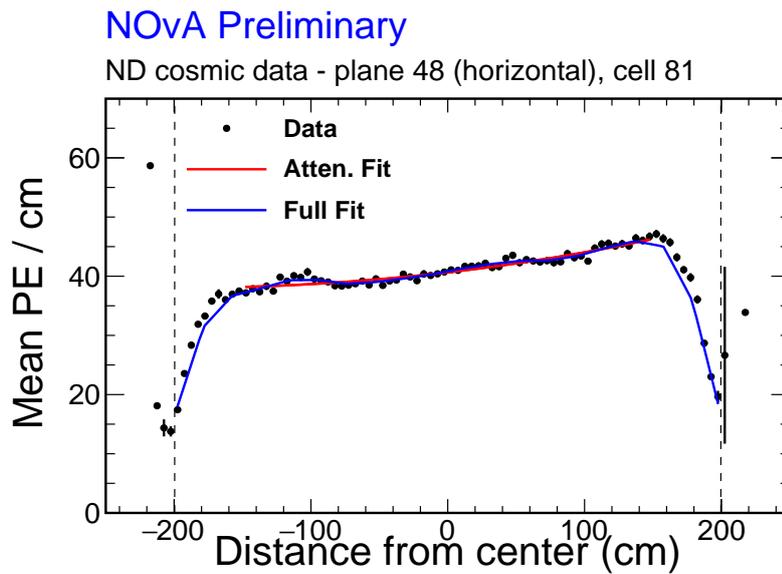


Figure 3.10: PE/cm vs W for a single cell at the Near Detector. A fitted function of the form in Equation 3.17 is shown in red. An additional fit, shown in blue, is done to account for the “roll-off” near the edges.

3.3.2 Absolute calibration

The absolute calibration takes the corrected PE/cm from the relative calibration and determines a scale factor to convert it into GeV. The absolute calibration also uses tri-cells and cosmic ray muons. However, the absolute calibration relies on stopping muons. Muons are in the minimum ionising particle (MIP) region between 100-200 cm from the end of the track. By only using stopping muons, we can ensure we only calibrate using the well-understood MIP portions of the muon track. By comparing energy depositions from data and MC (in corrected PE/cm and GeV respectively), a conversion from photoelectrons to energy is determined.

3.4 Summary

This chapter gives an overview of NOvA’s offline software, which is critical for interpreting collected data. The calibration and reconstruction techniques described enable high-level analysis of neutrino interactions. Monte Carlo simulation, among other things, allows direct comparison of measurements to model predictions. Additions I have made to NOvA’s software suite are described in Chapters 4 and 5. The offline software is integral

to the cross-section analysis described in Chapter 6.

Chapter 4

OPALProng: An improved prong view-matching algorithm

This chapter describes an extension to the prong reconstruction algorithm described in Section 3.2.3. This reconstruction algorithm was developed concurrently with the cross-section analysis described in Chapter 6. The reconstruction algorithm described in this chapter is not yet available in the full NOvA data and simulation sample. As a result, this work is not used in the cross-section analysis and is independent of the central focus of this thesis.

4.1 Motivation

As described in Section 3.2.3, the fuzzy k-means algorithm forms 2D prongs around a reconstructed vertex. This is done in the x and y detector views independently. To then form 3D prongs, a process known as view-matching is performed, whereby 2D prongs in the x view are matched with 2D prongs in the y view.

The view-matching algorithm has two notable issues. Firstly, it is a greedy algorithm¹, meaning that it makes the best decision it can at each step without regard for the best global solution. In this context, this means that, in making the best prong match in an event early on, we may be forced to make some bad prong matches later.

The second issue with the current view-matching algorithm is that it performs a 1-to-1 matching of prongs. This can lead to poor reconstruction in cases where two particle

¹Greedy is an established technical term in this context, rather than a descriptor.

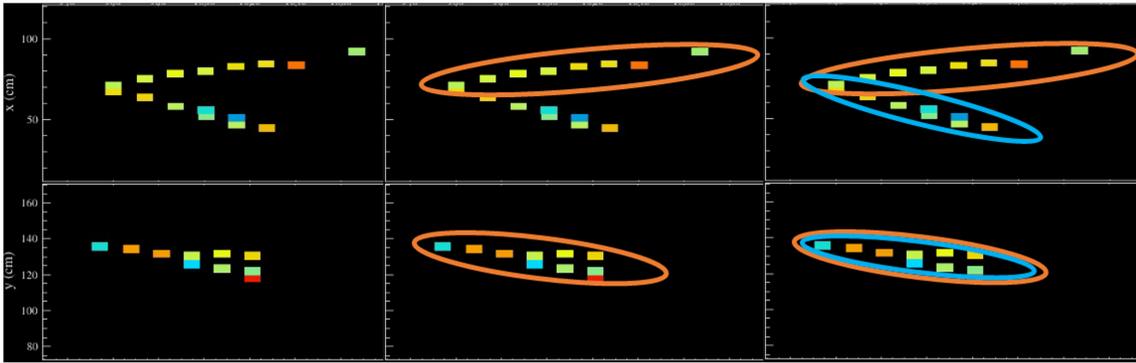


Figure 4.1: **Left:** An event display where two particle tracks are visible in the x view, but the tracks overlap in the y view. **Centre:** Example reconstruction under the standard view-matching algorithm. One prong is formed (orange), but we are forced to orphan one 2D prong. **Right:** Reconstruction possible with OPALProng. Both particles are reconstructed. Hits in the y view are shared between the orange and blue prongs.

tracks² are distinct in one detector view but overlap in the other. This can be seen in the centre of Figure 4.1, where some hits in the x view do not form a 3D prong. We say that the 2D prong composed of these hits is “orphaned”.

Overlapping **P**ArticLe (OPAL) reconstruction was created to address the two issues highlighted above. The first issue is addressed by evaluating the quality of view-matching on a global (whole event) level. The second issue is addressed by allowing a 2D prong in one view to be matched to two 2D prongs in the other view. A more in-depth explanation can be found in Section 4.2.

4.2 Methods

4.2.1 Evaluating prong match quality

Just as in the standard view-matching algorithm, a Kuiper test [71] is used to evaluate prong match quality. However, in the OPALProng framework, it is performed slightly differently. For each 2D prong, the cumulative energy fraction as a function of z is calculated (previously, this was calculated as a function of distance s along the prong). For a given pair of 2D prongs X_i and Y_j the cumulative energy fractions $E_{X_i}(Z)$ is shifted such that it begins at the same value of z as $E_{Y_j}(Z)$ ³.

²I am using the word “track” here to generically describe the path of a particle.

³This is to account for the fact that the z -components of the start point of x and y prongs will differ by at least one plane width.

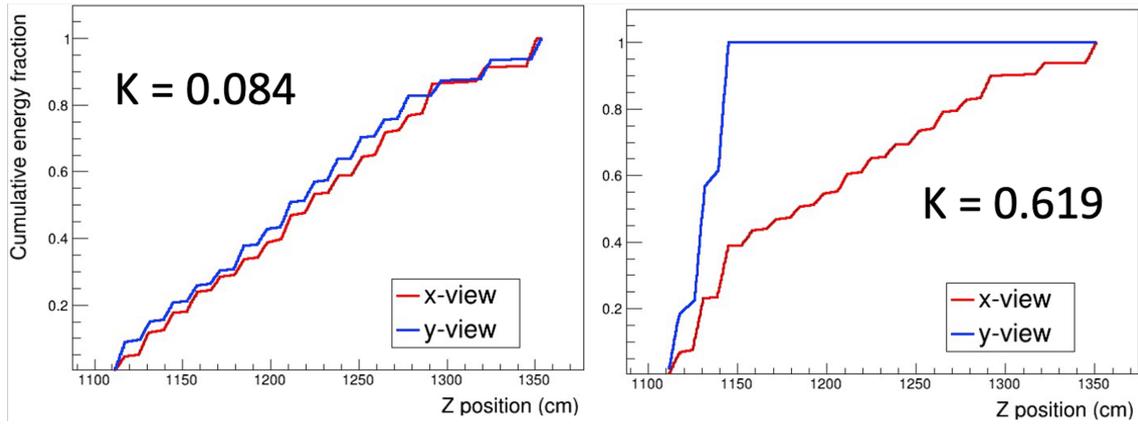


Figure 4.2: **Left:** Cumulative energy fraction of a pair of x and y prongs. The Kuiper score is low, so this is a good match **Right:** Cumulative energy fraction of a different pair of x and y prongs. The Kuiper score is high, so this is a bad match.

This shift can be described as follows

$$\begin{aligned}
 E_{X_i}(Z) &\rightarrow E'_{X_i}(Z) = E_{X_i}(Z + \Delta Z) \\
 \Delta Z &= Z_{Y_j}^0 - Z_{X_i}^0
 \end{aligned} \tag{4.1}$$

where Z^0 gives the z coordinate at the beginning of a prong. A score K is obtained by performing a Kuiper test on $E'_{X_i}(Z)$ and $E_{Y_j}(Z)$

$$\begin{aligned}
 K_{ij} &= D_{ij}^+ + D_{ij}^- \\
 D_{ij}^+ &= \max(E'_{X_i}(Z) - E_{Y_j}(Z)) \\
 D_{ij}^- &= \max(E_{Y_j}(Z) - E'_{X_i}(Z))
 \end{aligned} \tag{4.2}$$

The smaller the value of K_{ij} , the better the match is, as shown in Figure 4.2.

A Kuiper test is also performed for pairs of prongs in one view with a single prong in the other view, e.g. (X_i, X_j) and Y_k . The procedure is the same, but the cumulative energy fraction of the sum of the two prongs is used when performing the Kuiper test. In this way, it is possible to evaluate if the best solution involves, for example, matching one y prong to two different x prongs.

	y₁	NULL
x₁	0.967	1.00
x₂	0.143	1.00
(x₁, x₂)	0.148	
NULL	1.00	

Table 4.1: Kuiper score table for an example with two prongs in the x view and one prong in the y view. For every combination of prongs in one view paired with prongs in the other view, a Kuiper score is calculated. Some (nonsensical) combinations are never considered and hence are greyed out. The NULL row and column are included to allow prongs to be orphaned.

4.2.2 Finding a global solution

To find the best global solution, an aggregate metric is calculated for every possible way of matching all the prongs in the x view with all the prongs in the y view (let us call an allowed way of matching all prongs in an event a “solution”). First, a table of Kuiper scores is constructed for every match of one prong with another prong or pairs of prongs with another prong. An example is shown in Table 4.1. A row and column labelled NULL are included to allow the algorithm to leave some prongs orphaned. The NULL row/column is filled with a default score of 1.0. This corresponds to a bad Kuiper score to disincentivise orphaning prongs.

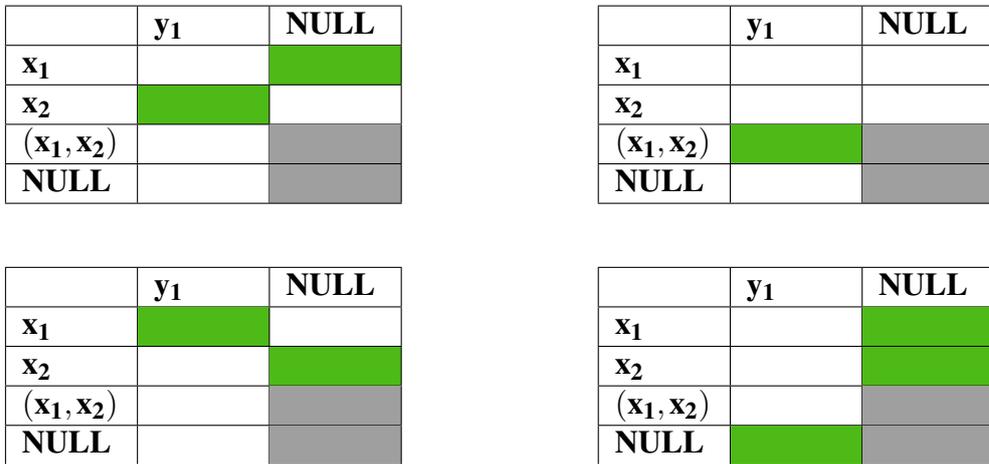


Figure 4.3: The four possible solutions when matching two x prongs with one y prong. Each table is one solution; boxes in green represent matches made

It is important to keep track of every possible allowed solution. Some solutions are nonsensical; for example, it is not possible to match a pair of x prongs with a pair of y prongs, nor is it possible to match X_i with Y_j and also match X_i with NULL (orphaning

X_i). So-called “match tables” are used to keep track of allowed solutions, as shown in Figure 4.3. A set of match tables is generated ahead of time and then referenced during reconstruction.

Each solution in the match tables is given a score equal to the mean of the Kuiper scores in the green boxes. The best global solution is the one which has the lowest mean Kuiper score. This is the solution we accept. For example, if we take Table 4.1 and test each of the match tables shown in Figure 4.3 starting with the top left and going clockwise, we find that the global metric is 0.572, 0.148, 1.00, and 0.984. The top right match table yields the lowest score, so we determine that prong y_1 should be matched with both x_1 and x_2 .

The match tables become extremely large for large numbers of x and y prongs. The number of rows (columns) N in a match table is given by

$$N = \frac{n(n-1)}{2} + n + 1 \quad (4.3)$$

where n is the number of x prongs (y prongs) in the event. The first term is the binomial coefficient $\binom{n}{2}$, i.e. one row (column) for each way of making a pair of prongs from the same view. The second term yields one row (column) for each prong by itself. The final term adds one extra row (column) for the NULL match. Due to the difficulty of generating match tables for large n , the preexisting view-matching algorithm is used for an event if the number of prongs in either view is greater than five.

It is also worth noting that we only attempt an overlapping match if a prong has three hits or more. This is because it is difficult to assess whether, for example, a match is better with a 1 hit overlap than without it.

4.2.3 Determining number of hits to overlap

If a short track overlaps with a long track, as seen in Figure 4.4, it is necessary to determine how many hits from the overlapping view are included in the resulting short 3D prong. We assert that the track will most likely be roughly as long in the y view as it is in the x view, hence if the prong ends at plane P_n in the x view, it should end at either plane P_{n-1} or P_{n+1} in the y view. To decide which plane to end at in the y view, we form one 3D prong ending

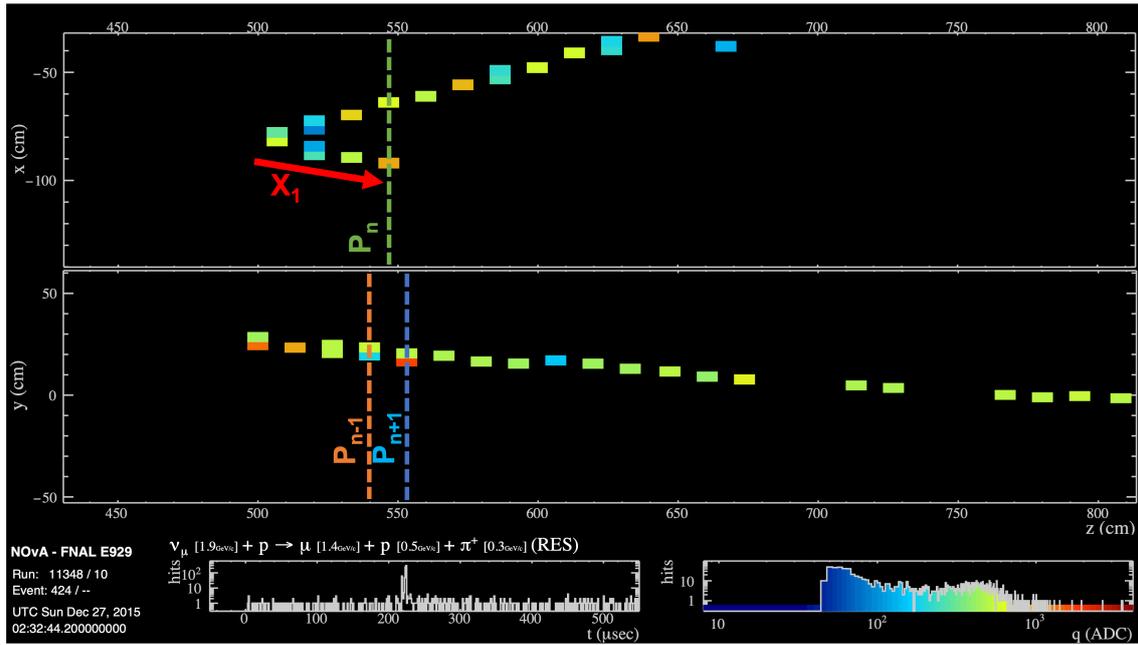


Figure 4.4: An event display of an interaction which will be reconstructed with overlapping prongs. It is necessary to determine how many hits in the y view will be included when making a 3D prong with the short 2D prong marked X_1

at P_{n-1} and another ending at P_{n+1} . A Kuiper score is calculated for each of these prongs. The prong with the smaller Kuiper score is saved, and the other prong is discarded.

4.2.4 Sharing hits

Now that overlapping prongs can be formed, it is necessary to share hits between prongs, i.e. for a given hit H_i , $X\%$ of the hit belongs to prong A, and $(100-X)\%$ of the hit belongs to prong B. This is done by weighting the hits within a prong.

To demonstrate how weights are calculated, consider an example. A 2D prong Y_i is matched with both prong X_j and X_k to form two 3D prongs, α and β respectively. The weight of the hit(s) in Y_i in the n -th plane $Y_{i,n}$ is given as follows

$$\begin{aligned}
 w_{n,\alpha} &= \frac{E_{\text{cal}}(X_{j,n-1}) + E_{\text{cal}}(X_{j,n+1})}{E_{\text{cal}}(X_{j,n-1}) + E_{\text{cal}}(X_{j,n+1}) + E_{\text{cal}}(X_{k,n-1}) + E_{\text{cal}}(X_{k,n+1})} \\
 w_{n,\beta} &= \frac{E_{\text{cal}}(X_{k,n-1}) + E_{\text{cal}}(X_{k,n+1})}{E_{\text{cal}}(X_{j,n-1}) + E_{\text{cal}}(X_{j,n+1}) + E_{\text{cal}}(X_{k,n-1}) + E_{\text{cal}}(X_{k,n+1})}
 \end{aligned} \tag{4.4}$$

where E_{cal} gives the calorimetric energy deposited by a prong in a given plane. In other

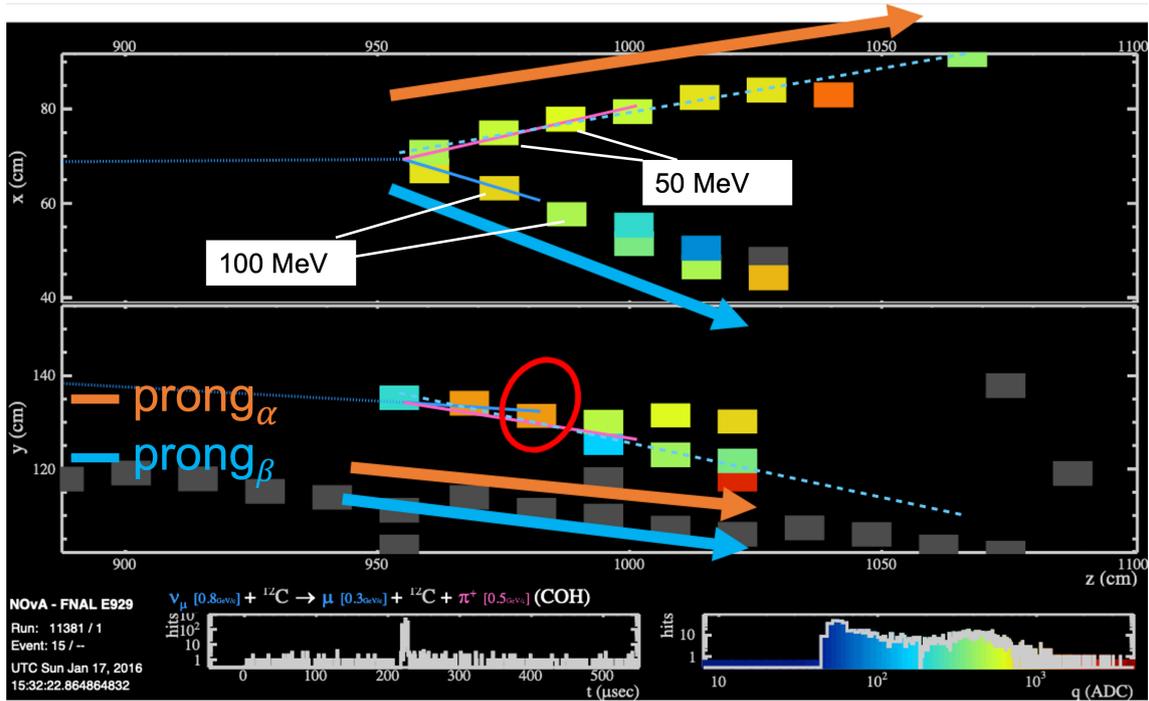


Figure 4.5: Toy example showing a hit in the y view (circled in red), which must be shared between two 3D prongs α and β . The upper prong in the x view deposits 50 MeV of calorimetric energy in both planes on either side of the hit to be weighted. The lower prong deposits 100 MeV in both of these planes. With reference to Equation 4.4 and Equation 4.6 we see that this hit will be shared with weights $w_{n,\alpha} = 0.33$ and $w_{n,\beta} = 0.67$.

words, the weight for prong α is given by the fraction of energy deposited by X_j in the two neighbouring planes. The philosophy here is that whichever prong deposits more calorimetric energy in the non-overlapping view (in this case, the x view) takes a larger fraction of the hit(s) $Y_{i,n}$ when forming a 3D prong. Note that Equation 4.4 satisfies the requirement

$$w_{n,\alpha} + w_{n,\beta} = 1.0 \tag{4.5}$$

The procedure outlined in Equation 4.4 can be further elucidated with a toy example as shown in Figure 4.5. Using the values from this example, we find

$$w_{n,\alpha} = \frac{50 + 50}{50 + 50 + 100 + 100} = \frac{100}{300} = 0.33$$

$$w_{n,\beta} = \frac{100 + 100}{50 + 50 + 100 + 100} = \frac{200}{300} = 0.67$$
(4.6)

The weights as calculated in Equation 4.4 must be included when calculating prong completeness and purity. These quantities are calculated with reference to the “leading particle” (LP), which can be understood as the true particle that the prong “belongs” to (i.e. the one which deposits the most energy in the prong). In the case of overlapping prongs, the LP is defined as the particle which deposits the most energy in the non-overlapping view. Completeness and purity are calculated as follows

$$\text{comp} = \frac{E_{LP} \text{ in prong}}{\text{Total } E_{LP} \text{ in detector}} = \frac{\sum_{\text{prong hits}} \min [(E_{\text{Total}} * w), E_{LP}]}{\sum_{\text{all hits}} E_{LP}}$$

$$\text{pur} = \frac{E_{LP} \text{ in prong}}{\text{Total } E \text{ in prong}} = \frac{\sum_{\text{prong hits}} \min [(E_{\text{Total}} * w), E_{LP}]}{\sum_{\text{prong hits}} E_{\text{Total}} * w}$$
(4.7)

where E_{Total} is the total deposited true energy in a given hit and E_{LP} is the true energy deposited by the LP in a given hit. This has the desired property that under-weighting gives high purity but low completeness, and over-weighting gives low purity but high completeness. Note that Equation 4.7 reduces to the preexisting way of calculating completeness and purity when all weights are equal to 1.0. Also note that if the “min []” terms in Equation 4.7 were not included, it would be possible to have purity/completeness greater than 1.0, which is nonsensical.

4.3 Evaluation

4.3.1 Results

Species	$N_{\text{OPAL}}/N_{\text{Regular}}$
muon	1.031
proton	1.097
piplus	1.116
piminus	1.093
photon	1.040
electron	1.036

Table 4.2: With reference to Figures 4.6-4.11, ratio of the number of prongs formed under the OPAL and regular view-matching algorithms

Shown in Figures 4.6-4.11 are an assortment of variables reconstructed under the “regular” (preexisting) and OPAL algorithms. All plots show only simulated events. Containment and fiducial cuts were used to make these plots. Note that Figure 4.7 is presented with a logarithmic y-axis scale. The containment and fiducial limits can be found in Equations 6.7 and 6.6, respectively, and were optimised for the cross-section analysis described in Chapter 6⁴. These figures show an increased number of prongs formed under the OPAL algorithm for all particle species.

Figures 4.8 and 4.9 show that the most improvement is seen at low energies and short prong lengths. These represent particles which previously may have partially overlapped with a more energetic particle in the event. Figure 4.10 shows that an improvement is seen for forward-going particles. The muon is frequently quite forward-going, so in the preexisting reconstruction paradigm other forward-going particles would often overlap.

In Figure 4.8 it can be seen that for large calorimetric energies the bin content is sometimes lower in the OPAL histogram than the “regular” histogram. Consider, for example, the muon plot in 4.8. It is incorrect to assume that some high energy muons are reconstructed under the preexisting algorithm but not under the OPAL algorithm. Rather, a change in the way calorimetric energy is calculated under the OPAL reconstruction causes a migration of prongs to lower calorimetric energy. Figure 4.12 shows that often the calorimetric energy is reduced under the OPAL reconstruction. The reason calorimetric

⁴To make these plots any reasonable fiducial and containment limits are acceptable. The fact that they were optimised for the cross-section analysis is not important.

energy is reduced can be elucidated with an example. Consider an event with a muon and a proton. The muon forms a long prong, and the proton overlaps with the muon in one of the detector views. Under the preexisting reconstruction, the proton is unreconstructed, and the muon prong contains additional calorimetric energy in the overlapping view coming from the proton. Under the OPAL reconstruction, both a muon and proton prong are formed, and the muon prong does not contain additional calorimetric energy as hits in the overlapping view are shared between the muon and proton prongs using hit weights as described in Section 4.2.4. As a result, the calorimetric energy of the muon prong is reduced to more accurately reflect the energy deposition of the muon.

Table 4.2 summarises the change in the number of prongs. The largest improvements are seen for protons and charged pion because OPAL does better than the regular algorithm for short prongs. Protons and charged pions tend to create shorter prong than other species, hence they see the largest improvements in reconstruction. Studies have shown that if a comparison is made of only the number of prongs shorter than 100 cm, then there is a comparable improvement in the number of prongs for most particle species.

4.3.2 Event Displays

Shown in Figures 4.13-4.15 are some comparisons of the OPALProng reconstruction and the preexisting reconstruction. An additional 3D prong is formed in each case when OPALProng is used. Note that there is currently no way of visually representing overlaps in the event display.

4.4 Summary

This chapter describes an improved algorithm for reconstructing prongs in NOvA. This algorithm makes prong matches based on the best global solution. Additionally, this algorithm allows a prong in one detector view to be matched with two prongs in the opposite view, thus allowing the reconstruction of overlapping particles. As a result, more particles of every type are reconstructed. This may allow for more events to be analysed, leading to smaller statistical uncertainties for many analyses. In addition, this algorithm improves NOvA's sensitivity to certain regions of phase space, for example forward-going pions.

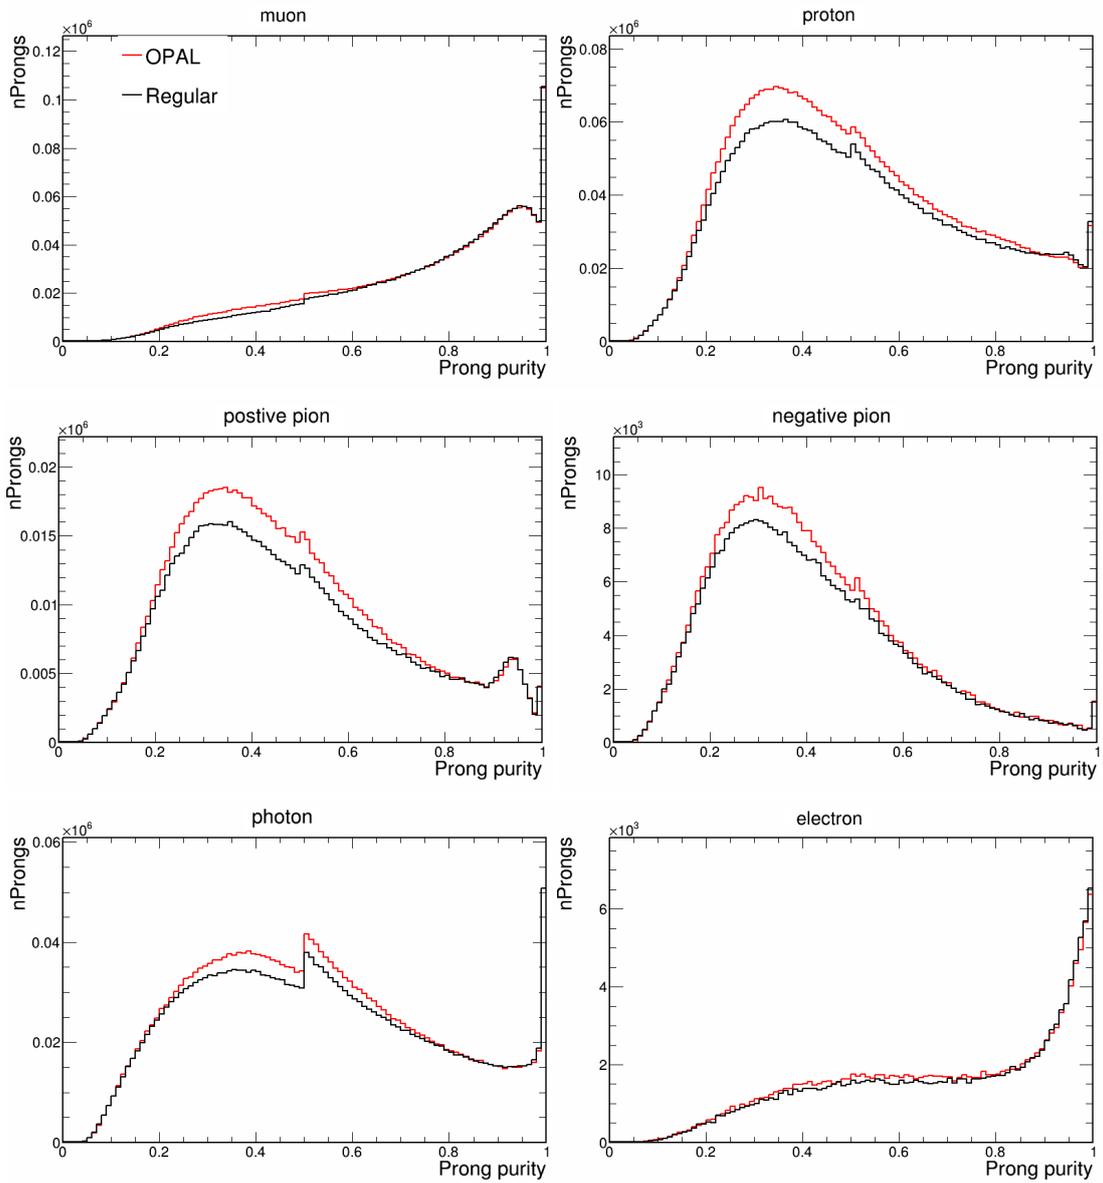


Figure 4.6: Prong purity in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms.

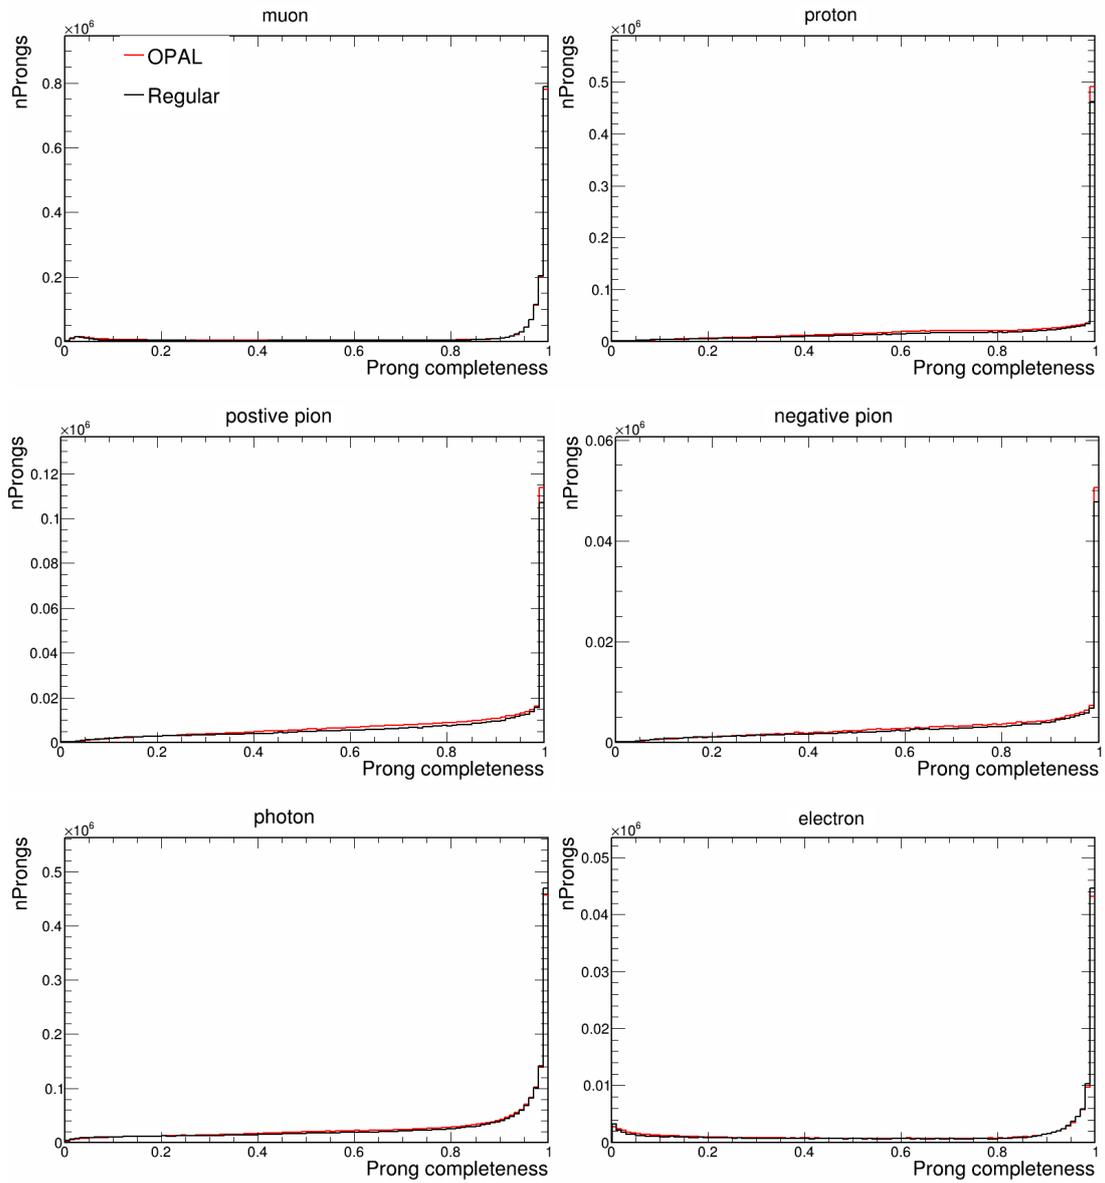


Figure 4.7: Prong completeness in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms.

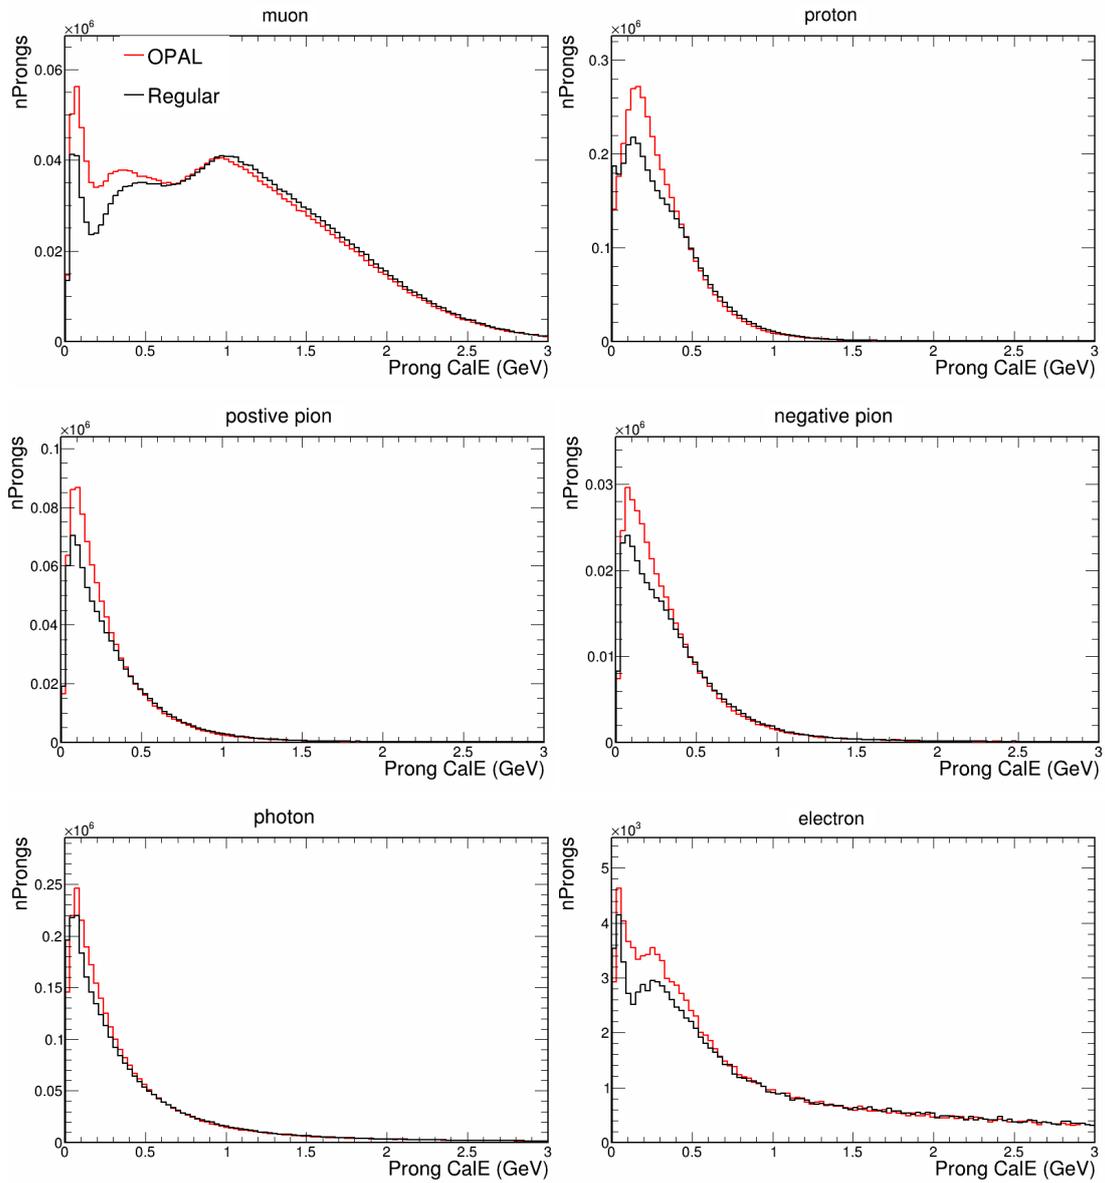


Figure 4.8: Prong calorimetric energy in simulated events, reconstructed using the “regular” (pre-existing) and OPAL reconstruction algorithms.

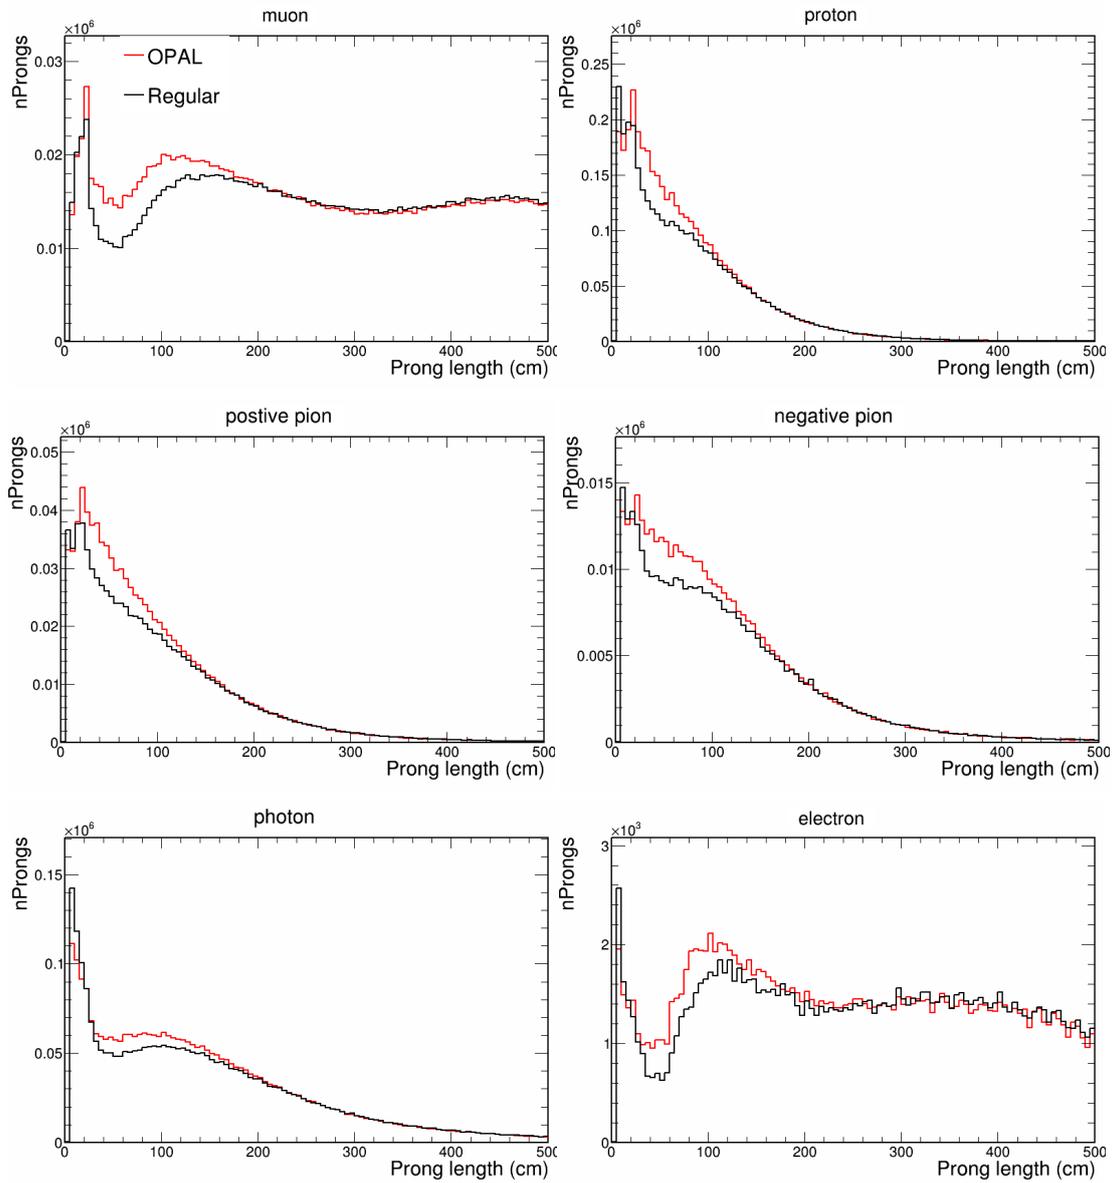


Figure 4.9: Prong length in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms.

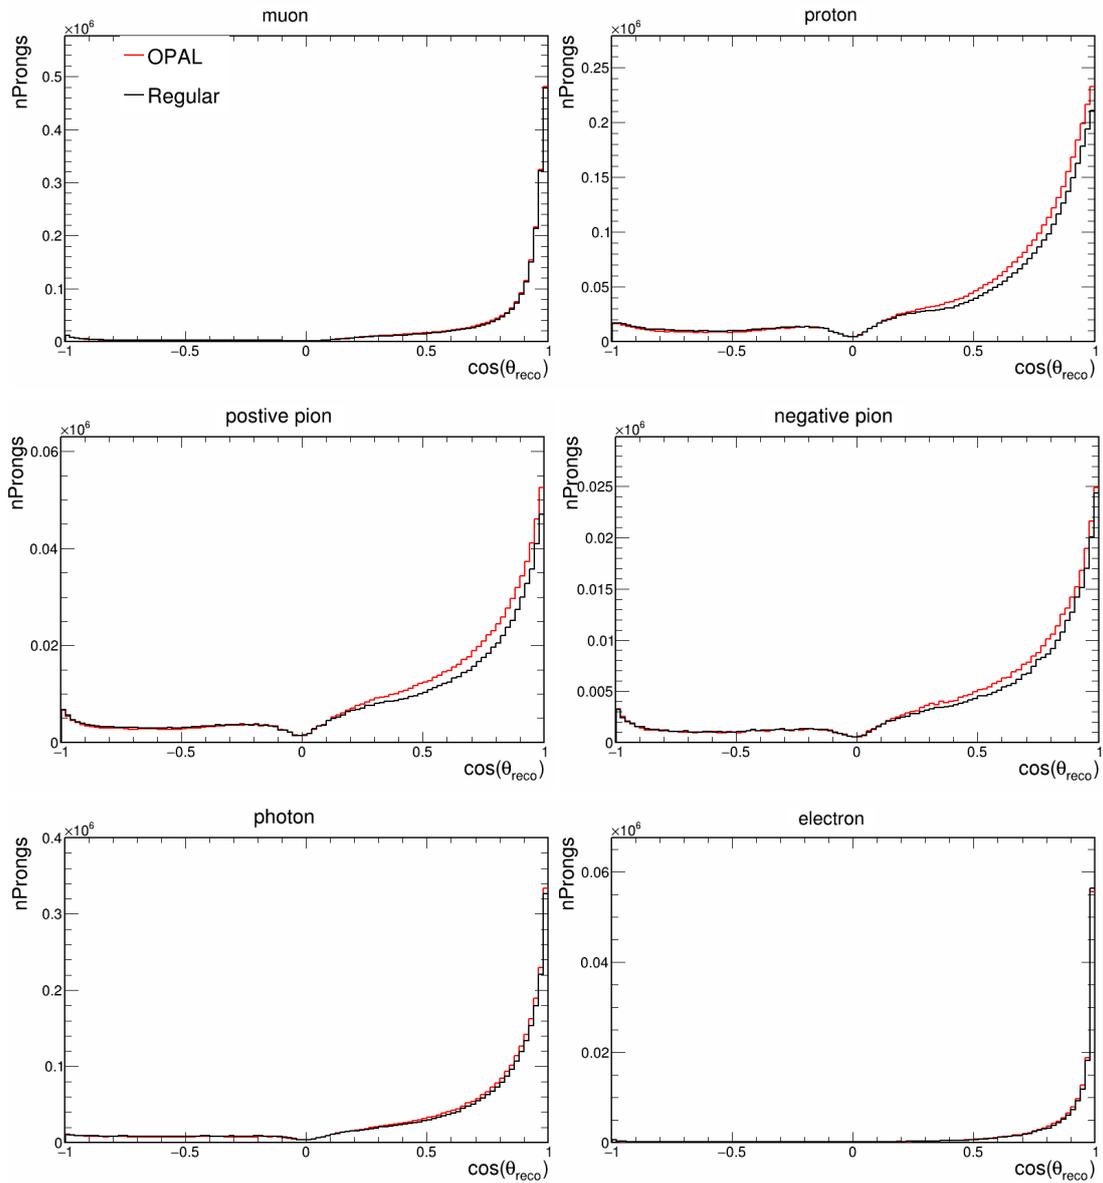


Figure 4.10: Cosine of reconstructed prong angle w.r.t. the beam in simulated events, reconstructed using the “regular” (preexisting) and OPAL reconstruction algorithms.

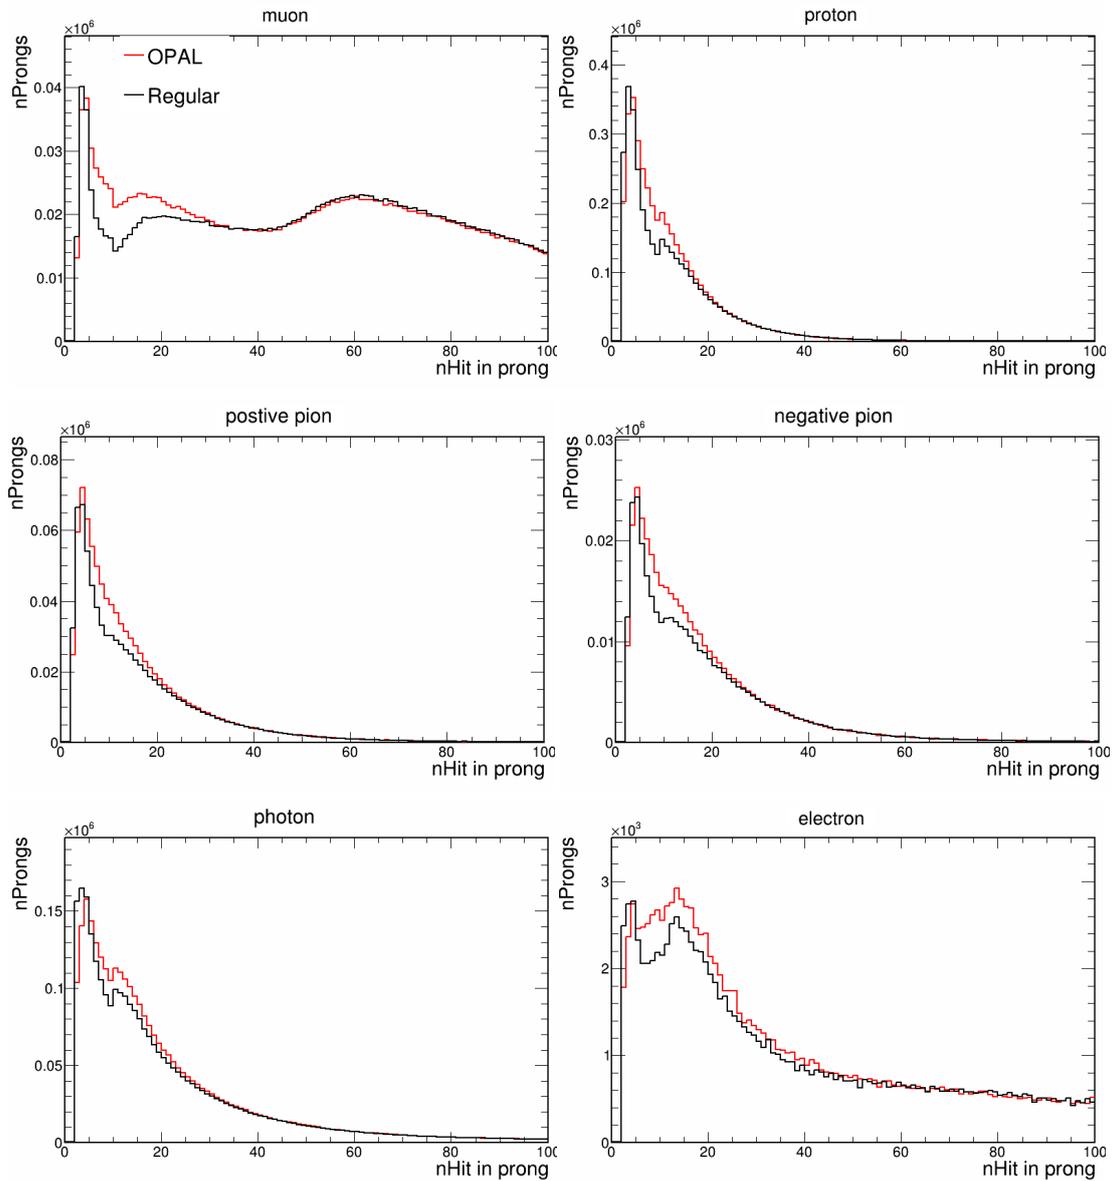


Figure 4.11: Number of hits in prong in simulated events, reconstructed using the “regular” (pre-existing) and OPAL reconstruction algorithms.

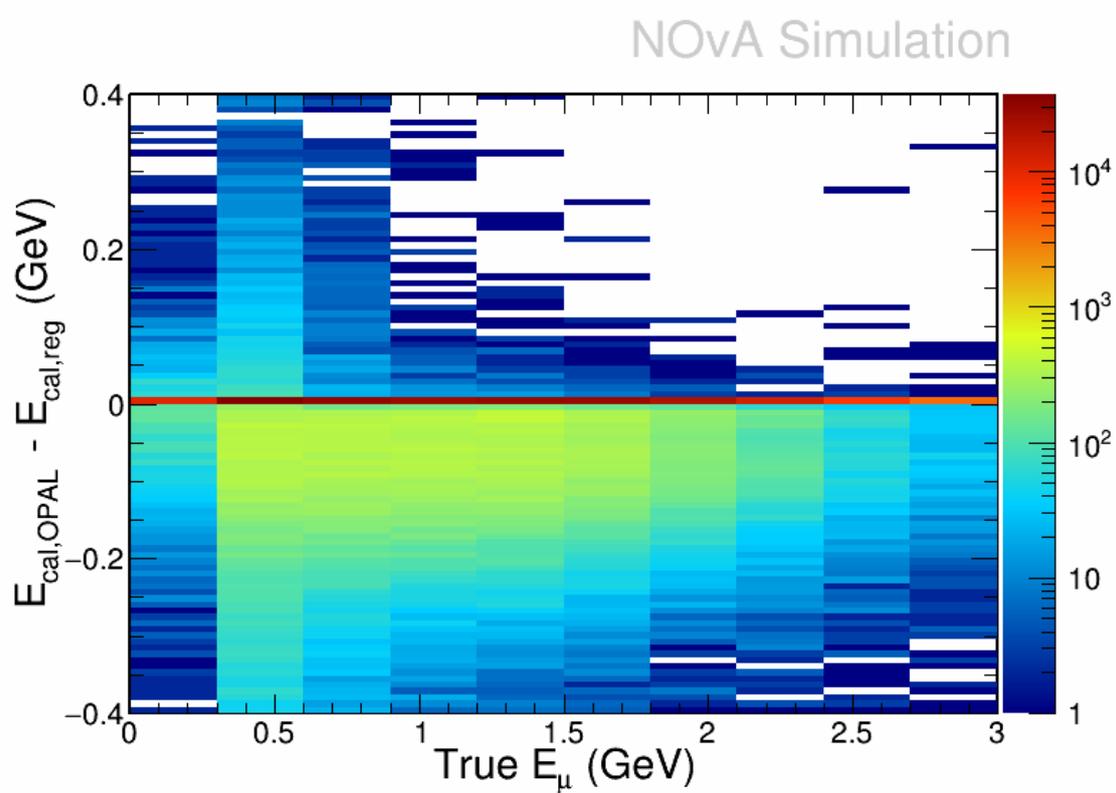


Figure 4.12: The difference between calorimetric energy calculated under the preexisting and OPAL reconstructions plotted vs true energy for a sample of muon prongs.

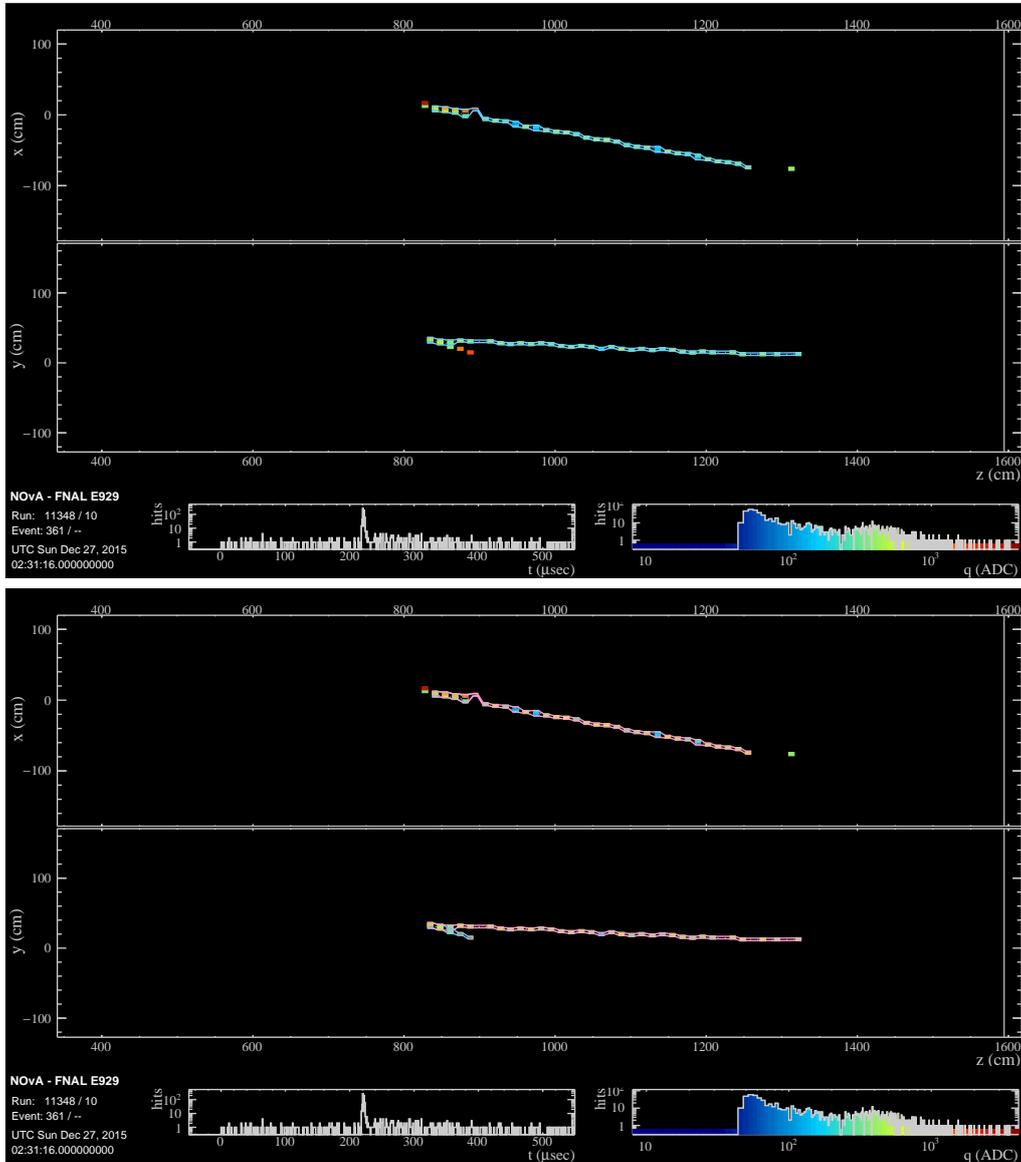


Figure 4.13: Event display showing 3D prongs made using regular reconstruction (top) and OPAL reconstruction (bottom). In the y view of the OPAL-reconstructed event an extra 3D prong can be seen, marked in blue. This prong is not visible in the x view as it is drawn behind the prong marked in pink.

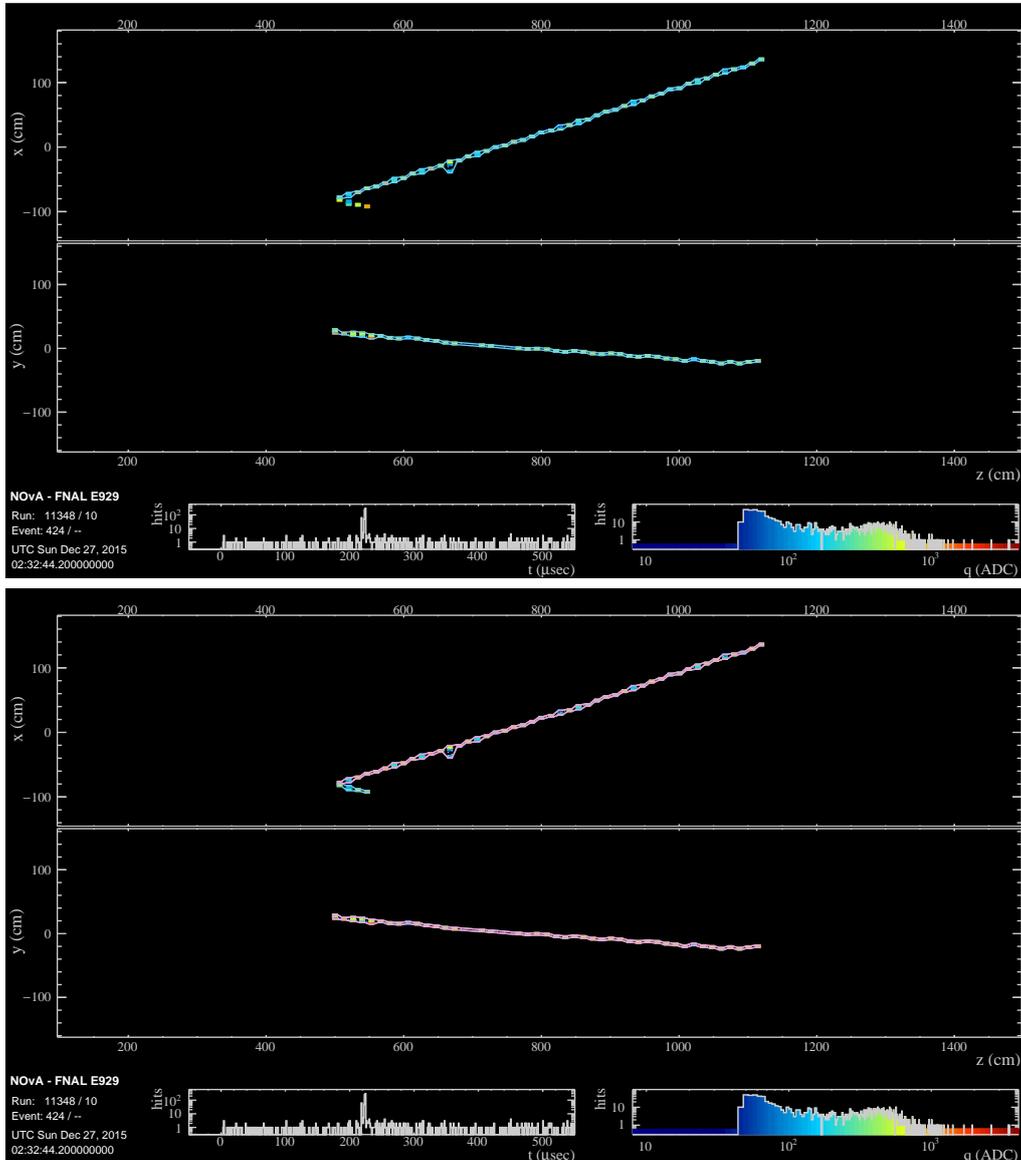


Figure 4.14: Event display showing 3D prongs made using regular reconstruction (top) and OPAL reconstruction (bottom). In the x view of the OPAL-reconstructed event an extra 3D prong can be seen, marked in blue. This prong is not visible in the y view as it is drawn behind the prong marked in pink.

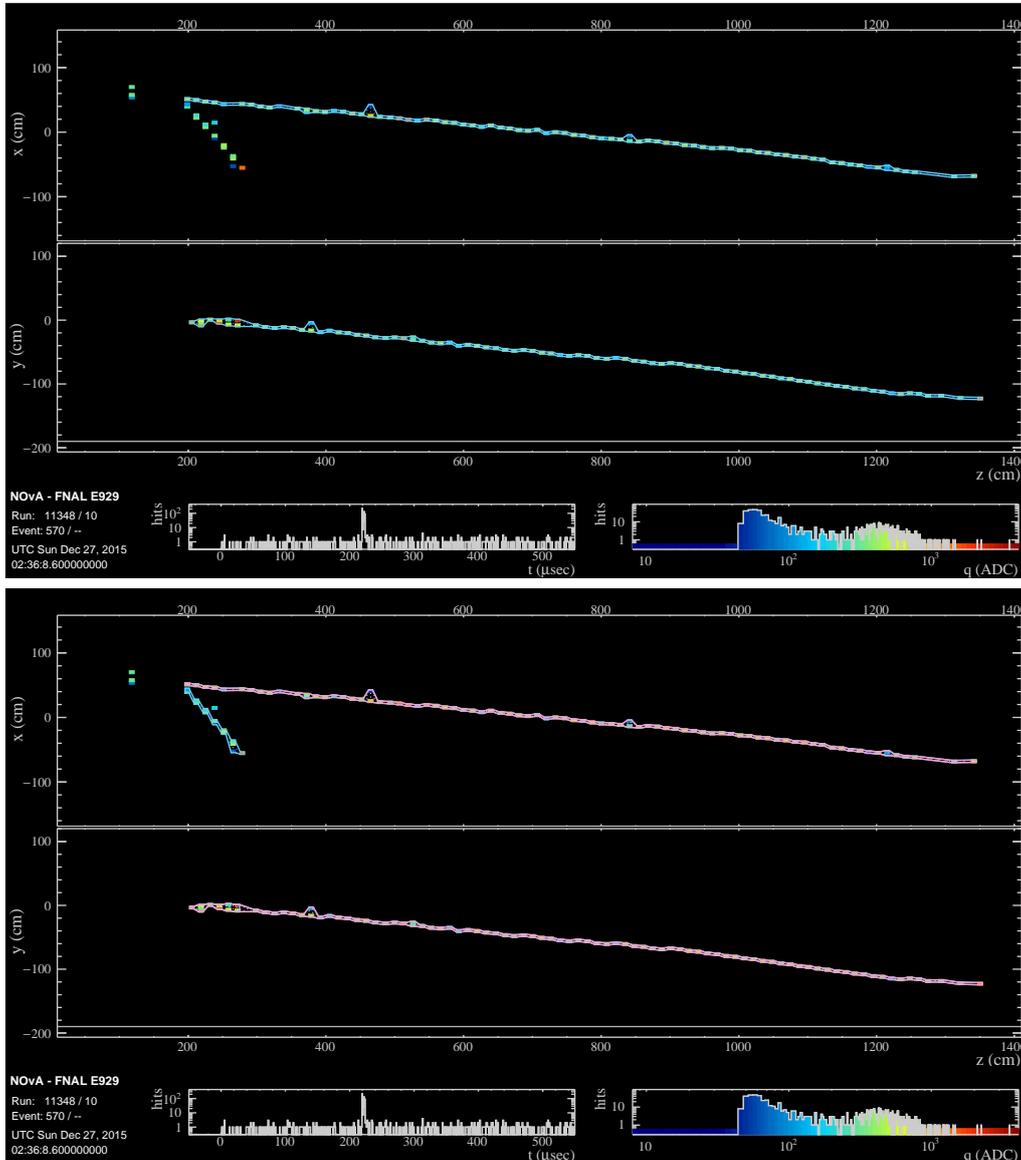


Figure 4.15: Event display showing 3D prongs made using regular reconstruction (top) and OPAL reconstruction (bottom). In the x view of the OPAL-reconstructed event an extra 3D prong can be seen, marked in blue. This prong is not visible in the y view as it is drawn behind the prong marked in pink.

Chapter 5

A systematic uncertainty on hadron scattering

Concurrently with developing the cross-section analysis outlined in Chapter 6, a new systematic uncertainty was implemented in the NOvA software framework. This systematic uncertainty is incorporated into the cross-section analysis, which is the central focus of this thesis.

Geant4Reweight [75] [76] is an externally developed software package which calculates event weights in order to replicate the effect of resimulating with altered scattering cross-sections for hadrons traversing the detector. This allows hadron scattering cross-section normalisations to be varied within an uncertainty determined from a fit to external data. The effects of varying the Geant4 physics can then be propagated to physics analyses done in NOvA.

Section 5.1 will briefly explain how Geant4Reweight works. The rest of this chapter will focus on NOvA-specific implementation and evaluation of results.

5.1 Geant4Reweight

In order to reproduce the effect of changing the scattering cross-section of a hadron, Geant4Reweight quantifies how much more or less probable a given particle's trajectory is under the modified cross-section compared to the nominal. For example, consider modifying the inclusive scattering cross-section of charged pions. A given pion travels a series of steps of size L_i with a scattering cross-section σ_i and finally scatters inelastically

(interacts). The probability of traversing a given step without interacting is

$$P_{i,\text{survive}} = e^{-\sigma_i L_i} \quad (5.1)$$

The probability of the entire trajectory is therefore given by

$$P_{\text{traj}} = \sigma_N \prod_{i=1}^N P_{i,\text{survive}} = \sigma_N e^{-\sum_{i=1}^N \sigma_i L_i} \quad (5.2)$$

where the factor σ_N is because the pion interacts at the very end of the final step¹. If the cross-section is modified, $\sigma \rightarrow \sigma'$, the weight this pion receives is given by

$$w = \frac{P'_{\text{traj}}}{P_{\text{traj}}} = \frac{\sigma'_N}{\sigma_N} \left(\frac{e^{-\sum_{i=1}^{N-1} \sigma'_i L_i}}{e^{-\sum_{i=1}^{N-1} \sigma_i L_i}} \right) \quad (5.3)$$

Now consider instead that the scattering cross-section for some exclusive channel is modified, e.g. pion charge exchange, $\pi^+ + N \rightarrow \pi^0 + N'$. The modification $\sigma_{\text{cex}} \rightarrow \alpha \sigma_{\text{cex}}$ is made. Pions receive a weight according to Equation 5.3, where the inclusive cross-section σ is now modified due to the modification of the exclusive component σ_{cex} . In addition, the weight for pions which truly undergo charge exchange receives an additional factor α , where α is the scaling factor applied to the charge exchange cross-section, as mentioned above.

The interested reader may refer to the Geant4Reweight paper [75] for a more in-depth explanation of Geant4Reweight. Note, however, that some equations written incorrectly in the paper have since been corrected [77].

5.2 Implementation in NOvA

This section gives a general explanation of how the systematic uncertainty is implemented in NOvA.

¹This can be understood as the probability of *not* surviving (i.e. interacting) during step N , $P_{N,\text{interact}} = 1 - P_{N,\text{survive}}$, in the limit that L_N tends to zero.

Name	Shorthand	Process
Pion absorption	abs	$\pi^\pm + N \rightarrow N' + X$
Quasi-elastic	inel	$\pi^\pm + N \rightarrow \pi^\pm + N' + X$
Charge exchange	cex	$\pi^\pm + N \rightarrow \pi^0 + N' + X$
Double charge exchange	dcex	$\pi^\pm + N \rightarrow \pi^\mp + N' + X$
Pion production	prod	$\pi^\pm + N \rightarrow n\pi + N' + X$

Table 5.1: Table of exclusive inelastic scattering channels for pions in Geant4Reweight. Here X represents any number of additional non-pion particles present in the final state

5.2.1 Generating weights

Geant4Reweight is used in NOvA to modify the scattering cross-sections for π^+ , π^- , and protons. Geant4Reweight also supports reweighting neutrons, however the weights generated for NOvA did not conserve the number of neutrino interactions². This problem could not be diagnosed in time, so neutron reweighting is disabled within NOvA. Nonetheless, a fit was done for the neutron cross-section, as described in Section 5.2.2.

For π^+ and π^- , weights are generated by varying the normalisation of each of the 5 exclusive inelastic scattering cross-sections independently³. The normalisation is controlled by a parameter which can be increased or decreased. A parameter setting of 1.0 represents an unchanged cross-section, and a parameter setting of 1.5 represents increasing the normalisation of the cross-section by 50%. A description of each of the exclusive parameters can be found in Table 5.1. There is only a single parameter for protons and neutrons, which controls the inclusive inelastic scattering cross-section.

For each event, a weight is calculated by varying one parameter while holding all the others at 1.0. For a given parameter, 6 weights are calculated, corresponding to parameter settings of 0.80, 0.90, 0.95, 1.05, 1.10, and 1.20. With 5 parameter for π^+ , this means $5 \times 6 = 30$ weights are calculated per event. There are a further 30 weights per event for π^- . Since protons only have one parameter, there are a further 6 weights per event. The weight for a user-requested parameter setting can be calculated at runtime by performing linear interpolation between the weights stored in simulation files.

²Geant4 is downstream of neutrino event generation (GENIE). Changes to Geant4 physics should not be able to change the number of simulated neutrino interactions.

³Note there is a sixth parameter which controls the normalisation of the inclusive scattering cross-section. This parameter is redundant since its effect can be reproduced by turning all 5 parameters controlling the exclusive cross-sections by the same amount.

Species	$\chi^2/\text{d.o.f.}$
π^+	118.9 / 48
π^-	55.8 / 25
n	9.9 / 5

Table 5.2: χ^2 per degree of freedom determined from fits to scattering data.

For an event containing multiple e.g. π^+ , the weight stored in the simulation files is the product of all the weights for each π^+ in the event, including those produced from secondary interactions, but excluding those from intranuclear scattering.

5.2.2 Fitting

A fit was performed to external scattering data to provide best fit values and uncertainties for each of the parameters. Since, at the time of writing, there is no scattering data for hadrons in the NOvA detector, scattering data in carbon⁴ was used instead. Fitting was done for each particle species independently. For charged pions, a fit was performed for the 5 exclusive parameter values. Note that inclusive scattering data additionally constrained this fit. Any parameters without data available were allowed to float during the fit with a 20% penalty term. For neutrons a fit was done for the inclusive parameter. For protons, no fit data was found at the time of writing.

Once the fit was completed, the fit value covariance matrix (and hence also the fit value uncertainties) was scaled by a factor χ^2/NDOF , where NDOF is the number of degrees of freedom in the fit. This was done to conservatively inflate the parameter uncertainties when confronted with a tension between data and prediction. See Table 5.2 for the χ^2 per degree of freedom.

Shown in Figures 5.1-5.3 are the Geant4 cross-sections and fit data for scattering in carbon. Note that interpreting fit quality by eye is difficult as the exclusive channels are all additionally constrained by the inclusive data, and hence are correlated. The sources of fit data can be found in the references [79–95]. These fit results, shown in Figures 5.4-5.6, are used to modify cross-sections in a homogeneous NOvA ND by varying parameter values randomly at runtime in a way that respects the correlations between parameters (see Section 5.3.1). The homogeneous NOvA ND is a fictional material with the same

⁴The NOvA near detector is 66% carbon by mass [78]

elemental composition as the NOvA ND, but with uniform composition. This allows us to aggregate the cross-sections of the many elements in the NOvA detector into a single cross-section which is modified during the reweighting. Since no data was found for proton scattering, the fit value is set to 1.00 ± 0.05 ⁵.

5.3 Usage

5.3.1 Multiverse systematic

As mentioned in Section 5.2.2, parameter values can be randomly varied in a correlated way. This is done using the covariance matrices determined during the fit. A “universe” is a collection of parameters which have all been shifted from the nominal parameter setting in this way. To create uncertainty bands on histograms, many universes are created, each with their own unique parameter shifts. An ensemble of universes is known as a multiverse. Using the event weights, a histogram can be made for each universe in the multiverse. Taking the RMS of the histograms allows an uncertainty band to be calculated. At runtime, the fit covariance matrix can be used to create the requested number of universes by doing correlated shifts of the parameters. See Figure 5.7 for an example. We see that proton and pion prong length are sensitive to this systematic, but calorimetric energy is not.

5.3.2 Varying parameters manually

Though the multiverse approach in Section 5.3.1 is the preferred method of using the weights, it is possible to vary the parameters manually. Figure 5.8 shows a sample of plots made in this way. We see that pion prong length is sensitive to the pion cross-section normalisation parameter, and proton prong length is sensitive to the proton cross-section normalisation parameter.

5.4 Summary

This chapter describes the implementation of a new systematic uncertainty in the NOvA offline software framework. Using the external package Geant4Reweight, weights are generated and stored in Monte Carlo simulation files. These weights reproduce the effect of re-running Geant4 with modified scattering cross-sections for charged pions and

⁵The uncertainty is an estimate, informed by the neutron fit value 0.975 ± 0.028 .

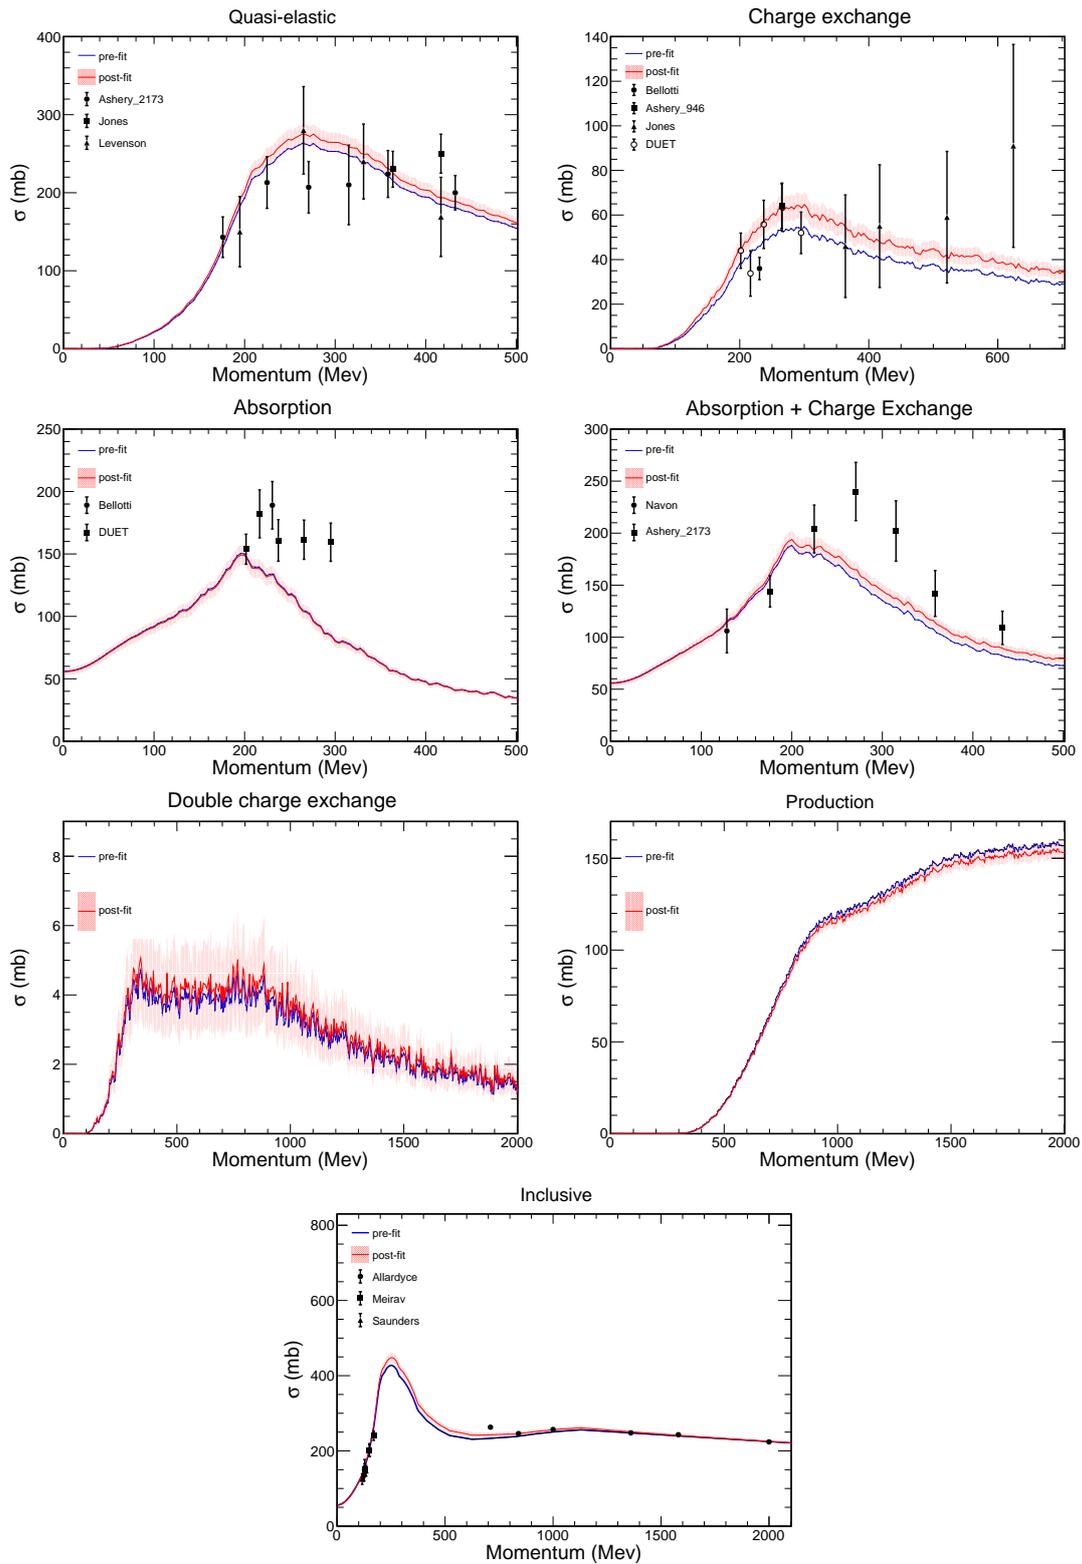


Figure 5.1: Inclusive and exclusive cross-sections for π^+ inelastic scattering in carbon. Pre- and post-fit predictions are shown, with fit data overlaid. No data was found for double charge exchange or pion production at the time of writing.

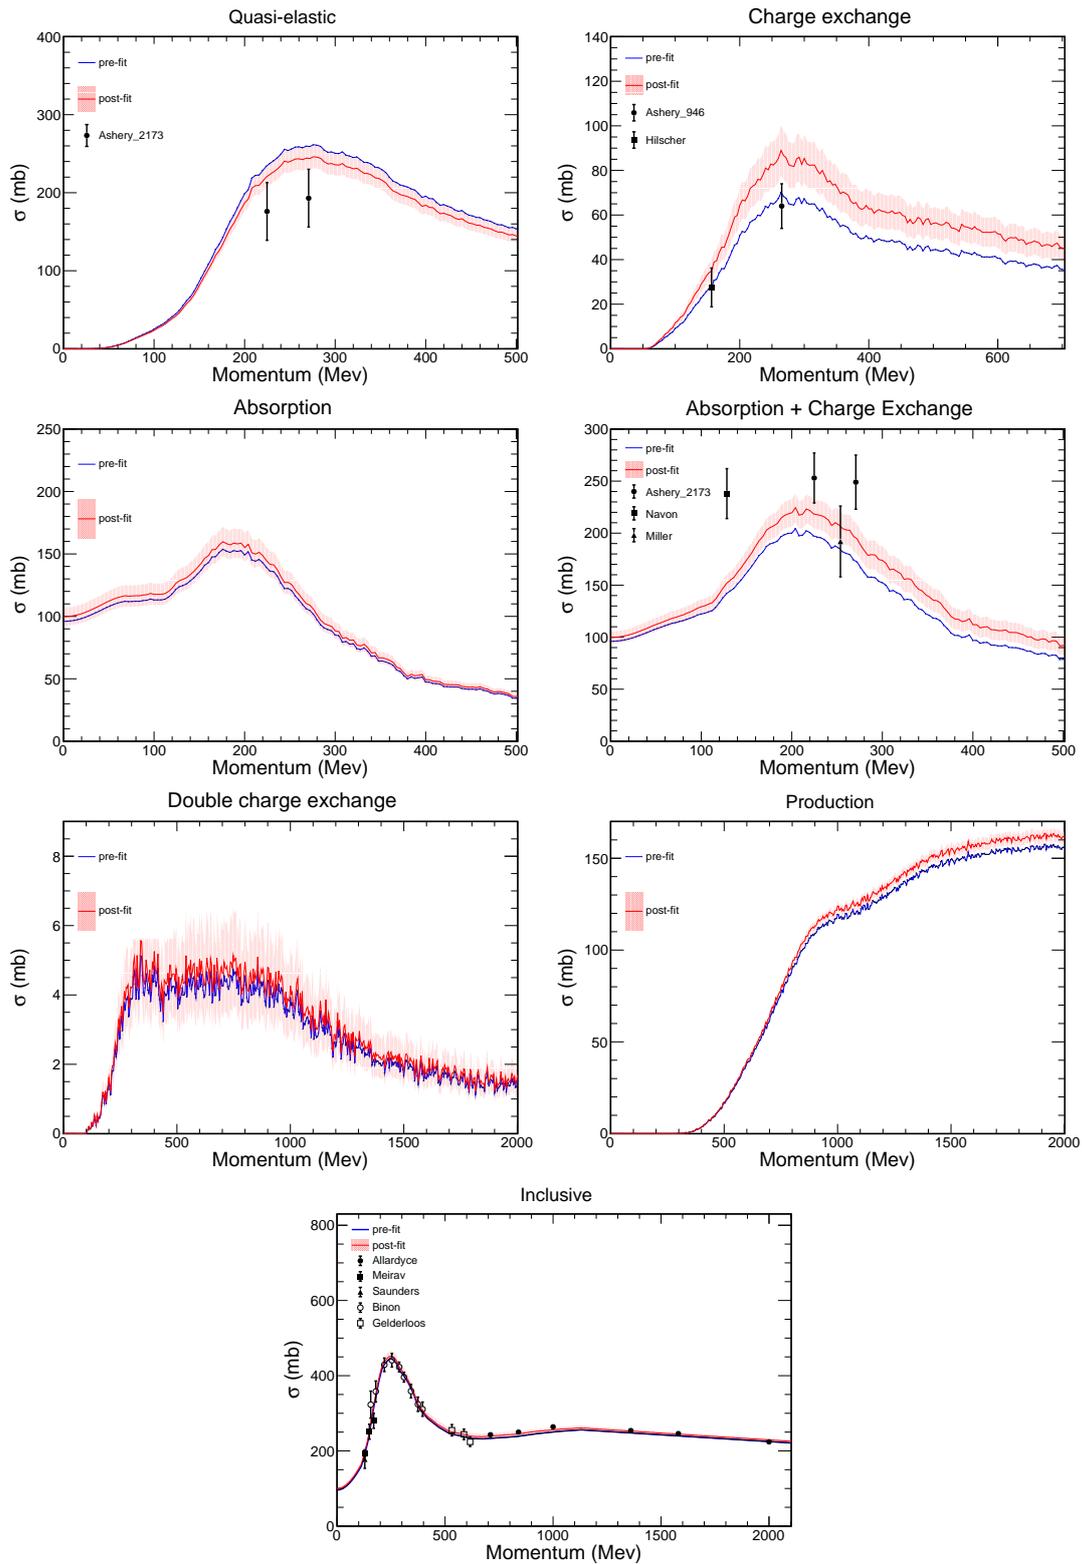


Figure 5.2: Inclusive and exclusive cross-sections for π^- inelastic scattering in carbon. Pre- and post-fit predictions are shown with fit data overlaid. At the time of writing, no data was found for double charge exchange, pion production, or pion absorption.

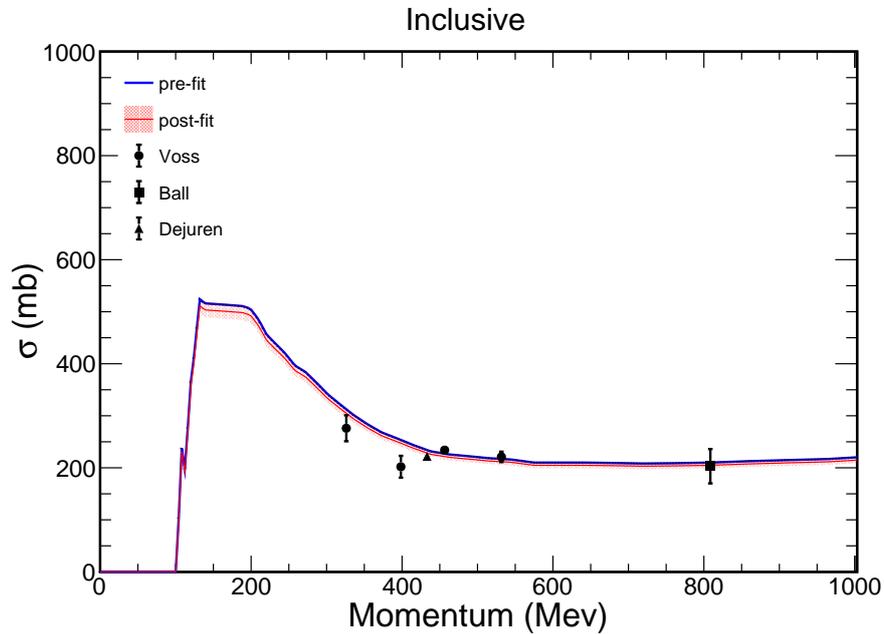


Figure 5.3: Inclusive cross-section for neutron inelastic scattering in carbon. Shown in blue is the pre-fit cross-section, and the post-fit cross-section with the error band is shown in red. Fit data is overlaid.

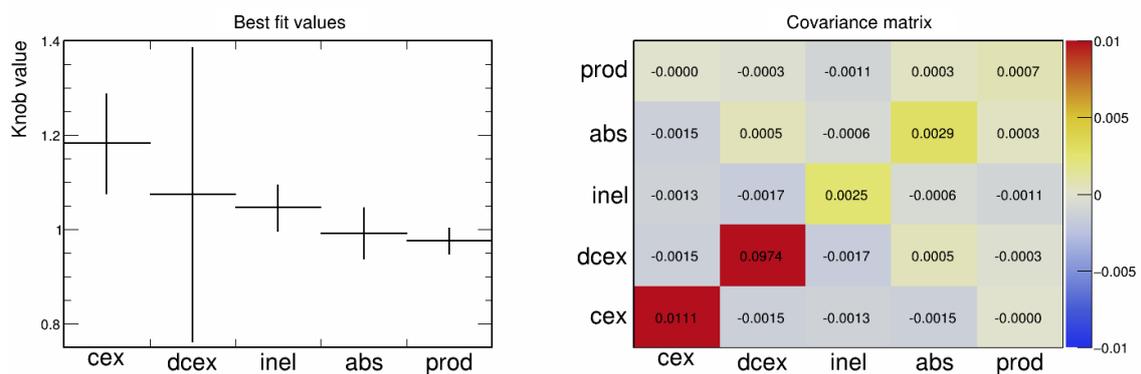


Figure 5.4: Results of fit for π^+ parameters. The best fit values and uncertainties are shown on the left, and the fit covariance matrix is shown on the right.

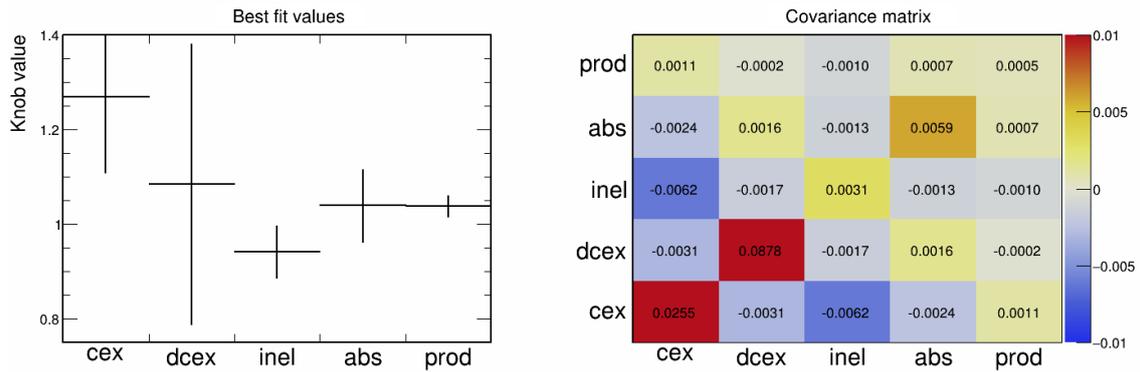


Figure 5.5: Results of fit for π^- parameters. The best fit values and uncertainties are shown on the left, and the fit covariance matrix is shown on the right.

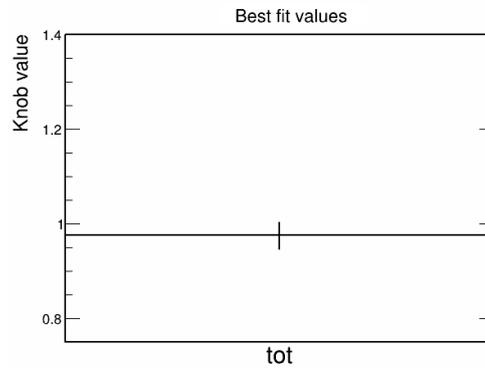


Figure 5.6: Result of fit for the neutron parameter. The best fit parameter value and uncertainty is shown.

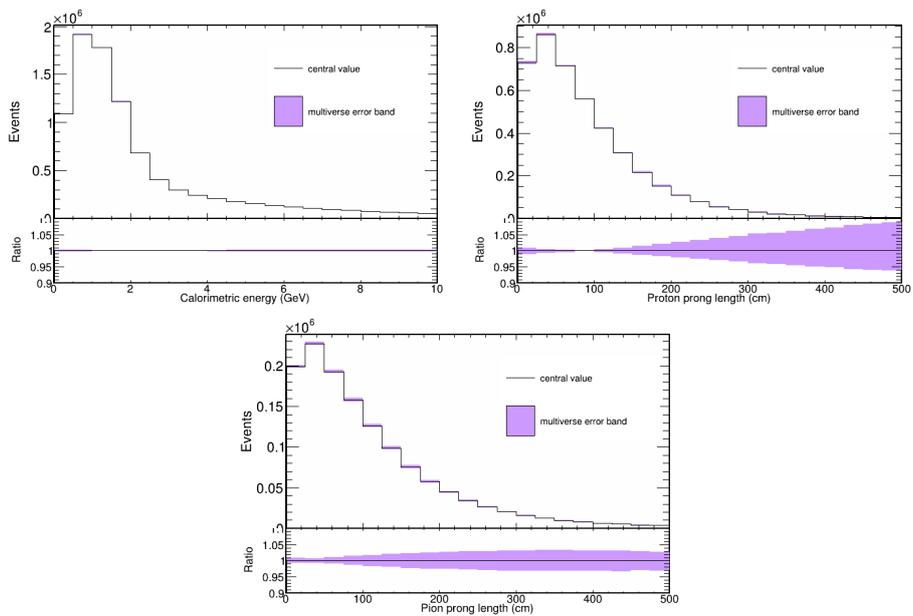


Figure 5.7: Simulated events with a multiverse error band drawn. The only cut applied here is that the true neutrino interaction vertex is inside the active detector limits

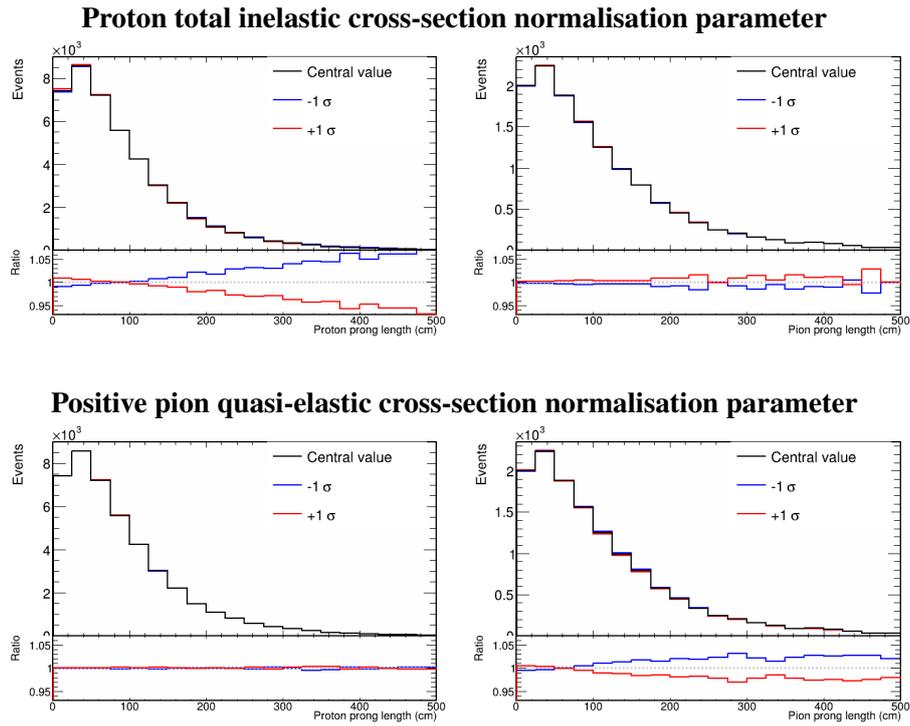


Figure 5.8: Simulated prong length histograms with shifts applied to proton and piplus parameters separately.

protons. A fit to external scattering data yields uncertainties and covariances for the cross-section parameters. This allows us to determine how much to vary parameters to generate uncertainty bands. By creating many universes, the correlation between parameters can be captured in the resulting uncertainty. This uncertainty is largest when looking at reconstructed pion and proton lengths, and smaller for other variables. This uncertainty is used throughout the cross-section analysis described in Chapter 6.

Chapter 6

Measurement of the ν_μ CC $1\pi^\pm$ cross-section

This chapter describes an analysis to measure the cross-section of $\nu_\mu + N \rightarrow \mu^- + \pi^\pm + X$, where X represents any other particles in the final state other than π^\pm . The cross-section is measured in bins of charged pion angle, using 13.8×10^{20} POT of FHC data.

6.1 Background and motivation

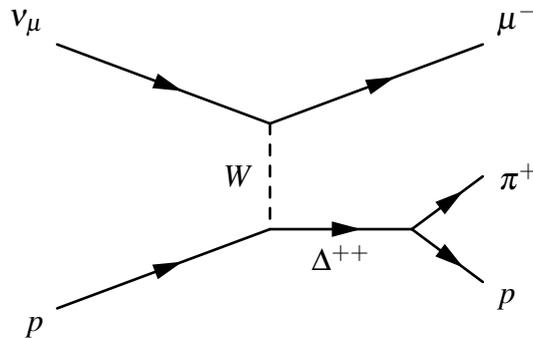


Figure 6.1: Diagram of resonant single pion production in a CC interaction on a proton.

As outlined in Section 1.5, knowledge of neutrino interaction cross-sections is a key ingredient in measuring neutrino oscillation parameters. In NOvA, neutrino energy estimation depends on the interaction topology. The energy estimation, in turn, impacts the measured values of neutrino oscillation parameters. Measuring cross-sections for different final state topologies reduces the uncertainty in oscillation measurements from cross-section modelling systematic uncertainties. This analysis measures a single-differential

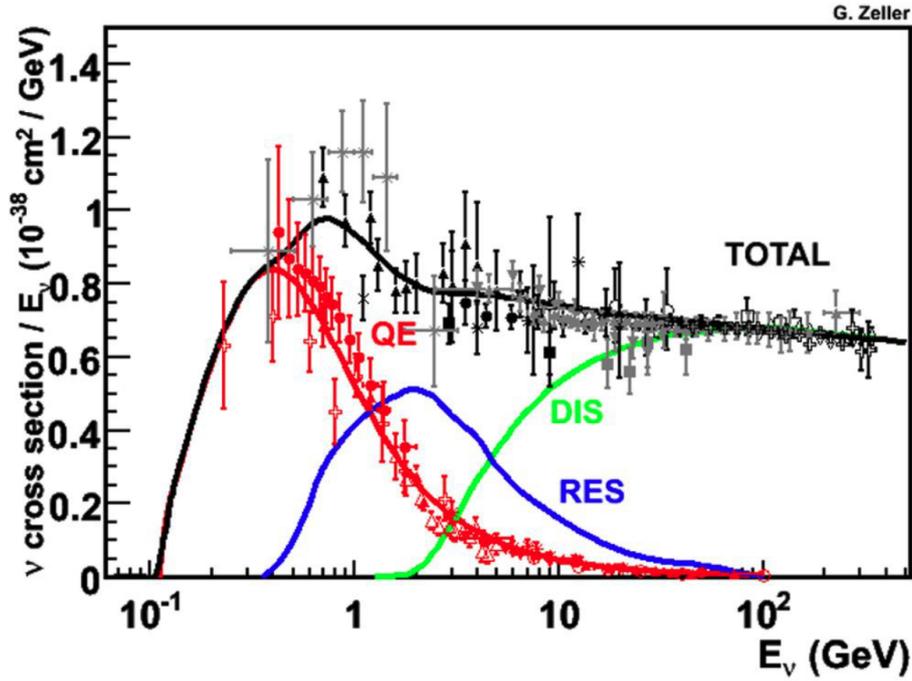


Figure 6.2: Charged-current neutrino interaction cross section as a function of energy. Figure from [7].

cross-section for single charged pion production in ν_μ CC interactions in the NOvA Near Detector.

This channel is sensitive to resonant pion production events (see to Section 1.4.1.2 for further details on resonant interactions). Figure 6.1 shows a Feynman diagram of a resonant interaction. Resonant interactions are the leading process at the peak energy of NOvA (~ 2 GeV), as seen in Figure 6.2. In addition, this measurement is a probe of final state interactions (FSI) in the nucleus, whereby events which initially do not contain a single charged pion become a $1\pi^\pm$ event through pion¹ absorption/production/charge-exchange in the nuclear medium.

6.1.1 Similar measurements

Other experiments have made measurements of the ν_μ CC $1\pi^\pm$ cross-section. These measurements were made in neutrinos beams of varying energies and on a range of nuclear targets. Table 6.1 shows a selection of ν_μ CC $1\pi^\pm$ measurements. Figure 6.3 shows a measurement of the ν_μ CC $1\pi^\pm$ cross-section made by MINERvA, which is reported in

¹Throughout this chapter “pion” will refer to π^+ or π^- unless stated otherwise.

Experiment	Year	Target	Citation
MiniBooNE	2011	CH ₂	[96]
MINERvA	2015	CH	[23]
T2K	2017	H ₂ O	[97]
T2K	2020	CH	[98]
MINERvA	2022	CH, C, H ₂ O, Fe, Pb	[99]

Table 6.1: Existing ν_μ CC $1\pi^\pm$ cross-section measurements.

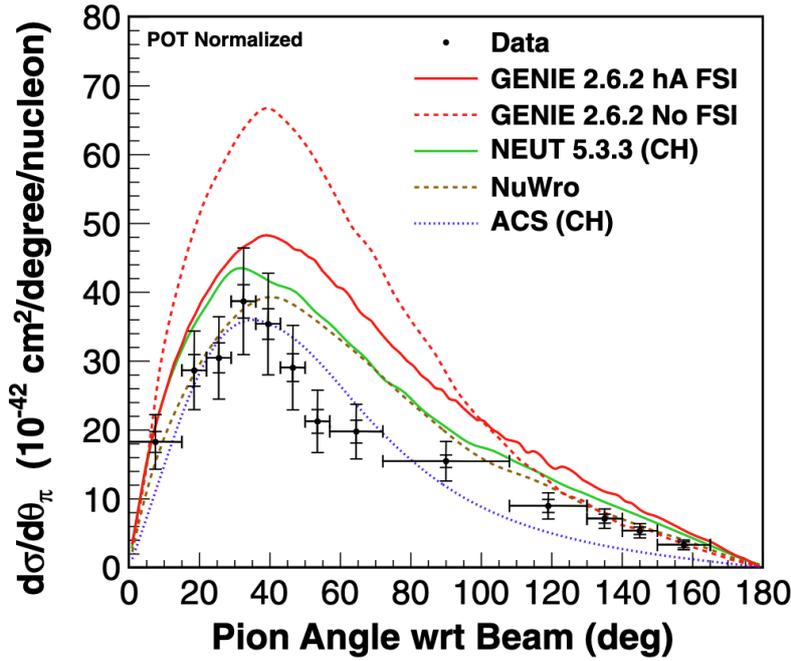


Figure 6.3: MINERvA measurement of CC $1\pi^\pm$ $d\sigma/d\theta_\pi$ compared to the GENIE, ACS, NEUT, and NuWro models. The inner (outer) error bars correspond to the statistical (total) uncertainties. Figure adapted from [23]. A portion of the original figure, which is not relevant to the current discussion, has been cropped out.

bins of pion angle, similar to the measurement made in this thesis. However, it should be noted that the average neutrino energy for the MINERvA measurement differs from the average neutrino energy at NOvA.

6.1.2 Analysis deliverable

This analysis reports a single-differential cross-section measurement in bins of pion angle with respect to the average neutrino beam direction. Pion kinematics are a probe of nuclear physics due to the final state interactions pions undergo as they traverse the nucleus.

Previous attempts to measure pion energy have shown significant bias, making it unsuitable as an analysis variable. Figure 6.4 shows a simple example of this, where

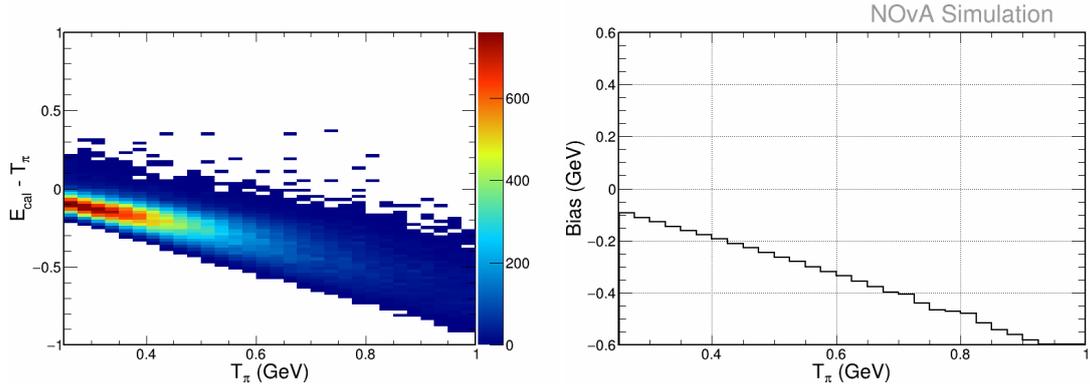


Figure 6.4: (Left) Plot of (reconstructed - true) pion energy vs true pion energy. Here calorimetric energy in the pion prong is used to estimate pion energy. (Right) Bias of pion energy estimation.

calorimetric energy is used as an energy estimator. For further details on how bias is calculated, see Section 6.4. Figure 6.4 shows a large bias, making unfolding in pion energy unreliable. Other more advanced pion energy estimators have been developed using machine learning methods. These estimators use several inputs to attempt to estimate pion energy yet still yield a large bias. Understanding the underlying reason for poor pion energy estimation is an area of ongoing active development in the NOvA collaboration. Potential issues include reconstruction failures causing pion prongs to miss some hits, as well as pion prongs including hits not created by the pion.

6.1.3 How to measure a cross-section

To measure the differential cross-section, the following equation is used

$$\left(\frac{d\sigma}{d\theta_\pi}\right)_\alpha = \frac{\sum_j U_{j\alpha} N_j^{\text{sel}} P_j}{\varepsilon_\alpha \Phi T (\Delta\theta_\pi)_\alpha} \quad (6.1)$$

The elements of Equation 6.1 are

- α Bin index in true θ_π
- j Bin index in reconstructed θ_π
- U Unfolding matrix, which maps reconstructed θ_π to true θ_π . It is determined using methods described in Section 6.7
- N^{sel} Number of events from data passing the selection outlined in Section 6.3

P Purity correction, determined from a template fit to data as described in Section 6.6. This term corrects for any background contamination

ε Efficiency correction, determined using Monte Carlo simulation

Φ The integrated muon neutrino flux, see Section 6.9

T Number of targets (nucleons) inside the fiducial volume, see Section 6.8

$(\Delta\theta_\pi)_\alpha$ Width of α -th θ_π bin

Efficiency and purity are defined as follows

$$\varepsilon = \frac{\text{Number of selected signal events}}{\text{Number of signal events}} \quad (6.2)$$

$$P = \frac{\text{Number of selected signal events}}{\text{Number of selected events}} \quad (6.3)$$

6.2 Analysis signal definition

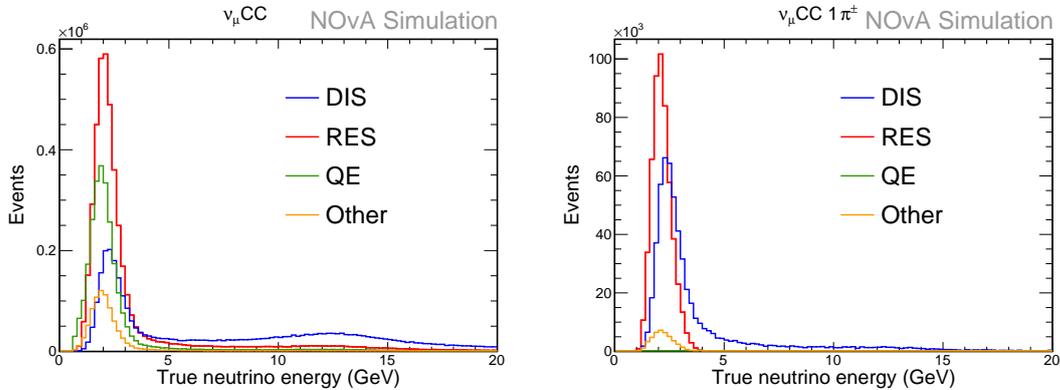


Figure 6.5: The number of signal events in simulation, broken down by interaction mode. (Left) A ν_μ CC inclusive signal definition, containing 1.1×10^7 events, of which 39 % are resonant interactions. (Right) A ν_μ CC $1\pi^\pm$ signal definition, containing 1.1×10^6 events, of which 46 % are resonant interactions.

This analysis targets events with a single charged pion in the final state. Some further kinematic constraints are applied to limit the analysis to regions of phase space that NOvA is sensitive to. Events which meet all the criteria are known as signal events. The signal for this analysis is defined as a neutrino interaction with

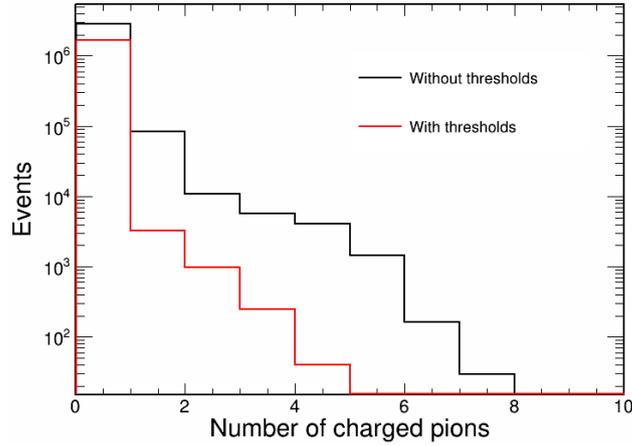


Figure 6.6: The number of pions in the final state in a sample of $\nu_\mu\text{CC}$ quasi-elastic interactions. Shown in black is the sample of all $\nu_\mu\text{CC}$ quasi-elastic interactions in the fiducial volume outlined in Section 6.3.3. For the histogram in red we further require that the muon and pions be within the kinematic thresholds outlined in Sections 6.2.2 and 6.2.1.

- Exactly one pion within the kinematic thresholds (see Section 6.2.1)
- Exactly one muon within the kinematic phase space (see Section 6.2.2)
- The true interaction vertex inside the fiducial volume (see Section 6.3.3)

Figure 6.5 shows that this signal definition is more sensitive to resonant pion production than an inclusive measurement. Note that the signal definition contains very few quasi-elastic interactions. From Figure 6.6 it can be calculated that 2.8% of all $\nu_\mu\text{CC}$ quasi-elastic interactions have a single pion in the final state. Furthermore, of the $\nu_\mu\text{CC}$ quasi-elastic interactions with a muon within the signal phase space (see Sections 6.2.2), only 0.2% have a single pion within the kinematic thresholds (see Section 6.2.1). Hence very few quasi-elastic events are found in the signal sample. Figure 6.7 shows an event display of a signal event.

6.2.1 Pion kinematic thresholds

It is difficult to reconstruct low-energy pions in the NOvA Near Detector as they may create hits in only one cell. In addition, prong reconstruction is difficult at angles close to 90° as the particle may only traverse one plane. A threshold on pion kinematics is added to the signal definition to ensure a result is only reported in regions of phase space that

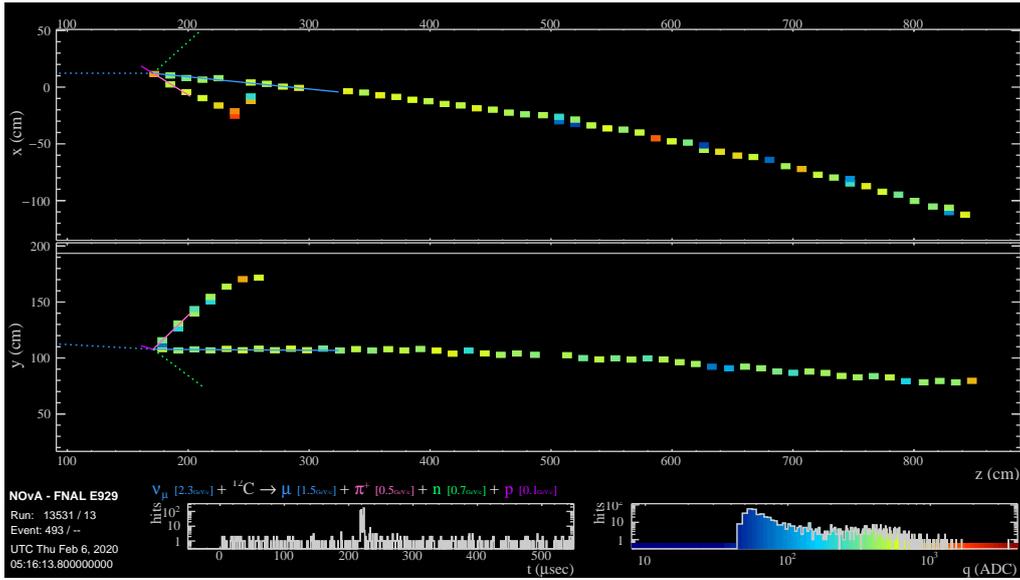


Figure 6.7: Event display of a simulated signal event.

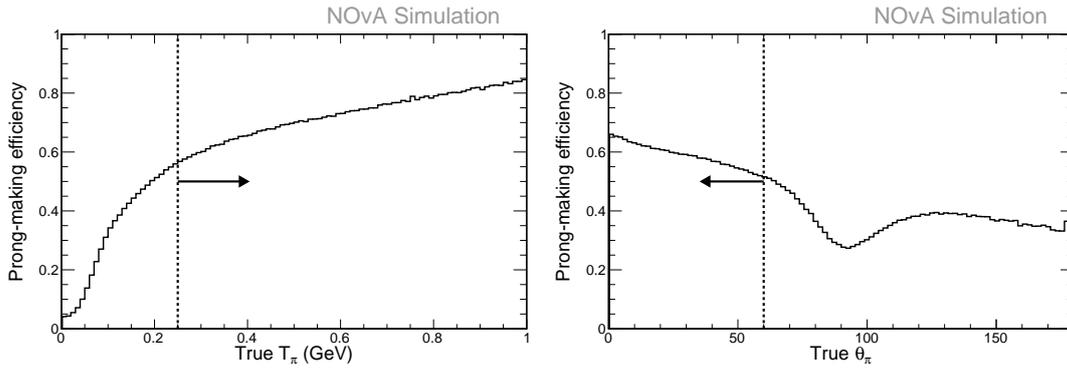


Figure 6.8: The fraction of pions which form a prong as a function of pion kinetic energy (left) and angle (right). The signal thresholds are shown with a dashed line. To make these plots, a simple ν_μ CC $1\pi^\pm$ signal definition is used, with no kinematic thresholds. Fiducial, containment, and quality selection cuts are applied.

the NOvA detector is sensitive to. The threshold on minimum pion energy and angle was determined by looking at the fraction of pions that form a prong as a function of energy and angle, as shown in Figure 6.8. The cut value is chosen based on two criteria. Firstly, the above threshold sample should have a high reconstruction efficiency. Secondly, the cut value should be some finite distance away from the region in which the efficiency changes rapidly. This is to avoid large changes in efficiency as bins migrate during unfolding². Cut

²Note that currently the measurement is not reported in bins of pion energy, and so no unfolding is done in pion energy. However it is hoped that an updated analysis may report a differential measurement in pion energy.

values which match these criteria were chosen by eye. As a result, this analysis requires that signal events contain exactly one pion with true kinetic energy greater than 250 MeV, and this pion should have a true angle w.r.t the beam direction less than 60° . The signal definition allows for multiple additional pions with energy less than 250 MeV, as there is no reliable way to reject these events.

6.2.2 Muon kinematic phase space

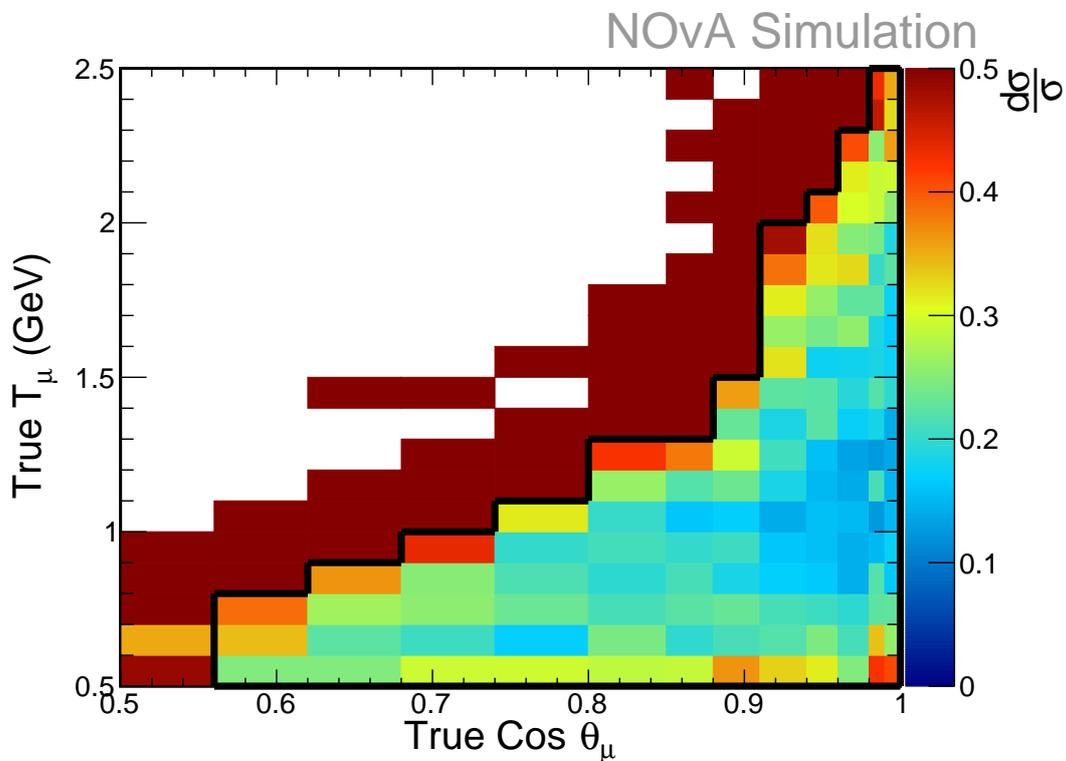


Figure 6.9: Fractional uncertainty on the cross-section in bins of muon kinematics. The muon phase space cut is indicated with a solid black border. The full signal and selection cuts are used to make this plot, excluding the muon phase space cut.

Muons with high energy or large angle are frequently uncontained, leading to large efficiency corrections. These efficiency corrections have large systematic uncertainties. To mitigate this, a muon phase space cut is added to the signal definition. A matching selection cut is also made using reconstructed energy and angle. The cut boundary was decided by

looking at the fractional uncertainty in the cross-section measurement, approximated as

$$\frac{d\sigma}{\sigma} \approx \sqrt{\frac{1}{N_{sel}} + \sum_{syst} \left(\frac{d\epsilon_{syst}}{\epsilon}\right)^2 + \left(\frac{dP_{syst}}{P}\right)^2} \quad (6.4)$$

which can be seen in Figure 6.9. The uncertainty on the efficiency and purity is comprised of a statistical and systematic component. The cut is defined as

$$\begin{aligned} 0.56 < \cos \theta_\mu < 0.62 & \quad \text{and} \quad 0.5 < T_\mu < 0.8 \\ 0.62 < \cos \theta_\mu < 0.68 & \quad \text{and} \quad 0.5 < T_\mu < 0.9 \\ 0.68 < \cos \theta_\mu < 0.74 & \quad \text{and} \quad 0.5 < T_\mu < 1.0 \\ 0.74 < \cos \theta_\mu < 0.80 & \quad \text{and} \quad 0.5 < T_\mu < 1.1 \\ 0.80 < \cos \theta_\mu < 0.88 & \quad \text{and} \quad 0.5 < T_\mu < 1.3 \\ 0.88 < \cos \theta_\mu < 0.91 & \quad \text{and} \quad 0.5 < T_\mu < 1.5 \\ 0.91 < \cos \theta_\mu < 0.94 & \quad \text{and} \quad 0.5 < T_\mu < 2.0 \\ 0.94 < \cos \theta_\mu < 0.96 & \quad \text{and} \quad 0.5 < T_\mu < 2.1 \\ 0.96 < \cos \theta_\mu < 0.98 & \quad \text{and} \quad 0.5 < T_\mu < 2.3 \\ 0.98 < \cos \theta_\mu < 1.00 & \quad \text{and} \quad 0.5 < T_\mu < 2.5 \end{aligned} \quad (6.5)$$

Note that the binning used here is the same as in the NOvA ν_μ CC inclusive analysis [100] so that the two analyses can be compared. See Section 6.4.2 for a study showing that this binning scheme is appropriate. The phase space determined here is a subset of ν_μ CC inclusive phase space.

6.3 Event selection

The event selection for this analysis is as follows

- Preselection
 - Quality cuts (more than 20 hits and 4 contiguous planes in the slice)
 - At least 2 large prongs in slice (see Section 6.3.2)

- A reconstructed vertex inside the fiducial volume (see Section 6.3.3)
- All prongs contained (see Section 6.3.4)
- Particle identification cuts (see Section 6.3.5)
- Reconstructed muon kinematic phase space cut (see Section 6.2.2)

Cut	Purity	Efficiency	Rel Eff
Preselection	0.089	0.16	0.16
Muon ID	0.25	0.059	0.36
Pion ID	0.44	0.038	0.65
Pion Rejection	0.47	0.026	0.67
Muon/Pion phase space	0.64	0.024	0.93

Table 6.2: Cutflow table showing purity, efficiency, and relative efficiency at each stage in the selection.

Cut	Signal	Outside μ/π^\pm phase space	$N_\pi = 0$	Non ν_μ CC	Non-fiducial	$N_\pi > 1$
Preselection	1.9e+05 (0.089)	7.6e+05 (0.36)	2.3e+05 (0.11)	8.2e+05 (0.39)	8.9e+04 (0.042)	5e+04 (0.024)
Muon ID	6.8e+04 (0.25)	1e+05 (0.38)	8.4e+04 (0.31)	9.4e+03 (0.035)	5.5e+03 (0.02)	2.5e+03 (0.0093)
Pion ID	4.4e+04 (0.44)	3.1e+04 (0.31)	1.6e+04 (0.16)	4.2e+03 (0.042)	2.7e+03 (0.027)	1.8e+03 (0.018)
Pion Rejection	3e+04 (0.47)	2e+04 (0.31)	9e+03 (0.14)	2.4e+03 (0.039)	1.9e+03 (0.031)	4.4e+02 (0.0069)
Muon/Pion phase space	2.8e+04 (0.64)	9e+03 (0.21)	4.8e+03 (0.11)	1.2e+03 (0.027)	5.2e+02 (0.012)	3.6e+02 (0.0082)

Table 6.3: The number of events from simulation at each step in the selection, broken down by signal and multiple background samples. Also shown is the fraction of the total number of events in each sample at each step in the selection.

The effect of the selection can be seen in Tables 6.2 and 6.3.

6.3.1 Background categories

The signal/background categories in Table 6.3 are defined as follows:

- **Signal:** An event matching all the signal definition criteria outlined in Section 6.2
- **Outside μ/π^\pm phase space:** Either the true muon kinematics are outside the phase space outlined in Section 6.2.2, or more likely there is no pion in the event with true energy greater than 250 MeV and true angle less than 60°
- **$N_\pi = 0$:** No π^\pm in the event
- **Non ν_μ CC:** Not a ν_μ CC interactions. Includes NC, $\bar{\nu}_\mu$ CC, and ν_e interactions
- **Non-fiducial:** The true neutrino interaction vertex is outside the fiducial volume
- **$N_\pi > 1$:** More than one π^\pm with $T_\pi > 250$ MeV

The largest background category is “Outside μ/π^\pm phase space” (events outside kinematic thresholds). This raises the question of whether the kinematic constraints should be relaxed. Doing so leads to a boost in purity but a drop in efficiency³. For example, suppose there was no pion energy threshold in the signal definition. In that case, the overall selection efficiency and purity are 2.1% and 70%, compared with the nominal efficiency and purity of 2.4% and 64%. The product of efficiency and purity is higher when a pion energy threshold is included, so this kinematic constraint is retained in the signal definition. It is also worth pointing out that identifying low-energy pions is difficult. In 40% of “Outside μ/π^\pm phase space” events, the prong identified as the pion was some other non-pion species (most often a proton). See Figure 6.19 for an example event display showing this.

6.3.2 Number of prongs

When performing particle identification to identify the muon and pion prongs, only prongs with at least 10 hits are considered. This is to ensure there are enough hits to reliably

³If thresholds in the signal definition were relaxed, there would be some selected events which were previously backgrounds as they were below threshold, but now are signal events. Hence purity is boosted. Conversely, many of the low energy pions would be unreconstructed, and so the signal selection efficiency would be reduced.

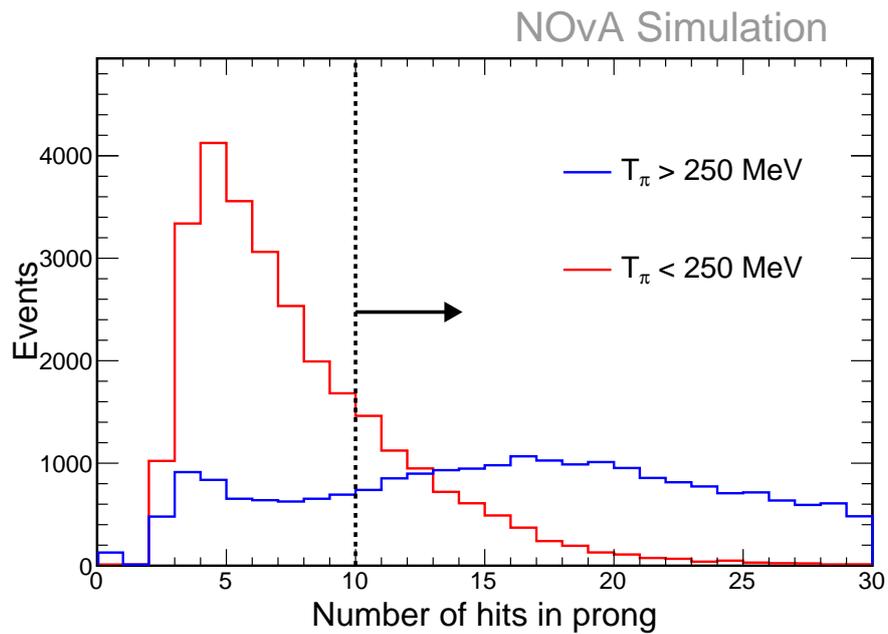


Figure 6.10: The number of hits in true pion prongs which are below (red) and above (blue) the pion energy threshold in the signal definition. To make these plots, a simple ν_μ CC $1\pi^\pm$ signal definition is used, with no kinematic thresholds. Fiducial, containment, and quality selection cuts are applied.

be able to perform particle identification, as well as reconstruction of kinematics. The threshold number of hits was determined by looking at the nHit distribution for pion prongs within the pion kinematic thresholds, as shown in Figure 6.10. No quantitative figure of merit was used to select a cut value. Rather the cut value was chosen by eye, as it excludes most of the below threshold pions while still retaining most of the above threshold pions. Since both a muon and a pion must be identified, we require at least two prongs in the event with more than 10 hits each.

6.3.3 Fiducial volume

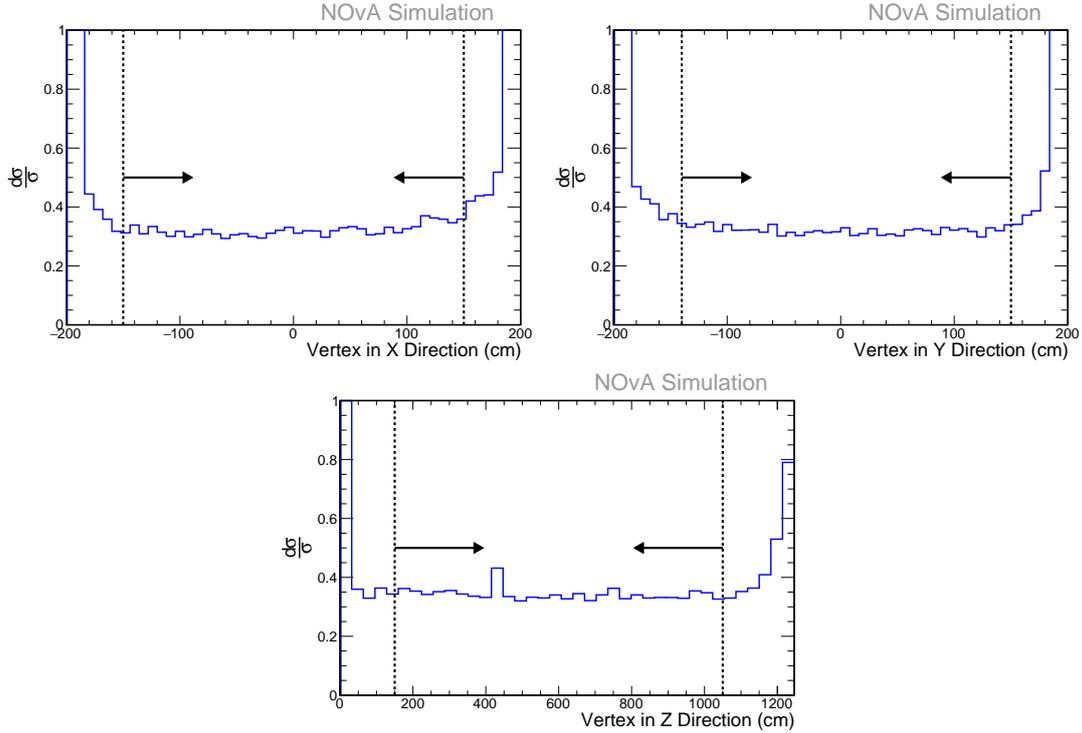


Figure 6.11: Fractional uncertainty on the integrated cross-section as a function of vertex x, y, and z components. To make this plot, a simple $\nu_\mu\text{CC}1\pi^\pm$ signal definition is used, with no kinematic thresholds and with true interaction vertex inside the detector. The selection cuts applied are an unoptimised muon identification cut, reconstructed vertex inside the detector, containment, and quality cuts.

A subset of the Near Detector volume is defined as the fiducial volume for this analysis. This is to guard against backgrounds from neutrino interactions occurring in the surrounding rock, which could be misidentified as neutrinos interacting inside the detector near an edge. The fiducial volume was determined by calculating the fractional uncertainty on the integrated cross-section as a function of vertex position in the x/y/z coordinates, as shown in Figure 6.11. Note that the origin of the coordinate system is approximately in the centre of the front face of the detector. The spike seen at $z \simeq 450$ cm is due to an understood bug in the simulation files used to make this plot. This bug is not present in the files used elsewhere in this thesis. The cut values were chosen to be the loosest cuts (i.e. closest to the detector extremities) in the region of minimum fractional uncertainty. Note that there are fluctuations in Figure 6.11 due to limited MC statistics. As a result, choosing the bin with the absolute minimum fractional uncertainty is not a good strategy. Rather,

there is some tolerance, and a round number is chosen as a cut value. The fiducial limits determined are

$$\begin{aligned} -150.0 < \text{Vertex } x \text{ (cm)} < 150.0 \\ -140.0 < \text{Vertex } y \text{ (cm)} < 150.0 \\ 150.0 < \text{Vertex } z \text{ (cm)} < 1050.0 \end{aligned} \quad (6.6)$$

Only events with a reconstructed interaction vertex inside this volume are selected. Furthermore, signal events must contain a true interaction vertex inside this volume.

6.3.4 Containment

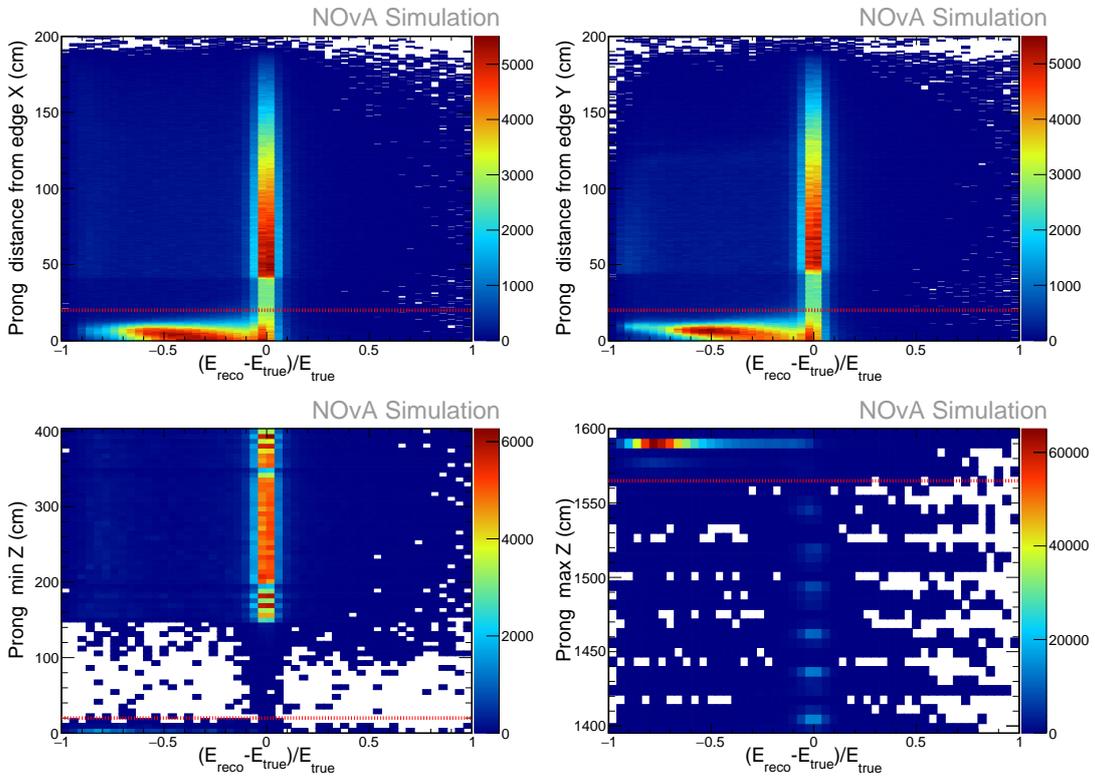


Figure 6.12: Muon energy resolution as a function of distance from the detector edge. The containment limits are marked with a red dashed line. These plots were made with fiducial, quality, and unoptimised muon ID cuts applied. When optimising, e.g. prong x containment limits, unoptimised containment cuts are applied in the y and z coordinates.

To ensure proper energy reconstruction is possible, prongs are required to be fully contained within the Near Detector. We require that prongs stop several centimetres before

reaching a detector wall to guard against exiting prongs where the hits nearest the detector edge are missed.

The containment limits were determined by studying how muon energy resolution degrades the closer the start/end point of the prong is to a detector edge. When the prong is near a detector edge, it was observed that there is a significant underestimation of muon energy. Hence we determine cuts by looking at the (reconstructed - true)/true energy plot as a function of distance from the detector edge in the x/y/z direction.

This study was done for muon prongs, shown in Figure 6.12. For the z-coordinate, the maximum and minimum cuts were determined independently, but for the x (y) coordinate, the maximum distance from the left and right (top and bottom) edges were determined simultaneously. Figure 6.12 shows that energy resolution is degraded when the prong is closer than 20 cm from the detector extremities in x and y, and also for prongs with $z > 1565$ cm or $z < 20$ cm. These cut values are marked with a red dashed line in the figure. Given the detector limits are at approximately ± 200 cm in x and y, this translates into containment limits of ± 180 cm. Similar studies for non-muon prongs yield the same limits; however, we additionally require that prongs stop before reaching the muon catcher, which starts at $z = 1275$ cm. Hence the containment limits are defined as follows

$$\begin{aligned}
 -180.0 \text{ cm} &< \text{Min}(\text{prong_start_x}, \text{prong_end_x}) \\
 180.0 \text{ cm} &> \text{Max}(\text{prong_start_x}, \text{prong_end_x}) \\
 -180.0 \text{ cm} &< \text{Min}(\text{prong_start_y}, \text{prong_end_y}) \\
 180.0 \text{ cm} &> \text{Max}(\text{prong_start_y}, \text{prong_end_y}) \\
 20.0 \text{ cm} &< \text{Min}(\text{prong_start_z}, \text{prong_end_z}) \\
 \textbf{Muon: } 1565.0 \text{ cm} &> \text{Max}(\text{prong_start_z}, \text{prong_end_z}) \\
 \textbf{non-Muon: } 1275.0 \text{ cm} &> \text{Max}(\text{prong_start_z}, \text{prong_end_z})
 \end{aligned} \tag{6.7}$$

6.3.5 Particle identification

Particle identification is done using the single particle trained prong CVN (see Section 3.2.4). Prong CVN is used to identify a muon and pion prong, and also to reject any

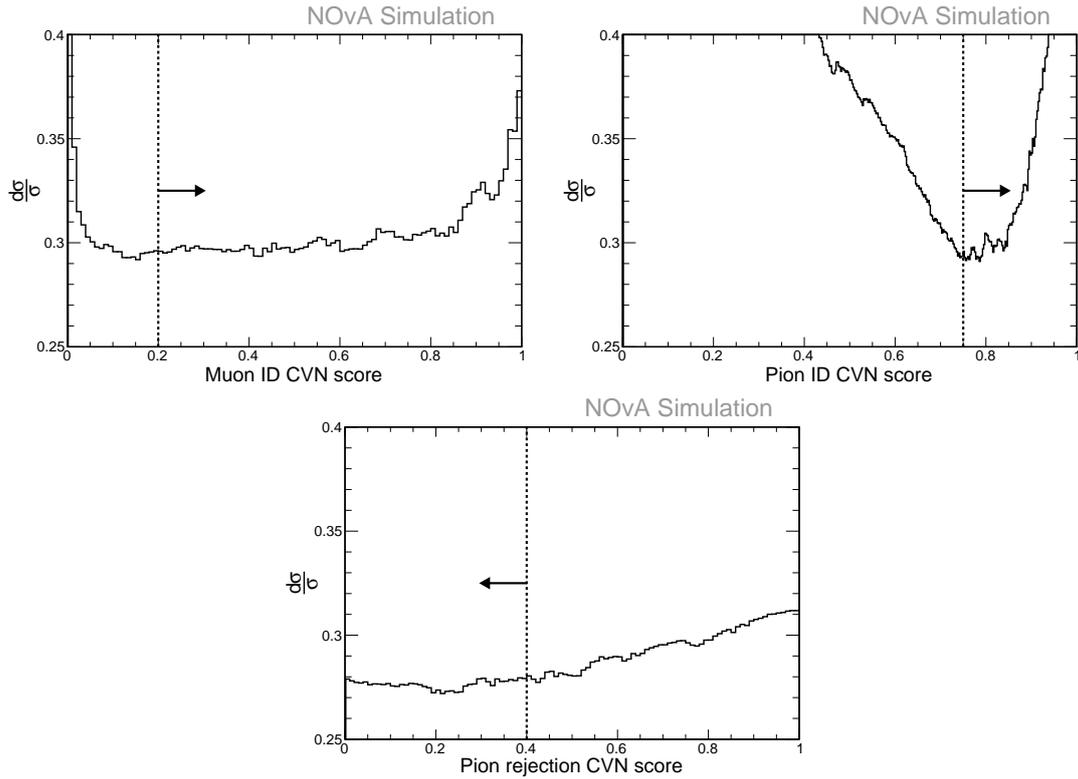


Figure 6.13: Fractional uncertainty on the integrated cross-section as a function of cut values. For muon and pion ID, we determine a minimum CVN score. For pion rejection, we determine a maximum CVN score. The cut values are marked with a dashed line. To make these plots, preselection cuts are applied. Here, the signal is ν_μ CC $1\pi^\pm$, with simple muon kinematic thresholds.

additional pions once the primary pion prong is found (to ensure $1\pi^\pm$). As mentioned in Section 6.3.2, prongs must have at least 10 hits before muon/pion identification is attempted. For pion rejection, we require at least 4 hits in the prong. Cut values were determined by looking at the fractional uncertainty on the integrated cross-section as a function of cut value.

Some interplay between these cuts is expected, e.g. the pion ID cut affects the location of the muon ID cut and vice-versa. Hence these cuts were optimised iteratively. First, the muon ID cut was optimised. Then with the muon ID cut applied, the pion ID cut was optimised. With both of these cuts applied, the pion rejection cut was optimised. Each of the cuts was then re-optimised by applying all the cuts from the previous round of optimisation, except for the one currently being optimised. This was done until all the cut values stabilised, which took 3 iterations of optimisation. The final fractional uncertainty

plots can be seen in Figure 6.13. Note that, similarly to in Section 6.3.3, there are statistical fluctuations in Figure 6.13, and hence we are free to choose a round number cut value which is loose while still in the region of minimum fractional uncertainty. Note that the prong identified as the muon is excluded when performing pion identification. Similarly, for pion rejection, the prongs identified as the muon and pion are excluded.

The cut values for muon ID, pion ID, and pion rejection are 0.2, 0.75, and 0.4, respectively.

6.4 Resolution and binning

Bin width is determined by looking at the bias and resolution of reconstructed quantities. We require that the bin width w_i satisfies the condition

$$w_i \gtrsim r_i + |b_i| \quad (6.8)$$

where r_i and b_i are the resolution and bias respectively in bin i . This is done to mitigate bin-to-bin migrations during unfolding.

For both pions and muons, the reconstructed angle is given by the angle with respect to the beam of the prong selected as the pion or muon. Muon energy is estimated using the track length in the active and muon catcher regions of the detector. This is done using preexisting spline-based energy estimators, which are used in the oscillation and ν_μ CC inclusive analyses.

To calculate the bias and resolution of some quantity X , a 2D histogram is made of $(X_{\text{reco}} - X_{\text{true}})$ vs X_{true} . A Gaussian distribution is fit in each slice of X_{true} . The mean and standard deviation of the fitted distribution gives the bias and resolution, respectively, at that value of X_{true} .

The muon kinematic binning scheme is less important than the pion binning scheme, since no unfolding is done in muon kinematics. Nonetheless, the muon kinematic resolution and bias are shown here for completeness.

Note that this thesis reports a result in bins of θ_π rather than $\cos \theta_\pi$. The reason for this is that measurements in $\cos \theta_\pi$ tend to put most of the events in a handful of bins near $\cos \theta_\pi = 1$, whereas reporting θ_π spreads the events out over several bins. In addition, θ_π

is more intuitive. Finally, it is trivial to convert from θ_π to $\cos \theta_\pi$. Note, however, that for muon kinematics the bins are defined in $\cos \theta_\mu$ space. This is done to match the binning scheme of the preexisting ν_μ CC inclusive analysis.

6.4.1 Pion angle

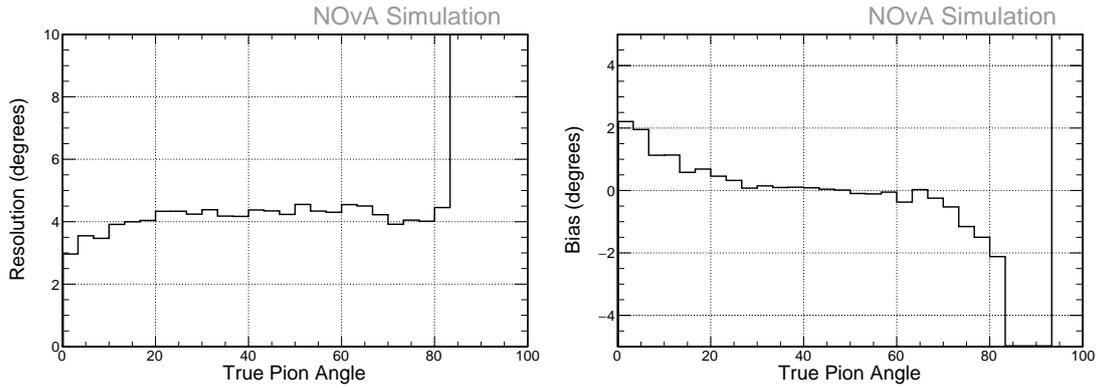


Figure 6.14: Resolution (left) and bias (right) of reconstructed pion angle, as a function of true pion angle.

With reference to Figure 6.14, the pion angle bin edges in degrees are defined as

$$[0, 6, 12, 18, 24, 30, 36, 42, 48, 54, 60] \quad (6.9)$$

6.4.2 Muon kinematics

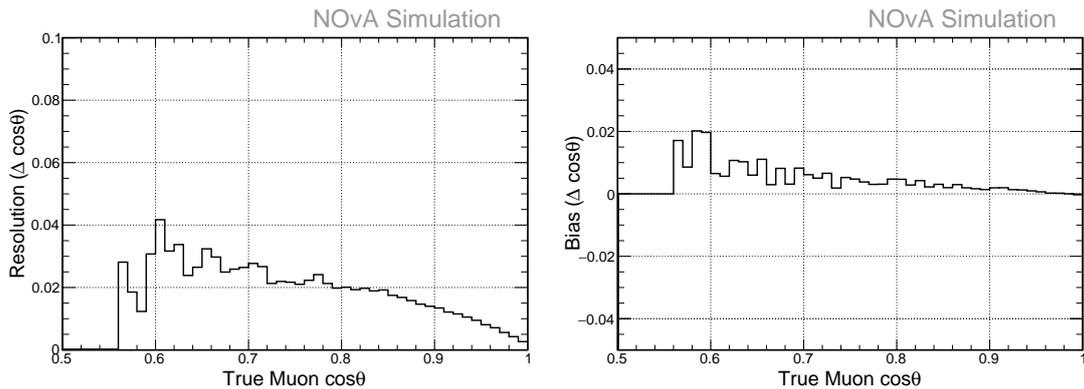


Figure 6.15: Resolution (left) and bias (right) of $\cos \theta_\mu$, as a function of true $\cos \theta_\mu$.

For muon kinematics the same binning as the ν_μ CC inclusive analysis is chosen, defined as follows:

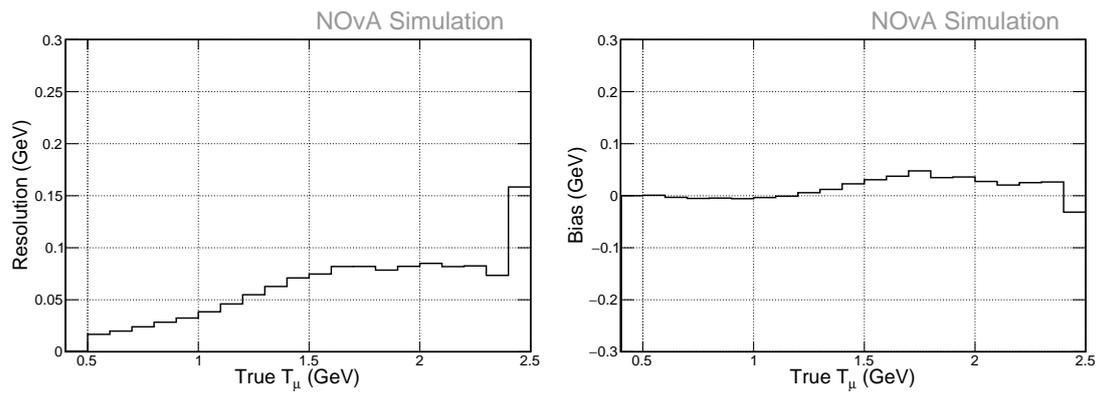


Figure 6.16: Resolution (left) and bias (right) of reconstructed muon kinetic energy, as a function of true muon kinetic energy.

- $\cos \theta_\mu$ - 13 variable-sized bins: 5 bins between 0.5-0.8, 1 bin between 0.8-0.85, 3 bins between 0.85-0.94, 2 bins between 0.94-0.98, and 2 bins between 0.98-1.0
- T_μ - 20 equally size bins between 0.5 and 2.5 GeV

Figures 6.15 and 6.16 show that this binning scheme is sufficient.

6.5 Systematic uncertainties

Name	Type
GENIE	reweight
PPFX	reweight
Beam transport	reweight
Geant4	reweight
Calibration shape	file-based
Calibration up/down	file-based
Light level up/down	file-based
Cherenkov up/down	file-based
Detector ageing	file-based

Table 6.4: All systematic uncertainties used in this analysis. Also listed is whether each is a reweight or file-based uncertainty.

Systematic uncertainties are quantified in one of two ways. The first approach involves re-running the NOvA simulation with adjusted detector properties. These are the so-called file-based systematics. The second approach involves applying weights to events in the nominal simulation to reproduce the effect of making some systematic shift. These are reweight systematics.

This section will give a brief description of each source of systematic uncertainty.

6.5.1 Cross-section modelling (GENIE)

The GENIE event generator [54] provides parameters which can be adjusted to change the simulation physics. NOvA uses the NOvARwgt framework [63] to apply weights to reproduce the effect of re-simulating with varied GENIE parameters. Many universes are created, where a universe is created by adjusting each of a set of parameters randomly within their respective uncertainties. This is the same as the process described in Section 5.3.1, but in this case all the parameters are treated as uncorrelated. By creating many universes, the interplay between different GENIE parameters can be fully captured. The uncertainty band is determined by taking the RMS of the universes.

6.5.2 Hadron production uncertainty (PPFX)

As described in Section 2.1, the NuMI neutrino beam is created by impinging a 120 GeV proton beam with a graphite target, producing hadrons [13]. These hadrons decay to muons and, more importantly, neutrinos. The Package to Predict the Flux (PPFX) [62] was created

to characterise the effect of uncertainty in hadron production at the target on the predicted neutrino flux. Weights are stored in the NOvA simulation files, which, when applied to an ensemble of events, replicate the effect of changing hadron production parameters. Similarly to the GENIE systematic, we create many universes and take uncertainty as the RMS.

6.5.3 Beam transport

In addition to the PPFX uncertainty, a collection of systematic beam transport uncertainties also affect the flux prediction. These arise due to uncertainties in the positioning and performance of the elements of the NuMI beamline. The uncertainties are

- ± 2 kA shift in the horn current used to focus hadrons
- ± 2 mm shift in the beam spot size
- ± 1 mm shift in the beam x and y position on the target
- ± 3 mm shift in the x and y position of horn 1 and horn 2
- ± 7 mm shift in the target z position
- Addition of a magnetic field in the decay pipe
- ± 1 mm shift in the thickness of the water layer cooling horn 1

These uncertainties are applied by reweighting events to the $\pm 1\sigma$ shifts and then conservatively taking the larger shift as the uncertainty.

6.5.4 Hadron reinteraction (Geant4)

As described in Chapter 5, the hadron reinteraction systematic (also known as the Geant4 reweighting systematic) uses event weights to replicate the effect of resimulating events with changed Geant4 physics for particles traversing the detector. Specifically, the weights reproduce the effect of making charged pion and proton inelastic scattering more or less likely. Similarly to the GENIE and PPFX uncertainties, this is implemented as a multiverse systematic (see Section 5.3.1 for an explanation of the term multiverse).

6.5.5 Calibration up/down

The calibration systematic characterises the uncertainty on the absolute detector calibration (see Section 3.3.2). Samples of Monte Carlo are created with a $\pm 5\%$ shift in the absolute detector calibration [101]. These samples define a $\pm 1\sigma$ error band on histograms.

6.5.6 Calibration shape

The detector calibration is different at the ends of cells compared to the middle of a cell (see Section 3.3.1). The relative calibration procedure does not fully capture this effect, so a calibration shape uncertainty is created, where Monte Carlo is resimulated with altered relative calibration to cover this uncertainty [101]. Only a $+1\sigma$ sample is created, which is then reflected across the nominal simulation to create a -1σ shift on histograms.

6.5.7 Light level up/down

The light level systematic is a two-sided uncertainty where events are resimulated with with a $\pm 5\%$ change in the amount of scintillation light produced [101].

6.5.8 Cherenkov up/down

The Cherenkov systematic is a two-sided uncertainty where events are resimulated with a $\pm 6.2\%$ change in the amount of Cherenkov light produced [101].

6.5.9 Detector ageing

To model the effect of detector degradation over time, a sample of Monte Carlo is made where the light model parametrisation decreases by 4.5% annually [102].

6.6 Data-driven purity estimation

This analysis uses a global template fitting method to estimate signal purity, as described in Section 6.6.2. However, to introduce the reader to the basic concepts, a more straightforward method will first be described in Section 6.6.1.

6.6.1 Template fit

The goal of template fitting is to obtain a data-driven estimate of purity. Signal and background templates are created from Monte Carlo simulation, in some variable⁴ that displays

⁴In this case, it is the output of a boosted decision tree. More details can be found in Section 6.6.3.1

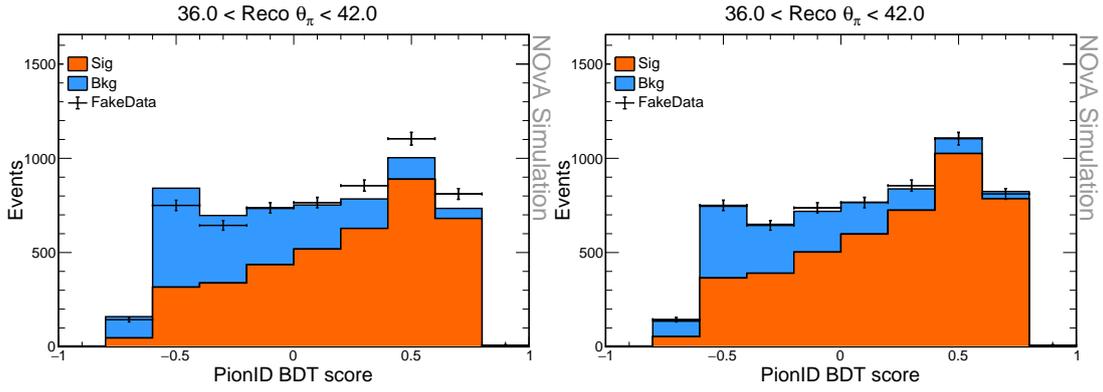


Figure 6.17: Example templates with fake data overlaid. (Left) pre-fit templates. (Right) post-fit templates

a shape difference between the templates, as shown in Figure 6.17. A χ^2 minimisation is performed to determine fit parameters θ_0 and θ_1 (which we can describe collectively as $\vec{\theta}$), which control the normalisation of the signal and background⁵ templates respectively. The χ^2 is formulated as follows

$$\chi^2 = (\vec{x} - \vec{\mu}(\vec{\theta}))^T \cdot V(\vec{\theta})^{-1} \cdot (\vec{x} - \vec{\mu}(\vec{\theta})) \quad (6.10)$$

where \vec{x} is the observed data, $\vec{\mu}$ is the total Monte Carlo simulation after modification by $\vec{\theta}$, and V is a covariance matrix describing uncertainties. The modified Monte Carlo is given by

$$\mu_i = \theta_0 \cdot N_i^{sig} + \theta_1 \cdot N_i^{bkg} \quad (6.11)$$

where N_i^{sig} and N_i^{bkg} are the number of events in template bin i of the signal and background templates, respectively, as predicted by the nominal simulation.

The covariance matrix can be expressed as the sum of a statistical and systematic component

$$V = V^{stat}(\vec{\theta}) + V^{syst}(\vec{\theta}) \quad (6.12)$$

⁵It is possible to define several background templates, e.g. “CC background”, “NC background”. However, one “Total background” template is used in this analysis. Hence that is how the fit process will be described

The statistical covariance is given by

$$V_{ij}^{stat} = \delta_{ij}\mu_i \quad (6.13)$$

As we fit for template normalisations, systematic samples are normalised to the integral of the nominal simulation. These are so-called ‘‘shape only’’ uncertainties. For a one-sided systematic shift $\vec{\mu}'(\vec{\theta})$ the covariance matrix is given by

$$V_{ij}^{syst} = (\lambda\mu'_i - \mu_i)(\lambda\mu'_j - \mu_j) \quad (6.14)$$

where λ is a normalisation factor

$$\lambda = \frac{\sum_i \mu_i}{\sum_i \mu'_i} \quad (6.15)$$

For a two sided systematic shift where $\mu^+(\vec{\theta})$ and $\mu^-(\vec{\theta})$ represent $+1\sigma$ and -1σ shifts respectively, the covariance is given by

$$V_{ij}^{syst} = \frac{1}{2}(\lambda^+\mu_i^+ - \mu_i)(\lambda^+\mu_j^+ - \mu_j) + \frac{1}{2}(\lambda^-\mu_i^- - \mu_i)(\lambda^-\mu_j^- - \mu_j) \quad (6.16)$$

For a multiverse systematic with N universes, the covariance is given by⁶

$$V_{ij}^{syst} = \frac{1}{N} \sum_{m=1}^N (\lambda^m \mu_i^m - \mu_i)(\lambda^m \mu_j^m - \mu_j) \quad (6.17)$$

The total systematic covariance matrix is given by the sum of the covariance matrices for each source of systematic uncertainty.

By minimising the χ^2 in Equation 6.10, the best fit parameters $\hat{\theta}$ are determined, from which an estimate of purity is determined

$$\hat{p} = \frac{\sum_i \hat{\theta}_0 \cdot N_i^{sig}}{\sum_i (\hat{\theta}_0 \cdot N_i^{sig} + \hat{\theta}_1 \cdot N_i^{bkg})} \quad (6.18)$$

⁶Note that the central value μ is known *a priori*, rather than being calculated as the mean of the universes. Hence Bessel’s correction is not needed here.

6.6.2 Global template fit

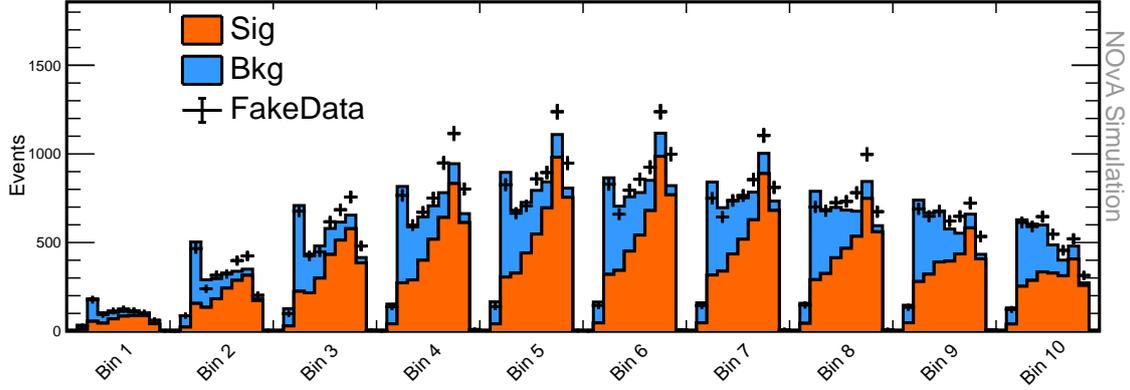


Figure 6.18: Concatenated fake data and pre-fit Monte Carlo as used in the global template fit. The x-labels “Bin 1” etc., are analysis bin labels (in this case, they are pion angle bins). i.e. Bin 1: ($0 < \theta_\pi < 6$), Bin 2: ($6 < \theta_\pi < 12$), etc. The fake data was made from MC by artificially increasing (decreasing) the amount of signal (background). Hence there is an excess (deficit) of events in the signal (background) rich regions.

The method outlined in Section 6.6.1 could be used to obtain an independent estimate of purity \hat{P}_j in each of the analysis bins j ⁷. However, this assumes that the purity in each analysis bin is independent of the purity in other analysis bins. This is not true, as systematic uncertainties can cause migrations between analysis bins. To account for these migrations, a global template fit is performed.

In a global template fit, we minimise

$$\chi^2 = (\vec{x}_g - \vec{\mu}_g(\vec{\theta}))^T \cdot V(\vec{\theta})^{-1} \cdot (\vec{x}_g - \vec{\mu}_g(\vec{\theta})) \quad (6.19)$$

where \vec{x}_g is the concatenation of the observed data in each analysis bin

$$\vec{x}_g = (\vec{x}^0, \vec{x}^1, \dots, \vec{x}^j) \quad (6.20)$$

Similarly, $\vec{\mu}_g(\vec{\theta})$ is the concatenation of the Monte Carlo simulation under the effect of the template normalisation parameters $\vec{\theta}$

$$\vec{\mu}_g(\vec{\theta}) = (\vec{\mu}^0, \vec{\mu}^1, \dots, \vec{\mu}^j) \quad (6.21)$$

⁷**Analysis bin:** bin we intend to report a cross-section measurement in.

where

$$\vec{\mu}^j = \theta_{2j} \cdot \vec{N}^{sig,j} + \theta_{2j+1} \cdot \vec{N}^{bkg,j} \quad (6.22)$$

Figure 6.18 shows concatenated fake data and Monte Carlo. The covariance matrix V is constructed in the same way as in Section 6.6.1, but the concatenated Monte Carlo is used instead.

The fit determines an estimate of purity in each analysis bin simultaneously. Once the best-fit parameters $\hat{\theta}$ have been determined, the estimate of purity is given by

$$\hat{P}_j = \frac{\sum_i \hat{\theta}_{2j} \cdot N_i^{sig,j}}{\sum_i (\hat{\theta}_{2j} \cdot N_i^{sig,j} + \hat{\theta}_{2j+1} \cdot N_i^{bkg,j})} \quad (6.23)$$

The post-fit uncertainty on the purity is obtained by creating 1000 universes from the fit parameter covariance matrix and then taking the RMS of the ensemble.

6.6.3 Implementation

6.6.3.1 Template variable

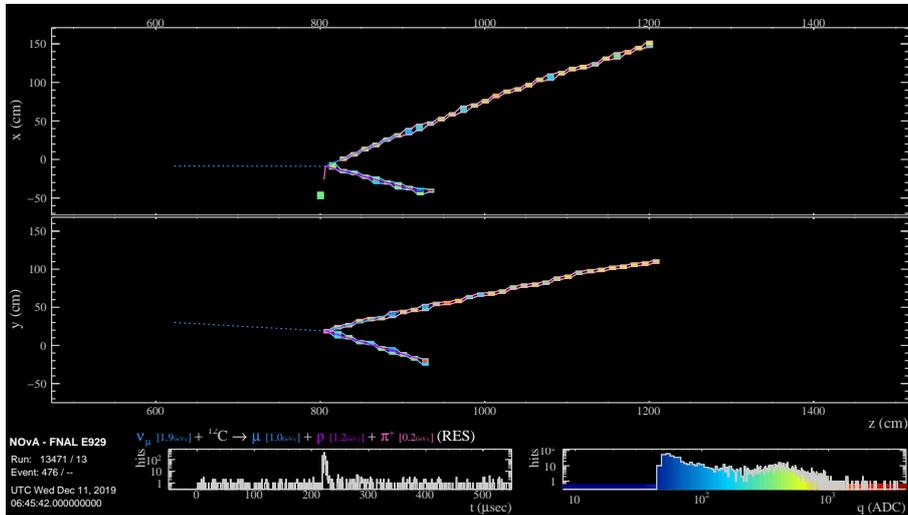


Figure 6.19: Selected background event (pion is below kinetic energy threshold in signal definition). We see that the prong identified by prong CVN as the pion prong is, in fact, a proton.

The ν_μ CC $1\pi^\pm$ analysis requires identifying a charged pion in the event. It was observed that a significant fraction ($\sim 58\%$) of simulated background events passing selection were events in which the prong identified as the pion (using prong CVN) was, in

fact, some other non-pion species. Figure 6.19 shows an event display showing one such simulated event. A boosted decision tree (BDT) was trained to mitigate these backgrounds.

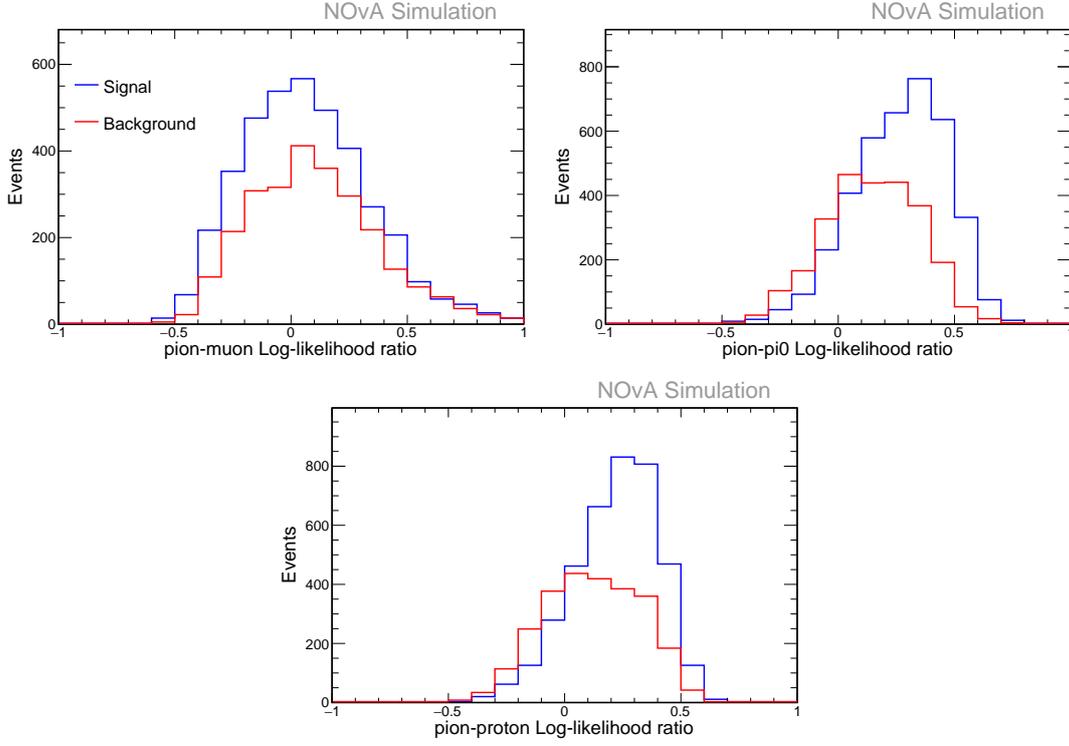


Figure 6.20: Training sample input variables to BDT.

As inputs, the BDT was given variables from the prong selected as the pion (using prong CVN). A pion-muon dE/dx log-likelihood ratio was derived by subtracting two existing electron log-likelihood ratios in the simulation files

$$\log \frac{L_\pi}{L_\mu} = \log \frac{L_e}{L_\mu} - \log \frac{L_e}{L_\pi} \quad (6.24)$$

Similar log-likelihood ratios were derived for protons and neutral pions. The BDT input variables are shown in Figure 6.20. Note that the pion-muon log-likelihood is centred near zero, as pions and muons have similar dE/dx profiles due to their similar masses. Note also that the signal and background histograms have a similar shape for the pion-muon log-likelihood for the same reason. Nonetheless the BDT output has better shape discrimination between signal and background when the pion-muon log-likelihood is used, hence it is retained as an input.

The BDT hyperparameters were optimised to maximise the integral of the ROC

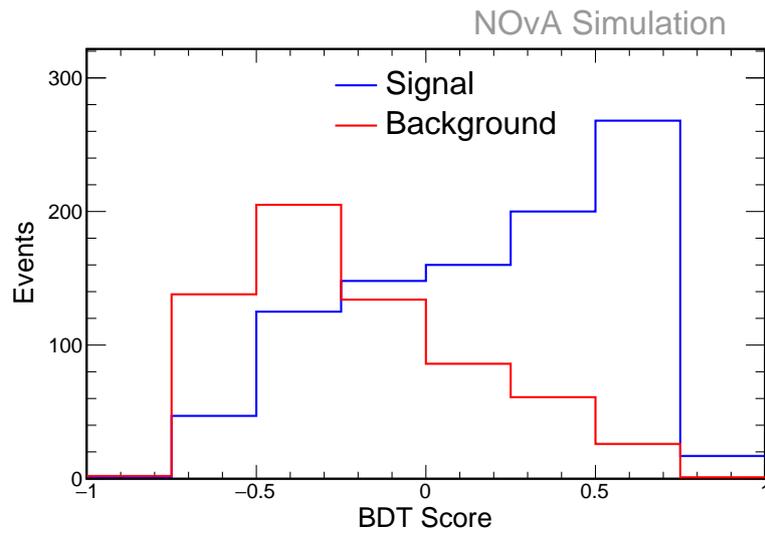


Figure 6.21: Output of BDT evaluated on the testing sample.

curve ⁸, while avoiding overtraining. An 80:20 train-test split was used on a sample of $\sim 3\%$ of the nominal Monte Carlo simulation, which was retired from further use. The BDT output evaluated on the testing sample is shown in Figure 6.21;

⁸The integral of the receiver operating characteristic curve is a standard figure of merit for binary classifiers

6.6.4 Fake data studies

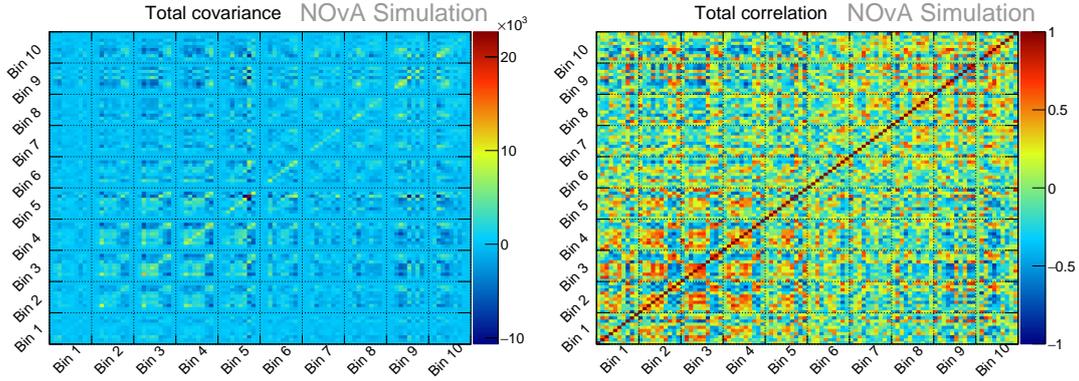


Figure 6.22: Total covariance (left) and correlation matrix (right).

In this section, the template fit is evaluated on shifted fake data. In Section 6.6.4.1, the robustness of the template fitting procedure is tested against artificial normalisation shifts applied to the fake data templates. In Section 6.6.4.2, the fitting procedure is tested on fake data made from an ensemble of systematically shifted GENIE universes. Figures 6.22-6.24 show the total covariance and correlation matrices, as well as them split into their systematic components. These Figures show which systematic uncertainties contribute the most to the overall template uncertainty. They also show the correlations introduced by systematic uncertainties. The covariance matrix calculation is outlined in Section 6.6.1. For a given covariance matrix M the correlation matrix C is defined as

$$C_{ij} = \frac{M_{ij}}{\sqrt{M_{ii}M_{jj}}} \quad (6.25)$$

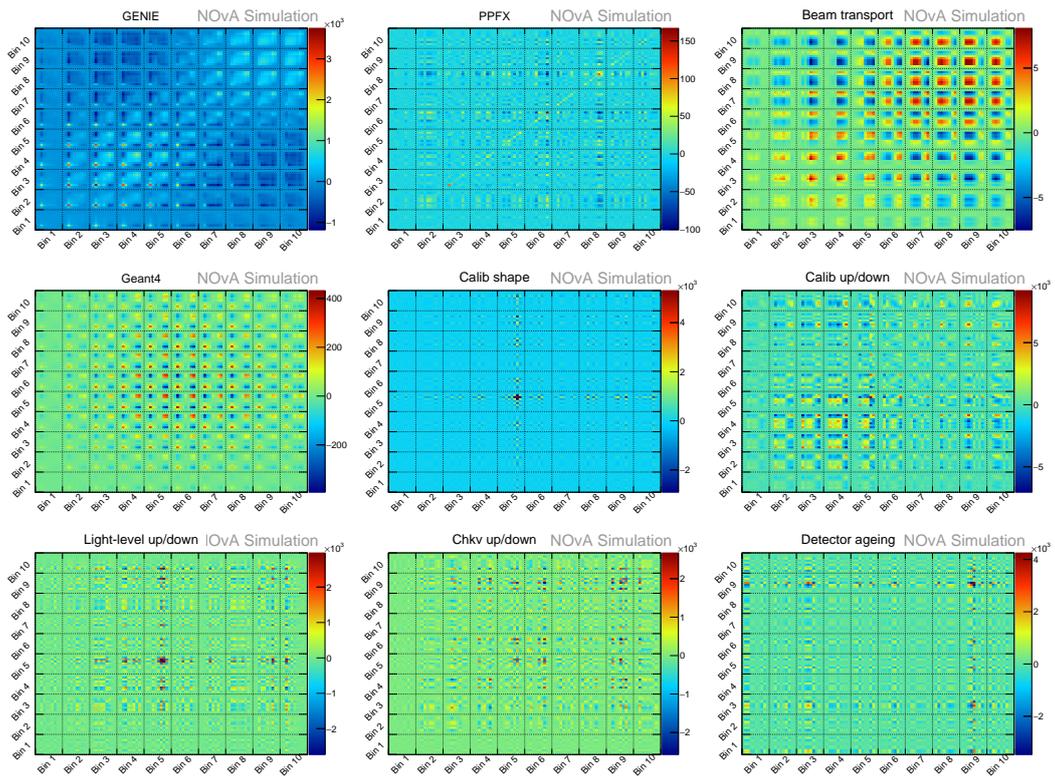


Figure 6.23: Each systematic component of the covariance matrix.

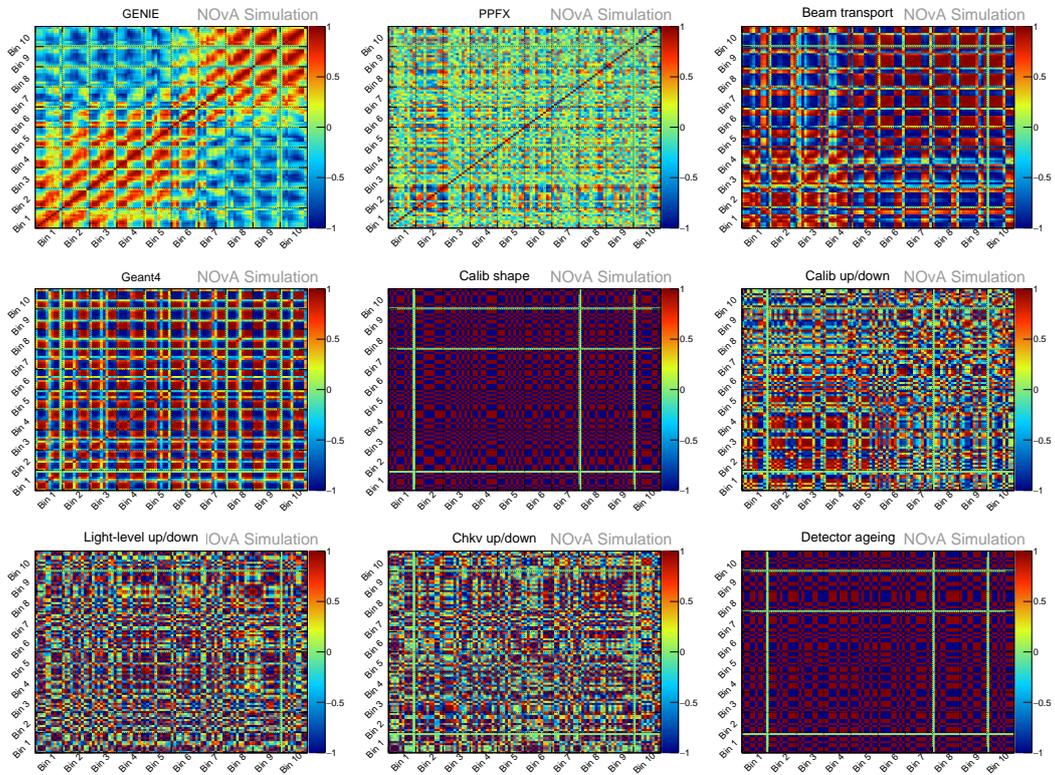


Figure 6.24: Each systematic component of the correlation matrix.

6.6.4.1 Template normalization shift

Here fake data is made by applying scale factors of 1.2 and 0.8 to the signal and background templates, respectively, and then applying random Poisson fluctuations. The pre- and post-fit templates from a single fake data universe can be seen in Figures 6.25 and 6.26. In this fake data universe, the χ^2 is reduced during the fit from 216.9 to 56.6, with 79 degrees of freedom. Figure 6.27 shows the pre- and post-fit purity estimates from this same fake data universe, with the true (fake data) purity overlaid. A χ^2 comparing true and pre-fit purity is 39.8, whereas, for true and post-fit purity, the χ^2 is 3.7, with 10 degrees of freedom. Hence we see that the post-fit estimate of purity is markedly better.

The fit is also evaluated on an ensemble of fake data universes. Each fake data universe is created by providing a different seed to the random number generator used to apply Poisson fluctuations. The results are shown in Figure 6.28, showing that the fit yields an unbiased purity estimate when a scale factor manually scales templates. Furthermore, the fit reduces the fractional uncertainty on the purity.

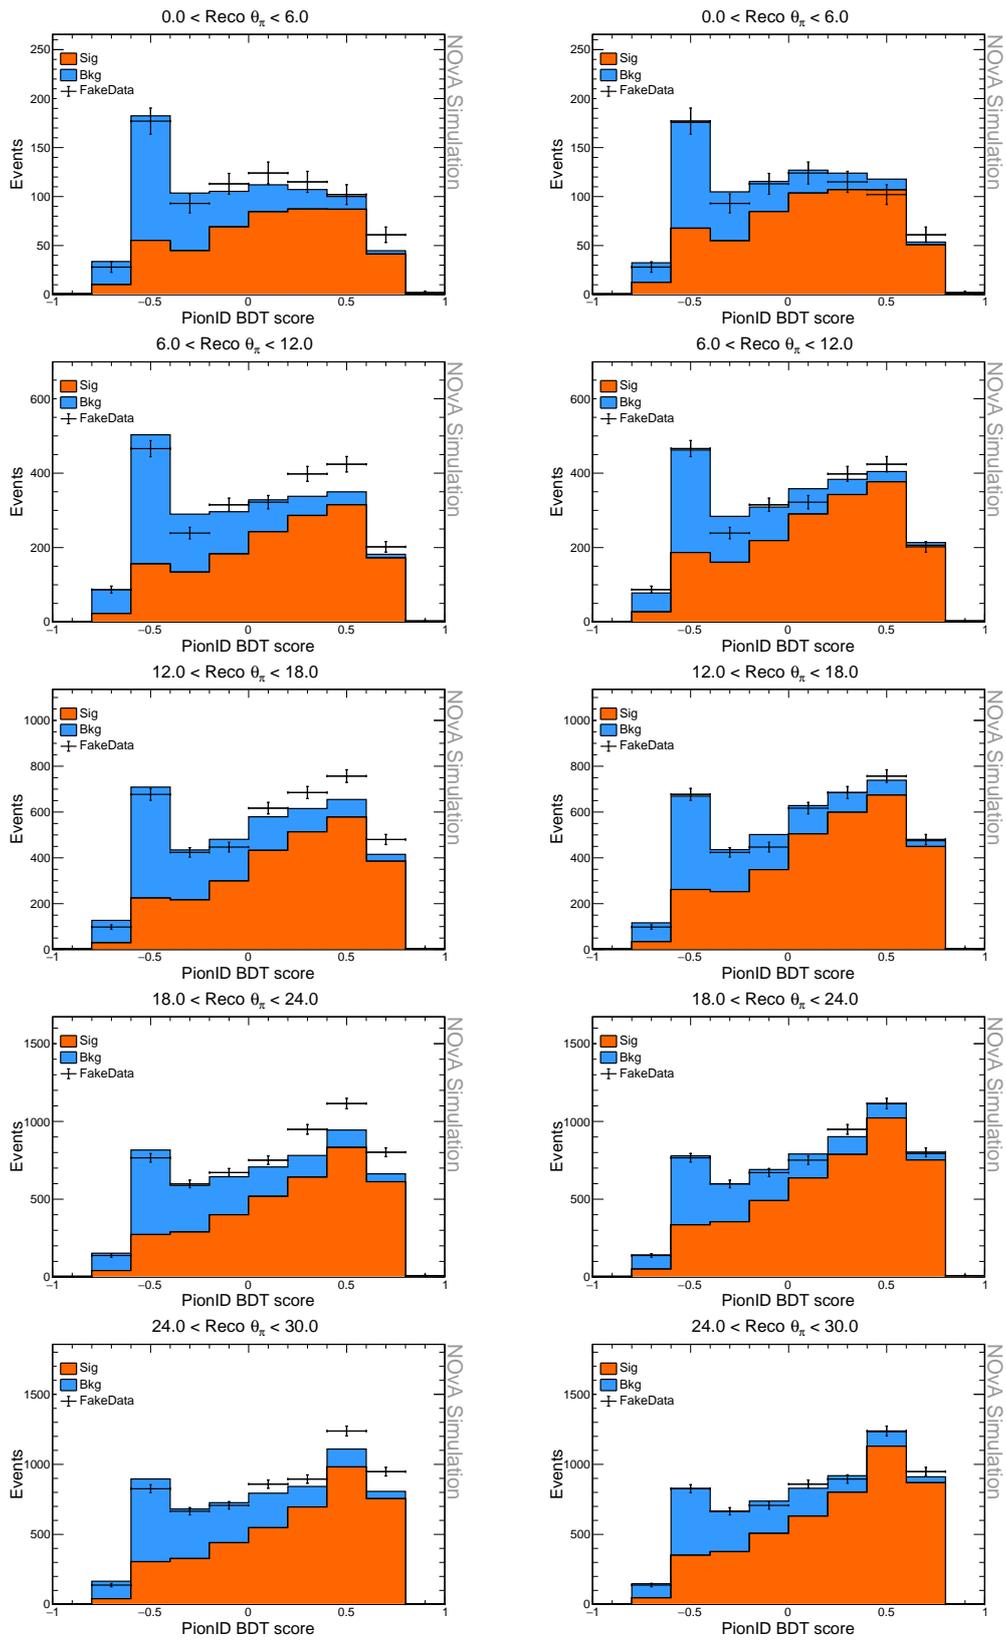


Figure 6.25: From a single fake data universe, pre-fit (left) and post-fit (right) templates with fake data overlaid. Shown here are the first five analysis bins.

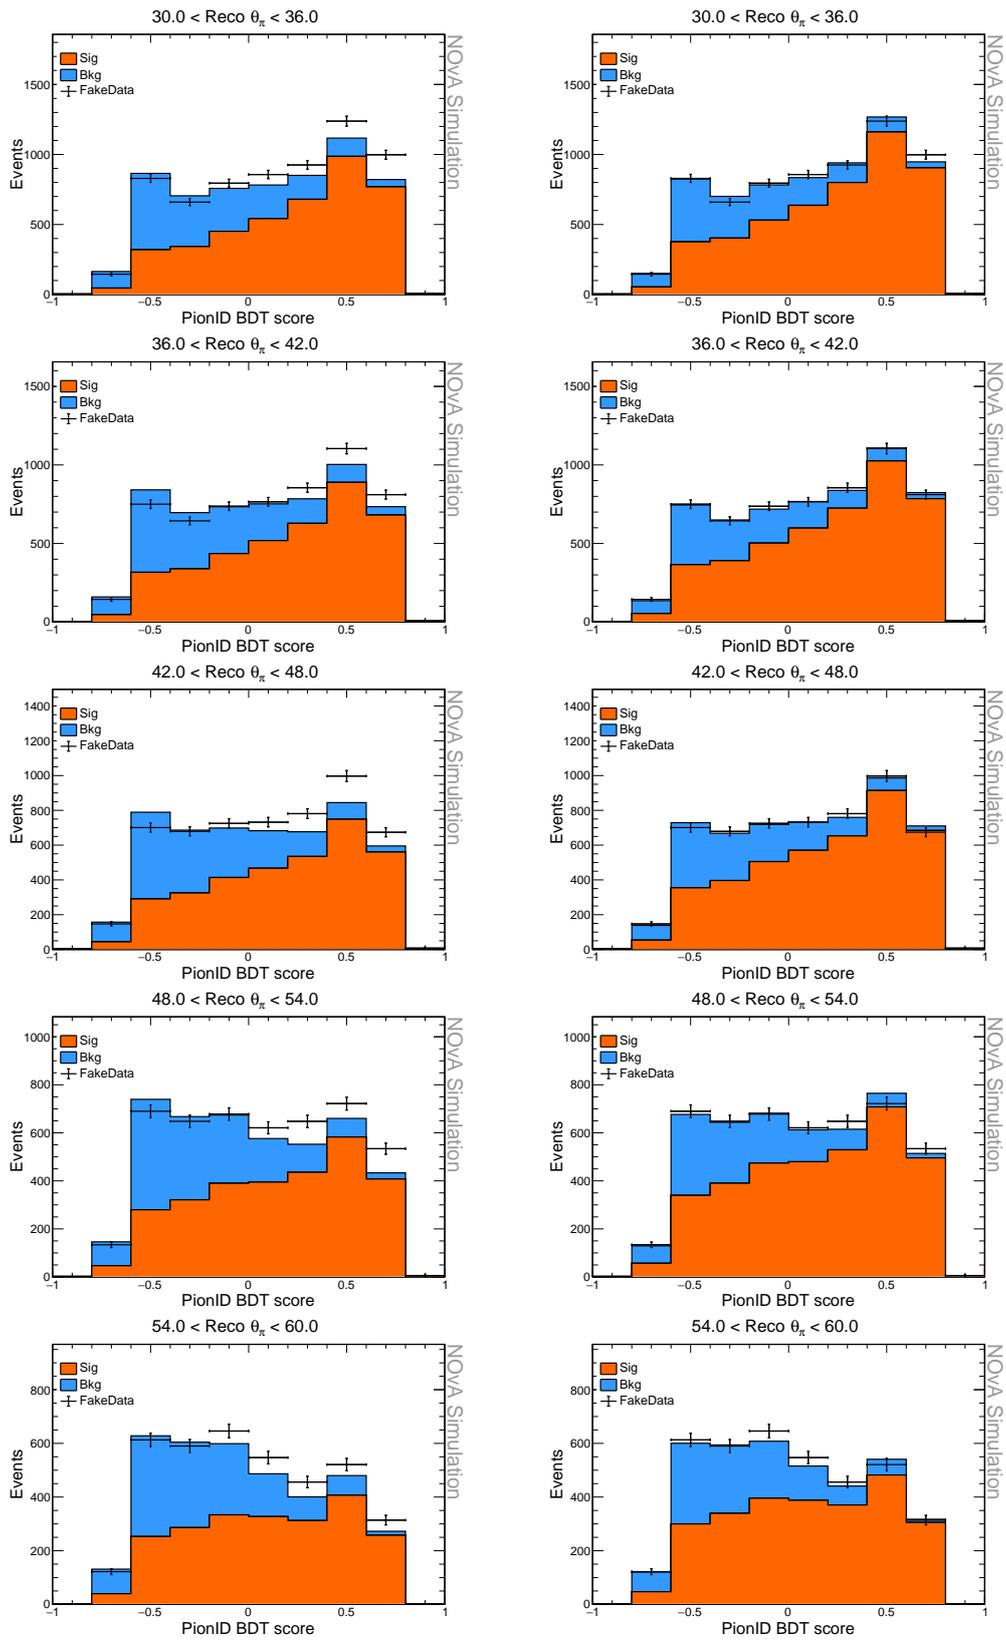


Figure 6.26: From a single fake data universe, pre-fit (left) and post-fit (right) templates with fake data overlaid. Shown here are the last five analysis bins.

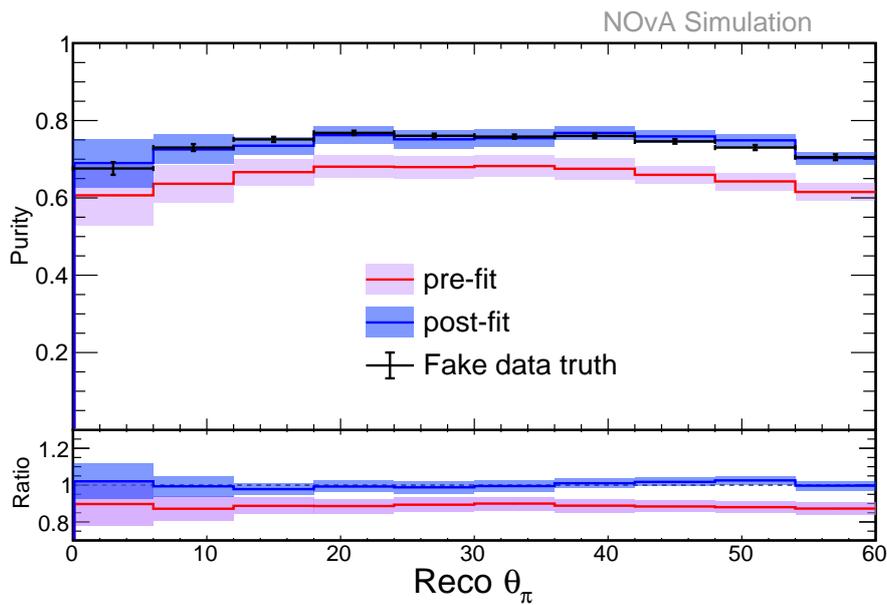


Figure 6.27: Pre- and post-fit purity estimates from a single fake data universe. The calculated χ^2 comparing true and pre-fit purity is 39.8, whereas, for true and post-fit purity, the χ^2 is 3.7, with 10 degrees of freedom.

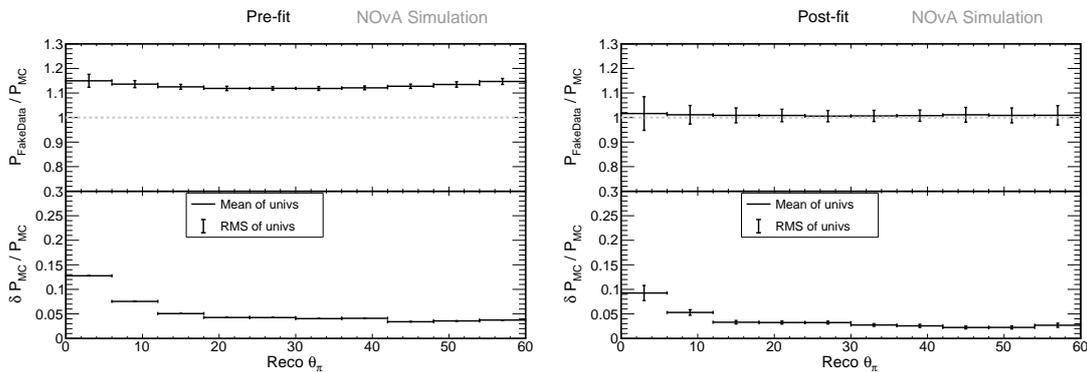


Figure 6.28: Comparison of pre-fit (left) and post-fit (right) purity estimates on an ensemble of 200 fake data universes, with scale factors of 1.2 and 0.8 applied to signal and background templates, respectively. (Top) Mean and RMS of data-MC purity ratio from each universe. (Bottom) Mean and RMS of fractional uncertainty on purity from each universe.

6.6.4.2 GENIE systematic shift

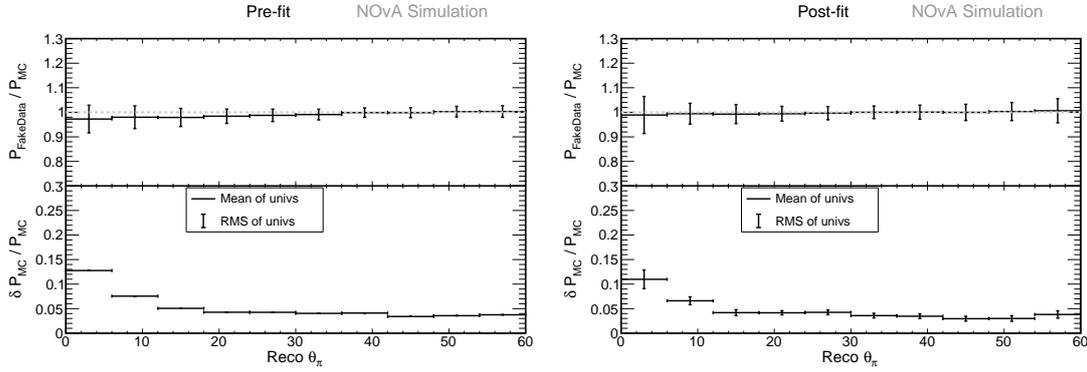


Figure 6.29: Comparison of pre-fit (left) and post-fit (right) purity estimates on an ensemble of 200 fake data universes, each made from a different systematically shifted GENIE universe. (Top) Mean and RMS of data-MC purity ratio from each universe. (Bottom) Mean and RMS of fractional uncertainty on purity from each universe.

To evaluate the robustness of the procedure under systematic shifts, the fit was performed on an ensemble of 200 systematically shifted fake data universes. The fake data in each universe is made from a different shifted GENIE universe, with random Poisson fluctuations applied. The results are shown in Figure 6.29. We find that the fit is robust against GENIE systematic shifts.

6.7 Unfolding

Experiments measure reconstructed quantities (e.g. reconstructed θ_π). However, to compare measurements made by different experiments, it is desirable to report cross-section measurements in bins of some true quantity (e.g. true θ_π). Detectors have finite resolution, so there is a smearing of truth quantities. We may write this as follows⁹

$$A\vec{y} = \vec{x} \quad (6.26)$$

where \vec{x} is the observed event count in reconstructed space bins, \vec{y} is the event count in truth space bins, and A is a smearing matrix characterising the detector's finite resolution. The goal of unfolding is to, in essence, attempt to invert the smearing matrix to obtain the

⁹In some sources, you may see $\vec{x} - \vec{b}$ on the RHS of Equation 6.26, where \vec{b} is the number of background events. Accounting for backgrounds is done via a purity correction as described in Section 6.6, but here we will assume this is already folded into \vec{x} .

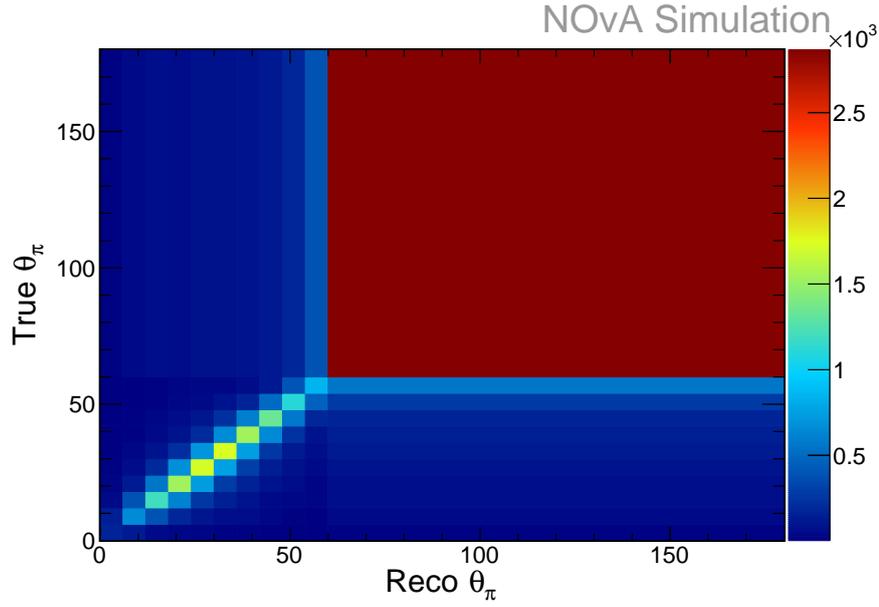


Figure 6.30: Pion angle smearing matrix. Note that the overflow is included, which appears as a large bin in the top right of the figure.

event count in truth space

$$\vec{y} = A^{-1}\vec{x} \quad (6.27)$$

In the presence of statistical fluctuations and imperfect modelling in our simulation, this matrix inversion is an ill-posed problem.

Many unfolding techniques exist; however, they usually have parameters which must be tuned. D’Agostini iterative unfolding is used in this analysis [103]. For this unfolding method, the parameter which must be tuned is the number of unfolding iterations. The Mean Squared Error (MSE) figure of merit is used to optimise the number of unfolding iterations [104]. MSE is defined as follows

$$MSE = \frac{1}{M} \sum_{i=1}^M \frac{V[\hat{y}_i] + b^2(\hat{y}_i, y_i)}{\hat{y}_i} \quad (6.28)$$

where V and b are the variance and bias respectively, and M is the number of bins. This figure of merit is optimised when it is minimised. To optimise the number of unfolding iterations, a sample of simulation independent from that used to create the smearing matrix is made by applying statistical fluctuations to the event counts in true space. This fluctuated

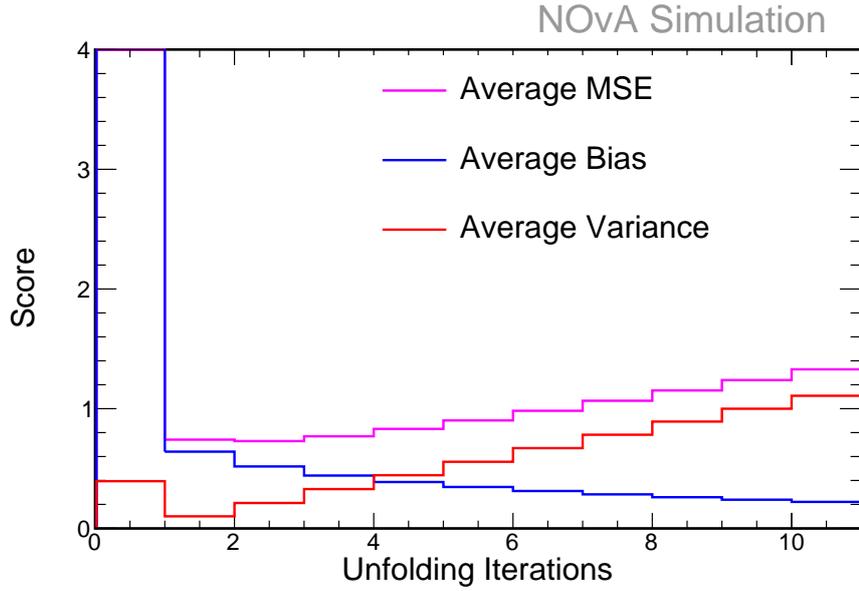


Figure 6.31: Average MSE as a function of the number of unfolding iterations.

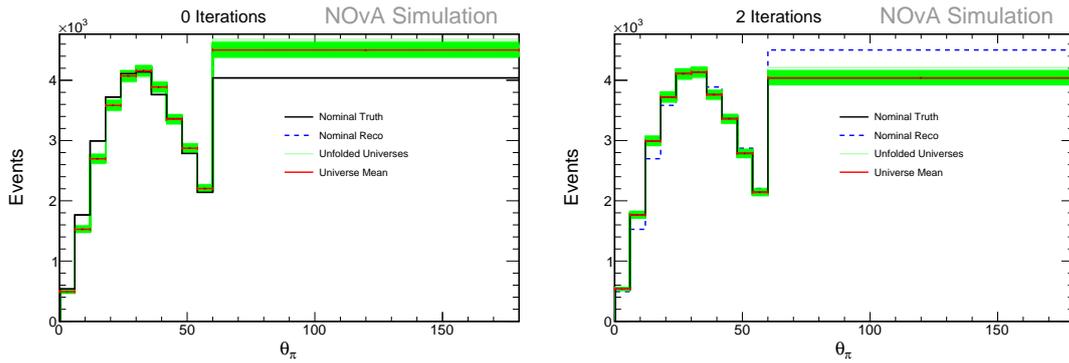


Figure 6.32: An ensemble of unfolded universes after zero iterations of unfolding (left) and two iterations of unfolding (right).

simulation is then smeared to reconstructed space

$$x_i = \sum_j \frac{A_{ij}}{\sum_i A_{ij}} y_j \quad (6.29)$$

Finally we unfold \vec{x} to yield $\hat{\vec{y}}$ and calculate the MSE for $N_{iter} = 1, 2, 3, \dots$. This process is done for many statistically independent universes, and the MSE is averaged

$$MSE_{avg} = \frac{1}{N_{univ}} \sum_{i=1}^{N_{univ}} MSE_i \quad (6.30)$$

where N_{univ} is the number of poisson universes ($N_{univ} = 1000$ here). Note that it is

important to include overflow bins during unfolding to allow for migration in/out of the sample. The smearing matrix can be seen in Figure 6.30. The MSE figure of merit is shown in Figure 6.31, from which we determine two iterations to be optimal. Figure 6.32 shows an ensemble of unfolded histograms for zero and two unfolding iterations.

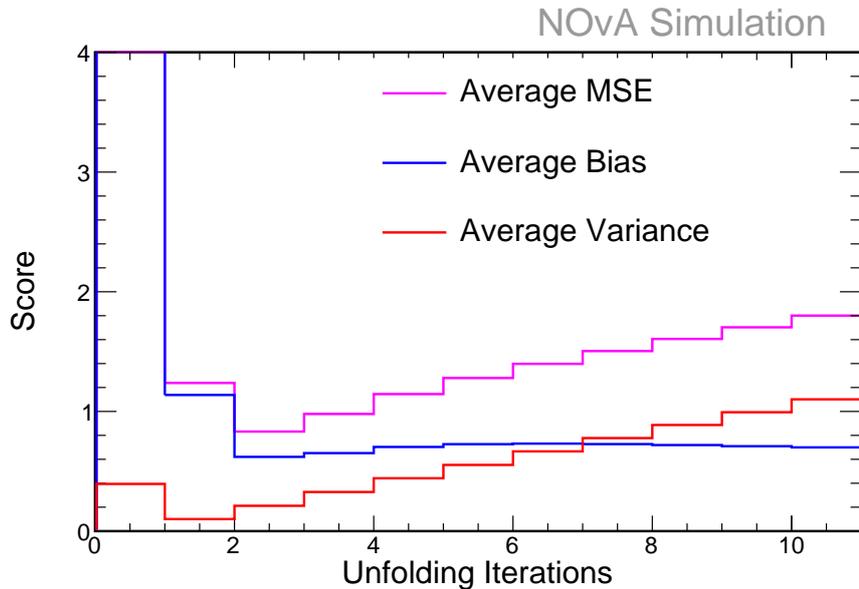


Figure 6.33: Average MSE as a function of the number of unfolding iterations. Here fake data is made by applying random shifts to the GENIE cross-section modelling parameters, but the matrix used for unfolding is made from the nominal simulation.

In order to test how the unfolding responds to systematic variations, the MSE study was repeated, but this time using fake data made by applying random shifts to the GENIE parameters. The optimal number of iterations was still found to be two, as shown in Figure 6.33.

As a check, the Monte Carlo simulation was unfolded without any Poisson fluctuations applied. This is known as a closure test and is shown in Figure 6.34. The unfolded and true distributions exactly coincide, as expected.

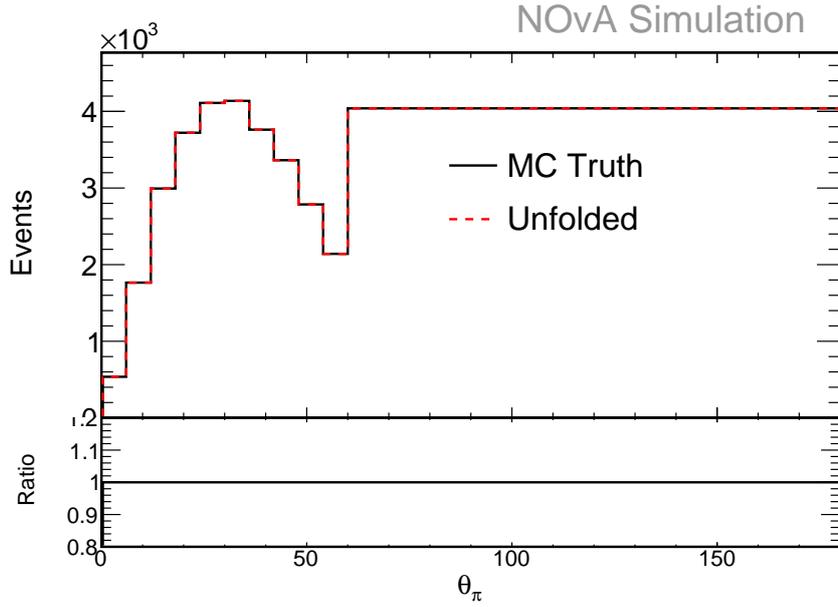


Figure 6.34: Closure test of unfolding.

6.8 Target counting

To enable comparisons of measurements made in different detector materials, cross-section measurements are usually reported per nucleon. To count the number of nucleons in the fiducial volume, first the fiducial mass is calculated. In simulation an ensemble of N points ($N = 10^6$ here) are placed randomly inside the fiducial volume V_{fid} . At each point, the material at that point is retrieved from the geometry simulation. Assuming a uniform density ρ in a small volume around this point, the mass is calculated for this small volume and added to the total. This can be expressed as follows

$$M_{\text{fid}} = \sum_{i=1}^N \frac{V_{\text{fid}}}{N} \rho(\text{mat}(x_i, y_i, z_i)) \quad (6.31)$$

where $\text{mat}(x, y, z)$ returns the material type at a point in the detector. The composition of the fiducial volume by material and element can be seen in Tables 6.5 and 6.6, respectively. The uncertainties arise due to uncertainties on the detector composition determined during a survey.

Once the fiducial mass is calculated, the number of nucleons is given by

$$N_{\text{nuc}} = M_{\text{fid}} N_A \quad (6.32)$$

Material	Density [g/cm ³]	Mass [kg]	Fraction of Total	Uncertainty [kg]
Air	1.19e-3	2.59	3.33e-5	<0.01
Glue	0.98	639.19	8.22e-3	31.96
PVC	1.49	28973	0.37	463.57
Scintillator	0.86	48102	0.62	288.61

Table 6.5: Composition of the fiducial volume, broken down by material.

Element	Z	Mass [kg]	Fraction of Total	Uncertainty [kg]
H	1	8289	0.107	64.97
C	6	51490	0.663	426.0
N	7	19.69	2.534e-4	0.145
O	8	2387	0.031	43.38
Na	11	2.060	2.651e-05	0.033
S	16	76.36	9.825e-4	1.291
Cl	17	12780	0.164	205.5
Ca	20	20.87	2.686e-4	0.334
Ti	22	2560	0.033	40.95
Sn	50	94.57	1.217e-3	1.513

Table 6.6: Composition of the fiducial volume, broken down by element.

where N_A is Avogadro's number. Using this method, the fiducial mass was found to be 77712 ± 481 kg, corresponding to $4.678 \pm 0.032 \times 10^{31}$ nucleons.

6.9 Flux estimation

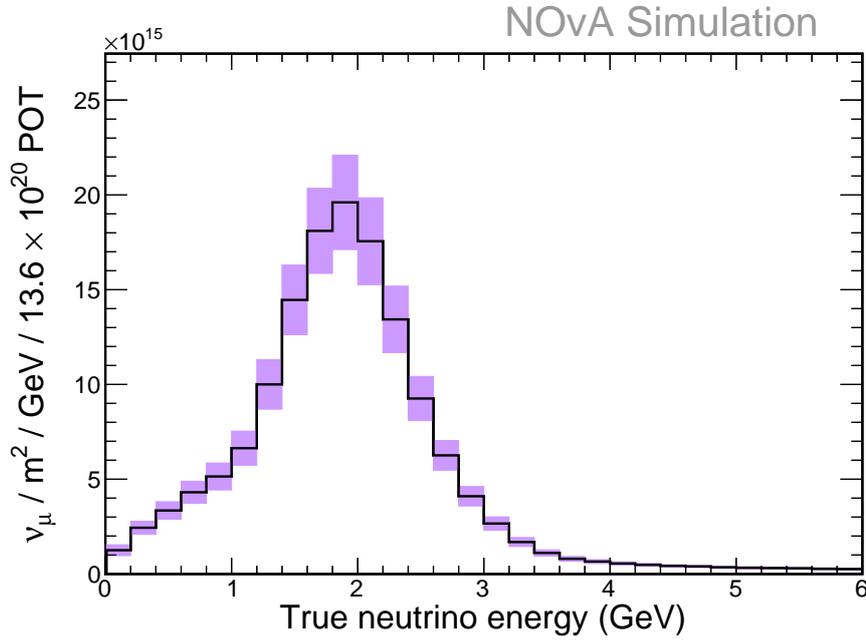


Figure 6.35: Predicted ν_μ flux with associated systematic error bars.

The Near Detector is an angular target of a finite size for the NuMI beam. As a result, the shape of the neutrino flux varies significantly across the face of the detector. The flux is therefore dependent on the fiducial volume. This flux is predicted using Monte Carlo methods. Consider the number of true neutrino interactions in the fiducial volume

$$N(E_\nu) = \Phi(E_\nu)\sigma(E_\nu)N_{\text{nuc}} \quad (6.33)$$

Rearranging Equation 6.33, we find an expression for the flux

$$\Phi(E_\nu) = \frac{N(E_\nu)}{\sigma(E_\nu)N_{\text{nuc}}} \quad (6.34)$$

Hence, a prediction of the flux can be obtained by reweighting true interactions by the inverse of the cross-section and dividing by the number of nucleons. To avoid the need to average over different nuclei and interaction modes, only ν_μ Neutral-current quasi-elastic interactions on carbon are used for this procedure.

Figures 6.35 and 6.36 show the estimated flux and fractional uncertainty, respectively. Note that the measurement reported in this analysis is flux-integrated. Using the method

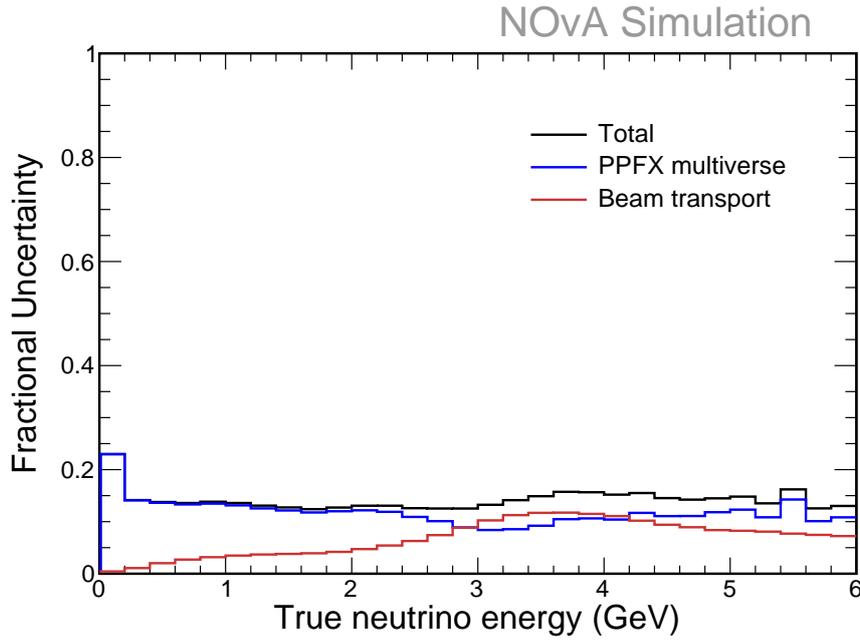


Figure 6.36: Fractional uncertainty on flux, broken down by source of systematic uncertainty

outlined above, the predicted integrated ν_μ flux is $1.59 \pm 0.18 \times 10^{17}/\text{m}^2/\text{GeV}/13.6 \times 10^{20}\text{POT}$.

6.10 Validation of analysis framework

Combining all the elements of the analysis and using Equation 6.1, a cross-section measurement can be done using fake data to test the analysis framework.

In addition to a cross-section measurement, a covariance matrix is also calculated. For a given source of systematic uncertainty, the covariance matrix of the cross-section σ is determined by calculating a shifted cross-section σ' using a systematically shifted response matrix for unfolding and a systematically shifted efficiency correction. For the beam transport and PPFX uncertainties, a systematically shifted value of the flux Φ is also used to calculate σ' . The covariance matrix from each systematic source is then calculated in the same way as in Equations 6.14 - 6.17.

Uncertainty from the purity correction is accounted for by varying the template fit parameters in a correlated way using the parameter covariance matrix determined during the fit. Similarly to in Section 5.3.1, many universes are created (1000 universes are used here), each with a different calculated purity. The cross-section is calculated in each

universe, and the ensemble of calculated cross-sections is used to determine a cross-section covariance matrix.

6.10.1 In-out test

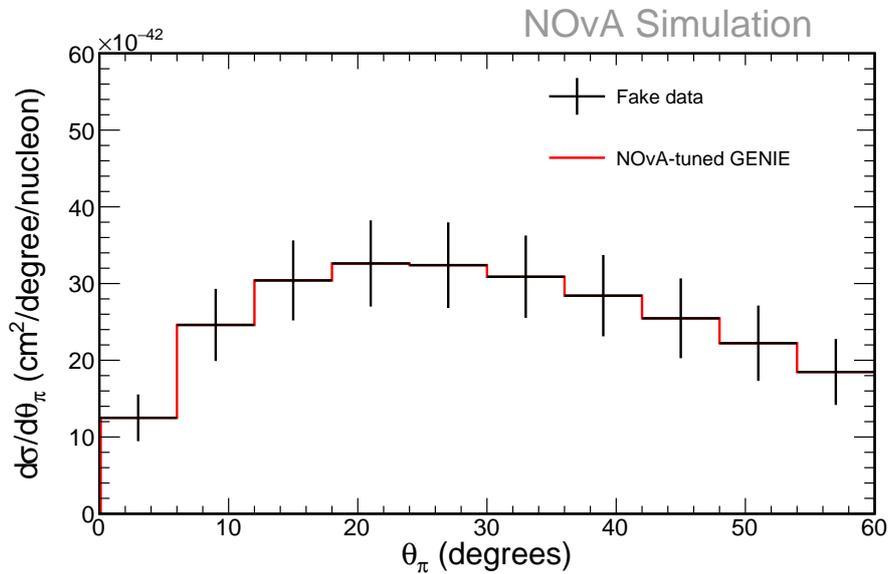


Figure 6.37: Cross-section calculated using Asimov fake data and full uncertainty calculation. Overlaid is the predicted true cross-section from the simulation.

Figure 6.37 shows a cross-section calculated using fake data identical to the nominal simulation used for unfolding, efficiency correction, etc. As expected, the calculated cross-section exactly matches the prediction from the NOvA simulation.

6.10.2 Poisson fluctuated fake data

Figure 6.38 shows a cross-section measurement made using Poisson fluctuated fake data. We see that the cross-section calculation is robust against statistical fluctuations. Figure 6.39 shows a breakdown of the sources of uncertainty. Note that the “Template fit” term encapsulates:

- All sources of systematic uncertainty on purity
- Statistical uncertainty on the number of selected signal events
- Uncertainties coming from unfolding

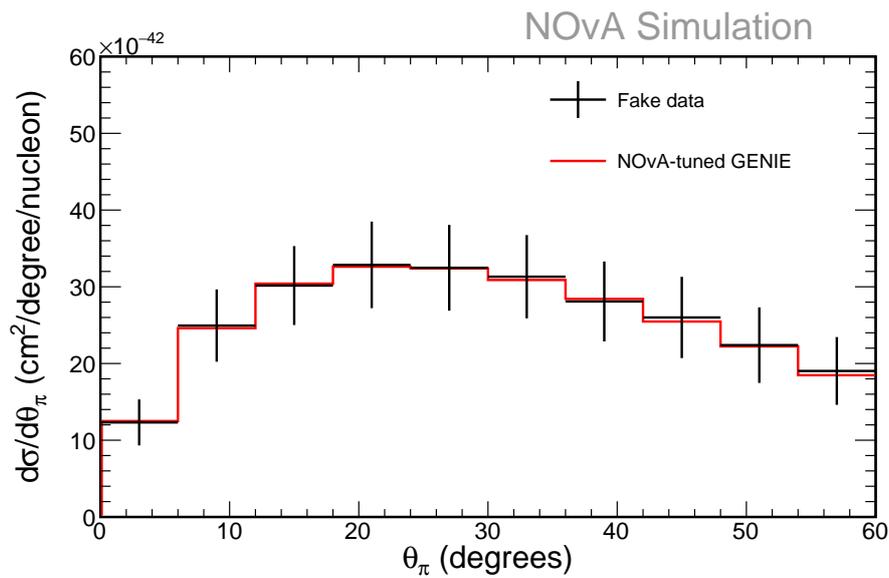


Figure 6.38: Cross-section calculation using fake data with Poisson fluctuations applied. Error bars include all statistical and systematic uncertainties. Overlaid is the predicted true cross-section from the simulation.

The “MC statistical” uncertainty in Figure 6.39 is the statistical uncertainty on the efficiency due to the finite size of the Monte Carlo sample.

Figure 6.40 shows the total covariance and correlation matrices. The various components of the covariance matrix are shown in Figure 6.41.

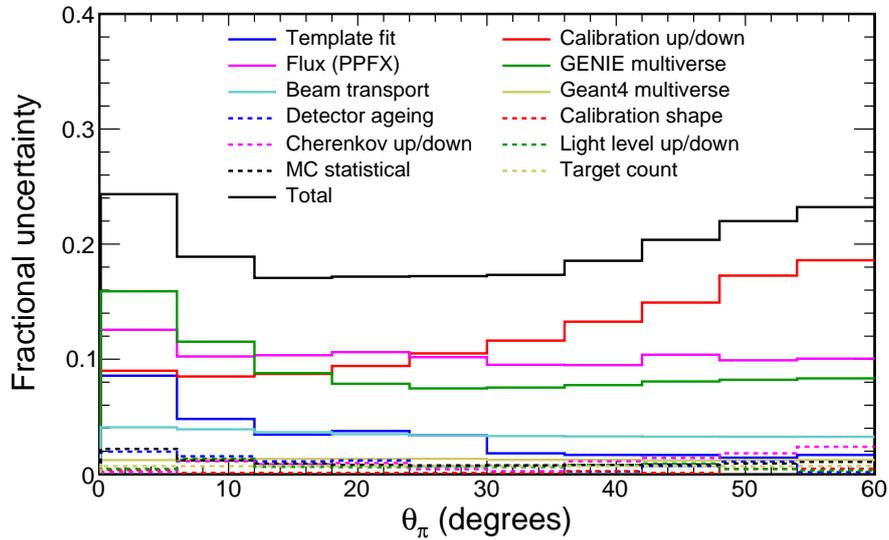


Figure 6.39: Fractional uncertainty on the cross-section measurement, broken down by sources of uncertainty.

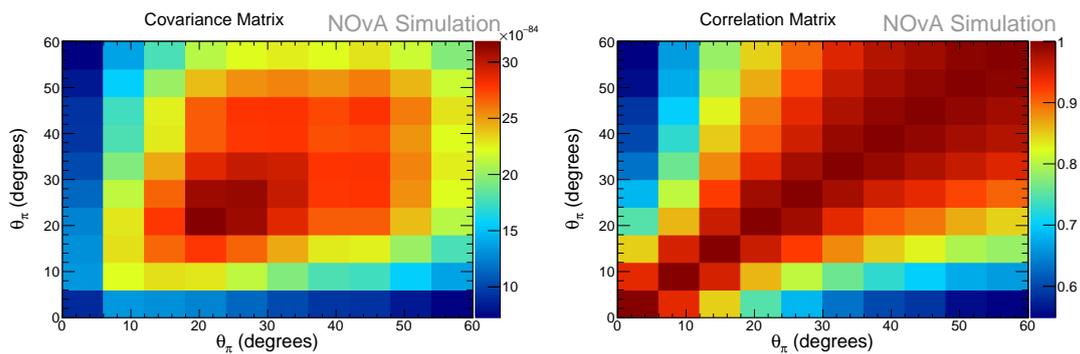


Figure 6.40: Total covariance (left) and correlation (right) matrices for the fake data cross-section measurement.

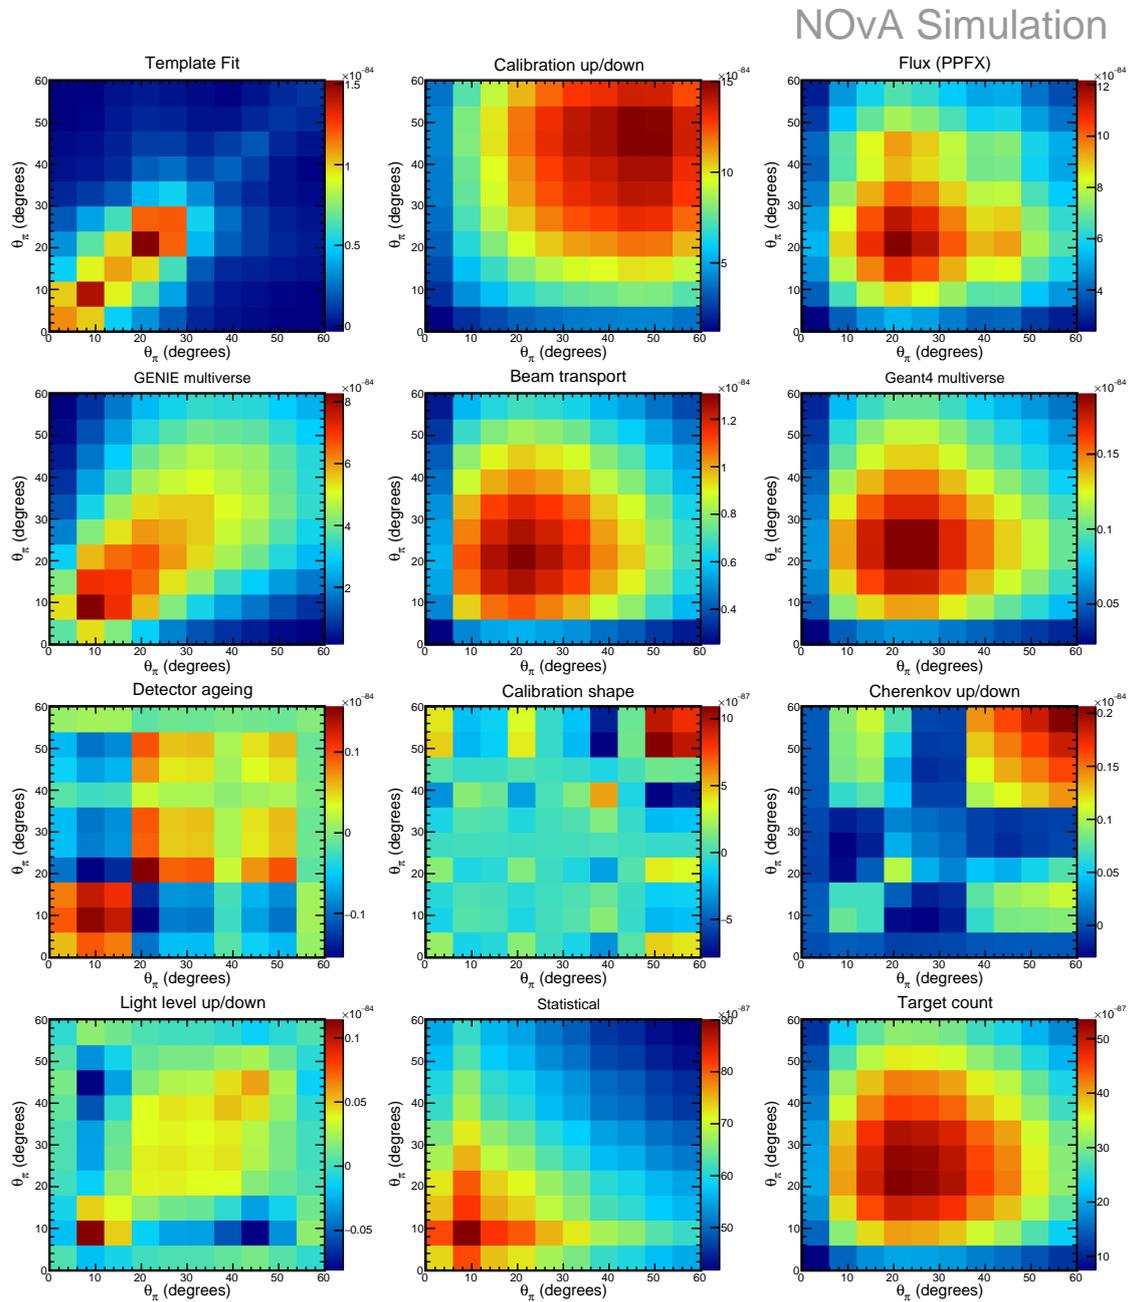


Figure 6.41: Covariance matrix for each source of uncertainty.

6.10.3 Untuned GENIE fake data

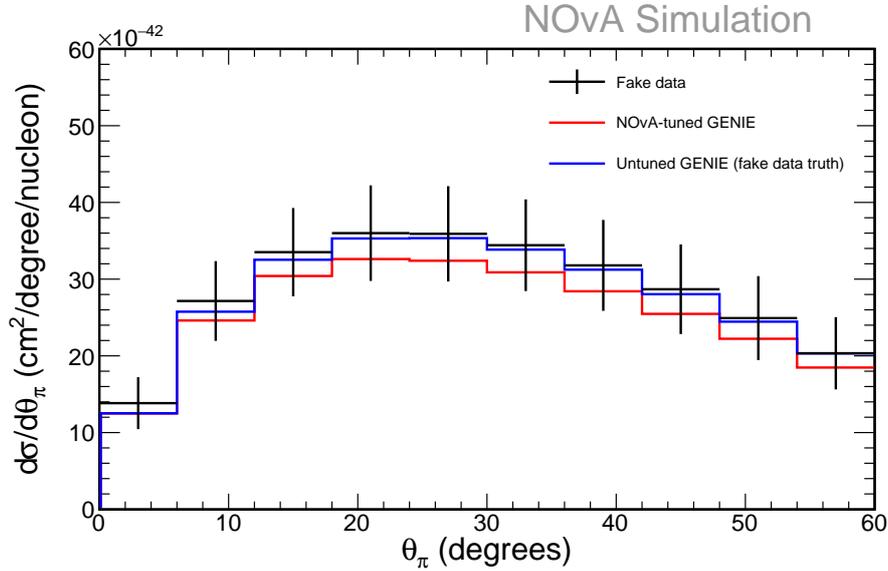


Figure 6.42: Cross-section calculation using fake data made from untuned GENIE simulation. Error bars include all statistical and systematic uncertainties. Overlaid is the predicted true cross-section from tuned and untuned simulation.

Figure 6.42 shows a cross-section calculation using fake data made from GENIE simulation without the NOvA MEC and FSI tunes applied. The calculated cross-section is consistent with the untuned true cross-section rather than being biased towards our nominal (tuned) simulation.

6.11 Summary

This chapter describes an analysis to measure the cross-section of ν_μ Charged-current interactions in the NOvA Near Detector with a single charged pion in the final state. This analysis is designed to target regions of phase space that NOvA is sensitive to. Selection cuts are chosen to minimize uncertainty on the cross-section measurement. A data-driven method of purity estimation is used to provide a better-motivated way of accounting for background contamination. The analysis framework has been tested to ensure it is not biased towards NOvA's base simulation model. The results of the template fit and cross-section measurement on Near Detector data are shown in Chapter 7.

Chapter 7

Results and conclusions

This chapter describes the results of the cross-section analysis described in Chapter 6. This analysis is the first to report a cross-section measurement in the NOvA Near Detector for ν_μ charged-current interactions with a single charged pion in the final state. The analysis was done using 13.8×10^{20} POT of data collected between 2014 and 2020.

7.1 Template fit

This section shows the results of the template fit method described in Section 6.6. The pre and post-fit templates are shown in Figures 7.1 - 7.2. The purpose of the template fit is to obtain a data-driven estimate of purity, which is shown in Figure 7.3. During the fit, the χ^2 was reduced from 217.25 to 144.38 for 79 degrees of freedom. Note that, to make the MC and data coincide, some shape changes to the templates are likely required. However the template fit is unable to change the template shapes. This may explain why the χ^2 is large.

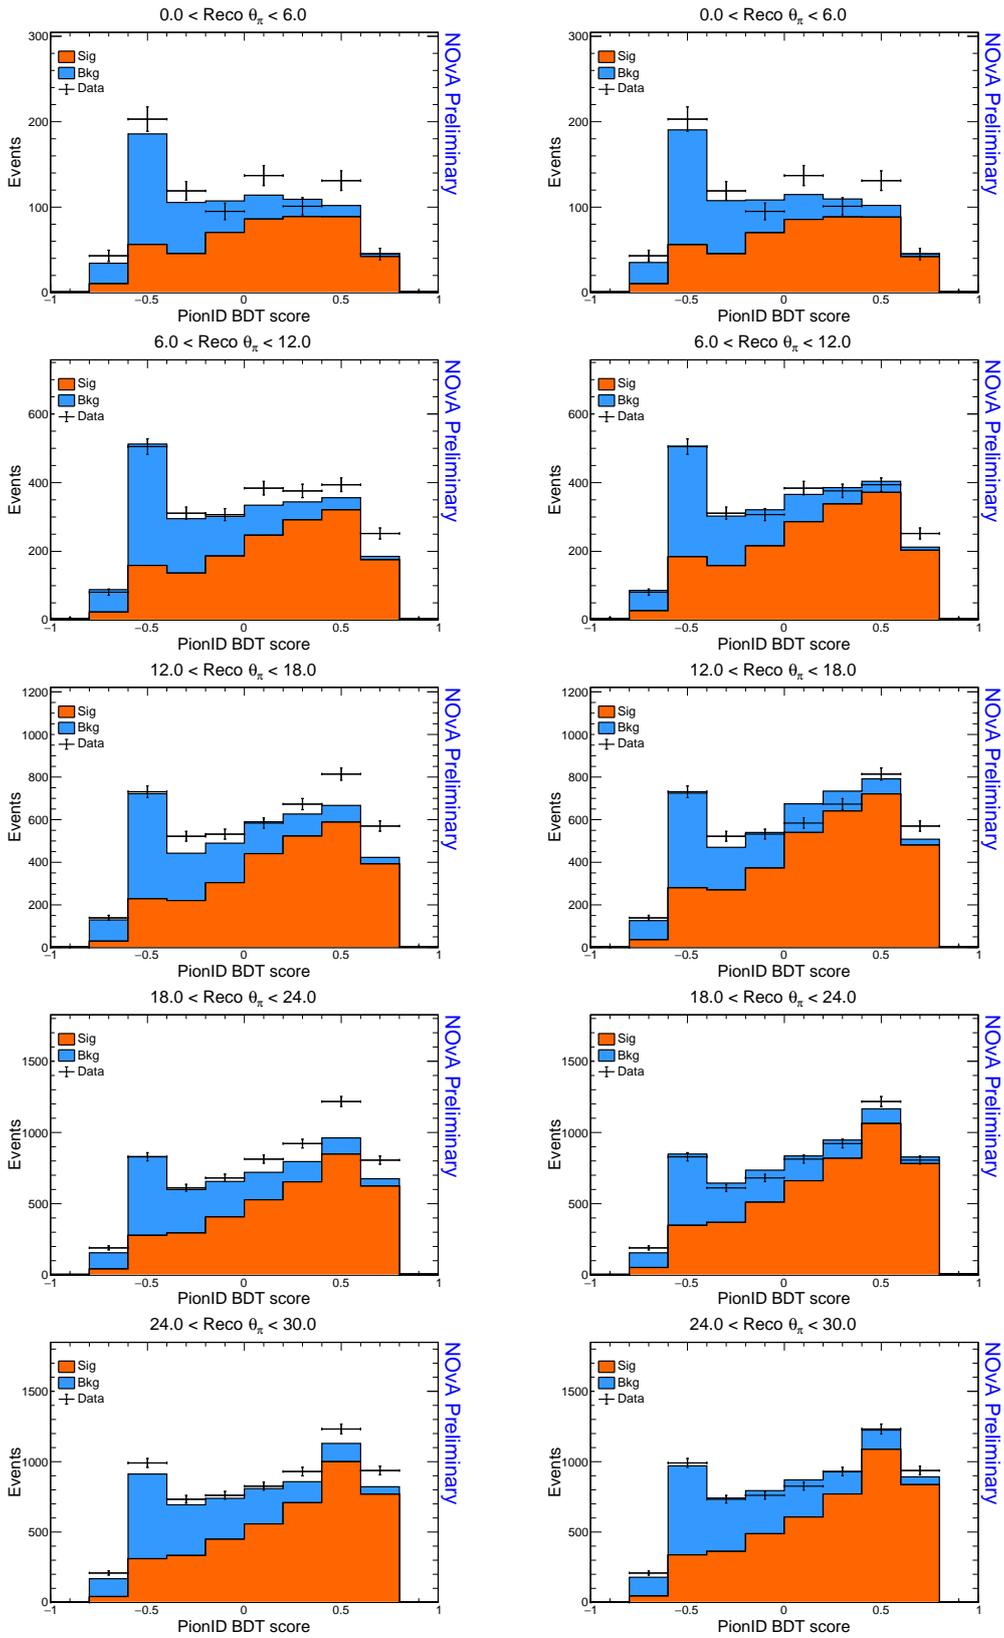


Figure 7.1: Pre-fit (left) and post-fit (right) templates with data overlaid. Error bars represent statistical uncertainty. Shown here are the first five analysis bins.

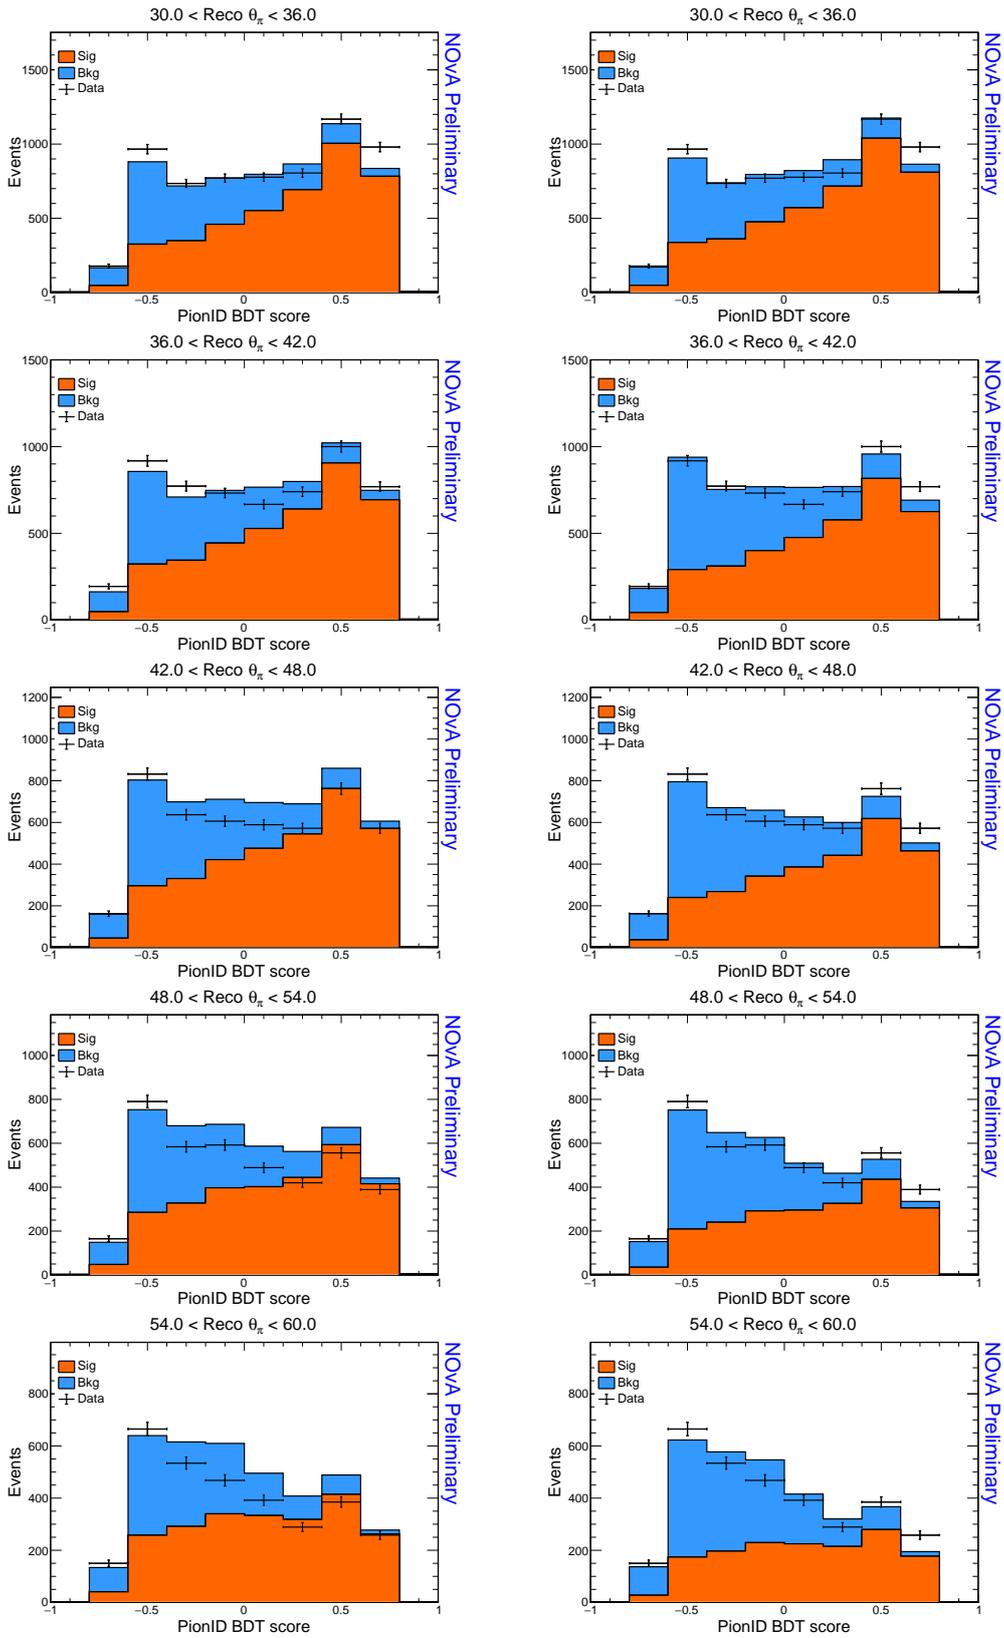


Figure 7.2: Pre-fit (left) and post-fit (right) templates with data overlaid. Error bars represent statistical uncertainty. Shown here are the last five analysis bins.

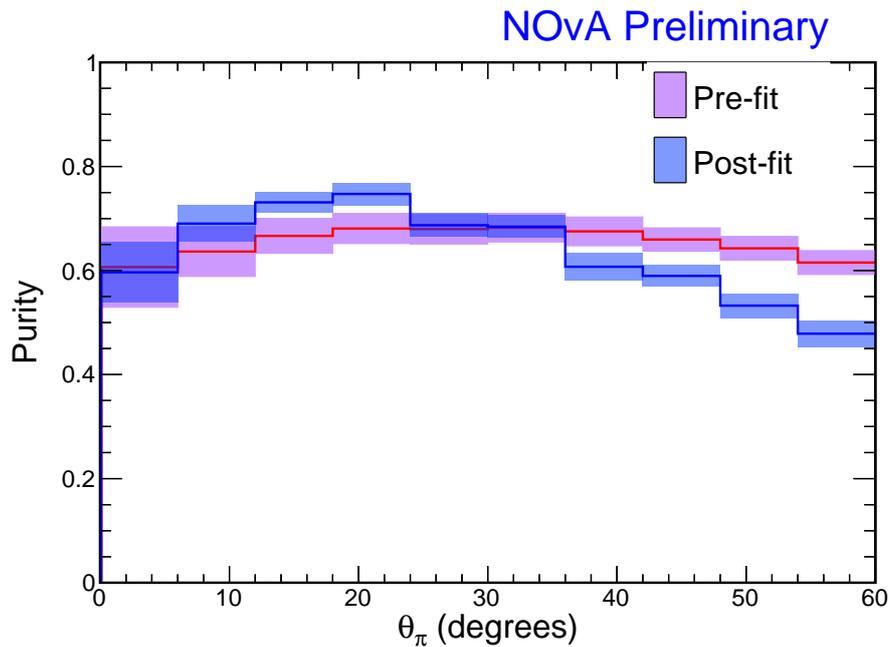


Figure 7.3: Pre and post-fit estimates of purity. Error bars include both statistical and systematic contributions.

7.2 Differential cross-section measurement

Using Equation 6.1, a differential cross-section measurement was made, as shown in Figure 7.4. A χ^2 test statistic comparing the measured and predicted cross-sections was calculated, yielding 51.82 for 10 degrees of freedom. Note that the error bars in Figure 7.4 are all positively correlated, and so estimating the χ^2 by eye is not possible. The total covariance and correlation matrix of the measurement can be seen in Figure 7.5.

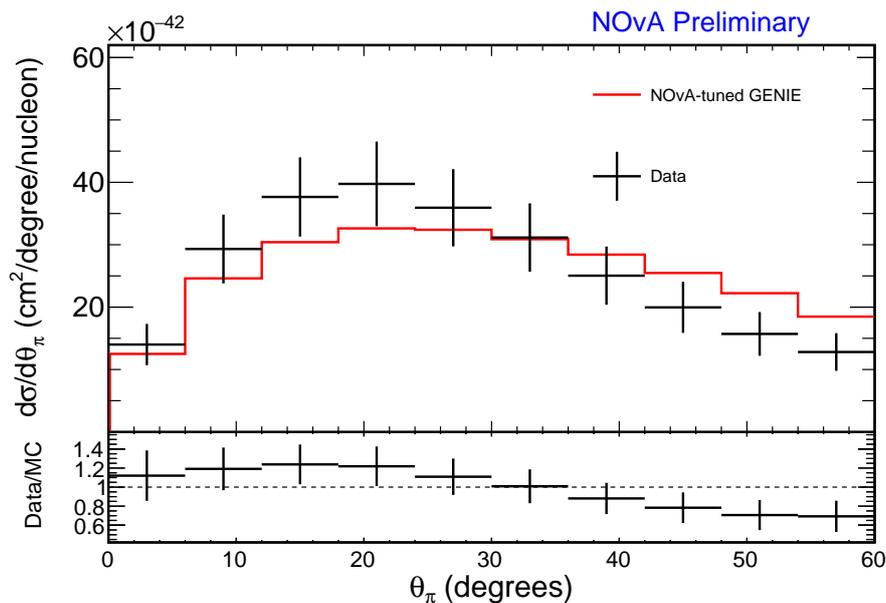


Figure 7.4: Differential cross-section measurement with the NOvA GENIE prediction overlaid. Error bars represent uncertainty from unfolding, statistical, and systematic sources.

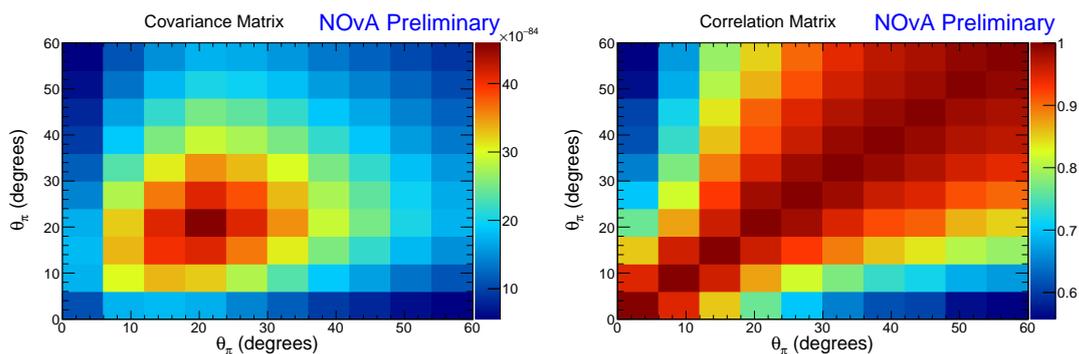


Figure 7.5: Total covariance (left) and correlation (right) matrices for the cross-section measurement.

7.3 Conclusions

The measurement of the differential cross-section $d\sigma/d\theta_\pi$ shown in Figure 7.4 shows that charged pions produced in $\nu_\mu \text{CC}1\pi^\pm$ neutrino-nucleus interactions are more forward-going (lower θ_π) than predicted in GENIE simulation.

The analysis outlined in this thesis uses many cutting-edge analysis techniques, from machine learning to template fitting. Further work might focus on developing selections with smaller systematic uncertainties, to reduce the systematic uncertainty on the measurement. The analysis presented here may serve as the basis for another

measurement of the $\nu_{\mu}CC1\pi^{\pm}$ interaction cross-section, reported for different observables. Possibilities include muon kinematics or the angle between the pion and muon. Such measurements would provide even more detailed information about where the predictions of current models are in tension with observations.

This cross-section measurement will aid the event generator community in creating the next generation of neutrino-nucleus interaction models. This, in turn, will allow the current and next generation of neutrino oscillation experiments to make more precise measurements of the properties of the neutrino by reducing the size of their cross-section modelling uncertainties.

Bibliography

- [1] U. of Zurich, “Overview of the Standard Model,” <https://www.physik.uzh.ch/en/researcharea/lhcb/outreach/StandardModel.html>, accessed: 2023-06-22.
- [2] C. D. Ellis and W. A. Wooster, “The average energy of disintegration of radium E,” *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 117, no. 776, pp. 109–123, 1927.
- [3] J. Link, P. Huber, and A. Haghigat, “The curious history of neutrinos and nuclear reactors,” *Nuclear News*, p. 5867, Dec 2020.
- [4] T. S. Electroweak, H. F. Groups, A. Collaboration, D. Collaboration, L. Collaboration, O. Collaboration, S. Collaboration, L. E. W. Group *et al.*, “Precision electroweak measurements on the Z resonance,” *Physics Reports*, vol. 427, no. 5-6, pp. 257–454, 2006.
- [5] T. Kajita *et al.*, “Atmospheric neutrino results from Super-Kamiokande and Kamiokande: Evidence for ν_μ oscillations,” *Nuclear Physics B-Proceedings Supplements*, vol. 77, no. 1-3, pp. 123–132, 1999.
- [6] “Mass hierarchy figure,” 2017. [Online]. Available: <http://www.staff.uni-mainz.de/wurmm/juno.html>
- [7] J. A. Formaggio and G. P. Zeller, “From eV to EeV: Neutrino cross sections across energy scales,” *Rev. Mod. Phys.*, vol. 84, pp. 1307–1341, Sep 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.84.1307>
- [8] D. Casper, “The nuance neutrino physics simulation, and the future,” *Nuclear Physics B-Proceedings Supplements*, vol. 112, no. 1-3, pp. 161–170, 2002.

- [9] A. Aguilar-Arevalo, C. Anderson, A. Bazarko, S. Brice, B. Brown, L. Bugel, J. Cao, L. Coney, J. Conrad, D. Cox *et al.*, “First measurement of the muon neutrino charged current quasielastic double differential cross section,” *Physical Review D*, vol. 81, no. 9, p. 092005, 2010.
- [10] T. Golan, J. T. Sobczyk, and K. M. Graczyk, “Modeling nuclear effects in NuWro Monte Carlo neutrino event generator,” Ph.D. dissertation, PhD thesis, University of Wroclaw, 2014.
- [11] R. Gran, “Model uncertainties for Valencia RPA effect for MINERvA,” 2017.
- [12] P. Coloma, P. Huber, C.-M. Jen, and C. Mariani, “Neutrino-nucleus interaction models and their impact on oscillation analyses,” *Physical Review D*, vol. 89, no. 7, p. 073015, 2014.
- [13] P. Adamson, K. Anderson, M. Andrews, R. Andrews, I. Anghel, D. Augustine, A. Aurisano, S. Avvakumov, D. Ayres, B. Baller *et al.*, “The NuMI neutrino beam,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 806, pp. 279–306, 2016.
- [14] “FCS still photos,” <https://mod.fnal.gov/mod/stillphotos/2014/0000/14-0007-01D.hr.jpg>, accessed: 2023-05-02.
- [15] A. Sutton, “Domain generalization with machine learning in the NOvA experiment,” Ph.D. dissertation, PhD thesis, University of Virginia, 2016.
- [16] D. Ayres, G. Drake, M. Goodman, J. Grudzinski, V. Guarino, R. Talaga, A. Zhao, P. Stamoulis, E. Stiliaris, G. Tzanakos *et al.*, “The NOvA technical design report,” 2007.
- [17] “FCS still photos,” <https://mod.fnal.gov/mod/stillphotos/2014/0000/14-0091-02D.hr.jpg>, accessed: 2023-06-03.
- [18] G. Davies and J. Vasel, “A technical note on cosmic overlays for NOvA analyses,” NOvA internal note 23378, Tech. Rep., 2017.

- [19] J. Wolcott, “Cross section models and uncertainties in NOvA,” NOvA internal note 58591, Tech. Rep., 2022.
- [20] A. Rodriguez and A. Laio, “Clustering by fast search and find of density peaks,” *science*, vol. 344, no. 6191, pp. 1492–1496, 2014.
- [21] M. Baird, “Global vertex reconstruction beginning with a modified hough transform,” NOvA internal note 8241, Tech. Rep., 2012.
- [22] E. Niner, “Observation of electron neutrino appearance in the NuMI beam with the NOvA experiment,” Ph.D. dissertation, PhD thesis, Indiana University, 2015.
- [23] B. Eberly, L. Aliaga, O. Altinok, M. B. Sazo, L. Bellantoni, M. Betancourt, A. Bodek, A. Bravar, H. Budd, M. Bustamante *et al.*, “Charged pion production in ν_μ interactions on hydrocarbon at $\langle E_\nu \rangle = 4.0$ GeV,” *Physical Review D*, vol. 92, no. 9, p. 092008, 2015.
- [24] J. Chadwick, “Intensitätsverteilung im magnetischen spectrum der β -strahlen von radium B + C,” *Verhandl. Dtsc. Phys. Ges.*, vol. 16, p. 383, 1914.
- [25] W. Pauli, “Pauli letter collection: letter to Lise Meitner,” Tech. Rep., 1930.
- [26] E. Fermi, “Versuch einer theorie der β -strahlen. i,” *Zeitschrift für Physik*, vol. 88, no. 3-4, pp. 161–177, 1934.
- [27] H. Bethe and R. Peierls, “The neutrino,” *Nature*, vol. 133, no. 3362, pp. 532–532, 1934.
- [28] B. Pontecorvo, “Inverse beta process,” *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.*, vol. 1, pp. 25–31, 1991.
- [29] L. W. Alvarez, “A proposed experimental test of the neutrino theory,” 1949.
- [30] C. L. Cowan Jr, F. Reines, F. Harrison, H. Kruse, and A. McGuire, “Detection of the free neutrino: a confirmation,” *Science*, vol. 124, no. 3212, pp. 103–104, 1956.

- [31] G. Danby, J. M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger, "Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos," *Physical Review Letters*, vol. 9, no. 1, p. 36, 1962.
- [32] M. L. Perl, G. Abrams, A. Boyarski, M. Breidenbach, D. Briggs, F. Bulos, W. Chinowsky, J. Dakin, G. Feldman, C. Friedberg *et al.*, "Evidence for anomalous lepton production in e^+e^- annihilation," *Physical Review Letters*, vol. 35, no. 22, p. 1489, 1975.
- [33] K. Kodama, N. Ushida, C. Andreopoulos, N. Saoulidou, G. Tzanakos, P. Yager, B. Baller, D. Boehnlein, W. Freeman, B. Lundberg *et al.*, "Observation of tau neutrino interactions," *Physics Letters B*, vol. 504, no. 3, pp. 218–224, 2001.
- [34] R. Davis Jr, D. S. Harmer, and K. C. Hoffman, "Search for neutrinos from the sun," *Physical Review Letters*, vol. 20, no. 21, p. 1205, 1968.
- [35] B. T. Cleveland, T. Daily, R. Davis Jr, J. R. Distel, K. Lande, C. Lee, P. S. Wildenhain, and J. Ullman, "Measurement of the solar electron neutrino flux with the Homestake chlorine detector," *The Astrophysical Journal*, vol. 496, no. 1, p. 505, 1998.
- [36] V. Gribov and B. Pontecorvo, "Neutrino astronomy and lepton charge," *Physics Letters B*, vol. 28, no. 7, pp. 493–496, 1969.
- [37] A. Abazov, O. Anosov, E. Faizov, V. Gavrin, A. Kalikhov, T. Knodel, I. Knyschenko, V. Kornoukhov, S. Mezentseva, I. Mirmov *et al.*, "Search for neutrinos from the sun using the reaction $\text{Ga}_{71}(\nu_e, e^-)\text{ge}_{71}$," *Physical Review Letters*, vol. 67, no. 24, p. 3332, 1991.
- [38] K. S. Hirata, T. Kajita, T. Kifune, K. Kihara, M. Nakahata, K. Nakamura, S. Ohara, Y. Oyama, N. Sato, M. Takita *et al.*, "Observation of b 8 solar neutrinos in the kamiokande-ii detector," *Physical review letters*, vol. 63, no. 1, p. 16, 1989.
- [39] P. Anselmann, W. Hampel, G. Heusser, J. Kiko, T. Kirsten, E. Pernicka, R. Plaga, U. Rönn, M. Sann, C. Schlosser *et al.*, "Solar neutrinos observed by GALLEX at Gran Sasso," *Physics Letters B*, vol. 285, no. 4, pp. 376–389, 1992.

- [40] M. Sanchez, W. W. Allison, G. Alner, D. Ayres, W. Barrett, M. Goodman, T. Joffe-Minor, and J. Thron, "Observation of atmospheric neutrino oscillations in Soudan 2." Argonne National Lab., IL (US), Tech. Rep., 2003.
- [41] Q. R. Ahmad, R. Allen, T. Andersen, J. Anglin, J. Barton, E. Beier, M. Bercovitch, J. Bigu, S. Biller, R. Black *et al.*, "Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory," *Physical review letters*, vol. 89, no. 1, p. 011301, 2002.
- [42] L. Wolfenstein, "Neutrino oscillations in matter," *Physical Review D*, vol. 17, no. 9, p. 2369, 1978.
- [43] S. Mikheyev and A. Y. Smirnov, "Resonant amplification of ν oscillations in matter and solar-neutrino spectroscopy," *Il Nuovo Cimento C*, vol. 9, pp. 17–26, 1986.
- [44] F. Hasert, S. Kabe, W. Krenz, J. Von Krogh, D. Lanske, J. Morfin, K. Schultze, H. Weerts, G. Bertrand-Coremans, J. Sacton *et al.*, "Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment," *Nuclear Physics B*, vol. 73, no. 1, pp. 1–22, 1974.
- [45] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, "The GENIE neutrino monte carlo generator: physics and user manual," *arXiv preprint arXiv:1510.05494*, 2015.
- [46] C. Llewellyn Smith, "Neutrino reactions at accelerator energies," *Physics Reports*, vol. 3, no. 5, pp. 261–379, 1972. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0370157372900105>
- [47] R. Smith and E. Moniz, "Neutrino reactions on nuclear targets," *Nuclear Physics B*, vol. 43, pp. 605–622, 1972. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0550321372900405>
- [48] A. Gil, J. Nieves, and E. Oset, "Many-body approach to the inclusive (e, e') reaction from the quasielastic to the Δ excitation region," *Nuclear*

- Physics A*, vol. 627, no. 4, pp. 543–598, dec 1997. [Online]. Available: <https://doi.org/10.1016%2Fs0375-9474%2897%2900513-7>
- [49] D. Rein and L. M. Sehgal, “Neutrino-excitation of baryon resonances and single pion production,” *Annals of Physics*, vol. 133, no. 1, pp. 79–153, 1981. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0003491681902426>
- [50] C. Berger and L. Sehgal, “Lepton mass effects in single pion production by neutrinos,” *Physical Review D*, vol. 76, no. 11, p. 113004, 2007.
- [51] W. Fabian and H. Arenhvel, “Pion exchange currents and isobar configurations in deuteron electro-disintegration below pion threshold,” *Nuclear Physics A*, vol. 258, no. 3, pp. 461–479, 1976. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0375947476904851>
- [52] J. Nieves, I. R. Simo, and M. V. Vacas, “Two particle–hole excitations in charged current quasielastic antineutrino-nucleus scattering,” *Physics Letters B*, vol. 721, no. 1-3, pp. 90–93, 2013.
- [53] G. Megias, T. Donnelly, O. Moreno, C. Williamson, J. Caballero, R. González-Jiménez, A. De Pace, M. B. Barbaro, W. M. Alberico, M. Nardi *et al.*, “Meson-exchange currents and quasielastic predictions for charged-current neutrino C_{12} scattering in the superscaling approach,” *Physical Review D*, vol. 91, no. 7, p. 073004, 2015.
- [54] C. Andreopoulos, A. Bell, D. Bhattacharya, F. Cavanna, J. Dobson, S. Dytman, H. Gallagher, P. Guzowski, R. Hatcher, P. Kehayias *et al.*, “The GENIE neutrino monte carlo generator,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 614, no. 1, pp. 87–104, 2010.
- [55] A. Bodek and U. Yang, “Higher twist, ξ w scaling, and effective LO PDFs for lepton scattering in the few GeV region,” *Journal of Physics G: Nuclear and Particle Physics*, vol. 29, no. 8, p. 1899, 2003.

- [56] M. Judah, “Measurement of the inclusive electron-neutrino charged-current cross section in the NOvA Near Detector,” Ph.D. dissertation, PhD thesis, Colorado State University, 2019.
- [57] B. Behera, “Measurement of the double differential inclusive muon neutrino charged-current cross-section in the NOvA Near Detector,” Ph.D. dissertation, PhD thesis, Indian Institute of Technology Hyderabad, 2018.
- [58] O. Benhar, P. Huber, C. Mariani, and D. Meloni, “Neutrino–nucleus interactions and the determination of oscillation parameters,” *Physics Reports*, vol. 700, pp. 1–47, 2017.
- [59] “Fermilab accelerator division records,” <https://www-bd.fnal.gov/FixedTargetPlots/today/ProtonPlots.html>, accessed: 2023-06-21.
- [60] M. Convery, M. Lindgren, S. Nagaitsev, and V. Shiltsev, “Fermilab accelerator complex: Status and improvement plans,” Fermi National Accelerator Lab.(FNAL), Batavia, IL (United States), Tech. Rep., 2018.
- [61] S. Agostinelli *et al.*, “Geant4 a simulation toolkit,” *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003. [Online]. Available: <https://inspirehep.net/files/6c9c0b62bbc8dc0401fca11a5fe5c87c>
- [62] L. Aliaga, M. Kordosky, T. Golan, O. Altinok, L. Bellantoni, A. Bercellie, M. Bétantcourt, A. Bravar, H. Budd, M. Carneiro *et al.*, “Neutrino flux predictions for the NuMI beam,” *Physical Review D*, vol. 94, no. 9, p. 092005, 2016.
- [63] J. Wolcott *et al.*, “NOvARwgt GitHub repository,” 2022. [Online]. Available: <https://github.com/novaexperiment/NOvARwgt-public>
- [64] K. Bays *et al.*, “NOvA 2020 cross-section tune tech note,” NOvA internal note 43962, Tech. Rep., 2020.
- [65] D. Pershey, J. Huang, and M. Judah, “TDSlicer technote,” NOvA internal note 27689, Tech. Rep., 2019.

- [66] R. C. Prim, “Shortest connection networks and some generalizations,” *The Bell System Technical Journal*, vol. 36, no. 6, pp. 1389–1401, 1957.
- [67] M. Gyulassy and M. Harlander, “Elastic tracking and neural network algorithms for complex pattern recognition,” *Computer Physics Communications*, vol. 66, no. 1, pp. 31–46, 1991.
- [68] M. Messier, “Vertex reconstruction based on elastic arms,” NOvA internal note 7530, Tech. Rep., 2012.
- [69] 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.
- [70] J. C. Dunn, “A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters,” 1973.
- [71] N. H. Kuiper, “Tests concerning random points on a circle,” in *Nederl. Akad. Wetensch. Proc. Ser. A*, vol. 63, no. 1, 1960, pp. 38–47.
- [72] 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, no. 09, p. P09001, 2016.
- [73] L. Cremonesi *et al.*, “Summary of single particle cvn training for prod5.1,” NOvA internal note 51176, Tech. Rep., 2021.
- [74] W. S. Cleveland, “Lowess: A program for smoothing scatterplots by robust locally weighted regression,” *The American Statistician*, vol. 35, no. 1, p. 54, 1981.
- [75] J. Calcutt, C. Thorpe, K. Mahn, and L. Fields, “Geant4Reweight: a framework for evaluating and propagating hadronic interaction uncertainties in Geant4,” *Journal of Instrumentation*, vol. 16, no. 08, p. P08042, 2021. [Online]. Available: <https://arxiv.org/pdf/2105.01744.pdf>

- [76] J. Calcutt *et al.*, “GeantReweight GitHub repository,” 2022. [Online]. Available: <https://github.com/NuSoftHEP/Geant4Reweight>
- [77] C. Sweeney, “Calculating weights in Geant4Reweight,” NOvA internal note 51223, Tech. Rep., 2021.
- [78] M. Strait, “Geometry changes tech note,” NOvA internal note 23132, Tech. Rep., 2018.
- [79] B. Allardyce, C. Batty, D. Baugh, E. Friedman, G. Heymann, M. Cage, G. Pyle, G. Squier, A. Clough, D. Jackson *et al.*, “Pion reaction cross sections and nuclear sizes,” *Nuclear Physics A*, vol. 209, no. 1, pp. 1–51, 1973.
- [80] D. Ashery, I. Navon, G. Azuelos, H. Walter, H. Pfeiffer, and F. Schlepütz, “True absorption and scattering of pions on nuclei,” *Physical Review C*, vol. 23, no. 5, p. 2173, 1981.
- [81] D. Ashery, D. Geesaman, R. Holt, H. Jackson, J. Specht, K. Stephenson, R. Segel, P. Zupranski, H. Baer, J. Bowman *et al.*, “Inclusive pion single-charge-exchange reactions,” *Physical Review C*, vol. 30, no. 3, p. 946, 1984.
- [82] W. P. Ball, “Nuclear scattering of 300 MeV neutrons,” 1952.
- [83] E. Bellotti, D. Cavalli, and C. Matteuzzi, “Positive-pion absorption by C nuclei at 130 MeV,” *Il Nuovo Cimento A (1965-1970)*, vol. 18, no. 1, pp. 75–93, 1973.
- [84] F. Binon, P. Duteil, J. Garron, J. Gorres, L. Hugon, J. P. Peigneux, C. Schmit, M. Spighel, J.-P. Stroot *et al.*, “Scattering of negative pions on carbon,” *Nuclear Physics B*, vol. 17, no. 1, pp. 168–188, 1970.
- [85] J. DeJuren and N. Knable, “Nuclear cross sections for 95-MeV neutrons,” *Physical Review*, vol. 77, no. 5, p. 606, 1950.
- [86] E. P. Guerra, S. Bhadra, S. Berkman, C. Cao, P. de Perio, Y. Hayato, K. Ieki, M. Ikeda, Y. Kanazawa, J. Kim *et al.*, “Measurement of σ_{ABS} and σ_{CX} of π^+ on

- carbon by the Dual Use Experiment at TRIUMF (DUET),” *Physical Review C*, vol. 95, no. 4, p. 045203, 2017.
- [87] C. Gelderloos, J. Brack, M. Holcomb, M. Keilman, D. Mercer, R. Peterson, R. Ristinen, and A. Saunders, “Reaction and total cross sections for 400–500 MeV π^- on nuclei,” *Physical Review C*, vol. 62, no. 2, p. 024612, 2000.
- [88] H. Hilscher, W.-D. Krebs, G. Sepp, and V. Soergel, “Charge exchange of 70 MeV π^- on c, al, cu and pb,” *Nuclear Physics A*, vol. 158, no. 2, pp. 602–606, 1970.
- [89] M. Jones, R. Ransome, V. Cupps, R. Fergerson, C. Morris, J. McGill, J. Zumbro, J. Comfort, B. Ritchie, J. Tinsley *et al.*, “Pion absorption above the Δ (1232) resonance,” *Physical Review C*, vol. 48, no. 6, p. 2800, 1993.
- [90] S. Levenson, D. Geesaman, E. Colton, R. Holt, H. Jackson, J. Schiffer, J. Specht, K. Stephenson, B. Zeidman, R. Segel *et al.*, “Inclusive pion scattering in the Δ (1232) region,” *Physical Review C*, vol. 28, no. 1, p. 326, 1983.
- [91] O. Meirav, E. Friedman, A. Altman, M. Hanna, R. Johnson, and D. Gill, “Total reaction cross sections of 50 and 65 MeV pions on nuclei,” *Physical Review C*, vol. 36, no. 3, p. 1066, 1987.
- [92] E. Miller, “Inelastic scattering of 150 MeV negative pions by carbon and lead,” *Il Nuovo Cimento (1955-1965)*, vol. 6, no. 4, pp. 882–911, 1957.
- [93] I. Navon, D. Ashery, J. Alster, G. Azuelos, B. Barnett, W. Gyles, R. Johnson, D. Gill, and T. Masterson, “True absorption and scattering of 50 MeV pions,” *Physical Review C*, vol. 28, no. 6, p. 2548, 1983.
- [94] A. Saunders, S. Høibråten, J. Kraushaar, B. Kriss, R. Peterson, R. Ristinen, J. Brack, G. Hofman, E. Gibson, and C. Morris, “Reaction and total cross sections for low energy π^+ and π^- on isospin zero nuclei,” *Physical Review C*, vol. 53, no. 4, p. 1745, 1996.

- [95] R. Voss and R. Wilson, “Neutron inelastic cross-sections between 55 and 140 MeV,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 236, no. 1204, pp. 41–51, 1956.
- [96] A. Aguilar-Arevalo, C. Anderson, A. Bazarko, S. Brice, B. Brown, L. Bugel, J. Cao, L. Coney, J. Conrad, D. Cox *et al.*, “Measurement of neutrino-induced charged-current charged pion production cross sections on mineral oil at $E_\nu \sim 1$ GeV,” *Physical Review D*, vol. 83, no. 5, p. 052007, 2011.
- [97] K. Abe, C. Andreopoulos, M. Antonova, S. Aoki, A. Ariga, S. Assylbekov, D. Autiero, S. Ban, M. Barbi, G. Barker *et al.*, “First measurement of the muon neutrino charged current single pion production cross section on water with the T2K near detector,” *Physical Review D*, vol. 95, no. 1, p. 012010, 2017.
- [98] K. Abe, R. Akutsu, A. Ali, C. Alt, J. Amey, C. Andreopoulos, L. Anthony, M. Antonova, S. Aoki, A. Ariga *et al.*, “Measurement of the muon neutrino charged-current single π^+ production on hydrocarbon using the T2K off-axis near detector ND280,” *Physical Review D*, vol. 101, no. 1, p. 012007, 2020.
- [99] A. Bercellie, K. Kroma-Wiley, S. Akhter, Z. A. Dar, F. Akbar, V. Ansari, M. Ascencio, M. S. Athar, L. Bellantoni, M. Betancourt *et al.*, “Simultaneous measurement of muon neutrino ν_μ charged-current single π^+ production in CH, C, H₂O, Fe, and Pb targets in MINERvA,” *arXiv preprint arXiv:2209.07852*, 2022.
- [100] M. Acero, P. Adamson, L. Aliaga, N. Anfimov, A. Antoshkin, E. Arrieta-Diaz, L. Asquith, A. Aurisano, A. Back, M. Baird *et al.*, “Measurement of the double-differential muon-neutrino charged-current inclusive cross section in the NOvA near detector,” *Physical Review D*, vol. 107, no. 5, p. 052011, 2023.
- [101] E. Catano-Mur, “Nova production 5.1 detector systematics executive summary,” NOvA internal note 53225, Tech. Rep., 2023.
- [102] O. Samoylov, “Nova production 5 detector simulation1 changes and systematics,” NOvA internal note 43935, Tech. Rep., 2020.

-
- [103] G. D'Agostini, "A multidimensional unfolding method based on Bayes' theorem," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 362, no. 2-3, pp. 487–498, 1995.
- [104] S.-K. Lin, "Iterative unfolding optimization with the mean squared error metric," NOvA internal note 40923, Tech. Rep., 2019.