Probing Nuclear Effects in Neutrino-Nucleus Scattering at the T2K Off-Axis Near Detector Using Transverse Kinematic Imbalances

Stephen Dolan

St Peter's College

University of Oxford



A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

Trinity Term, 2017

Probing Nuclear Effects in Neutrino-Nucleus Scattering at the T2K Off-Axis Near Detector Using Transverse Kinematic Imbalances

Stephen Dolan

Abstract

The precise characterisation of neutrino oscillations offers a range of auspicious opportunities in the search for physics beyond the standard model. In particular T2K and NOνA, as the current generation of accelerator-based long-baseline neutrino oscillation experiments, are just beginning to provide the first measurements of the complex phase that gives rise to CP-violation in the neutrino sector within the PMNS description of neutrino oscillations. However, the limiting systematic uncertainty in making such precision measurements of neutrino oscillations - and soon, for some measurements, the largest of all uncertainties - stems from a naivety of relevant ($\mathcal{O}(GeV)$) neutrino-nucleus interactions. These interactions are particularly difficult to understand since the bare neutrino-nucleon scatter is often obfuscated by a variety of poorly-understood 'nuclear effects', such as nucleon-nucleon correlations and hadronic re-interactions, which distort the outgoing particle kinematics. A better understanding of these nuclear effects is therefore essential to constrain the principle systematics associated with neutrino oscillation measurements.

This thesis details how the T2K experiment's off-axis near detector is used to provide a novel characterisation of nuclear effects. This is achieved through the first measurement of differential neutrino-nucleus scattering cross sections as a function of kinematic imbalances in the plane transverse to an incoming neutrino beam. The use of transverse kinematic imbalances are motivated by a series of recent phenomenological studies, which this thesis further develops, that illustrate how these observables can provide a new examination of nuclear effects at T2K. In order to extract the cross sections with minimal input model bias, a new cross-section extraction procedure, which minimises the use of the input model and facilitates a data-driven regularisation, is developed, validated and applied. Extensive comparison of the results to the predictions of a variety of neutrino-nucleus interaction simulations then enables a unique probe, and partial isolation, of some nuclear effects. This probe is expected to play a significant role in future neutrino-nucleus interaction model development, whilst also facilitating new constraints on the principle systematic uncertainties associated with neutrino oscillation measurements.

Statement of Originality

This thesis and the work presented in it are my own and were produced by me as a result of my own original research. Results and figures from published works by others have been clearly attributed, and this work has not been submitted for another qualification at this or any other university.

Despite this thesis being my own work, the result of any analysis performed as part of the T2K collaboration (whose papers have over 300 authors) should rightly be seen as an achievement of the entire collaboration and of other scientists whose work has been instrumental to even facilitate the possibility of making such analyses. This statement will detail which portions of this thesis are entirely my own original work and which are a review of the work of others.

Chapters 1 and 2 represent an overview of relevant theoretical and experimental works to provide appropriate context for the main analysis performed within this. These chapters largely represent a summary of work by others and all relevant documents have been cited. However, substantial elements of the interpretations of how neutrino interactions affect oscillation analysis (section 2.4) and the review of model dependence in existing analyses (in section 2.5) are my own.

Since they are not yet commonplace in neutrino physics, chapter 3 contains an overview of the observables that will be used in this work, including a review of existing works that have previously used them, which are duly cited. The exact characterisation of the observables is taken from a (cited) publication which was a joint work largely carried out Xianguo Lu, Luke Pickering and myself. This chapter also contains phenomenological studies to show the observables fitness for purpose. These are my own work but are heavily motivated by the studies in the aforementioned publication.

Chapter 4 contains an overview of the T2K experiment, whose data is analysed in this work. This is predominantly an overview of works performed by others, which are cited accordingly, whilst chapter 5 is a very brief overview of the strategy for the analysis in the later parts of this thesis and as such contains only my own work.

Chapter 6 details the event selection used in this work and the propagation of systematic uncertainties associated with it. Although I implemented the selection, almost all of the components have been adapted from the works of others, which have all been cited. In particular the selection used was based on the work of Laura Monfregola and adapted for

this event selection by Alexander Izmaylov, Pierre Bartet, Jiae Kim and myself. Although I produced the covariance matrices to characterise the associated systematic uncertainties, this was largely an error propagation procedure (i.e. I did not directly implement most of the systematics, I only propagated the uncertainties). The slight exception to this is the neutrino parent decay point systematic uncertainty, which David Coplowe, Alexander Izmaylov and I developed.

The first portion of chapter 7 (7.1) contains the cross-section extraction analysis methods used in this analysis. Although the likelihood fitter used is based on the work of others (predominantly the work of Sara Bolognesi), the implementation of it in this thesis has undergone substantial development from its original form. In particular the regularisation and the uncertainty propagation procedure are all my own work. At its core, the fitter uses the MINUIT numerical minimiser, which is suitably cited. The application of the cross-section extraction methods to the primary analysis within this thesis (as detailed in section 7.2) is my own work, but the software used to calculate event weights from model systematic variations is not, but this is cited.

The fake data studies presented in chapter 8; the results and interpretations presented in chapter 9; the evaluation and summary given in chapter 10; and the contents of the appendices are all my own work.

Acknowledgements

Any doctorate awarded for work with the T2K collaboration makes use of the combined work of many excellent scientists: some to build and calibrate the detectors, others to offer invaluable advice and many more to do the various bits of work that make analyses possible. For me perhaps the most influential are those that kept me motivated when I was lost in the depths of impossible `ROOT` errors or illogical fitter results. This thesis is at least partially a product of our ridiculously poorly-organised road/ski trips, our time spent at various pubs and izakayas and, of course, our visits to Shinchan's and karaoke. As such, I would like to thank all of the students, postdocs and academics who have made the last 4 years so enjoyable!

In particular I would like to thank all of my fellow students at Oxford - Raj, Kirsty, David, Will, Pete, Aiden, Tom, Tom, Charles, Abe and Zac for their friendship and collaboration over the past four years. I would also like to thank Luke, without whom I would likely still be writing macros in notepad, Dave for keeping me from becoming delusional, and Trevor for limiting the amount of offal in my meals. A special thanks is certainly due for Raj, for answering the constant barrage of statistics questions with seemingly unlimited patience, whilst also often being the instigator of many of the most memorable moments of my DPhil.

I certainly would not have been able to complete this work without the advice and support of my supervisor, Alfons, who has allowed me the freedom to work on whatever I wanted, whilst always being there for support. I would also like to thank Xianguo, Dave and Giles for their collaboration and excellent analysis advice.

The analysis in this thesis evolved through countless meetings with the T2K cross section group, who I would like to thank for all the wonderful feedback and encouragement. I would especially like to thank Kendall, Sara, Callum and Tom for their invaluable guidance. I count myself lucky to have ended up working as part of such a supportive team.

I would also like to thank all my family and friends. I owe so much to my parents, brother and sister for their constant support throughout my DPhil. I also thank Mike, Harley, Tim and Fran for their Mojito and White Russian mixing skills.

Finally, I would like to specially thank Emily for her continual love and support in the face of a stressed and unreasonable boyfriend. There is no way I could have managed this without you.

Contents

Introduction **1**

1 Neutrino physics **5**
1.1 A brief history of neutrino physics 5
1.2 Neutrino oscillation theory 7
1.2.1 Matter effects in neutrino oscillations 9
1.3 Modern long baseline neutrino oscillation experiments 10
1.4 Future sensitivity to CP-violation 12

2 Neutrino-nucleus interactions **13**
2.1 The charged-current quasi-elastic interaction 14
2.2 Nuclear effects in neutrino-nucleus scattering 16
2.2.1 Fermi motion 16
2.2.2 Final state interactions 18
2.2.3 Nucleon-nucleon correlations and the random phase approximation . 19
2.3 Other interaction modes 20
2.3.1 Resonance production 21
2.3.2 Coherent scattering 22
2.3.3 Deep inelastic scattering 23
2.3.4 Neutral-current elastic scattering 23
2.4 Neutrino interactions in oscillation analysis 23
2.4.1 Lepton kinematics approach - T2K 24
2.4.2 Calorimetric approach - NOνA 26
2.4.3 Neutrino interactions in oscillation analysis summary 28
2.5 Making ν_μ CCQE-like cross-section measurements 29
2.5.1 Flux-integrated and flux-unfolded differential cross sections 29
2.5.2 Challenges in characterising CCQE interactions 30
2.5.3 Previous CCQE and CCQE-like cross-section measurements 31
2.6 Summary 34

3 Single-transverse kinematic imbalance (STKI) 37
3.1 The single transverse variables (STV) . . . 38
3.2 Investigating the impact of nuclear effects on STKI . . . 39
3.2.1 Fermi motion . . . 40
3.2.2 Final state interactions . . . 40
3.2.3 2p-2h nucleon-nucleon correlations . . . 42
3.2.4 The nucleon axial mass . . . 43
3.2.5 The random phase approximation . . . 44
3.2.6 Prospects of probing nuclear effects in CC0π+Np interactions at ND280 using STKI . . . 45

4 The T2K experiment 47
4.1 The T2K neutrino beam . . . 48
4.1.1 J-PARC proton accelerator and neutrino beamline . . . 48
4.1.2 Off-axis approach . . . 49
4.1.3 Neutrino flux predictions . . . 50
4.2 ND280: The T2K near detector . . . 52
4.2.1 Time projection chambers . . . 54
4.2.2 First fine-grained detector . . . 57
4.2.3 Data simulation . . . 60
4.3 T2K data availability . . . 62

5 Analysis strategy 65

6 CC0π+Np selection 67
6.1 Selection . . . 68
6.1.1 Summary of cuts . . . 69
6.1.2 ECal π^0 veto cut . . . 72
6.2 Summary of selected events . . . 72
6.2.1 Backgrounds . . . 76
6.3 Efficiency and phase-space constraints . . . 76
6.3.1 Choosing the phase-space restrictions . . . 77
6.3.2 Total efficiency and purity . . . 79
6.4 Reconstructing the single-transverse variables . . . 79
6.4.1 Expected STV distributions . . . 85
6.5 Sideband samples . . . 91
6.5.1 Summary of sideband cuts . . . 91
6.5.2 Summary of selected sideband events . . . 92
6.5.3 Formation of the STV in the sideband samples . . . 92
6.6 Systematic uncertainties in the selection . . . 96

6.6.1 Sources of detector systematic uncertainties . . . 96
6.6.2 Total detector systematic uncertainty . . . 99

7 Binned likelihood fitting as a cross-section extraction technique 103
7.1 General overview of binned likelihood fitting . . . 103
7.1.1 Template fitting as an unsmearing method . . . 105
7.1.2 Ill-posed problems and the role of regularisation . . . 108
7.1.3 From a fit result to a flux-integrated cross section . . . 111
7.2 Implementation of a likelihood fitter . . . 112
7.2.1 Analysis binning . . . 112
7.2.2 Regularisation implementation . . . 114
7.2.3 Systematic parameters in the Fit . . . 115
7.2.4 Fitting mechanism and interpretation of errors . . . 119
7.2.5 Additional systematics for cross-section extraction . . . 121
7.2.6 Input simulation for fitting and associated model dependence . . . 126

8 Cross-section extraction validations and sensitivity 129
8.1 Systematic uncertainties in an Asimov fit . . . 129
8.2 Validating the analysis framework . . . 133
8.2.1 NEUT fake data . . . 134
8.2.2 GENIE fake data . . . 145
8.2.3 NuWro fake data . . . 154
8.2.4 NEUT 6D fake data . . . 163
8.2.5 NEUT 6D without 2p-2h fake data . . . 172
8.2.6 Fake data studies summary . . . 180

9 Cross-section results and interpretations 183
9.1 Real data results . . . 183
9.2 Simulation comparisons and further interpretations . . . 193

10 Evaluation 207

A Event generator implementations 211

B Unregularised results 213

C Outlook 215

Glossary 219

Bibliography 225

List of Figures

1.1 Feynman diagrams to illustrate matter effects in neutrino interactions. . . . 10
1.2 Published CP violation sensitivity curves from the LBNE (left) and HK (right) experiments. . . . 12

2.1 NUANCE predicted neutrino scattering cross sections overlaid with available data taken up to 2012. Figure taken from [39]. . . . 14
2.2 A Feynman diagram of a CCQE interaction. . . . 16
2.3 A comparison of widely used spectral functions. Figure adapted from [63]. . 18
2.4 A visualisation of FSI in a cascade model. Figure adapted from [54]. . . . 19
2.5 The RPA correction factor ($\sigma_{CCQE}/\sigma_{RPA}$) for CCQE neutrino interactions on carbon. . . . 20
2.6 The CCQE double differential cross section in Muon angle and energy according to several models plotted alongside data from the MiniBooNE experiment [71]. Figure taken from [72]. . . . 21
2.7 A Feynman diagram of a resonance production interaction. . . . 22
2.8 The neutrino energy smearing for different interaction modes seen at the T2K far detector. . . . 25
2.9 The change in preference for values of δ_{cp} for fake data with the same true value, but different 2p-2h models. Figure adapted from [80]. . . . 27
2.10 The reconstructed hadronic energy (E_{hadd}) for selected ν_μCC interactions in NOνA's near detector compared to a GENIE simulation. Figure taken from [82]. 28
2.11 Measurements of the ν_μ and $\overline{\nu}_\mu$ CCQE scattering cross sections per nucleon as a function of neutrino energy compared to predictions from the NUANCE simulation. Figure taken from [87]. . . . 32
2.12 The double differential cross section measured in available energy (E_{avail}) across six regions of q_3 measured by MINERνA. Figure taken from [107]. . . 35

3.1 An illustration of a CC0π+1p interaction with nuclear effects to provide a schematic definition of the three single transverse variables. Figure taken from [63]. . . . 38
3.2 The effect of different spectral functions to describe Fermi motion on the STV, formed from CC0π+Np events that are accessible at ND280. . . . 41

3.3 The effect of FSI on the STV, formed from CC0π+Np events that are accessible at ND280. 41
3.4 The effect of FSI on the STV, formed from full phase-space CCQE events. . 42
3.5 The relative STV shapes of CCQE and 2p-2h contributions of CC0π+Np events that are accessible at ND280. 43
3.6 The effect of the nucleon axial mass on the STV, formed from CC0π+Np events that are accessible at ND280. 44
3.7 The simulated Q^2 distribution of CC0π+Np events that are accessible at ND280. 45
3.8 The effect of RPA on the STV, formed from CC0π+Np events that are accessible at ND280. 45

4.1 The T2K experiment. 48
4.2 A schematic of the T2K neutrino beam production. 48
4.3 The neutrino energy from two-body pion decay as a function of the parent pion momentum for different off-axis angles. Taken from [115]. 50
4.4 The neutrino flux from the T2K beam for different off-axis angles, shown along side the muon neutrino oscillation probability. Taken from [111]. . . . 51
4.5 The (unoscillated) off-axis neutrino flux prediction at ND280 in FHC and RHC mode, broken down by neutrino type. Taken from [122]. 52
4.6 The fractional error on the muon neutrino flux in FHC mode. Adapted from [122]. 53
4.7 The T2K near detector complex. 53
4.8 An exploded view of the ND280 off-axis near detector labelling each sub-detector. Adapted from [26]. 55
4.9 An event display showing what appears to be a DIS event starting in FGD1 and where the ensuing fragments are tracked through ND280's tracker region. Adapted from [123]. 55
4.10 A simplified exploded view of one of the ND280 TPCs to illustrate its key components. Taken from [26]. 56
4.11 The reconstructed energy loss (dE/dx) as a function of reconstructed momentum in the TPCs. Taken from [128]. 58
4.12 A demonstration of how momentum by range can be calculated for an FGD-only track. Taken from [128]. 60
4.13 Deposited energy v.s. range for particles stopping in FGD1. Taken from [123]. 61
4.14 The variation of off-axis angle from the neutrino beam centre across ND280. Adapted from [132]. 63
4.15 Number of protons on target (POT) accumulated by T2K over time, broken down by whether T2K was running in neutrino (FHC) or antineutrino (RHC) mode, shown alongside the J-PARC proton beam power. 63

6.1 A simple breakdown of the different signal samples used in this analysis. . . 69
6.2 An overview of selection cuts for each sample. 69
6.3 The true muon momentum of events passing the cuts for each signal sample, stacked by true event topology. 74
6.4 The true muon angle with respect to the incoming neutrino direction ($cos(\theta_\mu)$) of events passing the cuts for each signal sample, stacked by true event topology. 74
6.5 The true proton momentum of events passing the cuts for each signal sample, stacked by true event topology. 75
6.6 The true proton angle with respect to the incoming neutrino direction ($cos(\theta_p)$) of events passing the cuts for each signal sample, stacked by true event topology. 75
6.7 The efficiency under a CC0$\pi + Np$ signal definition and the true distribution of proton and muon momentum and angle. 80
6.7 Cont. 81
6.8 Selection efficiencies as function of true muon and proton kinematics. 82
6.9 Selection efficiencies after proton phase-space constraints are applied, as a function of true muon kinematics. 82
6.10 Selection efficiencies after proton and muon phase-space constraints are applied, as function of true muon and proton kinematics. 83
6.11 Selection efficiency and purity after p phase-space constraints (excluding out of phase-space events) as a function of cut in each sample. 84
6.12 The schematic of the procedure to reconstruct the incoming neutrino direction from the neutrino PDP and the FGD1 vertex position. 84
6.13 The STV formed from the true kinematics of selected events broken down by true signal topology. 86
6.14 The STV formed from the true kinematics of selected events broken down by true signal interaction mode 87
6.15 The STV formed from the reconstructed kinematics of selected events broken down by true signal topology. 87
6.16 The STV formed from the reconstructed kinematics of selected events broken down by true signal interaction mode. 88
6.17 The row normalised smearing matrices for the STV. These indicate ND280's ability to accurately reconstruct the STV. 88
6.18 The selection efficiency as a function of STV bin is shown for both GENIE and NEUT. The binning is that which will be used for cross section extraction. 89
6.19 The signal portion of selected events broken down by interaction mode. The binning is that which will be used for cross section extraction. 90
6.20 A simple breakdown of the different sideband samples used in this analysis. The number of events expected in each sample within available ND280 data is indicated. 93

6.21 An overview of the selection cuts for each of the two sideband samples. . . . 93
6.22 The true proton and muon kinematics of events passing the cuts for the CC1π^+ sideband sample, stacked by true event topology. 94
6.23 The true proton and muon kinematics of events passing the cuts for the CCOther sideband samples, stacked by true event topology. 94
6.24 The STV formed in the CC1π^+ sideband sample from the true kinematics of selected events, stacked by true event topology. 95
6.25 The STV formed in the CCOther sideband sample from the true kinematics of selected events, stacked by true event topology. 95
6.26 A schematic of the detectors relevant to this analysis alongside a list of the corresponding systematics. 97
6.27 The relative standard deviation of the variation in the selected event rate in reconstructed STV bins after many variations of the neutrino mean parent decay point. 99
6.28 The relative standard deviation of the variation in the selected event rate in reconstructed STV bins after 500 variations of all detector systematics. . . . 101

7.1 An illustration of the reconstructed templates of true analysis bins and of the potential ill-posedness of unsmearing problems. 107
7.2 A toy data study to demonstrate the effect of changing the strength of regularisation on the fit result. 110
7.3 A demonstration of the L-curve method to extract an optimal regularisation strength 112
7.4 The covariance matrix used to construct the penalty term for detector systematic parameters, for each single transverse variable. 115
7.5 The covariance matrix used to construct the penalty term for model parameters. 118
7.6 The ND280 flux prediction and binning for systematic variation alongside the covariance matrix used to construct the penalty term for flux parameters. . . 119
7.7 The error in each STV truth bin due to variation of the efficiency with the cross-section model parameters. 123
7.8 The effect of the nucleon FSI mean free path dial on δp_T compared to the effect of turning nucleon FSI on/off in a NuWro simulation. 124
7.9 The error in each STV truth bin due to variation of the efficiency with the nucleon FSI MFP model parameter. 124
7.10 Covariances to describe the systematic uncertainty on the efficiency, flux normalisation and number of targets, formulated in each analysis bin of the STV. 125

8.1 The differential cross-section in the three STV obtained using an Asimov fit where the input simulation is NEUT 6B. 130

8.2 The contribution of each class of uncertainty on the CC0π+Np template weights in true STV analysis bins following an Asimov fit representative of real data POT. 131
8.3 The contribution of each class of uncertainty on the CC0π+Np template weights in true STV analysis bins following an Asimov fit representative of real data POT in which the sideband regions are excluded. 131
8.4 The contribution of each class of uncertainty on the CC0π+Np template weights in true STV analysis bins following an Asimov fit using the full simulation statistics ($\sim 5\times$ data POT). 132
8.5 The truth (or effective truth) distributions for each of the fake data used. . . 134
8.6 The L-curve used to determine the optimal regularisation parameters for each fit to NEUT fake data. 136
8.7 The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NEUT fake data. 137
8.8 A comparison of the pre and post-fit simulation in reconstructed bins from a fit in each of the STV (separately) to NEUT fake data. 138
8.9 The differential cross section in the three STV and the associated correlation matrices obtained from separate fits to NEUT fake data. 139
8.10 The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to NEUT fake data, to all test simulations. . 140
8.11 The differential cross section in the three STV obtained from a fit to NEUT fake data with various regularisation strengths, alongside the χ^2 of the results relative to each fake data set. 143
8.12 The differential cross section in the three STV obtained from a fit to NEUT fake data using, regularisation strengths deemed reasonable by the L-curve method, with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set. 144
8.13 The L-curve used to determine the optimal regularisation parameters for each fit to GENIE 2.8.0 fake data. 145
8.14 The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to GENIE 2.8.0 fake data. 146
8.15 A comparison of the pre and post-fit simulation in reconstructed bins from a fit in each of the STV (separately) bins to GENIE 2.8.0 fake data. 147
8.16 The differential cross section in the three STV and the associated correlation matrix obtained from separate fits to GENIE 2.8.0 fake data. 148
8.17 The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to GENIE 2.8.0 fake data, to all test simulations. 149

8.18 The differential cross section in the three STV obtained from a fit to GENIE 2.8.0 fake data for various regularisation strengths, alongside the χ^2 of the results relative to each fake data set. 152
8.19 The differential cross section in the three STV obtained from a fit to GENIE 2.8.0 fake data using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set. 153
8.20 The L-curve used to determine the optimal regularisation parameters for each fit to NuWro 11q fake data. 155
8.21 The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NuWro 11q fake data. 156
8.22 A comparison of the pre and post-fit input simulation in reconstructed bins from a fit in each of the STV (separately) bins to NuWro 11q fake data. . . . 157
8.23 The differential cross section in the STV and the associated correlation matrices obtained from fits to NuWro 11q fake data. 158
8.24 The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to NuWro 11q fake data, to all test simulations. 159
8.25 The differential cross section in the three STV obtained from a fit to NuWro 11q fake data for various regularisation strengths, alongside the χ^2 of the results relative to each fake data set. 161
8.26 The differential cross section in the three STV obtained from a fit to NuWro 11q fake data using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set. 162
8.27 The L-curve used to determine the optimal regularisation parameters for each fit to NEUT 6D (bug-fixed) fake data. 164
8.28 The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NEUT 6D (bug-fixed) fake data. 165
8.29 A comparison of the pre and post-fit input simulation in reconstructed bins from a fit in each of the STV (separately) bins to NEUT 6D (bug-fixed) fake data. 166
8.30 The differential cross section in the STV and the associated correlation matrix obtained from a fit to NEUT 6D (bug-fixed) fake data. 167
8.31 The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to NEUT 6D (bug-fixed) fake data, to all test simulations. 168
8.32 The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data for various regularisation strengths, alongside the χ^2 of the results relative to each fake data set. 170

8.33 The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set. 171
8.34 The L-curve used to determine the optimal regularisation parameters for each fit to NEUT (bug-fixed) fake data with the 2p-2h contribution removed. . . . 172
8.35 The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NEUT 6D (bug-fixed) fake data with the 2p-2h contribution removed. 173
8.36 A comparison of the pre and post-fit input simulation in reconstructed bins from a fit in each of the STV (separately) to NEUT (bug-fixed) fake data with the 2p-2h contribution removed. 174
8.37 The differential cross section in the STV and the associated correlation matrix obtained from a fits to NEUT 6D fake data with the 2p-2h contribution removed. 175
8.38 The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to NEUT 6D fake data without a 2p-2h contribution, to all test simulations. 176
8.39 The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data with its 2p-2h contribution removed with various regularisation strengths, alongside the χ^2 of the results relative to each fake data set. 178
8.40 The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data with its 2p-2h contribution removed and using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set. 179
8.41 The χ^2 goodness of fit of the extracted differential cross-section from a fit in each of the STV (separately) to each fake data set (which one is indicated by the legend) relative to every fake data set. 181

9.1 The L-curves obtained from many fits to the real data. 184
9.2 The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to real data. 185
9.3 A comparison of the pre and post-fit simulation in reconstructed bins from a fit in each of the STV (separately) to real data. 186
9.4 The differential cross section in the three STV and the associated correlation matrix obtained from (separate) fits to real data. The extracted cross-section is the primary result of this thesis. 187

9.5 The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to real data, to all test simulations. 188
9.6 The differential cross section in the three STV obtained from a fit to real data with various regularisation strengths, alongside the χ^2 of the results relative to each fake data set. 191
9.7 The differential cross section in the three STV obtained from a fit to real data with the regularisation strengths deemed plausible by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set. 192
9.8 A comparison of analysis results to predictions from NEUT 5.3.2.2 using $M_A^{QE} = 1.03$ and an SF nuclear model, split by interaction mode. 194
9.9 A comparison of analysis results to predictions from NEUT 5.3.2.2 using $M_A^{QE} = 1.03$ and an RFG+RPA nuclear model, split by interaction mode. This is similar to the starting point for T2K oscillation analyses. 194
9.10 A comparison of analysis results to predictions from GENIE 2.12.4 using $M_A^{QE} = 0.99$ and an RFG nuclear model, split by interaction mode (note that the default GENIE implementation contains no 2p-2h). 195
9.11 A comparison of analysis results to predictions from GENIE 2.12.4 using $M_A^{QE} = 0.99$ and an RFG nuclear model with the GENIE empirical 2p-2h model enabled, split by interaction mode [41]. This is similar to the starting point for NOνA oscillation analyses. 195
9.12 A comparison of analysis results to predictions from NuWro11 using $M_A^{QE} = 1.00$ and an LFG nuclear model, split by interaction mode. 196
9.13 A comparison of analysis results to predictions from NuWro11 using $M_A^{QE} = 1.00$ and an SF nuclear model, split by interaction mode. 196
9.14 A comparison of analysis results to predictions from GiBUU 2016 using $M_A^{QE} = 1.00$ and an LFG nuclear model, split by interaction mode. 197
9.15 A comparison of analysis results to predictions from a selection of different generator predictions from this section. 198
9.16 A comparison of analysis results to predictions from NEUT 5.3.2.2 using various nuclear models. 199
9.17 A comparison of analysis results to predictions from NuWro 11 using various nuclear models. 200
9.18 A comparison of analysis results to predictions from NEUT 5.3.2.2 using an SF nuclear model with various different FSI strengths. 201
9.19 A comparison of analysis results to predictions from NEUT 5.3.2.2 using an RFG+RPA nuclear model with various different FSI strengths. 202

B.1 The unregularised differential cross section in the STV and the associated correlation matrices obtained from fits to real data. 214

List of Tables

2.1 The fractional uncertainty on the number of expected events at the T2K far detector, split by error source and by sample. 26

6.1 The proportion of true signal topologies before and after the pi^0 veto cut alongside the proportion of CC0π and non-CC0π events rejected by the cut. 72
6.2 The proportion of each true signal topology in each reconstructed sample alongside the expected number of events in available data. 73
6.3 The proportion of each true interaction mode in each reconstructed sample alongside the expected number of events in available data. 73
6.4 The breakdown of various backgrounds to the CC0π+Np selection, separated by selection failure mode. 76
6.5 The proportion of each true signal topology in each reconstructed sideband sample alongside the expected number of events in available data. 92
6.6 The proportion of each true interaction mode in each reconstructed sideband sample alongside the expected number of events in available data. 92
6.7 The relative detector systematic error on the total event rate in each sample from sources contributing greater than 1% to a single proton signal sample. . 100

7.1 The δp_T analysis binning. 113
7.2 The $\delta\phi_T$ analysis binning. 114
7.3 The $\delta\alpha_T$ analysis binning. 114
7.4 A summary of all parameters included in the fit. 119

8.1 The size of each term in each maximum likelihood fit to NEUT fake data. . . 138
8.2 The size of each term in each maximum likelihood fit to GENIE 2.8.0 fake data. 147
8.3 The size of each term in each maximum likelihood fit to NuWro 11q fake data. 157
8.4 The size of each term in each maximum likelihood fit to NEUT 6D (bug-fixed) fake data. 166
8.5 The size of each term in each maximum likelihood fit to NEUT fake data with the 2p-2h contribution removed. 174

9.1 The size of each term in each maximum likelihood fit to real data. 186

A.1 A summary of the simulations used for event generation at ND280. 212
A.2 A summary of the event generators used for analysis sensitivity studies and interpretation of the analysis results. 212

Introduction

> *"Forty-two!" yelled Loonquawl. "Is that all you've got to show for seven and a half million years' work?"*
> *"I checked it very thoroughly," said Deep Thought, "and that quite definitely is the answer. I think the problem, to be quite honest with you, is that you've never actually known what the question is."*
>
> *Douglas Adams, The Hitchhiker's Guide To The Galaxy*

How did the Universe come to be? It is a wonderful property of scientific research that such a seemingly unknowable question can be broken down into manageable chunks (each of which can have a satisfactory unambiguous - potentially non-integer - answer).

One such chunk asks the question: Why is the Universe made up of only matter? It is generally assumed that in the very early Universe there was equal amounts of matter and antimatter but observations suggest that there is very little antimatter in the current Universe. In 1967 Andrei Sakharov proposed a set of three necessary conditions to explain this matter-antimatter asymmetry, one of which requires the violation of CP-symmetry (CP standing for charge parity) [1]. Although such CP-violation is allowed within the Standard Model of Particle Physics and has been observed in the mixing of quarks, the measured asymmetry is not sufficient to explain the matter-antimatter asymmetry in the Universe.

Despite a lack of sufficient CP-violation in the Standard Model, relatively recent observations have provided a tantalising possibility of CP-violation in the neutrino sector. The Standard Model requires neutrinos to be massless but measurements made at Super-Kamiokande and the Sudbury Neutrino Observatory around the turn of the millennium - recognised by the 2015 Nobel prize in physics - led to the discovery of neutrino oscillations, which shows that neutrinos do have mass. Although mass-less neutrinos are offered no possibility for CP-violation in the Standard Model, the physics of massive neutrinos is quite different: the current widely-used theory of neutrino oscillations, parametrised by the PMNS matrix and the three neutrino masses, allows for large CP-violation to be characterised by one of its nine parameters, the phase δ_{CP}.

So what is δ_{CP}? This is a manageable question (and one for which forty-two (degrees) would be an acceptable answer), in fact the current generation of long-baseline neutrino oscillation experiments - the T2K and NOνA experiments - are on the verge of providing the first answers, albeit with a large uncertainty. For example, with the data that the T2K

experiment is approved to collect, and reasonable best-case values of the other neutrino oscillation parameters, T2K will have sensitivity to show that δ_{CP} is not zero at 90% confidence limit or better, providing the true δ_{CP} is within the range $50^\circ < \delta_{CP} < 130^\circ$ (for an inverted neutrino mass hierarchy) [2].

T2K operates by accelerating protons into a fixed target to produce a ~ 0.6 GeV predominantly muon neutrino beam. The beam is then measured almost immediately after production at a suite of near detectors, and then again 295 km away at Kamioka by the Super-Kamiokande far detector. Neutrino oscillations are manifest in the changing flavour content of the beam and many of the neutrino oscillation parameters can be extracted by analysing how this changes as a function of neutrino energy. Typically the predominant cause of systematic uncertainty in a measurement of δ_{CP} by T2K, and indeed on all neutrino oscillation measurements using GeV-scale beams (as is required by all long-baseline experiments), is from a lack of knowledge regarding neutrino scattering processes and the dynamics of the nuclei off which they scatter.

This thesis presents a measurement of neutrino scattering off a hydrocarbon nuclear target at the T2K off-axis near detector (ND280) using novel observables specifically designed to probe, or facilitate future probes of, the effects that cause the largest uncertainty in neutrino oscillation measurements. The observables utilise kinematic imbalance in the plane transverse to the incoming neutrino (single-transverse kinematic imbalance - STKI) to investigate these 'nuclear effects', such as the initial momentum distribution of nucleons inside a nucleus (the Fermi motion); the re-scattering of exiting hadrons inside the nuclear medium following a neutrino interaction (final state interactions) and the prevalence of correlations between nucleons inside a nucleus that cause incoming neutrinos to effectively scatter off a bound state of multiple nucleons (nucleon-nucleon correlations).

The results form the first differential cross-section measurement as a function of these observables and allow a novel characterisation of these nuclear effects. In particular the provision of an improved understanding of Fermi motion and nucleon-nucleon correlations will directly reduce the systematic error on neutrino energy reconstruction at T2K, whilst a better knowledge of final state interactions will enable a more precise estimation of future neutrino-nucleus scattering cross-sections of exclusive interaction modes. These nuclear effects are currently poorly understood so, to minimise dependence on nuclear models, this measurement of STKI is presented as a fiducial cross-section of all charged-current muon neutrino scatters off hydrocarbon that produce one muon, at least one proton and no pions in the final state (CC0π+Np interactions). Some elements of this work have previously been reported in [3, 4, 5, 6].

This thesis begins by setting the historical, phenomenological and experimental context to motivate this measurement of STKI, before presenting a specially developed event selection

and method of cross-section extraction that ensures minimal model-dependence in the final result.

To achieve this, chapter 1 gives a very brief history and current status of relevant neutrino physics, whilst also summarising the current understanding of neutrino oscillations within the PMNS framework. It also discusses current and future planned neutrino oscillation experiments, providing preliminary details of how they are likely to suffer limitations from a limited knowledge of neutrino scattering.

Chapter 2 then summarises the dynamics of charged-current quasi-elastic scattering processes when an entire nucleon is resolved at one vertex, which are the predominant component of CC0π+Np interactions. Following this, the current understanding of how nuclear effects can distort the final state particle kinematics is discussed before an outline of the physics of inelastic and neutral-current neutrino scattering processes is presented. The role of uncertainties related to neutrino-nucleus scattering on neutrino oscillation analyses is then reviewed. This chapter finishes with a review of previous measurements of CC0π interactions.

Chapter 3 overviews the observables that characterise STKI. Previous measurements of these observables are discussed and series of phenomenological studies are shown to demonstrate that they offer an excellent probe of nuclear effects.

Chapter 4 details the T2K experiment, where particular focus is placed upon the components used directly in the measurement of STKI (the tracking region of ND280) and the simulation of fake data at ND280.

Chapter 5 uses the previous chapters to further motivate the extraction of a cross-section as a function of variables that characterise STKI and to outline an analysis strategy to probe nuclear effects with minimal model dependence.

Chapter 6 explains the selection and categorisation of events to be used for measuring STKI and constraining backgrounds. Uncertainties in the detector modelling, and the propagation of simulated particles through ND280, are summarised.

Chapter 7 describes the cross-section extraction method to be used in the analysis: a regularised binned maximum likelihood template fitter. Firstly the general principles of the method are detailed, focussing on how systematics can be handled and how Tikhonov regularisation facilitates a data-driven trade-off between bias to the input simulation and resolving the ill-posedness of cross-section extraction problems. Following this the exact implementation of the fitter for measuring STKI is specified.

The latter portion of this thesis first validates the cross-section extraction methodology with comprehensive fake data studies, before applying it to the latest data from ND280 to obtain the primary result of this thesis.

To do this, chapter 8 exhibits these extensive validation studies to establish the robustness of the analysis framework. This includes a plethora of fake data studies with an emphasis on ensuring there is negligible bias to the input simulation.

Chapter 9 then shows the measured STKI as a fiducial differential cross-section of CC0π+Np neutrino-nucleus interactions (on a hydrocarbon target) at ND280. The complicated phenomenological state of the field makes extracting physics conclusions non-trivial, but this is facilitated through an extensive range of model comparisons. The implications of the results with regards to reducing the systematic uncertainties in oscillation analyses are also discussed.

Finally chapter 10 summarises the analysis, the results and what has been learnt about nuclear effects.

Although specialist terms will be defined the first time they are used, a glossary and list of acronyms is also provided prior to the appendices for the reader's reference.

It should also be noted that, although every attempt is made to cite publicly available sources of reference throughout this thesis, some citations necessarily refer to internal T2K technical notes. If required, these may be obtained by contacting the T2K Spokesperson and International Co-Spokesperson. Contact details can be found at `http://t2k-experiment.org/contact-us/`.

Chapter 1

Neutrino physics

Neutrinos are suspected to be the second most abundant particle in the universe (after photons) and yet remain some of the most elusive. This naivety of neutrino nature is not for lack of study, but instead because their light, leptonic and neutral character allows interactions only via the weak force, instilling in neutrinos a profound shyness when it comes to interacting with matter. Indeed, when Wolfgang Pauli first postulated the neutrino in 1930 [7], to remedy an apparent lack of energy and angular momentum conservation in β-decay, he believed he had done a "terrible thing" in suggesting the existence of something that "cannot be detected", wagering a case of champagne that they would never be found. It took 26 years for Pauli to loose his bet, but on 14 June 1956 Clyde Cowen and Frederick Reins sent Pauli a telegram telling him that they had "definitely detected neutrinos" using Project Poltergeist: an apparatus to observe inverse β-decay stimulated by neutrinos from a nuclear reactor [8]. Six decades later, research is on the brink of identifying whether CP-violation occurs in neutrinos, potentially facilitating a mechanism to explain the observed matter antimatter-asymmetry of the universe. This chapter will overview select aspects of neutrino physics, with a view to providing a context for the research conducted herein.

1.1 A brief history of neutrino physics

Pauli's proposed "desperate remedy" to restore energy conservation was almost immediately developed by Enrico Fermi, as part of his theory of β-decay, published in 1934 [9]. Although it would take another 22 years for Cowen and Reins to confirm the existence of Pauli's and Fermi's neutrino, progress in experimental neutrino physics after this was much more rapid. In 1962 the AGS neutrino experiment at Brookhaven national lab utilised the first neutrino beam to observe a second type of neutrino - the muon neutrino: ν_μ - which produces muons in charged-current interactions [10], as opposed to the antielectron neutrinos, $\overline{\nu}_e$, seen in Cowen and Reins's inverse β-decay experiment which produce positrons. When the τ lepton was discovered at SLAC in the mid-1970s [11], it was expected that this would also have an

associated tau neutrino (ν_τ). It took until the year 2000, but the ν_τ was finally discovered at Fermilab's DONUT experiment [12].

The discovery of the τ signified the existence of a third generation of matter, it is interesting to ask whether there are any more. A partial answer came in 2006 when the LEP experiments at CERN published a precision measurement of the Z-boson resonance. By measuring the total Z-boson decay width and subtracting the 'visible' widths (from decays to hadrons and charged leptons), the 'invisible' partial width is obtained. It is assumed that this invisible width comes from Z-boson decays into light neutrino species (i.e. those that are lighter than half the Z-boson mass) and that each species contributes equally. Since the Z-boson decay width for $Z \rightarrow \nu\overline{\nu}$ is calculable, the number of neutrino species making up the invisible width can be found. The LEP experiments measured this to a very high precision to determine that there are three light neutrino species that couple to the Z-boson[1] [13]. It should be noted that this result does not exclude extra heavy neutrino species or species that do not couple to the Z-boson (sterile neutrinos).

Following their discovery, it was realised that neutrinos might make an ideal probe of environments that other radiation cannot penetrate. In particular, Ray Davis and John Bahcall believed neutrinos could be used to study the nuclear reactions taking place beyond the surface of the sun to test Bahcall's solar model. To do this, in the late 1960s Davis and Bahcall set up the Homestake experiment to measure the (electron) neutrino flux from the sun. The results suggested the flux was two to three times smaller than the model predictions [14]. Most early attempts to explain this discrepancy proposed that the solar model predictions were incorrect or that there was something wrong with the experiment. However, by the early 1990s, observations of helioseismology agreed with Bahcall's predictions, whilst results from the SAGE [15], GALLEX [16] and (Super-)Kamiokande [17] experiments corroborated the Homestake observations. An alternative solution to this 'solar neutrino problem' had been first offered in 1957 [18] and elaborated in 1967 [19] by Bruno Pontecorvo. Pontecorvo had suggested that neutrinos could oscillate between flavour states as they propagate, thereby reducing the electron neutrino flux at the Earth but maintaining the total neutrino flux. Pontecorvo had even suggested that 'from the point of view of [neutrino oscillation] detection possibilities, an ideal object is the sun' [19]. In 2002 the SNO experiment was able to directly test the neutrino oscillation hypothesis. SNO was able to measure both the electron neutrino flux (through charged-current interactions) and the total neutrino flux (through-neutral current interactions). The results showed, once again, that the electron neutrino flux was about a third of the solar models predictions, but that the total neutrino flux was in good agreement [20]. This was strong evidence that Pontecorvo was correct, clearly demonstrating that solar neutrinos undergo flavour change.

[1] More accurately, the combined results from the four LEP experiments find that the number of light and active neutrino types is 2.984 ± 0.008 [13].

Additional evidence for neutrino oscillations had also been provided 4 years prior to the SNO result with the observation of neutrinos coming from cosmic ray induced showers. The Super-Kamiokande experiment measured the flux of these electron and muon 'atmospheric' neutrinos (separately) as a function of incoming neutrino direction, which acts as a proxy for the distance an incoming atmospheric neutrino has travelled since generation (i.e. the baseline). In principle, it is expected that the atmospheric neutrino flux would be predictable. Whilst results from Super-Kamiokande observed electron neutrinos to be in reasonable agreement with this expectation, the results indicated that the muon neutrino flux was very dependent on the incoming neutrino direction. These observations could be explained by the disappearance of muon neutrinos ($\nu_\mu \rightarrow \nu_\tau$), in a way that was entirely consistent with neutrino oscillations [21].

Together the results from SNO and Super-Kamiokande show robust evidence for neutrino oscillations. However, such oscillations are only possible for massive neutrinos and neither Fermi's originally proposed neutrino, nor the neutrinos in the modern standard model, allow neutrinos to have any. Therefore the discovery of neutrino oscillations by the Super-Kamiokande and SNO experiments demonstrate beyond standard model physics in showing neutrinos must have mass, a discovery which was awarded the 2015 Nobel prize in physics [22].

Since the SNO and Super-Kamiokande results, neutrino oscillations have been studied with increasing accuracy using well-understood neutrino fluxes from both nuclear reactors and accelerator based neutrino beams. Modern experiments are now making a precision characterisation of the PMNS neutrino oscillation model discussed in the next section (section 1.2).

1.2 Neutrino oscillation theory

Potecorvo's originally proposed neutrino oscillations [18] drew analogy to previously observed neutral kaon mixing ($K^0 \rightarrow \overline{K^0}$), suggesting that neutrino mass states are different than their flavour states. Although the process $\nu \rightarrow \overline{\nu}$ has not been observed, this idea formed the conceptual foundation for a full theory of neutrino oscillations developed by Maki, Nakagawa and Sakata in 1962 [23] to complement Pontecorvo's 1967 extension [19] and a publication in 1968 by Gribov and Pontecorvo which directly addressed the Homestake result and proposed neutrino oscillations as the solution [24].

The formalised theory of neutrino oscillations considers the complementarity of the flavour states of the weak interaction ($|\nu_\alpha\rangle$, $\alpha = e, \mu, \tau$), that neutrinos are always created

in, and the neutrino mass states ($|\nu_i\rangle$, $i = 1, 2, 3$), that neutrinos propagate in:

$$|\nu_\alpha\rangle = \sum_i U^*_{\alpha i} |\nu_i\rangle \ . \tag{1.1}$$

Off-diagonal terms in the leptonic mixing matrix, U, known as the PMNS matrix (Pontecorvo-Maki-Nakagawa-Sakata), result in neutrinos being created in a superposition of mass states thereby causing them to change their flavour state upon propagation. In other words, this mixing of states means that a neutrino created in some flavour state and observed sometime later will have a finite chance of being measured as a different flavour state.

Only considering the known active flavours of neutrinos, the PMNS matrix can be written as:

$$\begin{aligned} U &= \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \end{aligned} \tag{1.2}$$

using the concise notation that $c_{\alpha\beta} = cos(\theta_{\alpha\beta})$ and $s_{\alpha\beta} = sin(\theta_{\alpha\beta})$. The lower line of equation 1.2 imposes the condition that U should be unitary to write U as a product of four sub-matrices, parametrised by three 'mixing angles' ($\theta_{12}, \theta_{23}, \theta_{13}$), one Dirac CP-violating phase (δ_{CP}) and two Majorana CP-violating phases (α_1 and α_2).

The final matrix in equation 1.2, containing the Majorana CP-violating phases, has physical consequences only if neutrinos are Majorana particles (i.e. neutrinos are their own antiparticles). But even if neutrinos are Majorana, these phases do not effect neutrino oscillation probability since this depends on UU^*, where the Majorana phases cancel.

To proceed from the mixing of neutrino mass and flavour states to oscillation probabilities, the neutrino propagation, from its initial conception in a superposition of mass states, must be considered. If a neutrino is created in a particular combination of mass eigenstates at time $t = 0$, each mass state at time t can be written (in natural units - $\hbar = c = 1$) as:

$$|\nu_i(t)\rangle = e^{-i(E_i t - \vec{p}_i \cdot \vec{x})} \mid \nu_i(0)\rangle, \tag{1.3}$$

where E_i and $\vec{p}_i$ are the energy and momentum of the propagating neutrino and $\vec{x}$ is its current position relative to where it was at $t = 0$. In the ultra relativistic limit (practically applicable to all observed neutrinos since their mass (m_i) is so small such that $|\vec{p}_i| >> m_i$ is always true) this can be rewritten using:

$$\begin{aligned} E_i &= \sqrt{p_i^2 + m_i^2} \simeq p_i + \frac{m_i^2}{2p_i} \approx E + \frac{m_i^2}{2E}, \\ t &\approx L, \end{aligned} \tag{1.4}$$

where E is the neutrino's total energy and L is the distance it has travelled such that:

$$|\nu_i(L)\rangle = e^{-im_i^2 L/2E}|\nu_i(0)\rangle. \tag{1.5}$$

This illustrates how each neutrino mass state propagates with a different phase (depending on the mass), demonstrating the importance of the difference between the masses of each state in neutrino oscillations. From here the PMNS matrix can be used to write the probability of a neutrino having started in some flavour state α, being found in flavour state β at time t:

$$P_{\alpha\to\beta} = |\langle\nu_\beta(t)|\nu_\alpha\rangle|^2 = \left|\sum_i U^*_{\alpha i}U_{\beta i}e^{-im_i^2 L/2E}\right|^2 \tag{1.6}$$

This can be written as:

$$\begin{aligned} P_{\alpha\to\beta} = \delta_{\alpha\beta} &- 4\sum_{i>j}\mathrm{Re}(U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j})\sin^2\left(\frac{\Delta m^2_{ij}L}{4E}\right) \\ &+ 2\sum_{i>j}\mathrm{Im}(U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j})\sin\left(\frac{\Delta m^2_{ij}L}{2E}\right), \end{aligned} \tag{1.7}$$

where $\Delta m^2_{ij} \equiv m_i^2 - m_j^2$.

From equation 1.7 and 1.2 it can be seen that three flavour neutrino oscillations can be described by the mixing angles and CP-violating phase in the PMNS matrix alongside two mass differences to describe the difference in neutrino mass state masses. It is these parameters that modern neutrino oscillation experiments aim to measure.

In particular the measurement of δ_{CP} would represent a characterisation of the CP-violation in the neutrino-sector (within the PMNS model), which, as suggested in the introduction to this thesis, has far reaching consequences for our understanding of the matter-antimatter imbalance in the Universe.

1.2.1 Matter effects in neutrino oscillations

Up until now neutrinos have been considered as propagating without disturbance from matter, but in reality most neutrino experiments involve neutrinos propagating through the Earth. This can make a substantial difference to neutrino oscillation physics. There are two mechanisms for neutrinos to interact without the neutrino converting into a charged lepton: neutral-current scattering from an electron, proton or neutron or charged-current scattering of an electron and a neutrino, both illustrated in figure 1.1.

A full account of how exactly these matter effects alter neutrino oscillation probabilities is beyond the scope of this thesis (although it is well described in other sources, e.g. [25]). For the purposes of this work it is sufficient to say that, whilst vacuum oscillations are only

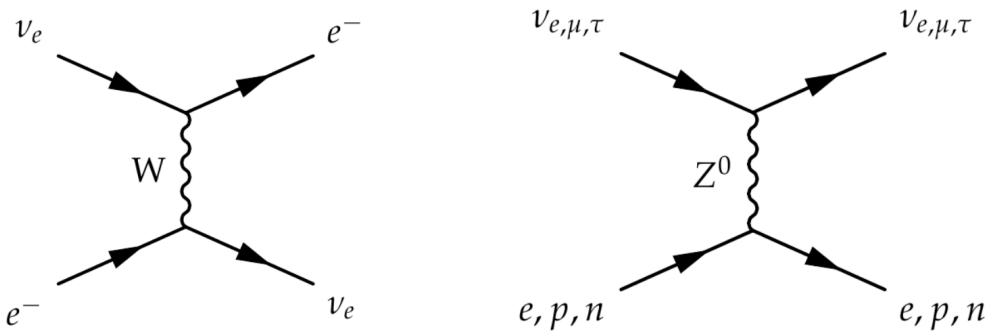


Figure 1.1: Feynman diagrams to illustrate matter effects in neutrino interactions. All neutrino flavours can undergo the neutral-current scattering (right) but only electron neutrinos can undergo charged-current scattering and still leave a neutrino in the final state (left).

sensitive to the square of the neutrino mass splitting, matter effects are sensitive to the signs of the mass splittings. Exploiting this in solar neutrino observations has allowed the determination that ν_2 is greater than ν_1, but it is not yet known whether ν_3 is the heaviest (a 'normal neutrino mass hierarchy') or the lightest (an 'inverted neutrino mass hierarchy'). Whether the neutrino mass hierarchy is normal or inverted is one of the most prominent unanswered questions in neutrino physics and the answer is expected to play a crucial role in determining the full particle physics model that explains neutrino masses and mixing. It may be possible for future neutrino oscillation experiments to utilise matter effects within the Earth to resolve the mass hierarchy.

Another important implication of matter effects in neutrino oscillations is that their impact is different for neutrinos and antineutrinos (since the CC interaction shown in figure 1.1 is not available to antineutrinos). This can lead to a measured CP violation $P(\nu_\alpha \rightarrow \nu_\beta) \neq P(\overline{\nu}_\alpha \rightarrow \overline{\nu}_\beta)$ which does not tell us anything interesting about matter-antimatter asymmetry at a fundamental level. It is therefore very important to account for matter effects when attempting to determine δ_{CP} to identify genuine neutrino-sector CP violation.

1.3 Modern long baseline neutrino oscillation experiments

As demonstrated in equation 1.7, the probability of neutrinos having oscillated away from their original flavour state depends on their energy and distance travelled. Modern accelerator-based long baseline neutrino oscillation experiments, such as T2K [26] and NOνA [27], exploit this by utilising a proton accelerator to produce a - predominantly muon - neutrino beam spread over a few GeV neutrino energy and observe the change in the flavour composition of

the beam over a few hundred kilometres, where neutrino oscillations are maximal. This allows measurements of muon neutrino disappearance (particularly sensitive to the oscillation parameters $sin^2(\theta_{23})$ and Δm^2_{23}) and electron neutrino appearance (particularly sensitive to the oscillation parameters $sin^2(\theta_{13})$ and δ_{CP}). These experiments consist of a far detector, positioned close to the oscillation maximum, and a near detector, positioned just after the neutrino beam production point, to constrain the properties of the neutrino beam. More details of the T2K experiment can be found in chapter 4.

Additionally, modern experiments are able to run with an antimuon neutrino beam, thereby gaining additional sensitivity to CP-violation. Since these experiments operate over long-baselines, the neutrino beams undergo significant matter effects (discussed in section 1.2.1), thereby also affording some sensitivity to the neutrino mass hierarchy.

The sensitivities of accelerator-based long baseline experiments complement other neutrino oscillation experiments. In particular modern experiments that measure the oscillation of antielectron neutrinos generated in nuclear reactors (such as Daya Bay [28, 29], Double Chooz [30, 31], RENO [32] and KamLAND [33]) are predominantly sensitive to θ_{12}, θ_{13} and Δm^2_{21}, thereby facilitating measurements of all known neutrino mixing parameters.

Although the values of the mixing angles and mass splittings are now known to the 5%-10% level, the mass hierarchy and δ_{CP} remain unknown and their determination is a principle goal of current and future long baseline neutrino oscillation experiments.

In order to extract oscillation parameters, long baseline neutrino experiments measure neutrino scatters of a particular neutrino flavour at their far detectors in bins of some observable(s), $\mathbf{\Theta}$ (often this is reconstructed neutrino energy). A Monte-Carlo simulation is used to predict the $\mathbf{\Theta}$ distributions for a particular set of oscillation parameters. These parameters are then varied until the predicted $\mathbf{\Theta}$ distributions are in best agreement with what is measured in order to obtain a best fit set of oscillation parameters. However, the measured $\mathbf{\Theta}$ spectrum depends on more than just the oscillation parameters, it is also affected by the detectors efficiency(ϵ), the incoming neutrino flux (Φ) and neutrino interaction cross sections (σ), such that:

$$N_{pred}(\mathbf{\Theta}) = \int \Phi(E_\nu)\sigma(E_\nu)P_{\nu_\alpha \to \nu_\beta}(E_\nu)\epsilon(E_\nu)U(E_\nu, \mathbf{\Theta})dE_\nu, \tag{1.8}$$

where $N_{pred}(\mathbf{\Theta})$ is the number of events as a function of the observable(s) $\mathbf{\Theta}$, E_ν is the neutrino energy and $U(E_\nu, \mathbf{\Theta})$ is the detector response function, describing how a true E_ν is reconstructed in $\mathbf{\Theta}$.

It is clear from equation 1.8 that in order to reliably measure the neutrino oscillation probability, $P_{\nu_\alpha \to \nu_\beta}(E_\nu)$, the unoscillated flux, the detector response and neutrino interaction cross-sections must be well understood. If any of these components are not well modelled

the extracted oscillation parameters may be biased whilst large uncertainties on these components will limit the sensitivity of the experiment.

1.4 Future sensitivity to CP-violation

Chapter 2 demonstrates that T2K and NOνA are nearing the point at which $N_{pred}(\mathbf{\Theta})$ is dominated by systematic errors for measurements of ν_μ disappearance measurements, the largest of which stem from neutrino interaction uncertainness. Although this may become the limiting factor in the precision of oscillation parameter measurements at T2K and NOνA, the problem is much more severe for the sensitivity of the next generation of accelerator-based long-baseline neutrino oscillation experiments: DUNE [34] and Hyper-Kamiokande (HK) [35]. Figure 1.2 demonstrates the expected δ_{CP} sensitivity of HK and DUNE as a function of exposure for different systematic uncertainties on $N_{pred}(\mathbf{\Theta})$. This demonstrates that uncertainties of 3% or lower (to be compared to current uncertainties of around 5-6%) are likely required to ensure a precise measurement of δ_{CP} that does not become systematics limited early in the experiments lifetime. As the principle systematic in current measurements of δ_{CP}, it will be essential to better understand the physical processes involved in neutrino-nucleus interactions in order to achieve this.

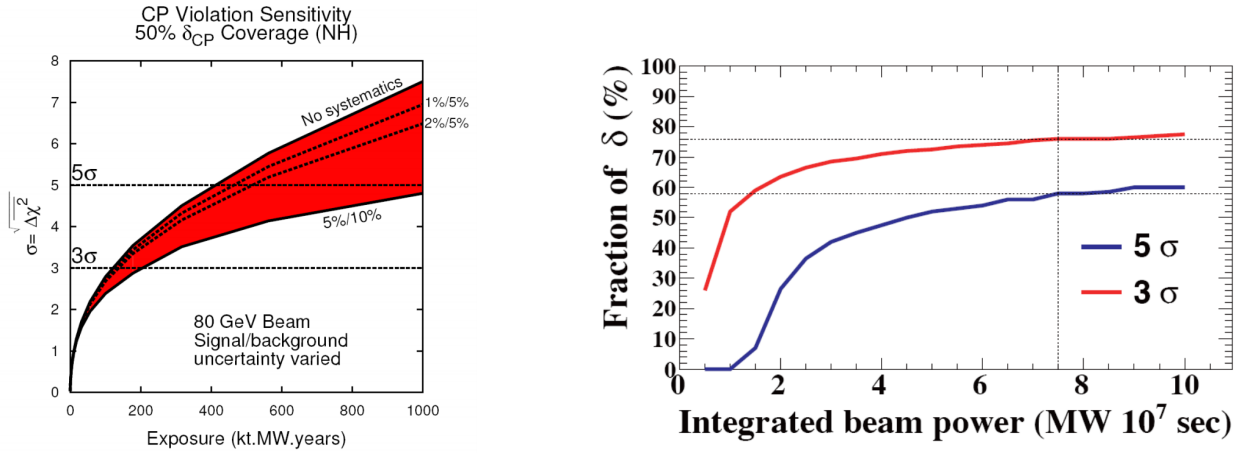


Figure 1.2: **Left:** Published CP violation sensitivity curves from the LBNE experiment (now renamed to DUNE) showing the predicted sensitivity to CP violation over 50% of the δ_{CP} parameter space, assuming a range of signal and background normalisation uncertainties, as a function of exposure [36]. **Right:** Published CP violation sensitivity curves from the HK experiment showing the predicted fraction of δ_{CP} parameter space for which a 3σ (red) or 5σ (blue) measurement of CP violation would be made for a given integrated beam power at HK assuming a total uncertainty of 3% on the far detector event rate [37]. Figure and caption adapted from [38].

Chapter 2

Neutrino-nucleus interactions

In the previous chapter, section 1.3 demonstrated that a sound understanding of neutrino-nucleus interactions is important for precise neutrino oscillation analyses. This chapter gives an introduction to neutrino-nucleus scattering, which is crucial to put the proceeding chapters in context, whilst building on section 1.3 to demonstrate why neutrino interaction uncertainties may soon become the limiting factor in oscillation analyses. Particular focus is placed on describing charged-current quasi-elastic (CCQE) interactions and nuclear effects, as it is predominantly these that limit T2K oscillation analyses and are investigated later in this thesis. A plot of measured neutrino scattering cross sections, as a function of incoming neutrino energy, is reproduced from [39] in figure 2.1 (this includes data taken up to 2012). This shows both the relevant neutrino energy scales for each interaction mode and the availability of measurements, which will be useful for reference within this chapter.

Throughout this thesis five different neutrino-nucleus interaction simulations (also known as event generators) are considered: NEUT [40], GENIE [41], NuWro [42], GiBUU [43, 44] and NUANCE [45]. Each uses Monte Carlo methods to simulate large productions of neutrino interactions based on models of the processes that will be discussed throughout this chapter. Where there is substantial difficulty in simulating a process (as is often the case for the simulation of nuclear effects), differences in the model implementation between the simulations will be discussed.

The chapter begins, in section 2.1, with a discussion of CCQE interactions, before the impact of nuclear-medium effects on such scatters is described in section 2.2. Following this, other neutrino interaction modes are considered in section 2.3. Section 2.4 then highlights the importance of understanding neutrino-nucleus scattering by providing a detailed description of its impact on oscillation analysis. In section 2.5, some of the challenges associated with extracting useful physics from measurements of CCQE-like neutrino-interactions are addressed and the confusing current state of the field, following modern measurements, is exhibited. Finally, in section 2.6, the chapter concludes with a summary of the current experimental and phenomenological challenges in measurements of neutrino-nucleus interactions with a view to motivating the primary analysis of this thesis.

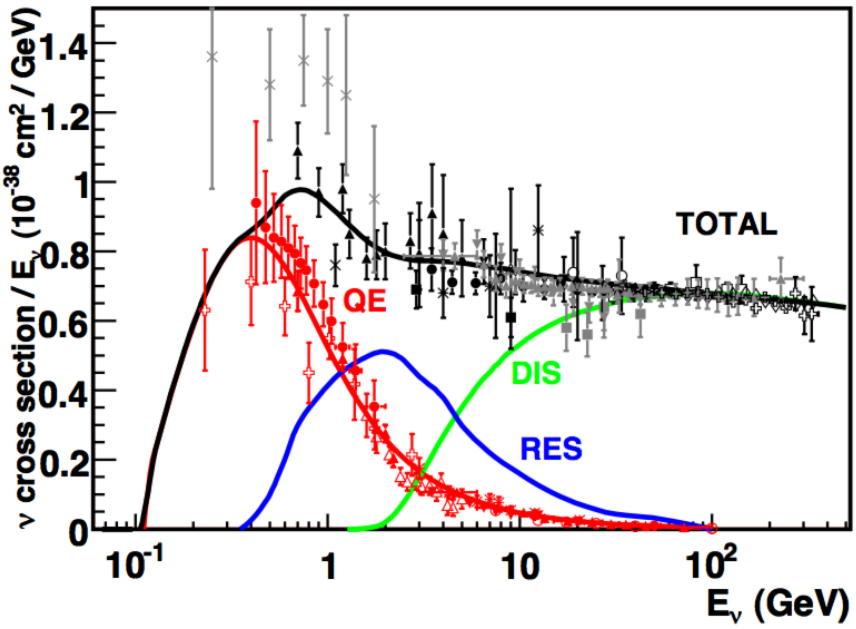


Figure 2.1: The total muon neutrino scattering cross section as a function of neutrino energy broken down by interaction mode as predicted by the NUANCE [45] neutrino interaction simulation. Available data up to 2012 is overlaid. Figure taken from [39].

2.1 The charged-current quasi-elastic interaction

At neutrino energies below $\sim$ 1.5 GeV, the primary way in which neutrinos interact with matter is via CCQE interactions (as shown in figure 2.1). In these interactions a W-boson is exchanged between a neutrino and nucleon, producing a charged lepton and a nucleon with altered isospin, as shown in figure 2.2. Since a nucleon is an extended object, it is very difficult to calculate a cross section for this interaction analytically. However, it is possible to parametrise the cross section in terms of parameters that can mostly be measured through β-decay and electron scattering experiments using the Llewellyn Smith model [46]. The differential cross section from this model is given as a function of the four momentum transfer squared (Q^2):

$$\frac{d\sigma}{dQ^2} = \frac{M^2 G_F^2 cos^2(\theta_c)}{8\pi E_\nu^2}\left(A(Q^2) \pm B(Q^2)\frac{s-u}{M^2} + C(Q^2)\frac{(s-u)^2}{M^4}\right). \tag{2.1}$$

Here the $\pm$ is for neutrinos/antineutrinos respectively, M is the nucleon mass, G_F is the Fermi coupling constant, θ_c is the Cabbibo angle, E_ν is the neutrino energy and s and u are the Mandlestam kinematic variables corresponding to the square centre of mass energy and the square four momentum respectively. For an outgoing lepton of mass m it can be calculated that: $s - u = 4ME_\nu - Q^2 - m^2$. A, B and C are functions of Q^2 built from the vector ($F_V^{1,2}$), axial vector (F_A) and pseudoscaler (F_P) nucleon form factors:

$$A(Q^2) = \frac{m_l^2 + Q^2}{M^2}\Big[(1+\tau)F_A^2 - (1-\tau)(F_V^1)^2 + \tau(1-\tau)(\xi F_V^2)^2 + 4\tau(\xi F_V^1 F_V^2) - \frac{m_l^2}{4M^2}\Big((F_V^1 + \xi F_V^2)^2 + (F_A + 2F_P)^2 - 4(1+\tau)F_P^2\Big)\Big], \tag{2.2}$$

$$B(Q^2) = \frac{Q^2}{M^2}F_A(F_V^1 + \xi F_V^2), \tag{2.3}$$

$$C(Q^2) = \frac{1}{4}\Big(F_A^2 + (F_V^1)^2 + \tau(\xi F_V^2)^2\Big), \tag{2.4}$$

where $\tau = Q^2/4M^2$, $\xi = (\mu_p/\mu_N - \mu_n/\mu_N) - 1$ and m_l is the outgoing lepton mass. μ_p and μ_n are the proton and neutron magnetic moments respectively, whilst μ_N is the nuclear magneton.

The form-factors themselves characterise the internal charge distribution of nucleons. Assuming a scatter off a stationary nucleon with an exponential charge distribution, $\rho(r) = \rho_0 e^{-Mr}$, the form factors take the form of the Fourier transform of this and can hence be described by a dipole: $F(Q^2) \propto (1 - Q^2/M_X^2)^{-2}$ [47], where M_X is some empirical parameter.

Since $F_V^{1,2}$ appears in the vector component of the electroweak nucleon current, its parameters can be accurately determined from observations of electro-magnetic interactions such as in electron scattering experiments. These observations have found that the dipole form works well for $Q^2 < 2.0$ GeV^2, but deviations are seen at larger Q^2 (see e.g. [48]). Extensions to $F_V^{1,2}$, which parametrise these deviations, are typically used by the neutrino scattering community. In this work the BBBA05 form factors are the most commonly used, which are defined and discussed in [49].

The pseudoscaler form factor can be rewritten in terms of Q^2 and the axial vector form factor [39] as:

$$F_P(Q^2) = \frac{2M^2}{Q^2 + m_\pi^2}F_A(Q^2), \tag{2.5}$$

whilst F_A itself is commonly assumed to have a dipole form such that it can be written:

$$F_A(Q^2) = \frac{F_A(0)}{(1 + Q^2/M_A^2)^2}, \tag{2.6}$$

and has two parameters, $F_A(0)$ and M_A, the nucleon axial mass. $F_A(0)$ can be found accurately from β-decay, whilst M_A is left as the only relatively unconstrained parameter in the Llewellyn Smith model. M_A has been measured by fits to bubble chamber $\nu_\mu - H_2$ and $\nu_\mu - D_2$ measurements as well as pion electroproduction data. A global fit to this data has found $M_A = 1.014 \pm 0.014$ GeV [50]. The determination of M_A in heavy target neutrino scattering experiments (originally at the MiniBooNE experience) has recently become a topic of interest in the neutrino scattering community, due to discrepancies seen with the bubble chamber results. This has become known as the 'MiniBooNE M_A puzzle' [51] and will be discussed further in section 2.5.

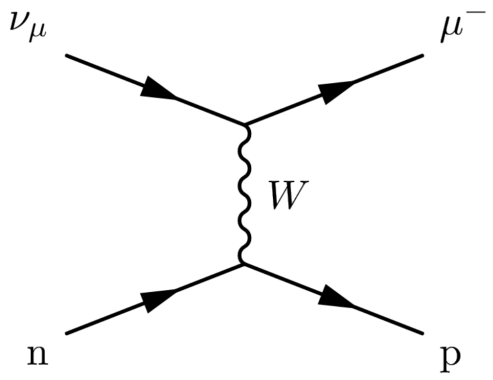


Figure 2.2: A Feynman diagram of a CCQE interaction.

2.2 Nuclear effects in neutrino-nucleus scattering

Section 2.1 describes CCQE interactions between neutrinos and free nucleons. Whilst this is sufficient to describe interactions on very light nuclear targets, such as in some older neutrino scattering experiments using bubble chambers, most modern experiments use heavy targets (typically carbon, oxygen or argon). Within these heavier targets there are additional 'nuclear-effects' to account for, which distort final state outgoing particles from what would be expected from a neutrino impinging on a free-nucleon.

These nuclear effects can broadly be split into initial state effects, including Fermi motion (FM), nuclear binding energy and multi-nucleon processes (MNP), and final state interactions (FSI). FM describes the momentum distribution of nucleons prior to any interaction whilst MNP describe correlations between nucleons and are thought to be dominated by '2p-2h' effects (these will be defined in section 2.2.3) for the neutrino energies relevant to T2K [52, 53]. FSI are outgoing hadron re-interactions inside the nuclear target. These processes can alter the kinematics of, or even entirely absorb, outgoing hadrons whilst also potentially ejecting additional hadrons from the subsequent nuclear-instabilities. A brief discussion of each of these effects and how they are typically modelled follows.

2.2.1 Fermi motion

As described above, FM is the initial state movement of nucleons inside the nucleus prior to any interaction, which is modelled as being isotropic. Such initial state nucleon momentum causes neutrino interactions with nucleons in a nuclear environment to undergo an unknowable event-by-event boost in the lab-frame. Since the nucleus is a complicated environment it is difficult to accurately predict the distribution of the initial state nucleon momenta for anything beyond the lightest of nuclei. Models that attempt to do this are collectively known as

spectral function models and those most commonly used in neutrino-interaction simulations are described below. A comparison of the FM predicted by each model in a carbon nucleus, produced using the NuWro [42] neutrino interaction simulation, is shown in figure 2.3.

A simple and commonly used spectral function model is the 'global' 'relativistic Fermi gas' (RFG) model (sometimes called global Fermi gas - GFG). Here nucleons within the nucleus are modelled as non-interacting fermions inside a pervading nuclear potential, such that all momentum states are filled from the ground state upwards. The highest momentum state has momentum p_F, known as the Fermi momentum, which depends on the the number of nucleons in a nucleus [54]. Overall this gives rise to the 'Fermi cliff' feature, manifest in the immediate cut off at p_F ($p_F \approx 230$ MeV/c for carbon) seen in figure 2.3.

The RFG spectral function treats all nucleons as feeling the same constant binding potential, whilst in reality this depends upon the local density of the nucleus. Such nuclear density alteration with nuclear radial position is known from electron scattering data [56], and can be used to construct a more sophisticated 'local Fermi gas' (LFG) spectral function. In an LFG the local nuclear density ($\rho(r)$) is used to build a nuclear potential that depends on the radial position of a nucleon within the nucleus (r) under the local density approximation [54, 55]. This effectively constructs a continuum of RFGs, each with a local Fermi Momentum $p_F(r)$ that it is dependent on the radial position of the nucleus such that $p_F(r) \propto (\rho(r))^{1/3}$. The overall effect of this is to smooth out the initial state nucleon momenta, as can be seen in figure 2.3. In this thesis the LFG models used are those implemented in NuWro 11q and GiBUU 2016 neutrino interaction generators, which are described in more detail in [42, 54] and [43] respectively.

Although an LFG model is more realistic than an RFG model, nucleons are still treated as non-interacting fermions (other than via the local nuclear potential). However, it is well known from electron scattering data that nucleon-nucleon interactions inside the nuclear medium can significantly alter the distribution of the initial state nucleon momenta [57, 58]. Models which incorporate these effects are known, confusingly, as spectral function (SF) models. In this thesis SF will refer to the SF model from Omar Benhar and collaborators [59]. This still utilises a shell model but considers modifications to the orbits due to two and three nucleon interaction potentials. Furthermore the SF model includes pairs of strongly repulsive nucleons known as short range correlations, which can have momentum above (the local or global) p_F, thereby giving the distribution of initial state nucleon momenta a high momentum tail [59, 54], as can be seen in figure 2.3.

Other, even more sophisticated, models of FM are available [43, 60, 61, 62] but - with the exception of [43] which is used in the GiBUU simulation - these are not implemented in neutrino interaction simulations.

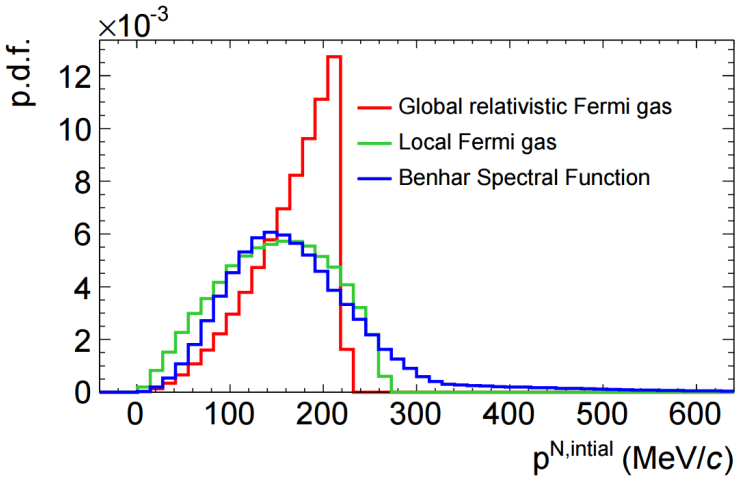


Figure 2.3: The initial nucleon momentum distribution within a Carbon nucleus as simulated by a global relativistic Fermi gas spectral function, local Fermi gas spectral function, and the Benhar spectral function. Figure adapted from [63].

2.2.2 Final state interactions

Following an interaction, the ejected hadrons must escape the nucleus they are created in before any measurement can be made. However, before this can happen it is possible for the hadrons to re-interact inside the nuclear medium where they can be absorbed, have their kinematics altered or stimulate additional nuclear emission (i.e. ejecting extra hadrons inside the nucleus). These re-interactions are known as final state interactions (FSI) and are extremely difficult to model or constrain with experimental data.

To attempt to model FSI, most neutrino interaction simulations use a cascade model. Within these models each hadron leaving the interaction vertex is treated independently and propagated through the nucleus in discrete steps of size based on the hadrons mean free path (MFP). The probability for a multitude of different interactions (e.g. elastic scatter, charge exchange, absorption) is calculated at each step based on the local nuclear density and, if necessary, such an interaction is simulated. This continues until the hadron leaves the nucleus. Outgoing mesons and nucleons are treated differently. A visualisation of FSI in a cascade model is shown in figure 2.4. The scale of the MFP and the relative probability for different FSI interactions are tuned from external hadron scattering data. Details of how this is done and more infomation on FSI implementation in neutrino-nucleus interaction simulations can be found in [54, 64, 65, 66].

It is worth noting that the GiBUU neutrino interaction simulation [44] uses an alternative approach, solving coupled hadronic transport equations to propagate all interaction products through the nuclear medium together. Although this method is substantially slower than

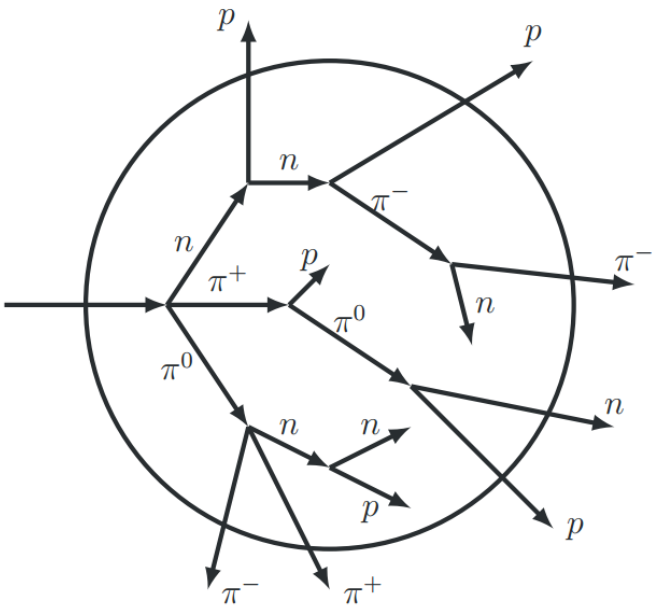


Figure 2.4: A visualisation of FSI in a cascade model. Figure adapted from [54].

a simpler cascade model, it provides a more sophisticated model of FSI, thereby allowing GiBUU to provide an important benchmark as the state of the art in FSI modelling.

2.2.3 Nucleon-nucleon correlations and the random phase approximation

Motivated by the MiniBooNE M_A puzzle, recent models of neutrino interactions aim to include the potential for neutrinos to scatter off a bound state of two or more nucleons [67, 68]. The presence of such nucleon correlations have previously been inferred in electron-nucleus scattering experiments and models to describe their effect on scattering cross sections have been developed (e.g. [69]). However, since it's unclear whether the axial vector component of the interactions will be properly covered in these models, these may not be representative for neutrino scattering [70].

New models for describing the role of multi-nucleon interactions in neutrino-nucleus scattering have recently been proposed by Nieves et al. [67] and Martini et al. [68]. In these models, CCQE events are referred to as one-particle one-hole (1p-1h) interactions but higher order two-particle two-hole (2p-2h) interactions with two nucleons are also considered. These 2p-2h interactions then leave two nucleons in the final state (although they are both still susceptible to absorption by FSI processes before leaving the nuclear medium).

In order to reproduce heavy-target neutrino scattering data, it has been found that both the Nieves and Martini models also need to include a nuclear screening effect, known

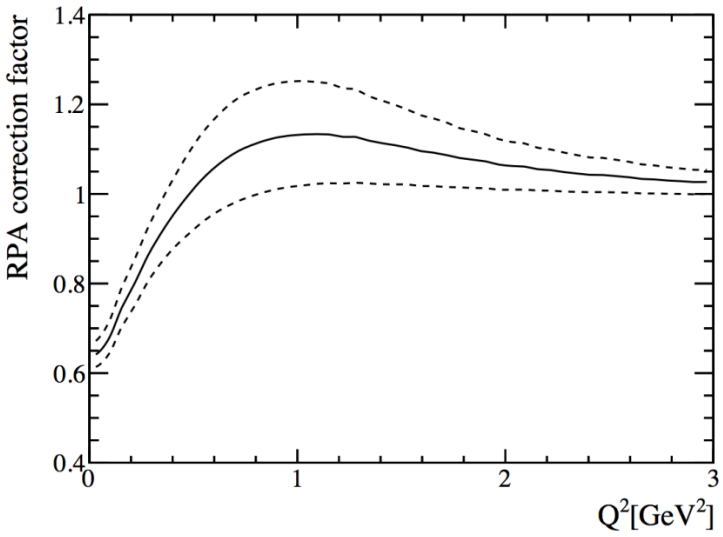


Figure 2.5: The RPA correction factor ($\sigma_{CCQE}/\sigma_{RPA}$) for CCQE neutrino interactions on carbon. This plot was generated using the T2K implementation of the Nieves RPA calculation in [67]. The dotted lines show approximate uncertainty bands.

as the random phase approximation (RPA) [67, 68]. RPA accounts for the impact of the nuclear medium on an electroweak propagator, which causes a Q^2 dependent alteration of the interaction cross-section. The RPA correction factor ($\sigma_{CCQE}/\sigma_{RPA}$) for CCQE interactions on carbon is shown in figure 2.5. Furthermore figure 2.6 illustrates the impact of multinucleon correlations, RPA and M_A on the agreement of the Nieves model with a CCQE measurement made by the MiniBooNE experiment (which has a hydrocarbon target), to illustrate the need for both 2p-2h and RPA if M_A is to be compatible with the early measurements described in section 2.1.

Most neutrino interaction simulations do not simulate full nucleon-nucleon correlations explicitly, but instead use either lookup tables, effective models or simplifications. In particular, no generators currently available include anything more than a very approximate model of the outgoing hadronic kinematics following 2p-2h interactions. It should also be noted that there is currently only circumstantial evidence that 2p-2h-like processes contribute significantly to measurements of neutrino interactions. It is hoped that new measurements - such as the one presented in this thesis - will be able to better characterise the 2p-2h contribution to neutrino interactions.

2.3 Other interaction modes

The discussion in this chapter has so far only considered charged-current neutrino-nucleus scattering which produces no mesons at the initial interaction vertex (although their emission

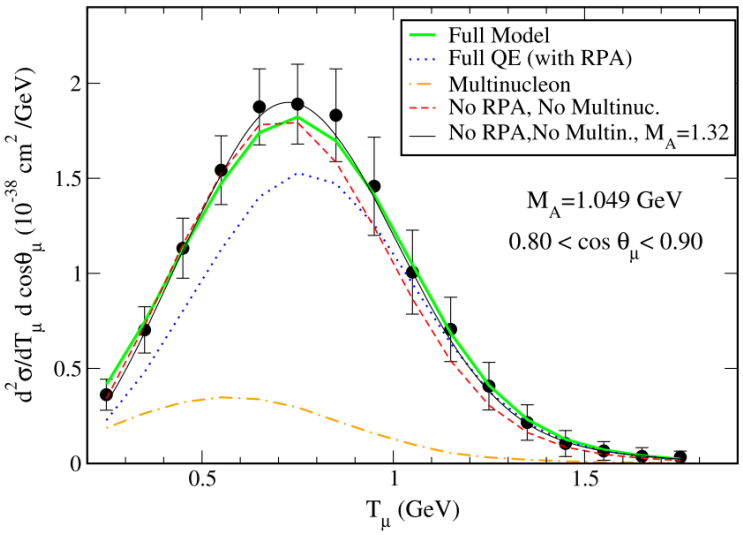


Figure 2.6: The CCQE double differential cross section in Muon angle and energy according to several models plotted alongside data from the MiniBooNE experiment [71]. It can be seen that data is fit well with an (enhanced) $M_A = 1.32$ GeV - the solid thin grey line - but is also compatible with a $M_A = 1.049$ GeV (compatible with results from bubble chambers) once multinucleon (2p-2h) and RPA effects are considered - the solid thick green line. Figure taken from [72].

can be stimulated by FSI processes). This section gives a very brief outline of other relevant neutrino interaction modes.

2.3.1 Resonance production

Once the centre of mass energy of a neutrino-nucleon scatter exceeds the mass of a delta baryon, it is possible that the interaction will excite a resonance state. For neutrinos with energies between about 1.5 GeV and 5 GeV this is the dominant interaction mode (as shown in figure 2.1). The subsequent delta baryon decay (still inside the nuclear medium) produces pions. Resonance production is allowed via charged- or neutral-current scattering in any channel that conserves charge. This leads to a wide variety of possible neutrino resonant pion production interactions, three charged current:

$$\nu_\ell + p \rightarrow \ell^- + p + \pi^+, \tag{2.7}$$

$$\nu_\ell + n \rightarrow \ell^- + p + \pi^0, \tag{2.8}$$

$$\nu_\ell + n \rightarrow \ell^- + n + \pi^+, \tag{2.9}$$

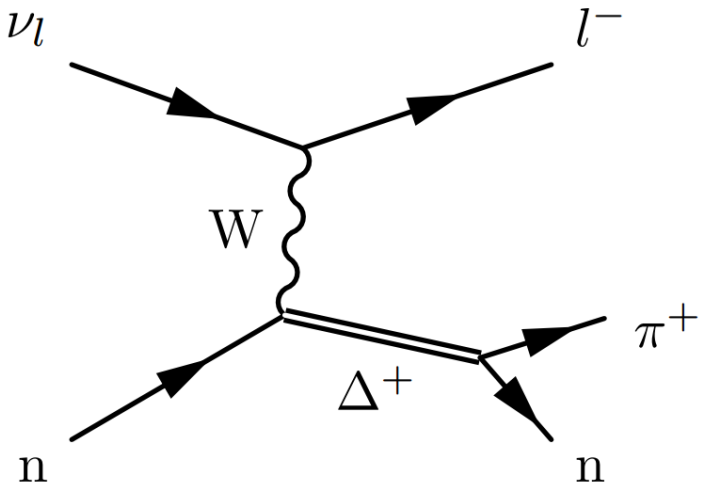


Figure 2.7: A Feynman diagram of a resonance production interaction.

and four neutral current:

$$\nu_\ell + p \rightarrow \nu_\ell + p + \pi^0, \tag{2.10}$$

$$\nu_\ell + p \rightarrow \nu_\ell + n + \pi^+, \tag{2.11}$$

$$\nu_\ell + n \rightarrow \nu_\ell + n + \pi^0, \tag{2.12}$$

$$\nu_\ell + n \rightarrow \nu_\ell + p + \pi^-. \tag{2.13}$$

Figure 2.7 shows an example of a resonance interaction.

Resonance production is typically described in neutrino interaction simulations using the Rein-Seghal model [73]. The calculations in this model are similar to those in the Llewellyn Smith CCQE model in that nuclear form factor are used to account for an extended nuclear target. These form factors are different than those seen in section 2.1 and have two parameters that cannot easily be determined from electron scattering: M_A^{RES} and C_5^A (the equivalent of C_5^A in the CCQE form factors is $F_A(0)$, which can be determined from β-decay).

2.3.2 Coherent scattering

In a coherent scatter a neutrino scatters, via a neutral- or charged-current, off the whole nucleus. The nucleus recoils but does not fragment, leaving it in the same final state as initial state. This is only possible at low Q^2 and this process has a very small interaction cross-section at neutrino energies relevant to accelerator-based long-baseline neutrino oscillation experiments. Upon interaction a single pion is produced, such that the possible coherent scattering modes for neutrinos are:

$$\nu_\ell + A \rightarrow \nu_\ell + A + \pi^0, \tag{2.14}$$

$$\nu_\ell + A \rightarrow \ell^- + A + \pi^+. \tag{2.15}$$

In neutrino interaction simulations coherent scattering is typically modelled with the Rein-Seghal coherent model [74, 75].

2.3.3 Deep inelastic scattering

At high enough neutrino energies, the incoming neutrinos become able to resolve individual quarks and interact via the exchange of a W or Z boson with them. This breaks apart the nucleon, producing a jet of hadrons in an interaction mode known as deep inelastic scattering (DIS). This is the dominant neutrino interaction mode for neutrinos with energy above about 10 GeV (as shown in figure 2.1).

Since neutrino DIS has long been used to validate the standard model and probe nuclear structure [39], there is a plethora of high energy (many GeV) neutrino data available, which has allowed a sound understanding of these DIS scatters. However, at the lower energies relevant to accelerator based long baseline neutrino oscillation experiments, there is very little data and it is not clear that DIS models can be extrapolated accurately to these energies.

2.3.4 Neutral-current elastic scattering

As the neutral current equivalent to CCQE interactions, neutrinos can undergo neutral-current elastic (NCE) scatters, typically ejecting a nucleon from the nucleus. The available NCE interaction modes are:

$$\nu_\ell + p \rightarrow \nu_\ell + p, \tag{2.16}$$

$$\nu_\ell + n \rightarrow \nu_\ell + n. \tag{2.17}$$

These are the dominant neutral-current interaction mode over the same energy range as CCQE is the dominant charged-current interaction mode. As the hadronic current is similar, they are modelled very similarly to CCQE interactions, see [76, 77] for more details.

2.4 Neutrino interactions in oscillation analysis

Equation 1.8 demonstrated that an understanding of neutrino interaction cross sections is essential to make accurate measurements of the neutrino oscillation parameters. This section discusses the effect of neutrino interactions on the sensitivity of neutrino oscillation experiments. It also demonstrates the possibility for a misunderstanding of neutrino interactions to bias the measurements.

This is discussed in the context of the two commonly used approaches to characterise the events used in neutrino oscillation analyses, one that uses only the outgoing lepton kinematics and another that additionally uses calorimetric information.

2.4.1 Lepton kinematics approach - T2K

For experiments with no sensitivity to elements of the hadronic component of neutrino interactions at the far detector, such as at T2K (see chapter 4 for more details), neutrino interactions must be characterised using limited information from the interaction. T2K's primary far detector event selections used for oscillation analysis attempt to identify events with one outgoing lepton and no pions or photons. The far detector is able to accurately identify the lepton and has different event samples for selected ν_μ and ν_e events. Because of T2K's relatively low neutrino beam energy ($\sim$ 1 GeV, see chapter 4), the majority of these events are expected to be CCQE interactions. Assuming the incoming neutrino scatters off a stationary target, in a CCQE event the neutrino energy can be fully reconstructed from only the outgoing lepton kinematics:

$$E_{rec}^{CCQE} = \frac{m_p^2 - (m_n - E_b)^2 + m_\ell^2 + 2(m_n - E_b)E_\ell}{2(m_n - E_b - E_\ell + p_\ell cos\theta_\ell)}, \tag{2.18}$$

where E_{rec}^{CCQE} is the reconstructed incoming neutrino energy, m_p and m_n are the proton and neutron masses, E_b is the nuclear binding energy of the neutron and p_ℓ, θ_ℓ, m_ℓ and E_ℓ are the momentum, the angle with respect to the incoming neutrino, the mass and the energy of the outgoing lepton respectively [78].

In order to obtain oscillation parameters, T2K relies on an input neutrino interaction simulation to map true neutrino E_{true} energy onto either the reconstructed neutrino energy, the measured lepton kinematics, or some combination of the two. This allows alterations of the oscillation parameters to be mapped onto measured quantities, such that they can be fit.

Since the nucleon is not really stationary but undergoes some Fermi motion (see section 2.2.1), the mapping from E_{rec}^{CCQE} in equation 2.18 is smeared (symmetrically, because Fermi motion is isotropic) by an amount that depends on the distribution of the Fermi motion. This is the first example of 'model dependence' in the oscillation analyses: the mapping from E_{rec}^{CCQE} to E_{true} depends on which Fermi motion model the input simulation uses, which in turn impacts the way in which the alteration of the oscillation parameters effect the reconstructed observables, which can alter the oscillation fit result.

There are also other, more severe, forms of model dependence in the T2K method. The selection of events with one visible lepton and no pions or photons does not only contain CCQE events. There are a significant number of resonance events (depending how often the

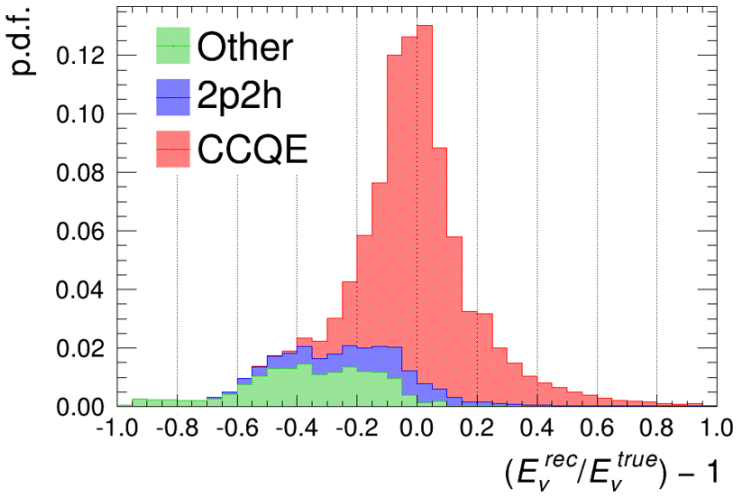


Figure 2.8: The 1D projection of the fractional mapping from E_{rec}^{CCQE} to E_{true} for simulated events selected at the T2K far detector to have one final state muon and no pions or photons, broken down by interaction mode. Zero on the x-axis corresponds to no bias. This plot was produced using the"NEUT 6B" simulation, which is defined in section 4.2.3 and further detailed in appendix A.

final state pions are below the momentum threshold to be seen in the far detector or are absorbed entirely by FSI) and also 2p-2h interactions. For events that are not CCQE, equation 2.18 does not give a representative reconstructed neutrino energy, this is demonstrated in figure 2.8, which shows the shape of a 1D projection of the mapping from E_{rec}^{CCQE} to E_{true}. From this it can be seen that there is not much average bias in the reconstruction of neutrino energy for CCQE events (just the spread from Fermi motion), whilst there is a substantial (and different) bias stemming from the presence of 2p-2h and other (mostly RES) events.

Since the input simulation includes the impact of some non-CCQE events being selected, this bias is accounted for in the propagation of the oscillation from E_{true} to the reconstructed observables. However, it is clear from figure 2.8 that if the input simulation has the wrong proportion of CCQE, 2p-2h and other events, or if the simulation incorrectly describes a particular interaction mode, the smearing will be described incorrectly. This would then bias the measurement of the oscillation parameters.

To account for these model dependencies, T2K introduces neutrino interaction nuisance parameters[1] (such as M_A and C_A^5) to the oscillation parameter fit. By including these parameters the fit is able account for the plausible variation in the models within the input simulation, whilst also allowing the near detector to constrain them. The fractional uncertainty on the number of expected events at the T2K far detector, following the near detector constraint, is shown in table 2.1. This illustrates that interaction uncertainties are already

[1]Alongside neutrino interaction nuisance parameters, T2K's oscillation analyses contain nuisance parameters to describe uncertainties in the neutrino flux and detector efficiency

the dominant systematic on current measurements of the oscillation parameters and will become the principle limitation if an improved understanding cannot be achieved.

	ν_μ selected events	ν_e selected events
Flux	3.6%	3.6%
Detector efficiency	3.9%	2.4%
Cross-section	4.1%	5.1%
Pion FSI	1.5%	2.5%
Total syst.	5.0%	5.4%
Statistical	8.6%	17.7%

Table 2.1: The fractional uncertainty on the number of expected events at the T2K far detector, split by error source (note that some of these uncertainties are highly correlated) [79].

Although the inclusion of nuisance parameters allows the oscillation analyses to accurately account for the uncertainty in a particular model, it does not necessarily account for the differences in predictions between models. Recent studies at T2K provided the oscillation analysis framework with fake data generated using the same oscillation parameters as the input simulation, but different neutrino interaction and nuclear models. Although it was found that the impact of changing models was not too severe considering current statistical errors, it was found that by the end of the proposed lifetime of T2K, there is significant potential for a large bias and/or unrealistic sensitivity in the extraction of the oscillation parameters [80]. In particular, it was found that altering the input CCQE/1p-1h and 2p-2h models could lead to an unrealistic increase in sensitivity for maximal δ_{cp}. This is demonstrated in the case where the input simulation is an implementation of the Nieves 2p-2h model [67] but the fake data was built using an implementation of the Martini 2p-2h model [68] in figure 2.9. This illustrates a much tighter exclusion of $\delta_{cp} = 0$ when using the Martini model. The extra sensitivity is not physical but instead comes from a bias to the input simulation that is not covered by the implemented nuisance parameters[2]. Overall this demonstrates the importance of understanding neutrino interactions in T2K-like oscillation analyses.

2.4.2 Calorimetric approach - NOνA

For experiments with more sensitivity to the hadronic component of neutrino interactions at the far detector, such as at the NOνA experiment, neutrino interactions can be characterised using the outgoing lepton kinematics in addition to calorimetric information [81, 82, 83]. Like T2K, NOνA is able to distinguish the outgoing lepton flavour (but not its charge) in charged-current neutrino interactions at the far detector and forms a sample of ν_μCC and

[2]T2K has recently added additional nuisance parameters to the oscillation fit to account for this.

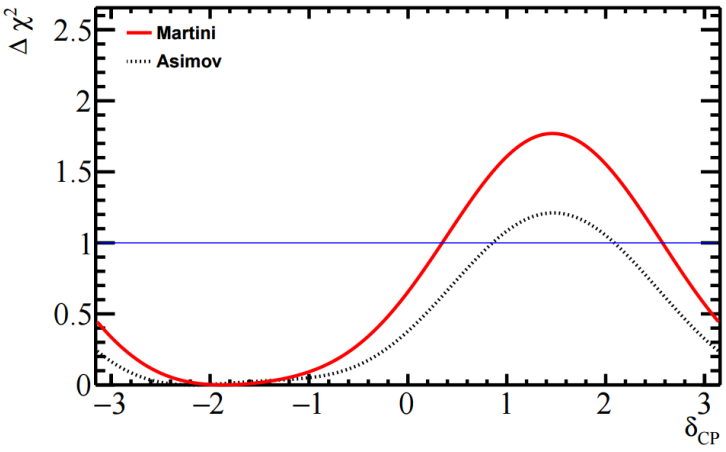


Figure 2.9: The resultant preference for values of δ_{cp} using an input simulation that was built using an implementation of the Nieves 2p-2h model [67] and fake data that was either itself (an 'Asimov' fit) or fake data that was built using an implementation of the Martini 2p-2h model ('Martini') [68]. Both fake data contained the same true value of the oscillation parameters, including $\delta_{cp} = -1.601$. Figure adapted from [80].

ν_eCC events. In each event in the CCν_μ sample the incoming reconstructed neutrino energy is found from:

$$E_{rec}^{\nu_\mu CC} = E_\mu + (E_{calo} - E_\mu) = E_\mu + E_{had}, \tag{2.19}$$

where $E_{rec}^{\nu_\mu CC}$ is the reconstructed neutrino energy, E_μ is the reconstructed muon energy and E_{had} is the estimated energy of the hadronic component of the interaction, formed from the sum of all energy deposits, E_{calo}. For events in the ν_eCC selection the reconstructed neutrino energy is just E_{calo}. Similarly to T2K, NOνA oscillation analyses then fit the reconstructed neutrino energy spectrum in order to obtain oscillation parameters, using a smearing matrix to go from the the E_{true} spectrum, that the oscillations can be applied in, to $E_{rec}^{\nu_\mu CC}$.

NOνA's method of neutrino energy reconstruction does not rely on assuming events are CCQE, like T2K's, but does have its own difficulties. NOνA's method is still susceptible to the event by event boost imparted on each interaction due to Fermi motion and the model dependence this entails (detailed in section 2.4.1), although NOνA's larger beam energy - around 2GeV - makes this a smaller effect, but more importantly NOνA has very limited sensitivity to energy deposited by neutrons. This makes the smearing from true to reconstructed neutrino energy heavily dependent on the model of neutron multiplicity and energy, which itself depends on, among other things, the poorly understood FSI and 2p-2h contributions to neutrino interactions.

The difficulties in accurately modelling the hadronic energy are manifest in the poor agreement that NOνA observed between data and simulation for their first muon-neutrino disappearance result [82], as illustrated in figure 2.10. The subsequent uncertainty that is

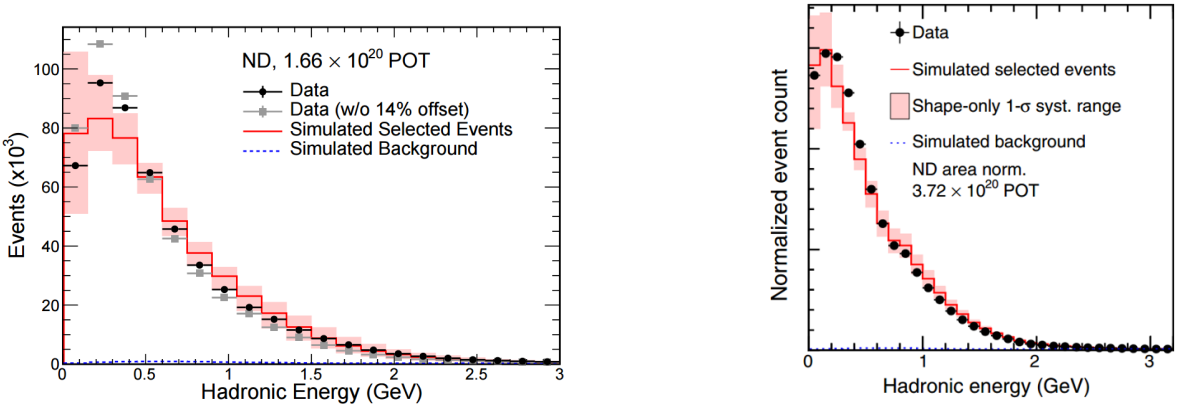


Figure 2.10: The reconstructed hadronic energy (E_{hadd}) for selected ν_μCC interactions in NOνA's near detector (which is a smaller version of the far detector) is shown alongside the input simulation prediction (from the GENIE neutrino-nucleus interaction simulation [41]) for (left) NOνA's 2016 muon-neutrino disappearance result [82] (with and without a normalisation offset to achieve better agreement), and for (right) NOνA's latest result [83]. Figure adapted from [82] and [83].

applied to cover the disagreement becomes NOνA's largest systematic error in their measurement of $sin^2(\theta_{23})$ and Δm^2_{32}. It should be noted that NOνA's latest disappearance result (also shown in figure 2.10) seems to observe better data-simulation agreement in hadronic energy, but that the hadronic energy uncertainty remains the largest systematic [83]. Overall, just as for T2K, NOνA's limiting systematic stems from neutrino interaction uncertainties.

2.4.3 Neutrino interactions in oscillation analysis summary

Sections 2.4.1 and 2.4.2 demonstrate that whichever method of characterising neutrino interactions is used, there is substantial scope for neutrino interaction uncertainties to limit experimental sensitivity to the oscillation parameters, or worse to bias the results or give an unrealistic sensitivity. Specifically, section 2.4.1 demonstrates that T2K's method is susceptible to models of the Fermi motion and the relative contributions on CC-non-QE events, but is not effected much by FSI. On the other hand, section 2.4.2 shows how NOνA's method does not directly rely on the correct modelling of the relative interaction mode contributions, but instead leans on the input simulation predictions of 2p-2h and FSI.

In either case it seems that, in order to make precision measurements of neutrino oscillations, a detailed understanding of both GeV-scale neutrino-nucleus scattering cross sections and nuclear effects are essential. This also highlights the importance and complementarity of having measurements from T2K and NOνA, since the experiments are susceptible to quite different potential model dependencies.

2.5 Making ν_μ CCQE-like cross-section measurements

The focus of this thesis is on the understanding of nuclear effects through a measurement of a flux-integrated fiducial differential cross section of ν_μ CC0$\pi + Np$ ($N > 0$) interactions: charge-current ν_μ scatters with one muon, no pions and at least one proton in the final state. This section aims to define exactly what this is and why it is a useful cross section to measure. In particular this section will discuss the merits of measuring a flux-integrated ν_μ CC0$\pi +$ Np topology cross section, over a (conceptually easier to understand) flux-unfolded CCQE interaction mode cross section. Additionally, previous relevant cross-section measurements will be discussed to overview the current state of the neutrino-nucleus scattering field.

2.5.1 Flux-integrated and flux-unfolded differential cross sections

In neutrino interaction differential cross-section measurements, no attempt is usually made to 'correct' for the incoming neutrino flux in each bin of the measurement. Instead, each bin is normalised by the total integrated flux, defining the 'flux-integrated' (sometimes called 'flux-averaged') cross section as:

$$\left(\frac{d\sigma_{flux\,integrated}}{dx}\right)_i = \frac{N_i^{signal}}{\epsilon_{i,sim}^{signal}\Phi N_{targets}^{FV}} \times \frac{1}{\Delta x_i}, \tag{2.20}$$

where i is the bin index, x_i is the differential variable, Δx_i is the bin width, Φ is the integrated flux, $N_{targets}^{FV}$ is the number of targets in the fiducial volume of the interaction target, N_i^{signal} is the number of signal events in the i^{th} bin of the differential variable and $\epsilon_{i,sim}^{signal}$ is the efficiency from the input simulation, computed in bins of true signal kinematics.

Because Φ is the integrated neutrino flux impinging on the experiment making the measurement, the result is experiment dependent. Indeed, when flux-integrated results are released they always include the shape of the incoming neutrino flux such that a model can be convoluted with the flux to be comparable to the result.

If differential measurements are to be compared between experiments directly, they must be measured as a 'flux-unfolded' cross section, which would need to be measured as a function of both the observable of interest (x_i) and the true neutrino energy. In general the true neutrino energy is not accessible from experimental observables and so forming a flux-unfolded cross section necessarily relies heavily on the input simulation to form, which may obscure any interesting physics content in the result. For this reason most modern measurements of differential neutrino scattering cross sections (including the one in this thesis) are flux-integrated.

2.5.2 Challenges in characterising CCQE interactions

Directly identifying CCQE interactions on a nuclear target is not practically possible, even with an ideal detector and event selection. This is because the nuclear effects discussed in section 2.2 distort any ability to differentiate between interaction modes on an event by event basis. For example: charged-current resonant production with pion absorption, via an FSI inside the nucleus, is effectively indistinguishable from a charged-current quasi-elastic scatter. Therefore, any attempt to directly measure CCQE interactions on a heavy nuclear target is likely to rely excessively on the (poorly understood) models of nuclear effects inside an interaction simulation, effectively regurgitating elements of said simulation and obscuring the interesting physics. For this reason, most modern experiments designed to measure neutrino interactions will mostly no-longer attempt to report CCQE measurements and will instead favour measurements of the CC0π (often called CCQE-like) topology.

In a true CCQE interaction off a stationary nucleon, the entire interaction can be fully characterised by the kinematics of the outgoing lepton (for example, the neutrino energy can be accurately calculated using equation 2.19). Indeed, previous measurements have attempted to measure unfolded (i.e. truth-level) differential cross sections of CCQE interactions as a function of a variable that is not directly accessible inside the detector, but can be calculated knowing the interaction is CCQE and using the assumption that the target nucleon is stationary (typically true neutrino energy E_ν or true four momentum transfer Q^2). However, this is not the case for CC0π measurements. Attempting to measure a CC0π cross section as a function of such variables would result in a smearing from reconstructed quantity to true quantity that cannot even by approximated without utilising models of poorly understood effects in an interaction simulation (effectively suffering the same model dependence as described in section 2.4.1). This can lead to significant propagation of model assumptions into the measured cross section. To avoid this, most modern experiments measure differential cross sections in quantities that are directly accessible in the detector. This includes final state particle kinematics but also observables known as $E_{\nu,QE}$ and Q^2_{QE}, where the differential cross section is measured as a function of the effective neutrino energy or four momentum transfer, as calculated using the true lepton kinematics of CC0π events when assuming a CCQE interaction.

A significant additional challenge in cross-section measurements stems from the limited acceptance of the detectors that are used. No currently operational detectors are able to see all low momentum outgoing particles and few have sensitivity to all angles. This creates another way for model dependence to enter analyses if the variables measured in a differential cross section do not well describe the detector acceptance. For example, consider an experiment that attempts to measure an unfolded single-differential ν_μCC0π cross-section in outgoing muon angle, but is not able to see a significant number of muons below some

momentum threshold. These 'out of acceptance' (OOA) low momentum muons are spread over the angular phase space in a way that the detector has no sensitivity to. Therefore when the angular distribution is unfolded to extract a truth-level cross section, there is significant dependence on the distribution of OOA events, which will have to be inferred from a model. There are a few different options to avoid such model dependence. The experiment could instead measure a multi-differential cross section in variables that can well describe the acceptance, in the case of a ν_μCC0π measurement a double-differential cross-section in muon momentum and angle is likely sufficient. Alternatively the experiment could opt to measure a 'fiducial' single differential cross-section in which the muon momentum is restricted to the phase space the detector is sensitive to.

2.5.3 Previous CCQE and CCQE-like cross-section measurements

Early measurements of the CCQE cross section (in the 1970s to early 1990s) came from hydrogen or deuterium bubble chambers. Following these measurements, the CCQE cross section could be accurately and consistently described using the theory outlined in section 2.1, with a value of $M_A \sim 1.0$ GeV [39].

However, more modern experiments use complex nuclei as a target, making nuclear effects much more important. The resultant modern picture that has built up of CCQE interactions on heavier nuclear targets from few-GeV neutrino beams has become quite complicated. In 2006 the K2K experiment fit M_A to reconstructed muon kinematics from interactions impinging on a water target, that left two or less tracks in the detector, to find $M_A = 1.2 \pm 0.12$ GeV [84]. Following this, in 2009, the NOMAD experiment, which has a carbon target, measured a total CCQE cross section and extracted a value of $M_A = 1.05 \pm 0.05$, consistent with the bubble chamber measurements [85]. However, in 2010, MiniBooNE measured a double differential CCQE cross section in muon kinetic energy and angle, alongside a single differential cross-section measured in Q^2_{QE} to find $M_A = 1.35 \pm 0.17$ GeV [71]. As outlined in section 2.1, this discrepancy became known as the MiniBooNE M_A puzzle, and, as discussed in section 2.2.3, the current thinking is that this arises from MiniBooNE not having included a 2p-2h contribution when extracting M_A, thereby causing it to act as an effective parameter which describes both CCQE and 2p-2h interactions simultaneously. This was demonstrated in figure 2.6, which illustrates a good fit to the MiniBooNE data without inflating M_A once the Nieves model of 2p-2h is included. This data, and more, is summarised in figure 2.11. It should be noted that MiniBooNE also found a larger value of M_A in their $\overline{\nu}_\mu$ when fitting a double differential distribution with an RFG spectral function, which is not obvious from figure 2.11 [86].

Some currently running experiments have also attempted to measure a CCQE cross-section. In 2014 the MINERνA experiment measured Q^2_{QE} to find a result consistent with

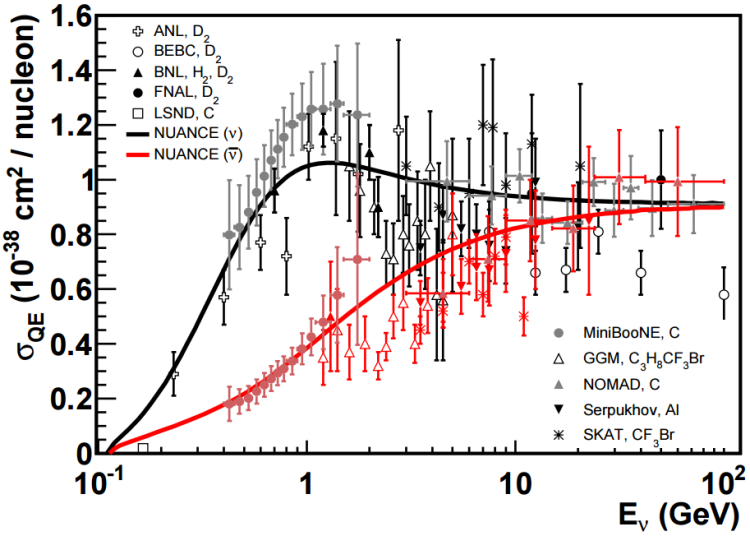


Figure 2.11: Measurements of the ν_μ (black) and $\overline{\nu}_\mu$ CCQE scattering cross sections per nucleon as a function of neutrino energy. The data is from experiments that utilise a variety of neutrino targets, which target is shown in the legend, from early bubble chamber experiments (refs [88, 89, 90, 91, 92]) to more recent experiments that use heavier nuclear targets (refs [71, 85, 86, 93, 94, 95, 96, 97]). The MiniBooNE M_A puzzle is apparent in the data. The data are shown alongside the predicted cross-section from the NUANCE simulation. Figure taken from [87].

$M_A \sim 1.0$ GeV, so long as an effective multi-nucleon correlation model (the transverse enhancement model - TEM [69]) was included [98]. MINERνA also noted an increase in energy deposit around the vertex relative to what was in their input simulation (which does not include any multi-nucleon correlations). This is consistent with the ejection of additional protons suggested by 2p-2h models (see section 2.2.3). In 2015 T2K also released a CCQE cross section from their ND280 off-axis near detector [99]. In this result there was explicitly no attempt to account for 2p-2h processes and M_A was extracted to be $1.26^{+0.21}_{-0.18}$ GeV ($1.43^{+0.28}_{-0.22}$ GeV) using the absolute measured (shape-only) 2D muon momentum and angle distribution. Soon after, T2K released a second CCQE result, this time using their on-axis INGRID detector [100]. This result extracted a separate cross section depending on whether one or two tracks were identified, a combined cross section was also shown. The result found better agreement with $M_A \sim 1.0$ GeV for the two track cross section and better agreement with an enlarged M_A for the one track and combined cross section. However, the result also explicitly showed that the measured cross section differed by 5%-10% when the spectral function in the input simulation was changed. Overall the INGRID CCQE result demonstrates the crux of the confusion over CCQE cross section measurements: the potential sensitivity of results to the input simulation and the cross section extraction method.

With the exception of the bubble chamber experiments (since measurements on hydrogen/deuterium are subject to no/little nuclear effects), all of the cross-section measurements

discussed so far suffer from the first issue discussed in section 2.5.2 : each attempts to correct for nuclear effects to access a signal the experiment is not directly sensitive to. Additionally, most measurements of differential cross sections in variables that do not characterise the detector acceptance (i.e. any single differential measurement or measurements of Q^2_{QE} or E_ν) make no attempt to restrict the phase space to what the detector is actually sensitive to (i.e. these measurements are full phase space rather than fiducial - this issue is discussed in section 2.5.2). Overall this illustrates that there is some, potentially large, model bias in each of the results. Indeed, recent attempts to simultaneously fit multiple CCQE results shows that there are large tensions between them, making the extraction of useful constraints or model preferences for oscillation analyses very difficult [101].

As discussed in section 2.5.2, to attempt to remedy this confusing situation, experiments are now measuring cross sections of interaction topologies rather than interaction modes directly (e.g. measuring CC0π rather than CCQE). Both T2K and MINERνA have recently published such topology-based measurements (it should also be noted that MiniBooNE provide the tools to extract a CC0π cross section from their CCQE result). In 2015 MINERνA published a measurement of a fiducial CC0$\pi + Np$ cross-section which showed a (relatively weak) preference for $M_A \sim 1.0$ GeV and no 2p-2h when using an RFG spectral function [102][3]. However, preliminary new data from MINERνA, which takes into account a revaluation of the MINERνA flux prediction [103] seems to prefer a larger 2p-2h contribution [104, 105]. In 2016 T2K published a double differential CC0π cross section in outgoing muon momentum and angle using ND280 [106]. Within this publication the CC0π result was extracted twice, once using a selection focussed on only forward going muons and once using a selection with some acceptance over the entire angular phase space. In the kinematic analysis bins that overlapped, excellent agreement between the results was shown, thereby helping validate the cross section extraction methods. The results were compared to predictions directly from the Nieves [67] and Martini [68] models, both with and without a 2p-2h contribution, and indicated a preference for requiring 2p-2h. However, it is not clear that agreement could not also be achieved by inflating M_A or altering RPA in the CCQE component of the result.

Although the T2K and MINERνA topology based results are an important step forward in removing the model dependence inherent in previous characterisations of CCQE interactions, it remains difficult to use them to draw firm physics conclusions that are useful to oscillation analyses. The primary reason for this is that the kinematic predictions of any theory compared to the measurements is subject to plausible alterations from many cross section and nuclear effects. Since the observables measured often offer limited separation

[3]Interestingly, of all the models fit, the preferred TEM model in MINERνA's CCQE analysis [98] was the most disfavoured. This further demonstrates the potential for the cross-section extraction method to alter the physics conclusions of a result.

power between these different effects, any attempt to extract information about a particular physics process largely degenerate with others. For example, an increased 2p-2h contribution can often be substituted with an increased M_A, less RPA suppression, or, in MINERνA's result that is sensitive to proton kinematics, an alteration to nucleon FSI. Therefore complementary measurements are required to facilitate further model separation and parameter determination, ideally using observables that can isolate different cross-section and nuclear effects.

One attempt to do this was recently published in 2016 by the MINERνA collaboration [107]. MINERνA measure a 2D distribution of the 'available energy' (a less model dependent proxy for the energy transfer from the neutrino to the hadronic system, obtained from the measured calorimetric energy deposited) and the three momentum transfer, q_3, of all charged-current interactions. The motivation for this is that it is expected that both RPA and 2p-2h have separable characteristic signatures in these kinematic variables. The results are shown as a CC-inclusive double differential cross section in figure 2.12, compared to the GENIE neutrino interaction simulation with an RFG spectral function. Taken at face value the results show that the data is in overall poor agreement with the input simulation but that the effect of RPA is well isolated and clearly required for improved agreement in the first two bins of q_3. However, the determination of q_3 relies on the input simulation to account for the impact of neutrons, which MINERνA cannot reliably detect. This was partially addressed when this result was presented at the NuInt 2015 conference, as it was demonstrated that the need for RPA in the early q_3 bins was present at the reconstructed level (where the data is not impacted by the input simulation in any way) [108]. It is acknowledged that the simulation is in too poor agreement with the data to make strong statements about the 2p-2h contribution. Additionally the potential to account for discrepancies by altering the spectral function used is not shown. Overall this analysis shows strong evidence that, when using an RFG spectral function, an RPA-like suppression of the CCQE contribution is required.

2.6 Summary

Neutrino oscillation analyses requires an accurate identification of the incoming neutrino energies. For $\mathcal{O}$(GeV) neutrino beams this is difficult to achieve since, with current neutrino energy reconstruction techniques, each reaction mode may be reconstructed with some bias which is, on all nuclei other than hydrogen, further obfuscated by various nuclear effects (such as final state interactions, Fermi motion and 2p-2h processes), which are poorly understood. For this reason, the accuracy of neutrino oscillation analyses are limited by the understanding of both reaction mode cross sections and the role of nuclear effects. This problem is made more complicated since nuclear effects distort any ability to differentiate between interaction

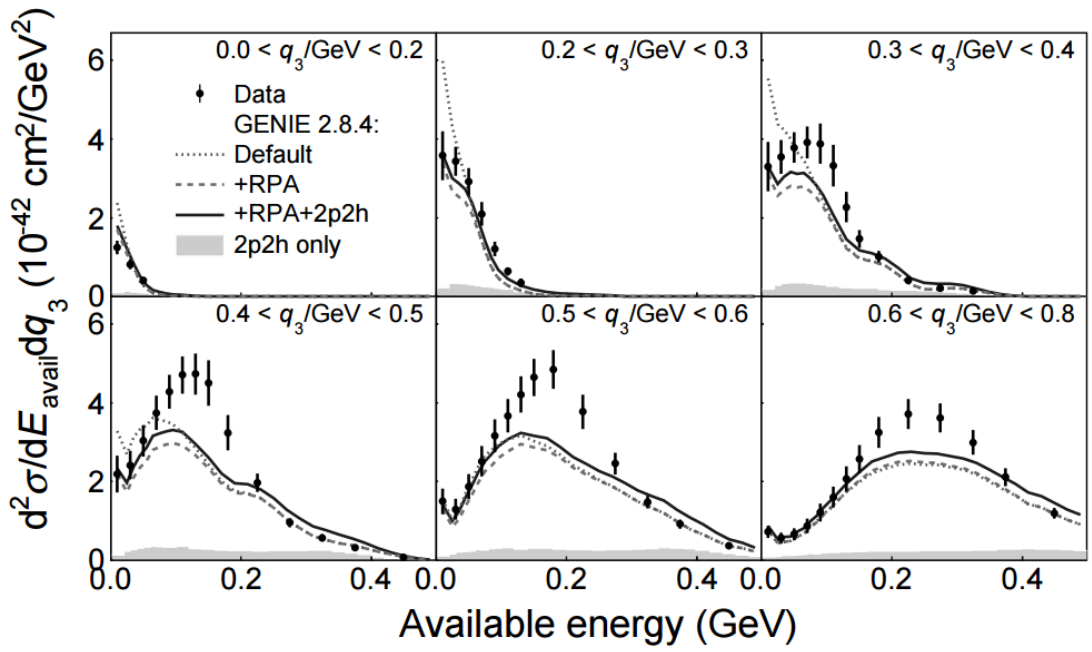


Figure 2.12: The double differential cross section measured in available energy (E_{avail}) across six regions of q_3 measured by MINERνA. The result is compared to the predictions of the GENIE neutrino interaction simulation with and without RPA and where the 2p-2h contribution is shown separately. Figure taken from [107].

modes (e.g. charged-current resonant production with pion absorption inside the nucleus is effectively indistinguishable from CCQE scattering). Current experimental efforts have been focused on lepton-kinematics based variables, which are largely degenerate with alterations to various nuclear and cross-section effects, making the role of interaction mode, nuclear effects and neutrino energy very difficult to untangle.

The highly convoluted nature of this problem is clearly demonstrated by the contention surrounding the MiniBooNE experiment's CCQE cross-section measurement which reported an excess over what was expected [71]. Assuming a sound measurement, this could simply be explained by a larger CCQE cross section (parametrised by a larger M_A) or alternatively by MiniBooNE having not fully considered correlations between nucleons. The degeneracy in the solution is the root of the problem.

MINERνA's CC-inclusive result, shown in figure 2.12, represents an interesting attempt to disentangle various cross-section freedoms and nuclear effects using observables that are specifically sensitive to different effects in different regions of their distributions. The analysis presented in this thesis aims to complement this measurement by using other observables that offer similar isolation of nuclear effects.

Chapter 3

Single-transverse kinematic imbalance (STKI)

Chapter 2 has demonstrated that current measurements of CCQE interactions and the CC0π topology struggle to untangle the various different nuclear effects from each other and from plausible alterations to the CCQE cross section. This degeneracy at least partially stems from the fact that measuring only the lepton kinematics in a CC0π interaction utilises only part of the final state. If both the outgoing lepton and hadron's momenta are accounted for then, in the absence of nuclear effects, some balance between the lepton and hadron kinematics is expected (i.e. one should be calculable from the other). Any imbalance observed for CCQE interactions is therefore a direct consequence of nuclear effects.

In this thesis a novel set of observables are used to directly probe such an imbalance in measurements of CC0π+Np interactions. These observables exploit the 'single transverse kinematic imbalance' (STKI)[1] between the outgoing lepton and highest momentum proton in the plane transverse to the incoming neutrino. The motivation for using STKI to study nuclear effects was summarised in a 2016 publication which demonstrated it to act as a very powerful probe of nuclear model and FSI in a full phase space measurement of an exclusive interaction mode [3], but - prior to the analysis contained in this thesis - STKI had yet to be measured beyond the reconstructed-level (where it is very difficult to use for model discrimination).

The purpose of this chapter is to first overview the observables that characterise STKI, in section 3.1, before showing that their ability to probe nuclear effects extends to measures of a CC0π+Np topology in a restricted kinematic phase space that is realistically accessible to the ND280 off-axis near detector of the T2K experiment (as is used in this thesis), in section 3.2. Additionally, section 3.2 also builds on the studies in [3] by including an investigation into the extent to which STKI can probe 2p-2h and RPA, and whether any sensitivity is degenerate with plausible alterations to FSI and the spectral function. The

[1] the 'single' part of the name distinguishes STKI from double transverse kinematic imbalance, presented in [109].

chapter concludes by discussing the prospects for using STKI to probe nuclear effects using ND280.

3.1 The single transverse variables (STV)

Without nuclear effects in a CC0π+1p topology (which therefore contains only CCQE interactions), the momentum of the outgoing lepton, projected into the plane transverse to an incoming neutrino, is exactly equal and opposite to that of the proton. However, once nuclear effects and/or other topologies are considered, the transverse momenta are no longer balanced, this is STKI. This imbalance can be fully characterised, by a set of three observables - the 'single transverse variables' (STV). The construction of these variables, from the kinematics of the outgoing lepton and proton, is shown schematically in figure 3.1 and discussed further in [3]. Quantitatively, these variables are defined as:

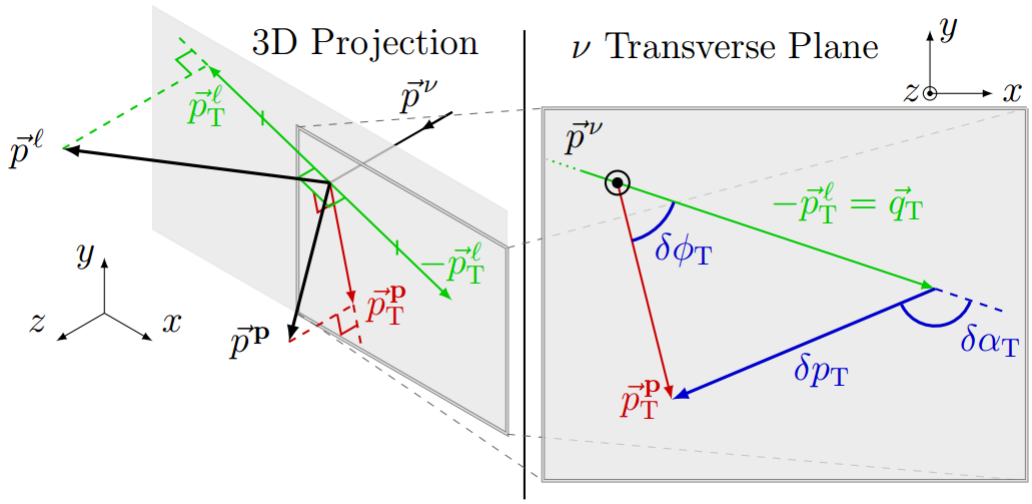


Figure 3.1: An illustration of a CC0π+1p interaction with nuclear effects to provide a schematic definition of the three single transverse variables, $\delta\phi_{\mathrm{T}}$, δp_{T} and $\delta\alpha_{\mathrm{T}}$. Here $\vec{p}^{\,\nu}$ is the incoming neutrino momentum whilst $\vec{p}^{\,\ell/\mathrm{P}}$ is outgoing lepton/proton momentum. The T subscript donates projection into the plane transverse to an incoming neutrino. Figure taken from [63].

$$\delta p_T \equiv |\vec{p}_T^{\,\ell'} + \vec{p}_T^{\,N'}|, \tag{3.1}$$

$$\delta\alpha_T \equiv \frac{-\vec{p}_T^{\,\ell'} \cdot \delta\vec{p}_T}{p_T^{\ell'}\delta p_T}, \tag{3.2}$$

$$\delta\phi_T \equiv \arccos\frac{-\vec{p}_T^{\,\ell'} \cdot \vec{p}_T^{\,N'}}{p_T^{\ell'}p_T^{N'}}. \tag{3.3}$$

δp_T is the overall three momentum imbalance in the plane transverse to the incoming neutrino. More transversely balanced interactions are characterised by small δp_T. In neutrino

scattering, δp_T has previously been utilised to investigate neutrino scattering by the Big European Bubble Chamber and NOMAD, both of these had significantly higher energy beams than T2K [110, 85].

$\delta\phi_T$ is the angular transverse imbalance measured as the angle between $-\vec{p}_T^{\,\ell'}$ and the direction of δp_T. Transverse balance is characterised by $\delta\phi_T = 0$. $\delta\phi_T$ has previously seen use as a selection variable by NOMAD and INGRID [85, 100], whilst MINERνA has measured the reconstructed (pre detector-effect corrected) coplanarity angle, $\varphi = \pi - \delta\phi_T$, in a CC0$\pi$ selection and found it to be in good agreement with a GENIE simulation [102]. Although $\delta\phi_T$ potentially contains less information about the transverse imbalance than δp_T, it only requires angular information to reconstruct and is therefore more accessible and may be reconstructed with a higher resolution in experiments.

$\delta\alpha_T$ can provide information regarding whether the effect causing the transverse imbalance is accelerating or decelerating the final state hadron. For constant $\delta\phi_T$, figure 3.1 shows that $\delta\alpha_T < 90^\circ$ corresponds to a boosted proton transverse momentum, whereas when $\delta\alpha_T > 90^\circ$ the proton transverse momentum has been reduced. In the absence of transverse imbalance $\delta\alpha_T$ is undefined. $\delta\alpha_T$ was proposed for the first time in [3] and has not, to the authors knowledge, been measured prior to the analysis presented within this thesis.

3.2 Investigating the impact of nuclear effects on STKI

It has previously been shown that the STV can act as excellent probes of the spectral function model and FSI in a full phase space measurement of an exclusive interaction mode [3, 63]. However, as discussed in section 2.5, exclusive interaction modes cannot be identified in a model independent way, and practical limitations of detectors limit the phase space that is accessible. This section aims to specialise the studies in [3, 63] to investigate the sensitivity to nuclear effects for a realistically measurable signal definition at ND280, whilst also extending the study to cover sensitivity to 2p-2h and RPA.

In this section the NuWro 11q [42] and "NEUT 6D" (defined in section 4.2.3 and further detailed in appendix A) neutrino interaction simulations are used to generate muon-neutrino interactions on a carbon target for various configurations of nuclear effects. An analysis of these events is then made to facilitate an investigation of the impact of different nuclear effects on the STV distributions formed from a CC0π+Np selection with the following constraints on the muon/leading proton momentum ($p_{\mu/\mathrm{p}}$) and angle ($\theta_{\mu/\mathrm{p}}$), such that they are measurable at ND280: $p_\mu > 250$ MeV/c, $p_\mathrm{p} > 450$ MeV/c, $cos(\theta_\mu) > -0.6$ and $cos(\theta_\mathrm{p}) > 0.4$. These phase space limitations are the same as will be used in subsequent chapters to define the signal definition for cross-section extraction. These are chosen by examining the proportion

of all signal events selected (selection efficiency) in each of the proton and muon kinematic variables and will discussed further in chapter 6.

3.2.1 Fermi motion

As discussed in section 2.2.1, the FM described by the spectral function imparts an isotropic event-by-event boost to the interacting system. To investigate the effect of different spectral functions, NuWro is used to generate phase space restricted CC0π+Np events using RFG, LFG and SF models, without any FSI or 2p-2h. Probability density functions (PDF) in the STV are built from these events and are shown in figure 3.2.

In the absence of other nuclear effects, δp_T is, by construction, the transverse projection of the FM. This gives δp_T a relatively unambiguous sensitivity to the spectral function model. By comparing the δp_T PDF in figure 3.2 to the PDF of the raw Fermi momentum in figure 2.3, it is clear that δp_T can distinguish the key differences between RFG spectral function and the others. In particular the Fermi cliff drop off at the Fermi momentum (here $p_f \approx 230$ MeV) is clearly identifiable in δp_T. There is also some limited separation between the LFG and SF models at the high and low extremes of δp_T. The former stems from the short range correlations (SRC) in the SF model that are not in LFG model but, subsequent sections will demonstrate that this long δp_T tail caused by SRC resemble the 2p-2h component, which is not surprising considering they are very similar physical processes.

As discussed in section 3.1, $\delta\alpha_T$ characterises the apparent acceleration/deceleration of the final state nucleon by nuclear effects. In the absence of other nuclear effects, any apparent acceleration comes from the shift of the lab frame caused by the FM. Since FM is isotropic, this gives a flat $\delta\alpha_T$ distribution, as seen in figure 3.2.

Figure 3.2 shows that $\delta\phi_T$ is rather less sensitive to the spectral function model. The larger typical values of the FM in an LFG and SF model seen in figure 2.3 correspond to a shallower peak and longer tail of the $\delta\phi_T$ PDF, but this is a weaker projection of the sensitivity contained in δp_T.

3.2.2 Final state interactions

FSI will change the kinematics of outgoing hadrons whilst potentially ejecting additional particles from the nucleus or stimulating pion emission, as detailed in section 2.2.2. To investigate the impact of FSI on the STV, NuWro is used to generate phase space restricted CC0π+Np events with and without FSI present using an LFG spectral function and no 2p-2h. Probability density functions in the STV are built from these events and are shown in figure 3.3, the inlays show close-up of the δp_T and $\delta\phi_T$ tails.

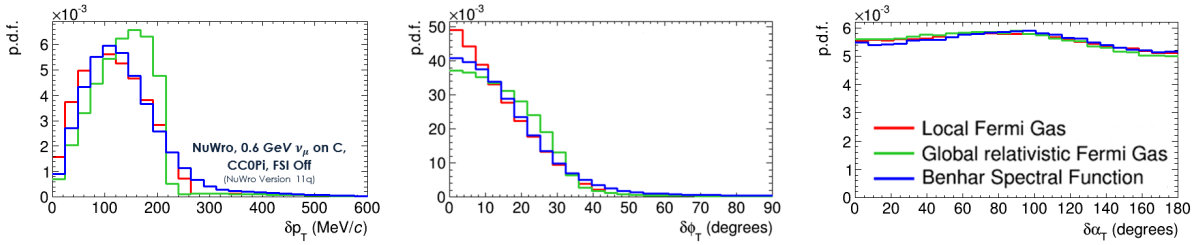


Figure 3.2: A NuWro 11q simulation of 0.6 GeV muon neutrinos on a carbon target is generated for three widely used spectral functions (global relativistic Fermi gas (RFG), local Fermi gas (LFG), and Benhar spectral function (SF) [59]). FSI effects are not simulated. Events containing one muon, no pions and at least one proton are selected, providing they pass the phase space limitations discussed in section 3.2. These events are used to build probability density functions in the STV.

FSI is expected to increase the degree of transverse imbalance in these events (the exception being when FSI coincidently counteracts the imbalance from another source). This is characterised by a large transverse momentum and angular imbalance, leading to the increase in the relative number of events in the tails of δp_T and $\delta\phi_T$, as shown in figure 3.3.

In general the addition of FSI causes final state protons to have a lower momentum, having exchanged some with the nucleus. This is manifest in the increase in the relative number of events seen at $\delta\alpha_T > 90°$ for exclusive interaction modes in figure 3.3.

Comparing the effect of FSI for restricted phase space CC0π+Np events in figure 3.3 to the full phase space CCQE events shown in figure 3.4, it is clear that there is a severe reduction in sensitivity. This is not too surprising since FSI typically reduces the momentum of final state nucleons substantially such that the harsh restriction of proton momentum in the phase space restrictions removes a large amount of the events that are characteristic of FSI. This demonstrates that improved experimental resolution to low momentum protons would help characterise nucleon FSI, as is expected to be provided by future liquid argon detectors.

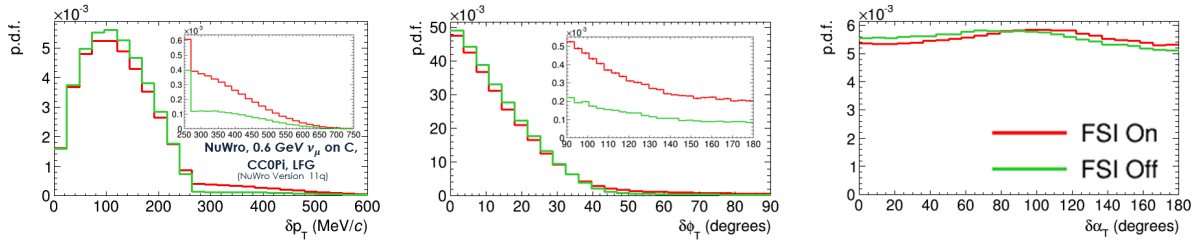


Figure 3.3: A NuWro 11q simulation of 0.6 GeV muon neutrinos on a carbon target using an LFG spectral function is generated with FSI turned on and off. Events containing one muon, no pions and at least one proton are selected, providing they pass the phase space limitations discussed in section 3.2. These events are used to build probability density functions in the STV. The inlays show a zoom and/or shift in the x axis to demonstrate that the δp_T and $\delta\phi_T$ tails remain sensitive to FSI.

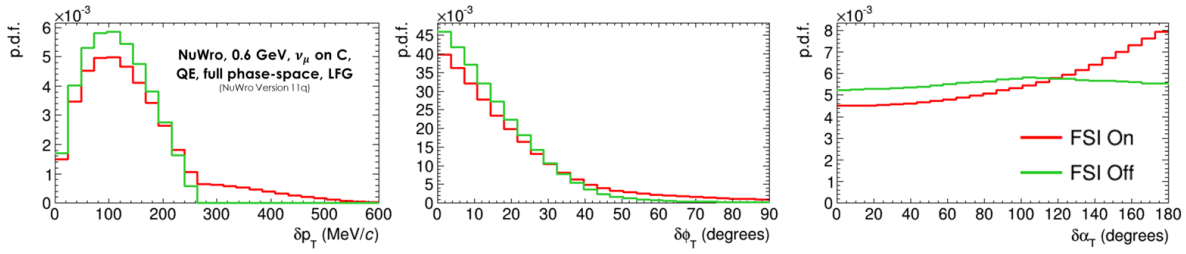


Figure 3.4: A NuWro 11q simulation of 0.6 GeV muon neutrinos on a carbon target using an LFG spectral function is generated with FSI turned on and off. Full phase space CCQE events are used to build probability density functions in the STV.

3.2.3 2p-2h nucleon-nucleon correlations

For 2p-2h interactions an interpretation of the STV is more challenging. Although the interaction is with a bound state of two nucleons (as discussed in 2.2.3), the STV are reconstructed using only one final state proton. Since now only part of the final state hadronic system is being considered, even without FSI and FM, there is always a transverse imbalance. To investigate the impact of 2p-2h on the STV, NuWro is used to generate phase space restricted CC0π+Np events using an LFG spectral function with a FSI contribution. Probability density functions in the STV are built from these events for the CCQE component and the 2p-2h component (separately) using two 2p-2h models available within NuWro. These are shown in figure 3.5.

As discussed in section 2.2.3, it is important to note that the 2p-2h models shown here are the NuWro implementation of the Nieves and Martini models and so are not fully representative of the theory described by the authors in [52, 53]. As such, the predictions shown in figure 3.5 should only be taken to be roughly indicative of the shape of a 2p-2h contribution. The specific model differences shown should not be assumed to be dependable.

Similarly to FSI, 2p-2h is expected to almost always increase transverse momentum imbalance, but with a smaller expected contribution at low δp_T, since in 2p-2h events there will always be a missed transverse momentum contribution from the second nucleon. This is observed in the δp_T PDF in figure 3.5.

A corresponding transverse angular imbalance leads to events spread over $\delta\phi_T$ such that 2p-2h may become a very prominent contribution in the tail, as shown in figure 3.5. Although it is very unlikely to find very small δp_T from a 2p-2h event, events in which the two outgoing nucleons are co-linear or back-to-back could give small $\delta\phi_T$.

Since, even before FSI, the proton used to form the STV has only a portion of the momentum transferred to the hadronic system (as the other correlated nucleon has the rest), it has undergone an apparent deceleration leading the 2p-2h contribution to be concentrated at high $\delta\alpha_T$, shown in figure 3.5.

So long as the shape of the 2p-2h contribution the STV in figure 3.5 is reasonably representative, as the arguments above suggest, it can be seen that, with reasonable constraints on FM, FSI and non-CCQE components of the selection it may be possible to extract a 2p-2h contribution from a measurement of the STV. Although the small FSI contribution shown in section 3.2.2 limits the sensitivity to FSI, since FSI and 2p-2h share the same regions of the STV phase space, it potentially increases sensitivity to 2p-2h.

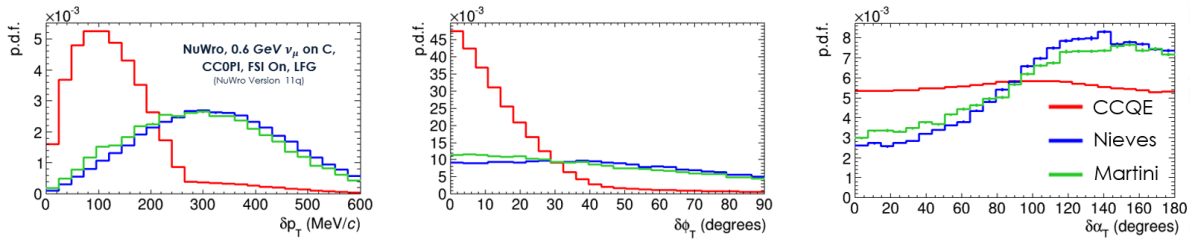


Figure 3.5: A NuWro 11q simulation of 0.6 GeV muon neutrinos on a carbon target with an LFG spectral function is generated with two widely used 2p-2h models [52, 53]. FSI effects are simulated. Events containing one muon, no pions and at least one proton are selected, providing they pass the phase space limitations discussed in section 3.2. These events are used to build probability density functions in the STV of the CCQE and 2p-2h contribution (using two different models) separately (i.e. there is no information about the relative normalisation of the CCQE and 2p-2h components in these plots).

3.2.4 The nucleon axial mass

To investigate the impact of altering M_A on the STV, NuWro is used to generate phase space restricted CC0π+Np events present using an LFG spectral function with an FSI contribution for different values of M_A. Probability density functions in the STV are built from these events and are shown in figure 3.6.

As discussed in chapter 2, for variables formed from only the outgoing lepton kinematics (as is typically used), changing M_A tends to alter the shape such that it is difficult to distinguish an altered M_A from a 2p-2h contribution in the data from an individual experiment. Conversely, figure 3.6 shows that the shape of the STV is almost entirely independent of M_A, thereby making nuclear effects much easier to distinguish. For example, a MiniBooNE-like excess in the measurement of CC0π events using the STV would not be able to be simultaneously explained by a larger M_A and multi nucleon correlations; the former would manifest only as a normalisation shift whilst the latter would exhibit an additional significant shape change.

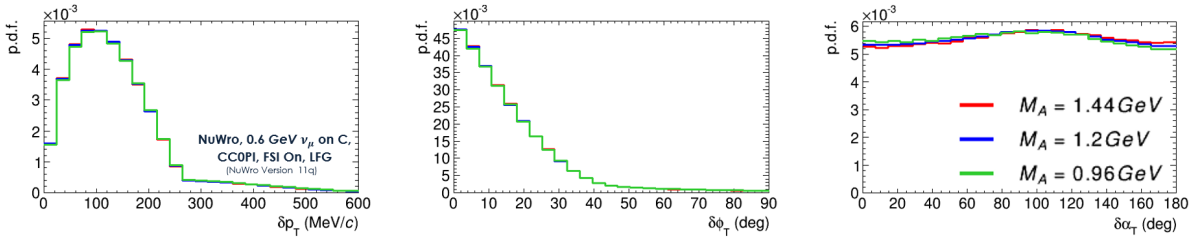


Figure 3.6: A NuWro 11q simulation of 0.6 GeV muon neutrinos on a carbon target with an LFG spectral function is generated with three different values of the nucleon axial mass (M_A). FSI effects are simulated. Events containing one muon, no pions and at least one proton are selected, providing they pass the phase space limitations discussed in section 3.2. These events are used to build probability density functions in the STV.

3.2.5 The random phase approximation

Section 2.2.3 illustrated that RPA can cause a Q^2 suppression or enhancement of the CCQE cross section, having a particularly large effect at low Q^2. To investigate the impact of RPA on the STV, NEUT is used[2] to first assess the Q^2 distribution of phase space restricted CC0π+Np events compared to those without such restrictions. This is shown in figure 3.7. Following this, NEUT is used to generate phase space restricted CC0π+Np events using an RFG spectral function with an FSI contribution with and without RPA. This is shown in figure 3.8.

Comparing the Q^2 distribution of CC0π+Np events, in figure 3.7, with the impact of RPA as a function of Q^2 shown in figure 2.5, it can be seen that a smaller impact from RPA is expected for the phase space restricted events. Figure 3.8 then demonstrates that RPA, as implemented in NEUT, has very little impact on the shape of the STV, with the possible exception of a small increase in the relative number of events at large $\delta\alpha_T$. Overall this shows that, within NEUT, RPA is expected to cause only a small normalisation shift across all STV. On the one hand this severely limits any sensitivity to RPA, but on the other allows further isolation of the nuclear effects the STV are more sensitive to.

Although this study to assess the impact of RPA on the STV is informative, it is important to note that its conclusions are very dependent on the specific RPA implementation within NEUT, which is tuned to allow good agreement with measurements of outgoing lepton kinematics and therefore may not be appropriate to describe the specifics of the outgoing hadronic kinematics encapsulated in the STV.

[2]NEUT is used to study RPA, as its implementation is more accessible than other available simulations.

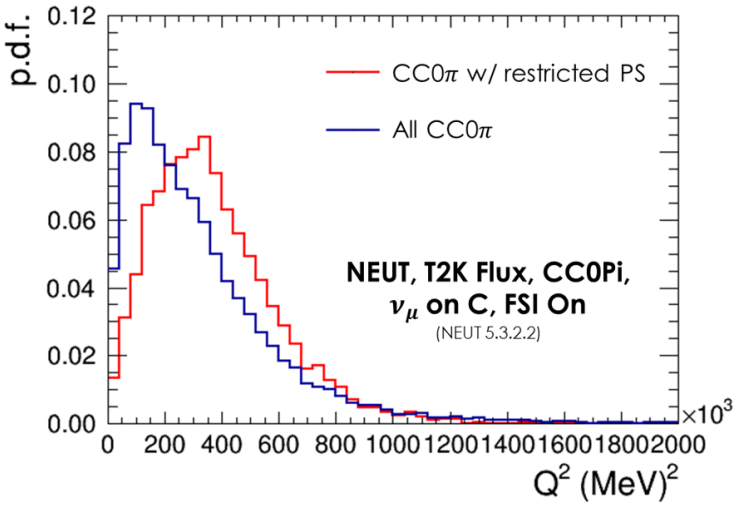


Figure 3.7: A NEUT 6D simulation of muon neutrinos generated with the T2K flux [26], on a carbon target with an RFG spectral function, is generated. FSI effects are simulated. Events containing one muon, no pions and at least one proton are selected. These events are used to build probability density functions in Q^2 to evaluate the relative expected impact of RPA for events all such events and only those that pass the phase space limitations discussed in section 3.2.

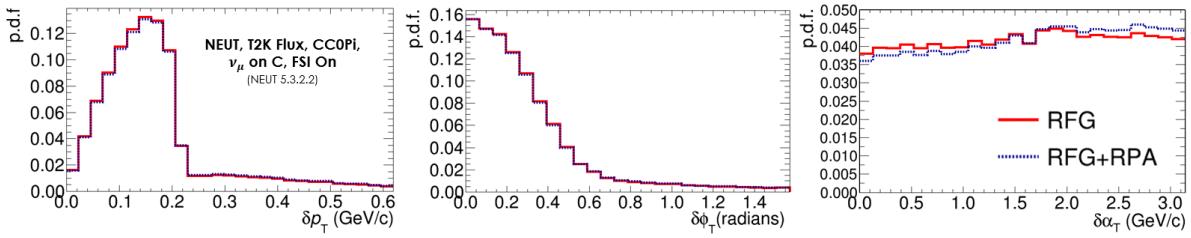


Figure 3.8: A NEUT 6D simulation of muon neutrinos generated with the T2K flux [26], on a carbon target with an RFG spectral function, is generated with and without RPA. FSI effects are simulated. Events containing one muon, no pions and at least one proton are selected, providing they pass the phase space limitations discussed in section 3.2. These events are used to build probability density functions in the STV.

3.2.6 Prospects of probing nuclear effects in CC0π+Np interactions at ND280 using STKI

The studies in section 3.2 demonstrate that, even in the restricted phase space easily accessible to ND280, STKI offers a powerful new probe of nuclear effects. The shape of the observables are independent of M_A and RPA, but are potentially sensitive to nucleon FSI, 2p-2h and FM, the latter of which is separated from the two former within the STV phase space. This may allow a firm resolution of the MiniBooNE M_A puzzle, whilst also facilitating an interesting partial isolation of FM from 2p-2h and FSI. This is complementary to existing measurements with enhanced sensitivity to RPA (MINERνA's CC-inclusive measurement [107]) or those that currently struggle to disentangle M_A and 2p-2h. Moreover,

since ND280's acceptance requires relatively high momentum protons, it is expected that the impact of nucleon FSI should be small and may potentially be suitably constrained from electron-scattering data to allow a separation from 2p-2h.

Overall the primary advantages of using STKI to probe nuclear effects is that the size and shape of different STV features correspond unambiguously to particular combinations of nuclear effects and the STV are directly measurable from reconstructed kinematics. Simple measured outgoing particle kinematics (such as muon or proton momentum and angle) are equally as accessible but require models to interpret, whilst other variables with similar nuclear effect isolation potential (e.g. Q^2 or q_3 and energy transfer) require inference from a neutrino interaction simulation to form from detector measurements.

Chapter 4

The T2K experiment

T2K [26] is a long-baseline neutrino oscillation experiment which aims to characterise the mixing of muon neutrinos with other neutrino flavours. T2K consists of a predominantly muon (or antimuon) neutrino beam produced at the J-PARC facility [111]; the INGRID and ND280 near detectors placed $\sim$ 280 m downstream from the target used to produce the beam; and the Super-Kamiokande (SK) water Cherenkov far detector located a further 295 km away. The layout of the T2K experiment is summarised in figure 4.1. Both SK and the ND280 near detector are placed 2.5° off-axis from the beam centre, allowing exposure to a relatively narrow band neutrino flux with a peak energy of 0.6 GeV. T2K first received beam in late 2009 and has been taking physics data since early 2010.

T2K's design goals were to measure the previously last unknown PMNS mixing angle (θ_{13}), alongside making precision measurements of Δm^2_{23} and $sin^2(\theta_{23})$, which have all been successfully completed, although accuracy continues to improve. T2K is also now moving beyond these goals and is placing constraints on δ_{CP}: T2K's most recent flagship result, published in January 2017 and updated in April 2017, is able to exclude the CP conservation hypothesis ($\delta_{CP} = 0, \pi$) at a 90% C.L for the first time (for any experiment) [112].

Alongside measurements of neutrino mixing angles and mass differences, T2K has made some world-leading neutrino-nucleus cross-section measurements, some of which have already been discussed in chapter 2. Beyond this T2K has also completed sterile neutrino searches [113] and made measurements of neutrino time of flight between ND280 and SK [114].

This section will overview the components of the T2K experiment, focussing on the tracking region of ND280 that is directly used in the subsequent chapters. This will begin with a brief discussion of the T2K neutrino beam production and the ensuing flux constraints in section 4.1. Following this, the T2K near detectors will be detailed in section 4.2, with focus on the off-axis ND280 near detector. Section 4.2 will also overview the different ways in which data is simulated at ND280. Finally, in section 4.3 T2K's currently available data will be discussed, alongside the scope for extra data availability in the near-term future. The T2K far detector, SK, will not be discussed further, since it is not an important component

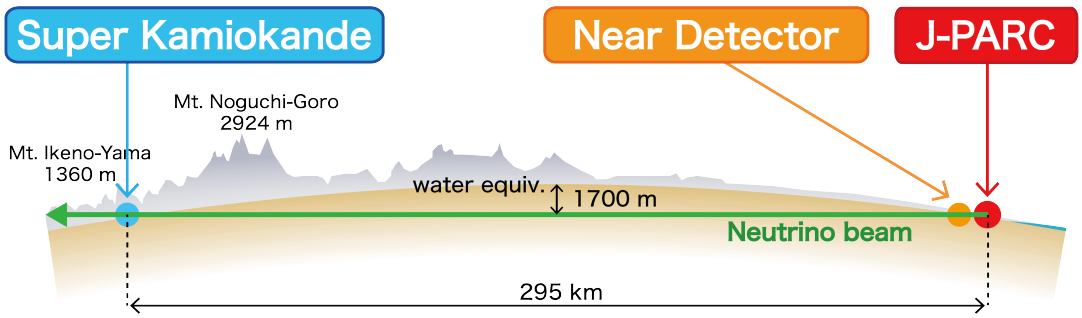


Figure 4.1: The T2K experiment.

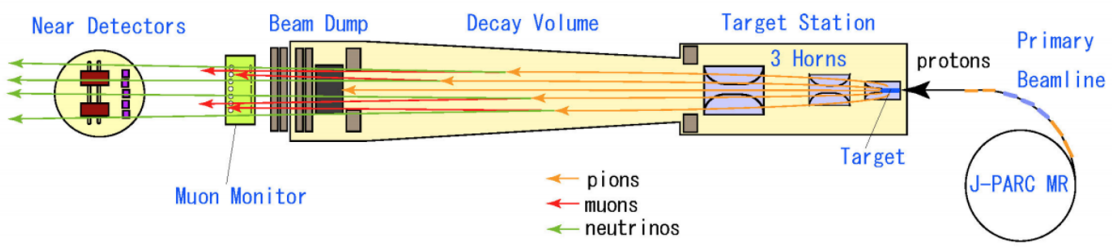


Figure 4.2: A schematic of the T2K neutrino beam production.

of the primary analysis in this thesis.

4.1 The T2K neutrino beam

The T2K neutrino beam [111] is produced at J-PARC in Tokai-mura on the East coast of Japan. It is designed to produce an intense quasi-monoenergetic beam of muon neutrinos or antimuon neutrinos, with minimal wrong-sign or electron neutrino contamination and a peak energy of ~ 0.6 GeV, such that there is maximal neutrino oscillations at SK. A schematic of the neutrino beam production is shown in figure 4.2. This section provides an overview of the production of the T2K neutrino beam; the off-axis approach to utilising neutrino beams and the prediction of the flux seen by ND280 and SK.

4.1.1 J-PARC proton accelerator and neutrino beamline

The J-PARC proton accelerator [116] consists of three components: a linear accelerator to accelerate a H^- beam to 400 MeV; a rapid cycling synchrotron, which begins with charge stripping foils to create a proton beam that is then accelerated to 3 GeV; and a main ring

(MR) synchrotron, which further accelerates some portion of the protons to 30 GeV. The neutrino beamline [26] begins by extracting 30 GeV protons from the J-PARC MR accelerator, before steering them towards SK and into a graphite target. This is called the primary neutrino beamline. The impinging protons stimulate the production of mesons (mostly pions and some kaons) in the target that are then, in the secondary neutrino beamline, charge selected and focussed into a decay volume by magnetic horns. By switching the polarity of the magnetic horns the charge of the mesons focussed by the horns is also reversed, thereby allowing the production of a predominantly neutrino or antineutrino beam. The neutrino or antineutrino beam configurations are often called 'forward horn current' (FHC) mode or 'reverse horn current mode' (RHC) mode respectively.

Whilst the focussed mesons are in the $\sim$ 96 m long decay volume, most decay to muon neutrinos (or antineutrinos) via the following processes:

$$\pi^+ \to \mu^+ + \nu_\mu \qquad \text{and} \qquad K^+ \to \mu^+ + \nu_\mu \qquad \text{(FHC)},$$

$$\pi^- \to \mu^- + \overline{\nu}_\mu \qquad \text{and} \qquad K^- \to \mu^- + \overline{\nu}_\mu \qquad \text{(RHC)}.$$

However, imperfect horn focussing allows some wrong-sign contamination and some of the subsequent muons also decay, alongside some of the kaons via another mode, to produce a small (anti)electron neutrino component to the beam:

$$\mu^+ \to e^+ + \nu_e + \overline{\nu}_\mu \qquad \text{and} \qquad K^+ \to \pi^0 + e^+ + \nu_e \qquad \text{(FHC)},$$

$$\mu^- \to e^- + \overline{\nu}_e + \nu_\mu \qquad \text{and} \qquad K^- \to \pi^0 + e^- + \overline{\nu}_e \qquad \text{(RHC)}.$$

The length of the decay volume is chosen to maximise the meson to muon neutrino conversion rate, whilst minimising the electron neutrino contamination in the beam. At the end of the decay volume there is a beam dump that stops all hadrons and muons below $\sim$ 5 GeV. Muons above 5 GeV that pass through the beam dump are able to enter the muon monitor (MUMON), which provides a real-time measure of the beam direction and intensity. The neutrinos pass through unimpeded and make up the T2K beam.

4.1.2 Off-axis approach

ND280 and SK are not placed in the centre of the neutrino beam, instead they sit 2.5° off-axis from it. This approach is to exploit the fact that, in the kinematics of the two-body pion decay used to produce the majority of the neutrinos in the beam, the outgoing neutrino energy depends only weakly on the parent pion momentum beyond some scattering angle dependent threshold. Quantitatively, considering energy and momentum conservation in two-body pion decay (in the lab frame), the outgoing neutrino energy (E_ν) can be written in terms of the pion energy (E_π) and momentum (p_π), the pion mass (m_π), the muon mass

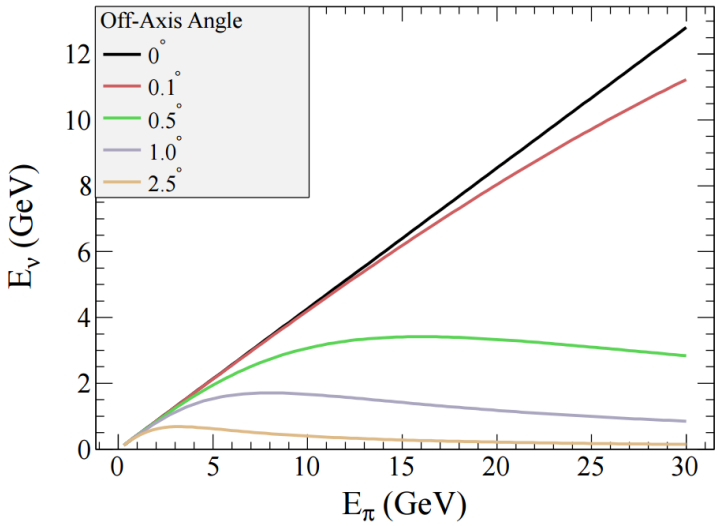


Figure 4.3: The neutrino energy from two-body pion decay as a function of the parent pion momentum for different off-axis angles (between the incoming pion and outgoing neutrino directions). T2K opts use a 2.5° off-axis angle. Taken from [115].

(m_μ), and the angle between the initial pion direction and outgoing neutrino direction (θ) as:

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - p_\pi cos(\theta))}. \tag{4.1}$$

Plotting $E_\nu(E_\pi)$ for various θ, as shown in figure 4.3, it can clearly be seen that there is a minimal dependence on E_π above some E_π, which depends on the scattering angle.

By utilising this off-axis effect, T2K is able to use a neutrino beam with a narrow spread of neutrino energies, but at the cost of a lower intensity beam. The beam configuration and off-axis angle is chosen to ensure a balance between intensity and beam spread whilst ensuring a peak energy of ~ 0.6 GeV, such that SK is at the oscillation maximum. This is illustrated in figure 4.4, which demonstrates the narrow neutrino energy at the oscillation maximum afforded by a 2.5° off-axis angle

4.1.3 Neutrino flux predictions

The T2K flux prediction is an essential part of the successful prediction of neutrino interaction rates at ND280 and SK, making it a key component of oscillation and cross-section analyses. The prediction of the neutrino flux begins at the end of the primary neutrino beamline, with a simulation of ~ 30 GeV protons impinging on a graphite target. A FLUKA2008 [117, 118] based simulation is used to model hadronic interactions in the target to produce kinematic predictions for the secondary mesons. These are then propagated through the secondary neutrino beamline using a GEANT3 [119] based simulation,

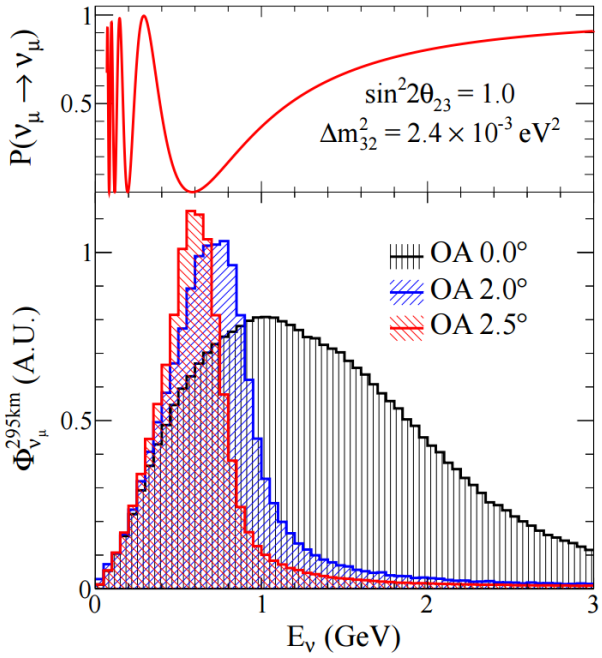


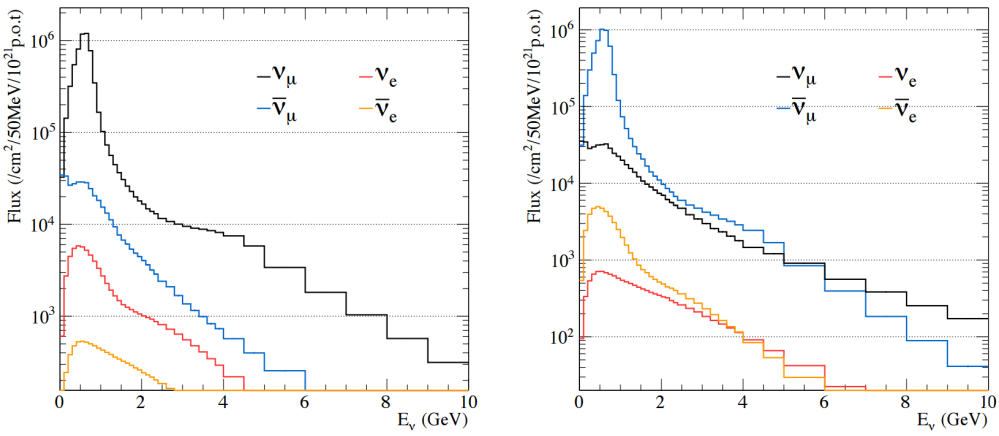
Figure 4.4: The (arbitrarily normalised) neutrino flux for different off-axis angles shown along side the muon neutrino oscillation probability to justify a peak neutrino beam energy of $\sim$ 0.6 GeV and an off-axis angle of 2.5°. Taken from [111].

thereby allowing a prediction of the subsequent neutrino flux. The latest T2K flux prediction, broken down by neutrino type, is shown in figure 4.5, demonstrating the very small beam contamination.

The simulation uses in-situ measurements of the primary proton beam, magnetic horn currents and the beam alignment and intensity (from MUMON and verified by INGRID) in order to tune the simulation. However, since it involves the simulation of multiple relatively low energy strong interactions, one of the most difficult aspects of the simulation is the hadronic interactions inside the graphite target, and this is challenging to constrain with measurements at J-PARC. To facilitate an improved understanding of such interactions, T2K makes use of the NA61/SHINE experiment at CERN.

NA61/SHINE is a dedicated hadron-production experiment at the CERN SPS which utilises a $\sim$30 GeV proton beam and a replica T2K target (alongside a thinner version of it) to make valuable measurements of the hadron production directly relevant to the T2K beam [120, 121]. The results are used to tune hadronic interaction models which then enables a tuning of the hadron production in the T2K flux simulation, providing valuable constraints on the flux prediction[1] [111]. To illustrate this, figure 4.6, shows the fractional uncertainty on the muon neutrino flux in FHC mode, broken down by uncertainty source, before and after the latest NA61/SHINE data is used to tune the simulation. It is important to note

[1]Currently only thin target results are used, analysis of the T2K replica target data is under way.



(a) FHC (b) RHC

Figure 4.5: The (unoscillated) off-axis neutrino flux prediction at ND280 in FHC and RHC mode, broken down by neutrino type. Taken from [122].

that, even prior to the latest tune, NA61/SHINE was still used to constrain the neutrino flux predictions, without this the flux uncertainties would be substantially larger than the previous ∼12% shown in figure 4.6.

4.2 ND280: The T2K near detector

The T2K near detectors are situated about 280 m away from the neutrino beam production point and are built with the primary purpose of constraining the neutrino interaction rate at SK to reduce systematic uncertainties in neutrino oscillation analyses. The near detector facility, shown in figure 4.7, contains both the INGRID on-axis detector and the ND280 off-axis detector.

INGRID [26] consists of a series of modules positioned to form a cross with the neutrino beam centre in the middle. Each module contains layers of scintillator and inactive iron, to allow charged particle tracking and ranging. Although INGRID's centre is on-axis, it covers a wide range of off-axis angles, thereby allowing it to fulfil its primary purpose as a monitor of the neutrino beam direction and profile. The INGRID modules can also be used to make neutrino cross-section measurements with a different but complementary flux than at ND280. In addition to the iron and scintillator modules shown in figure 4.7, INGRID also contains a 'proton module', situated in-between the central vertical and horizontal modules, which is a fully active polystyrene scintillator tracker, capable of high precision differential cross-section measurements.

The ND280 off-axis near detector is composed of a series of central specialised sub-

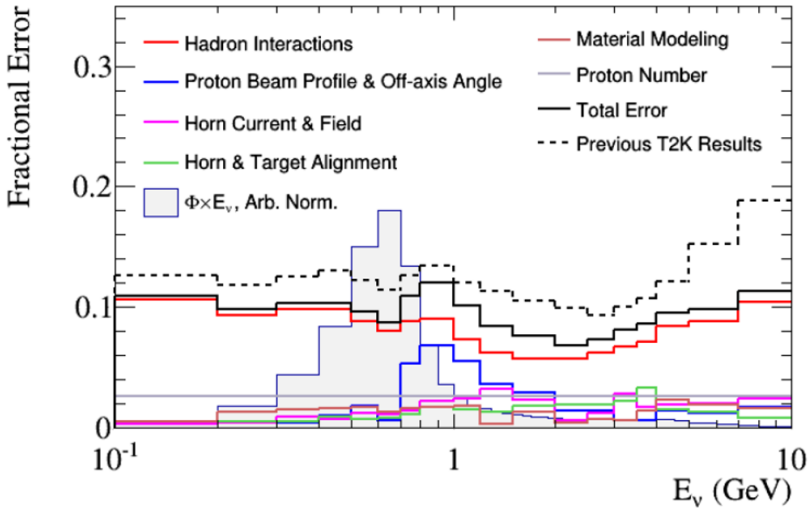


Figure 4.6: The fractional error on the muon neutrino flux in FHC mode, broken down by error source, is shown alongside the neutrino flux shape for comparison. The 'previous T2K results' dashed line refers to the flux prediction prior to the latest tune using the NA61/SHINE data (although this still included some constraints from NA61/SHINE). Adapted from [122].

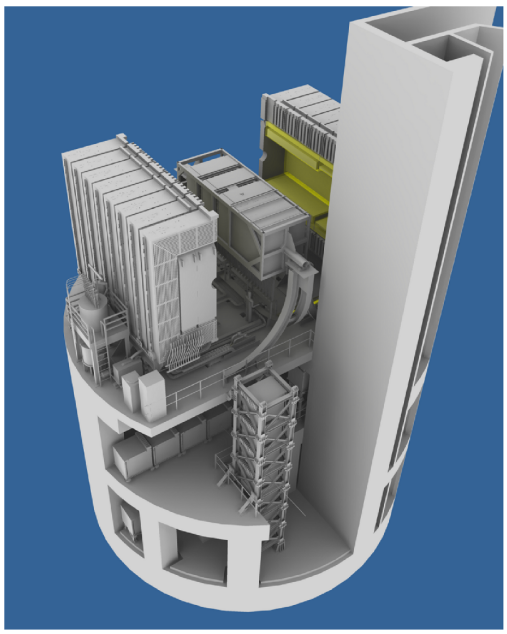


Figure 4.7: The T2K near detector complex. The top floor contains the ND280 off-axis near detector (with the magnet open to see the constituent sub-detectors), whilst the INGRID on-axis near detector is on the lower floor. The beam centre is within the area that the INGRID horizontal and vertical modules cross. Taken from [26].

detectors surrounded by an electromagnetic calorimeter (ECal), consisting of interleaved layers of lead and scintillator, which itself is contained within the (former UA1 experiment's) magnet, providing a 0.2 T dipole field [26]. The magnet is instrumented with the side range muon detector (SMRD), containing scintillator to tag high angle escaping muons from the central detectors or entering backgrounds. A schematic of ND280 is shown in figure 4.8.

The most upstream central sub-detector is the pi-zero detector (PØD), which consists of tracking planes of scintillating bars, interleaved with fillable water bags or lead/brass sheets. Downstream of the PØD is the tracker region, comprising of three time projection chambers (TPCs) and two fine grained detectors (FGD1 and FGD2). The FGDs are both made of finely segmented scintillating bars which provide a neutrino target mass and charged particle tracking, whilst FGD2 also contains inactive water layers. Events leaving the FGDs or the PØD can be tracked into the TPCs, which provide high-resolution tracking and thereby allow the curvature of charged particles to be used to make excellent measurements of charged particle momentum. As discussed in section 4.2.1, this can then be combined with a measurement of particle energy loss to produce a metric of charged particle identity.

As the principle components of ND280 used in the subsequent chapters, the TPCs and the most upstream FGD (FGD1) will be described in more detail in the following subsections alongside how events are simulated at ND280. To illustrate aspects of the tracker regions experimental resolution and reconstruction techniques, figure 4.9 shows an ND280 tracker event display.

4.2.1 Time projection chambers

ND280's gaseous argon TPCs [124] are designed to provide high-resolution tracking of charged particles to facilitate a measurement of their momentum and particle type. Each of ND280's three TPCs consists of an inner box containing an argon-based gas alongside a central cathode plane, allowing the TPC to be instrumented to set up a drifting electric field in the same orientation as the ND280 magnetic field. Readout planes are placed in the same orientation as the central cathode on the sides of the TPC. This is shown schematically in figure 4.10.

As charged particles pass through the gas in the TPC, ionisation electrons are produced which then drift away from the central cathode towards one of the two TPC readout planes in an electric field of around 280 V/cm. The maximum drift distance to a readout plane is approximately 90 cm. The charge and arrival time, which are readout by twelve 342 mm $\times$ 359 mm 'micromegas' modules [125] in each plane, are then combined to produce a 3D reconstructed image of the paths of the traversing charged particles. The micromegas charge-collection element ('pad') segmentation is 70 mm^2 (giving 1728 pads per micromegas

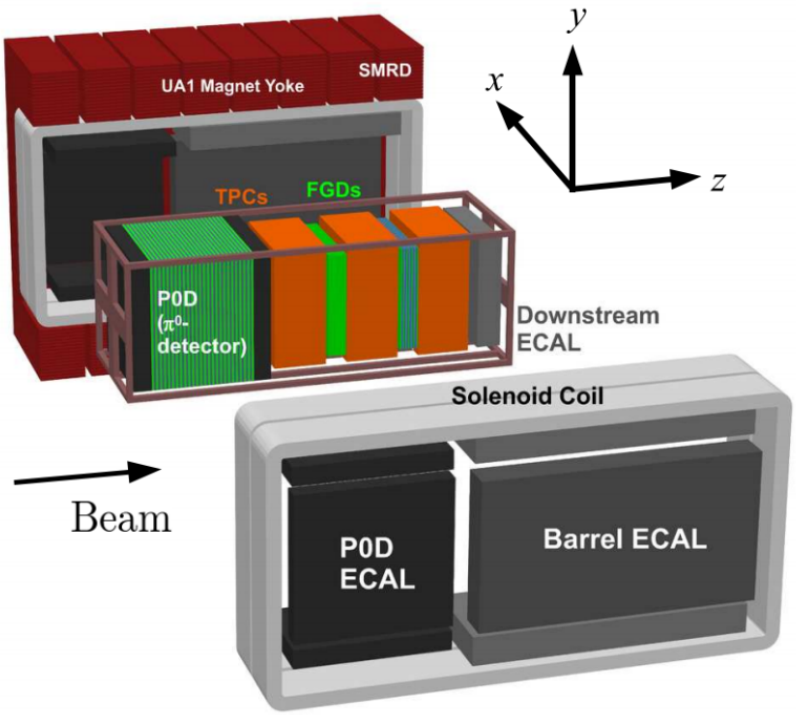


Figure 4.8: An exploded view of the ND280 off-axis near detector labelling each sub-detector. Adapted from [26].

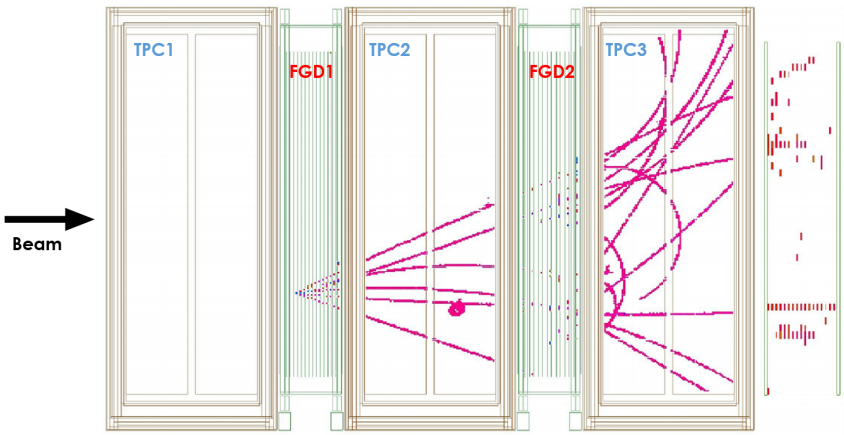


Figure 4.9: An event display showing what appears to be a DIS event starting in FGD1 and where the ensuing fragments are tracked through ND280's tracker region. Adapted from [123].

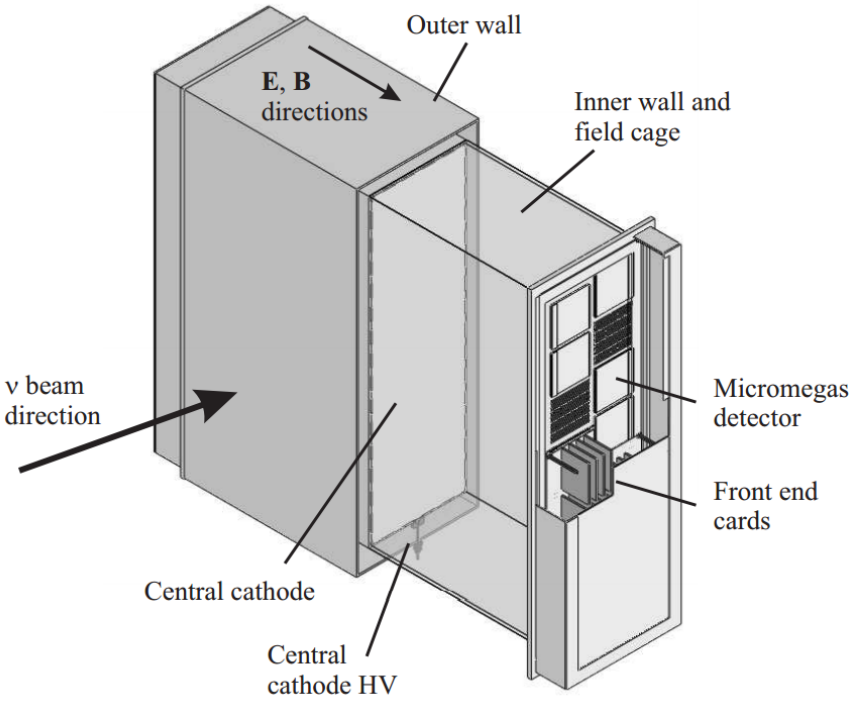


Figure 4.10: A simplified exploded view of one of the ND280 TPCs to illustrate its key components. Taken from [26].

module) with a sampling length of 700 mm, providing a point spatial resolution of around 0.7 mm. Upon reaching the readout plane, drifting elections enter an amplification region which applies a very large electric field (around 27 kV/cm) over a short distance (128 μm) such that they can unbind more elections, thereby creating an avalanche and providing a gain of about 1500.

The charge and time information is digitised and readout from the micromegas detectors using six front-end electronics cards, each containing four custom integrated circuits called 'AFTER', which shape and buffer the signals into 511 time-bin switched capacitor arrays. These front end cards then connect to a single front-end mezzanine board, which combines the signals and performs zero-suppression, before sending the subsequent signal for storage. The charge response of the AFTER electronics is set such that one count from the analogue to digital conversion (ADC) output is produced from around 170 electrons. The energy loss for a minimum ionizing particle (MIP) in the TPC gas is about 1.2 KeV/cm (shown in figure 4.11), corresponding to roughly 100 liberated primary electrons per cm. Overall this produces a signal of around 1000 ADC counts on each pad, to be compared with a noise level of 4-5 ADC counts (approximately 800 electrons equivalent) [124, 126]. This large signal to noise ratio affords the TPC an overall MIP detection efficiency of close to 100%.

The TPC high resolution readout allows an accurate measure of track direction. Fur-

thermore, the TPCs are within the ND280 magnetic field and therefore charged particles travelling through them produce curved tracks. By measuring the ensuing curvature of the tracks, the TPCs facilitate an accurate measure of charged particle momenta (and charge - by curvature direction). Figure 4.9 exhibits the TPCs tracking capabilities, being clearly able to distinguish charged tracks and curvatures. Combining the 0.2 T magnetic field with the aforementioned 0.7 mm point spatial resolution allows the TPCs to meet their momentum resolution goal of $\delta p_\perp / p_\perp < 0.1 p_\perp$ GeV [124] (where $p_\perp$ is the momentum perpendicular to the electric and magnetic fields in figure 4.10).

Beyond measuring the kinematics of charged particles, the TPCs are also capable of distinguishing between some particle types by measuring their energy loss (estimated from the amount of ionisation) in the gas as a function of their momentum (obtained from curvature). The energy loss of a relativistic charged particle characterises its velocity but is largely independent of its mass [127]. On the other hand, the momentum of a charged particle depends on both the particle's velocity and mass. As such, considering both gives a handle on the particles mass and therefore its type. This is clearly shown in figure 4.11, which demonstrates how the particles position in TPC reconstructed momentum-energy loss space is largely characteristic of its type and can therefore be used for PID (particle identification).

The TPC PID algorithm [124] exploits this by first utilising the truncated mean[2] of charged particle energy loss in the gas. The measured energy loss per unit length is then used to calculate 'pulls', $\delta_{E_i}(p)$: the number of standard deviations the measured energy loss is from what would be expected for a particular particle type, i, with momentum, p, as measured by the track curvature (the expectations are tuned from test beam data). The TPC MIP energy loss resolution is 7.8 ± 2% , which allows the probability of misidentifying muons as elections to be only 0.2% for $-1 < \delta_{E_e}(p) < 2$ where $p < 1$ GeV.

4.2.2 First fine-grained detector

Of ND280's FGDs [123], the most upstream fine-grained detector (FGD1) is a fully active polystyrene scintillator tracker, containing 5,760 9.61 mm × 9.61 mm × 1864.3 mm scintillator bars. These are arranged into 30 layers of 192 bars each, where each layer is aligned such that it is perpendicular to the incoming neutrino and the proceeding layer.

As charged particles pass through the FGD scintillator bars, they stimulate the emission of scintillation light. This light is then channelled down the bar, by a thin wavelength shifting fibre at its centre, to a multi-pixel photon counter (MPPC) [129, 130] (containing 667 pixels), which can measure the light intensity, a proxy for energy deposit, and time of arrival.

[2]The truncated mean is used to mitigate the impact of stochastic effects such as secondary interactions.

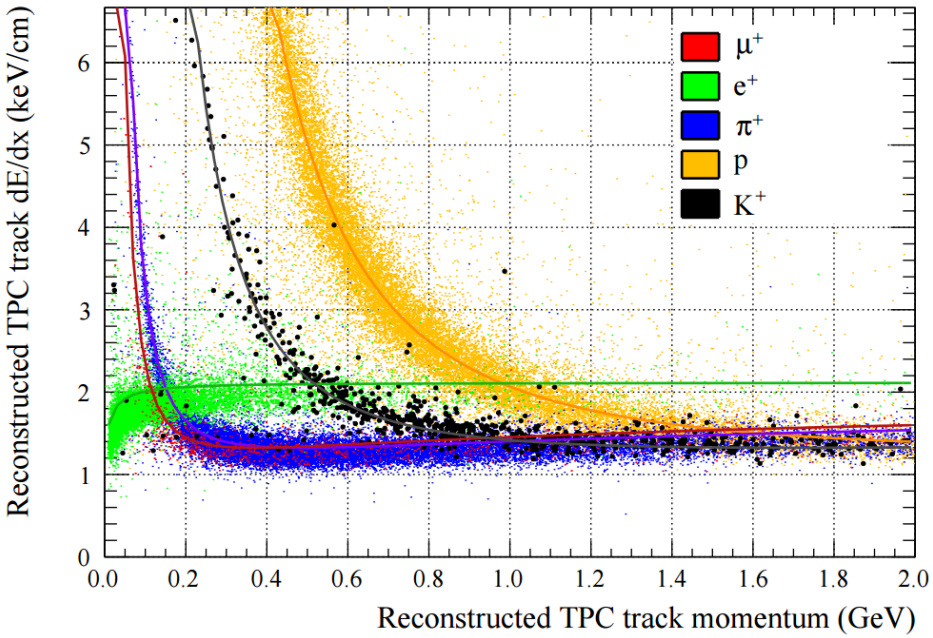


Figure 4.11: The reconstructed energy loss (dE/dx) as a function of reconstructed momentum in the TPCs. The curves show the expected distribution for various particle types (from calibration studies) whilst the scatter-plot shows the reconstructed distribution from a GENIE 2.8.0 neutrino interaction simulation at ND280 for different true particle types. This shows that, within a limited momentum region, the measurement of these observables is able to differentiate between some charged particle types. Taken from [128].

The intensity of the light can be measured by the number of photoelectrons (pe) recorded. A MIP passing through a scintillator bar with a path length of 1 cm deposits approximately 10 MeV of energy (as shown in figure 4.13) which corresponds to a 20-30 pe yield at the MPPC. Therefore the energy resolution of the FGDs is more limited from Poisson fluctuations (∼20%) than by electronics noise (which is around 1 pe) [123]. From a typical interaction it is expected that charge will be deposited in about 60 scintillator bars.

The signals from the MPPCs are amplified with a gain of around 5×10^5 before being digitised and readout using front-end boards containing the same AFTER integrated circuits discussed in section 4.2.1. The data from these boards can then be collected using specialised data collection cards and sent for storage. The charge response of the FGD AFTER electronics is such that 1 pe corresponds to around 40 ADC counts, meaning a MIP signal is expected to produce approximately 1000 ADC counts (to be compared with a noise level of around 40 ADC counts).

Primarily, FGD1 is designed to provide a target mass for neutrino interactions whilst supplying a reasonable degree of tracking. As a target, the fiducial volume of FGD1 contains approximately one tonne of target mass of which 86.1% is carbon, 7.35% hydrogen and the remaining contribution mostly from oxygen (3.70%) making it well-approximated as a CH target.

Because of the small cross-section of the scintillator bars, FGD1 is able to provide an accurate measure of outgoing track angle. Although the FGD's small profile means it is not able to measure particle curvature accurately enough to extract momentum, the detector is instead capable of determining momenta based on charged particle distance travelled ('momentum by range'), so long as the particle stops within FGD1. This is done by summing the energy deposits, from interaction vertex to stopping point, deemed to be associated with the charged particle track. This procedure is shown schematically in figure 4.12.

In a similar way to the TPC, FGD1 can also provide particle identification. However, rather than measuring energy loss and momentum, the FGD1 PID algorithm uses the total energy deposited and the length of the track ('track range'). Since, for a given energy, the range of a particle depends on its energy loss within the material, this gives a similar handle on PID as the TPC method. The distribution of FGD1 track range and energy deposited is shown in figure 4.13.

The FGD PID algorithm is similar to that used for the TPC, discussed in section 4.2.1. A pull, $\delta_{E_i}(L)$, is formed as the number of standard deviations the measured energy deposit is from what would be expected for a particular particle type, i, with track range, L, as measured by the FGD. As discussed above, the uncertainty in the measured energy loss for

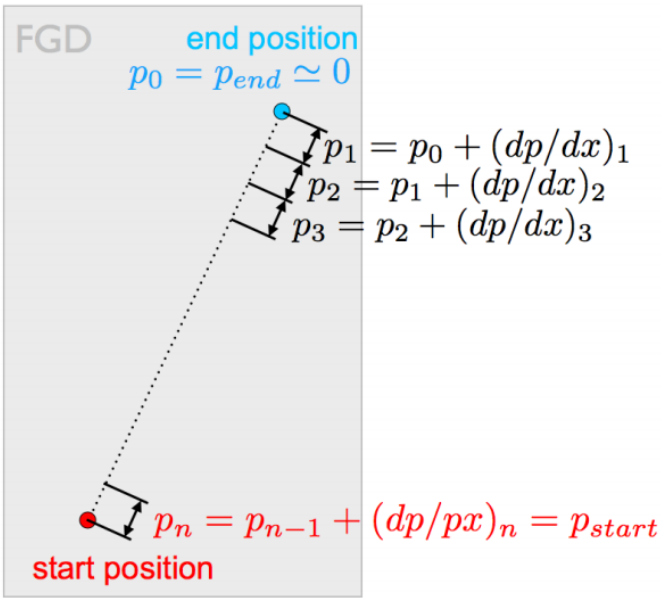


Figure 4.12: A demonstration of how momentum by range can be calculated for an FGD-only track. Taken from [128].

MIPs is dominated by the plausible statistical fluctuations in the 20-30 pe they produce in a scintillator bar.

4.2.3 Data simulation

In order to estimate backgrounds and detector efficiencies in any but the most simple event selections, it is essential to have a Monte Carlo simulation of data. This section will briefly discuss the fake data used by ND280, that will be used as an input to the analysis in the subsequent chapters.

ND280's event simulation begins with the sampling of interaction vertices throughout the detectors and magnet according to the relative masses of the different components of ND280 and the flux intensity at different off-axis angles (derived from the flux predictions discussed in section 4.1.3 and illustrated in figure 4.14). For example, most interactions are placed inside the ND280 magnet yoke since it makes up the majority of ND280's target mass and of the events placed in the FGDs there are more in the bottom left than the top right (if looking in the direction of the neutrino beam), since the former is at a more on-axis angle (as illustrated in figure 4.14) .

Following this placement of vertices, the NEUT [40] neutrino interaction simulation is used to generate a neutrino-nucleus interaction at each (where the probability of generating a particular interaction is proportional to the NEUT prediction for its cross section). The

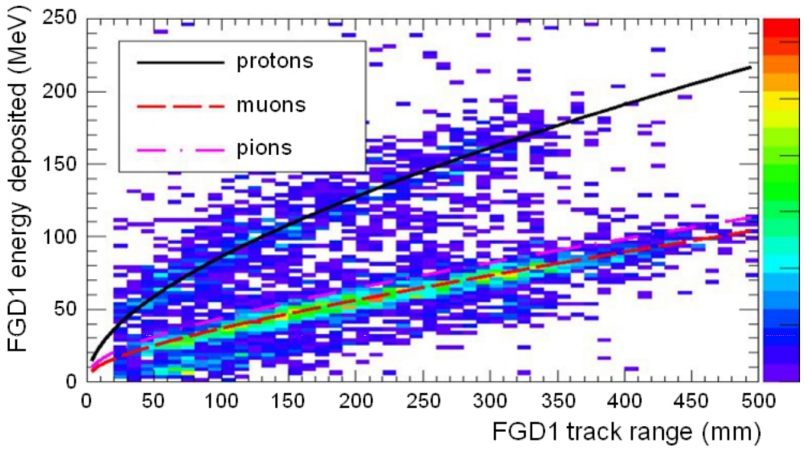


Figure 4.13: Deposited energy v.s. range for particles stopping in FGD1. The scatterplot is filled using stopping particles in neutrino beam data, whilst the curves show the simulated expectations for different particle types. This shows that, above some energy momentum threshold, the FGD is able to distinguish protons from muons and pions. Taken from [123].

default NEUT used by ND280 (known as "NEUT 6B" in this thesis) utilises NEUT version 5.3.2 with a SF model for CCQE interactions and an RFG spectral function for other interaction modes, all with $M_A = 1.21$ GeV. NEUT 6B simulates 2p-2h interactions based on the Nieves model, using an implementation described in [131]. NEUT 6B has a cascade FSI model, as described in 2.2.2 and does not simulate RPA by default. Unfortunately, the available ND280 NEUT 6B simulation contains a bug in which nucleon FSI are often generated as if they were on a much heavier target (typically iron) than is found in the regions of interest in common analyses (typically carbon or oxygen). Although this means the ND280 simulation predictions should not be taken to be representative of the true NEUT predictions, chapter 8 will demonstrate that the analysis in this thesis is robust against this issue.

As an alternative generator (and to guard against model dependence) there is also an ND280 simulation where events are generated using the GENIE neutrino interaction simulation. The ND280 GENIE 2.8.0 simulation uses an RFG nuclear model for all interaction modes with $M_A = 0.99$ GeV and does not simulate 2p-2h or RPA. The FSI implementation is not a full cascade model but is instead an empirical model that uses the total cross section for each possible nuclear process as a function of hadron energy. This is known as GENIE's 'hA' FSI model and more details can be found in [41].

In addition to the NEUT 6B and GENIE, there is also an ND280 simulation using the NuWro 11q neutrino interaction simulation and one using a slightly modified NEUT (known as NEUT 6D) simulation which uses NEUT version 5.3.2.2 and contains a fixed nucleon FSI simulation alongside the required information to study systematic errors from nucleon FSI. Unfortunately these simulations only generate events within the ND280 detectors and not

within the surrounding magnet, such that there are fewer events rejected in event selections because of activity stemming from out of target fiducial volume interactions. This gives these simulations an unrealistically large selection efficiency making them impractical as an input simulation but can still serve as very useful test fake data sets. NEUT 6D is generated with the same models as NEUT 6B, whilst the NuWro simulation uses an LFG spectral function; a different implementation of Nieves 2p-2h and a different FSI cascade model. Further details regarding the event generators used can be found in appendix A

Although it is useful to have these representative simulations of neutrino interactions, it is absolutely essential that any analysis that are dependent on poorly understood effects (which in neutrino scattering physics is almost all analyses) should not strongly depend on this input neutrino-nucleus interaction simulation, as chapter 2 has demonstrated to cause significant complications in the interpretation of previous measurements.

Once the neutrino-nucleus interaction simulation has produced an event, the final state particles (i.e. those that leave the nucleus) are propagated through ND280 using a GEANT4 [119] based simulation. This accounts for particle energy loss, curvature due to the magnetic field and secondary hadronic interactions. The GEANT4 tracks are then the input for the ND280 data acquisition and electronics simulation (elecsim), which simulates the detectors response to them. The reconstructed output from elecsim then looks identical to real data, save for it also containing the true properties of the interaction alongside those reconstructed by elecsim.

4.3 T2K data availability

As stated in the introduction to this chapter, T2K has been running since 2009 and, as of March 2017, has collected around 25% of its planned data. The amount of data acquired is usually referred to in terms of 'protons on target' (POT): the number of protons that have impinged on the target in the primary neutrino beamline. Figure 4.15 shows the number of POT collected over time, alongside the J-PARC proton beam power, thereby demonstrating that T2K has been accelerating in its data taking rate. Figure 4.15 also indicates that, as of March 2017, there is 12.83×10^{20} POT FHC data available but, since - as of April 2017 - the latest T2K run has not yet finished, the data from 'Run 8' is yet to be calibrated. Overall there is 6.68×10^{20} POT FHC data which is suitable for cross section analysis. The imminent availability of almost twice this much data means there is considerable scope to soon lower statistical errors in FHC cross section analyses.

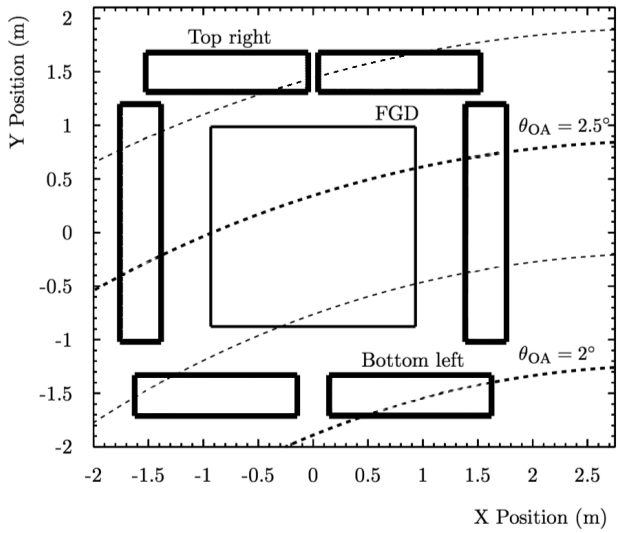


Figure 4.14: The variation of off-axis angle from the neutrino beam centre across ND280. The 'Bottom left' and 'Top right' labels are as seen if looking through ND280 in the beam direction. Adapted from [132].

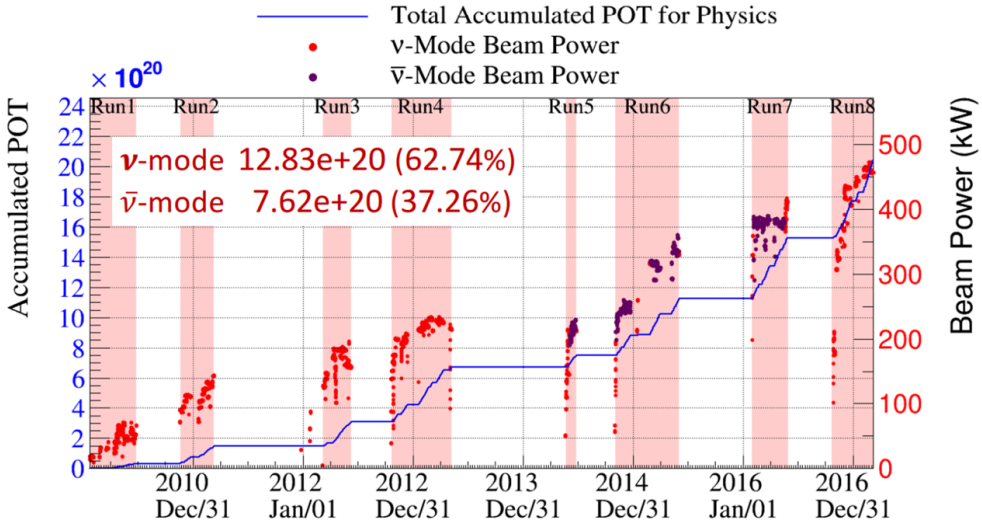


Figure 4.15: Number of protons on target (POT) accumulated by T2K over time, broken down by whether T2K was running in neutrino (FHC) or antineutrino (RHC) mode, shown alongside the J-PARC proton beam power. The beam intensity is measured using five current transformers in the primary neutrino beamline [111].

Chapter 5

Analysis strategy

The primary aim of this thesis is to provide a better characterisation of the physical processes that lead to the principle systematics in neutrino oscillation analyses, particularly the nuclear-effects discussed in chapter 2. To achieve this, a measurement of a flux-integrated, fiducial ν_μ CC0π+Np differential cross section on a hydrocarbon (CH) target is made as a function of each of the single-transverse variables, discussed in chapter 3, which may offer a novel probe of such processes. As highlighted in chapter 2, it is essential that the extracted cross-section is not biased by the input neutrino-nucleus interaction simulation. To minimise the scope for such potential bias particular care is made to avoid any substantial dependence on the input interaction simulation. Specifically:

- cross sections are measured as a function of observables (the STV) that are directly accessible in ND280 and do not require any model assumptions to build (unlike E_ν or Q^2);
- the cross sections are only formed in the region of proton and muon kinematic phase space that ND280 is directly sensitive to (the extracted cross section is fiducial);
- a flux-integrated cross-section is extracted, such that there is no attempt to understand the neutrino energy distribution in each bin of the result;
- the extracted cross section signal definition is defined in terms of particles that exit the nucleus such that there is no need to utilise a nuclear model to correct for FSI effects (i.e. a topology based signal definition is chosen - CC0π+Np - rather than one based on interaction mode - like CCQE);
- 'sideband' samples are used to assess that the implemented background models are able to well-describe the data within the freedom afforded to them in their systematic parameters;
- extensive fake data studies are made to ensure the cross-section extraction method is able to remain unbiased from the input simulation across a wide range of plausible real data scenarios.

The rest of the thesis will detail the cross-section extraction framework, before applying it to real data and offering an interpretation of the result. Various components of the ND280 off-axis near detector, described in chapter 4, are used to achieve this measurement: the FGD1 is used as a CH target, whilst both FGD1 and the TPCs will be used for tracking, measuring kinematics and particle identification.

Chapter 6 will present a selection of CC0π+Np events starting in FGD1, with a high efficiency and minimal backgrounds. This chapter will also overview the 'sideband' selection, which identifies events corresponding to the dominant backgrounds of the signal selection, in this case CC1π and CCNπ events. Following this, chapter 7 presents a method of going from this signal and sideband event selection to a cross-section measurement using a regularised binned likelihood template fit. Particular focus is placed on the role of regularisation and the scope for it to bias cross-section measurements. Chapter 8 then validates the cross-section extraction method and assess possible sensitivity to some nuclear effects. To do this numerous fake data studies are performed to demonstrate that the method is not biased to the input simulation and has a rigorous handling of uncertainties. Finally chapter 9 applies the cross-section extraction method to real data and assesses the implications of the results through extensive model comparisons.

Chapter 6

CC0π+Np selection

This section details the event selection used as an input to the cross-section analysis in this thesis. As discussed in the previous chapters, this selection will aim to extract events from neutrino interactions with a CH target that produce one muon, no pions and at least one proton in the final state (CC0π+Np interactions). The event selection used to achieve this is adapted from well-established selections that have been designed to facilitate near detector constraints in the oscillation analyses and to measure other cross sections. In particular, the selection and implementation of associated systematic uncertainties, detailed in [133] (and used as the selection for the first analysis in [106]) was used as a starting point for the event selection in this thesis.

Within this chapter, section 6.1 will overview the event selection, focussing on modifications or additions from [133]. This covers a restructuring of the cuts; the addition of multi-proton samples; a π^0 veto cut using the ECal and the the identification of the incoming neutrino direction using simulations of the secondary neutrino beamline. The distribution of resultant events is then described in section 6.2, including the breakdown of events in each sample by true topology and interaction mode in addition to the proton and muon kinematic distributions. Following this, section 6.3 describes the efficiency and purity of the selection (defined in equations 6.1 and 6.2) and shows that building constraints on the proton and muon kinematic phase-space into the signal definition can dramatically increase efficiency whilst allowing a measurement to be made using only events which ND280 is able to reliably reconstruct. Section 6.4 then describes how the STV are built from reconstructed kinematics and the associated potential for model dependence. Sideband samples, used to constrain the predominant backgrounds from the signal selection, are then defined in section 6.5. Finally section 6.6 describes the systematic uncertainties associated with the selection.

Unless otherwise specified, this chapter uses the default ND280 NEUT 6B simulation for neutrino interactions (described in section 4.2.3 and further detailed in appendix A), produced with 33.16×10^{20} POT ($\sim 5\times$ real available data POT). Since this is an analysis of a ν_μ cross section rather than $\overline{\nu}_\mu$, all simulations used are based on T2K operating in FHC mode.

Throughout this section it will be useful to assess the selection performance using the purity and efficiency metrics. These are defined as:

$$\text{Purity} = \frac{\text{Number of selected signal events}}{\text{Number of selected events}}, \tag{6.1}$$

$$\text{Efficiency} = \frac{\text{Number of selected signal events}}{\text{Total number of signal events}}. \tag{6.2}$$

6.1 Selection

Since this selection aims to extract neutrino interactions on CH, the FGD1 detector is used as a target. In addition to FGD1; FGD2, TPC2 and TPC3 are all used to characterise particle tracks. Events with one muon, at least one proton and no other identified particles are selected where the vertex is found to be within the FGD1 fiducial volume (FV). These events are split into four signal samples, divided based on which detectors are used for reconstructing the selected event and how many protons are detected. This is necessary as the detector response; detector acceptance; and the breakdown of interactions modes and topologies in each of these categorisations can be quite different. Overall, the following selection samples are used:

- μ TPC + p TPC : μ-like TPC track and p-like TPC track
- μ TPC + p FGD : μ-like TPC track and p-like FGD track
- μ FGD + p TPC : μ-like FGD track and p-like TPC track
- μ TPC + multi p : μ-like TPC track, highest momentum p-like track in the TPC and at least one more p-like TPC track or FGD track

Although other samples are possible (e.g. a muon-like FGD vertex + FGD proton(s)) these are found to have extremely low event rates (<30 expected in available data). Figure 6.1 shows a breakdown of all these samples and the simulated number of events in each within available data.

A list of cuts used to construct each sample can be found in Figure 6.2. The cuts in this selection are built heavily on that described in [133] and will be briefly summarised below. The major changes to the selection are the addition of a sample with multiple protons in the final state (μ TPC, multi p in figure 6.1) as well as the inclusion of a π^0 veto cut using the ECal. These are described in more detail in the following sections.

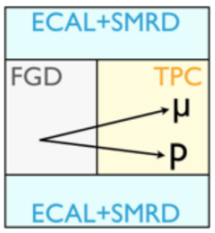
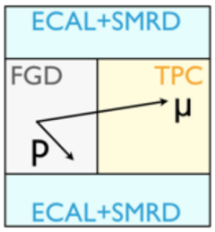
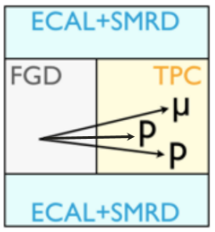
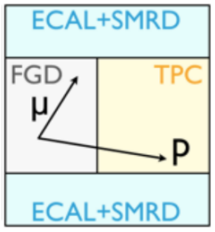
Sample	$\mu TPC + pTPC$	$\mu TPC + pFGD$	$\mu FGD + multi\ p$	$\mu FGD + pTPC$
				
Description and number of expected events	• Both μ and p candidates are tracked in TPC • **1677 events**	• μ candidate tracked in TPC • p candidate tracked in FGD only • **1222 events**	• μ candidate tracked in TPC • Multi p candidates, leading p in TPC • **154 events**	• μ candidate tracked in FGD only • p candidate tracked in TPC • **919 events**

Figure 6.1: A simple breakdown of the different signal samples used in this analysis, split depending on the detectors that were used to reconstruct the events and the number of protons identified. The number of events expected in each sample within available ND280 data is indicated.

<table>
<tr><td>$\mu TPC + pTPC$</td><td>$\mu TPC + pFGD$</td><td>$\mu TPC + multi\ p$</td><td>$\mu FGD + pTPC$</td></tr>
<tr><td colspan="4">Event Quality</td></tr>
<tr><td colspan="3">HMN Vertex</td><td>HMP Vertex</td></tr>
<tr><td colspan="3">Muon PID</td><td>Only Muon and Proton Tracks</td></tr>
<tr><td colspan="3">One Negative Track with Common Vertex</td><td>Leading Proton TPC Track</td></tr>
<tr><td colspan="3">Only Muon and Proton Tracks with Common Vertex</td><td>Long FGD Track</td></tr>
<tr><td colspan="3">No Michel Electron</td><td>FGD Muon PID</td></tr>
<tr><td colspan="3">ECal π^0 Veto</td><td>Common Vertex</td></tr>
<tr><td colspan="2">One Proton with Common Vertex</td><td rowspan="2">Multi Proton with Common Vertex</td><td>Stopping Muon</td></tr>
<tr><td>TPC Proton</td><td>FGD Stopping Proton</td><td>One Proton</td></tr>
</table>

Figure 6.2: An overview of selection cuts for each sample, these are detailed in section 6.1.1.

6.1.1 Summary of cuts

This section gives a summary of the cuts used in this analysis, many of these are ported from [133] and are further detailed therein.

Event quality This cut reads a data quality flag stored in real data to determine whether the detector and beam were functioning correctly. For simulated data this cut is always passed.

HMN/P vertex To accurately establish a vertex position in FGD1 it is best to have a track that enters, and can be identified by, the TPC. The vertex position is preferentially set using a muon track (since these leave very clean signals in the tracker). For this reason this cut looks for the highest momentum negative (HMN) track in the TPC (usually the muon), and uses its combined TPC and FGD tracking information to set the vertex position. If there is no negative track in the TPC this cut uses the highest momentum positive track in the TPC (HMP). If there are no TPC tracks associated with the FGD this cut is failed.

HMN vertex based cuts

Muon PID This cut is passed if HMN track is determined to look muon-like based on the TPC PID method outlined in section 4.2.1. Note that muon-like means the track is likely to be a muon or a negative pion, since the TPC cannot reliably distinguish between these.

One negative track with common vertex This cut is failed if there are extra tracks (in addition to the HMN track) identified as being negatively charged in the TPC, which appear to have a similar vertex position to the HMN vertex.

Only muon and proton tracks with common vertex This cut is failed if there is a track in the TPC or FGD which is determined as being neither muon- or proton-like by their PID algorithms (described in 4.2.1 and 4.2.2) and that appears to have a similar vertex position to the HMN vertex.

No Michel electron This cut is failed if a Michel electron/positron is identified in the FGD. A Michel electron is the electron from a muon decay. In the samples which apply this cut, the muon is already identified as having left the FGD, so a Michel electron would likely come from a low energy pion decaying into a muon which would subsequently decay itself. A Michel electron/positron is identified by seeing a charge deposition consistent with an electron/positron in the FGD around one muon lifetime (2.2 μs) after the event start time.

ECal π^0 veto This cut is failed if the ECal identifies a charge deposition consistent with a photon from a π^0 decay. This cut is newly added to this selection and so its performance is validated in section 6.1.2.

One proton with common vertex This cut is passed if exactly one track, in either the TPC or the FGD, share a common vertex with the HMN track and is identified as being proton-like by the TPC or FGD PID algorithms.

TPC proton This cut is passed if the identified proton-like track is in the TPC.

FGD stopping proton This cut is passed if the identified proton-like track stops within the FGD.

Multi proton with common vertex This cut is passed if there is more than one proton-like track, at least one of these protons must be tracked by the TPC.

HMP vertex based cuts

Leading proton TPC track This cut is passed if the track used to form the HMP vertex is identified as being proton-like by the TPC PID algorithm.

Long FGD track This cut is passed if there is a track in the FGD above a certain length. This cut ensures the candidate muon track has enough FGD information to have a reasonable measure of its kinematics and also loosely helps separate muons from pions, as muons typically have a higher momentum in true CC0π interactions.

FGD muon PID This cut is passed if the long FGD track is identified as looking muon-like based on the FGD PID method. Note that muon-like means the track is likely to be a muon or a charged pion, since the FGD cannot reliably distinguish these.

Common vertex The muon candidate from the previous cut must share a similar vertex position to the HMP track.

One proton with common vertex This cut is failed if there are any more tracks, other than the muon candidate from the previous cut, with a similar vertex to the HMP track.

6.1.2 ECal π^0 veto cut

As described in section 6.1.1, this cut uses the ECal to identify charge deposits that are consistent with being from a π^0 decay photon. This cut was recently developed by the T2K collaboration for a CC1π^+ selection but can also be applied here in order to reduce π^0 backgrounds. The details of the cut's implementation and associated systematics are reported in [134]. To validate its use in this selection, Table 6.1 shows the effect of the number of selected and signal events in each sample alongside the contribution of π^0 events. From this table it is clear that the cut is actively rejecting a significant amount of non CC0π events with a minimal impact on CC0π events and is therefore working effectively.

	Topology				Event Rejection	
Cut Status	CC0π	CC1π^+	CC1π^0	Other	CC0π	Other
Before π^0 veto	84.7%	6.15%	3.98%	5.17%		
After π^0 veto	88.8%	5.71%	1.39%	4.10%	4.97%	33.8%

Table 6.1: The proportion of true signal topologies before and after the π^0 veto cut (see Figure 6.2) alongside the proportion of CC0π and non-CC0π events rejected by the cut. All topologies are encompassed such that numbers in each row of the topology column add to 100%.

6.2 Summary of selected events

In each sample it is helpful to consider the breakdown of events by both true final state topology and interaction mode, this is shown in table 6.2 and 6.3 respectively. The former table shows the purity of finding a CC0π+Np true topology (the sum of the second and third columns) and gives the breakdown of backgrounds in each sample. Table 6.3 confirms the expectation that the largest contribution to CC0π+Np events comes from CCQE and 2p2h contributions across most samples (within the simulated data). The exception to this is the multi-proton topology, which is dominated by RES events. The topology table demonstrates that few events in this sample actually have a CC1π^+ final state, so many of these RES events likely follow a CC0π+Np true topology and are subject to pion-absorption FSI, and subsequent proton emission, from the unstable nucleus.

In addition to the coarse breakdown of total selected events provided by the tables, it is also very useful to have an idea of the simulated breakdown of interaction modes in simple proton and muon kinematics (p, $cos(\theta)$). This is useful to provide a preliminary understanding of the detector acceptance and a preliminary assessment of whether the cuts have caused any unpredictable detector inefficiencies. These distributions are provided in

figures 6.3 to 6.6, which show the true proton and muon kinematics of events selected within each signal sample stacked by true interaction topology.

	Topology						
Sample	CC0π0p	CC0π1p	CC0πNp	CC1π^+	CCOther	Other	Events
μTPC+pTPC	0.01%	46.1%	36.1%	9.55%	4.23%	3.98%	1677
μTPC+pFGD	0.00%	50.1%	36.2%	8.01%	2.35%	2.40%	1222
μFGD+pTPC	0.00%	60.3%	23.8%	5.36%	2.60%	7.92%	919
μTPC+Np	0.00%	1.74%	64.0%	18.9%	8.70%	6.71%	154

Table 6.2: The proportion of each true signal topology in each reconstructed sample alongside the expected number of events in available data (predicted by NEUT). Note that here Np refers to a multi-proton final state.

	Interaction Mode					
Sample	CCQE	2p2h	CCRES	CCOther	Other	Events
μTPC+pTPC	59.8%	14.6%	18.9%	2.42%	4.06%	1677
μTPC+pFGD	67.3%	10.5%	17.5%	1.51%	3.23%	1222
μFGD+pTPC	77.1%	5.28%	8.80%	0.94%	7.90%	919
μTPC+Np	15.5%	17.9%	49.4%	10.3%	6.69%	154

Table 6.3: The proportion of each true interaction mode in each reconstructed sample alongside the expected number of events in available data (predicted by NEUT).

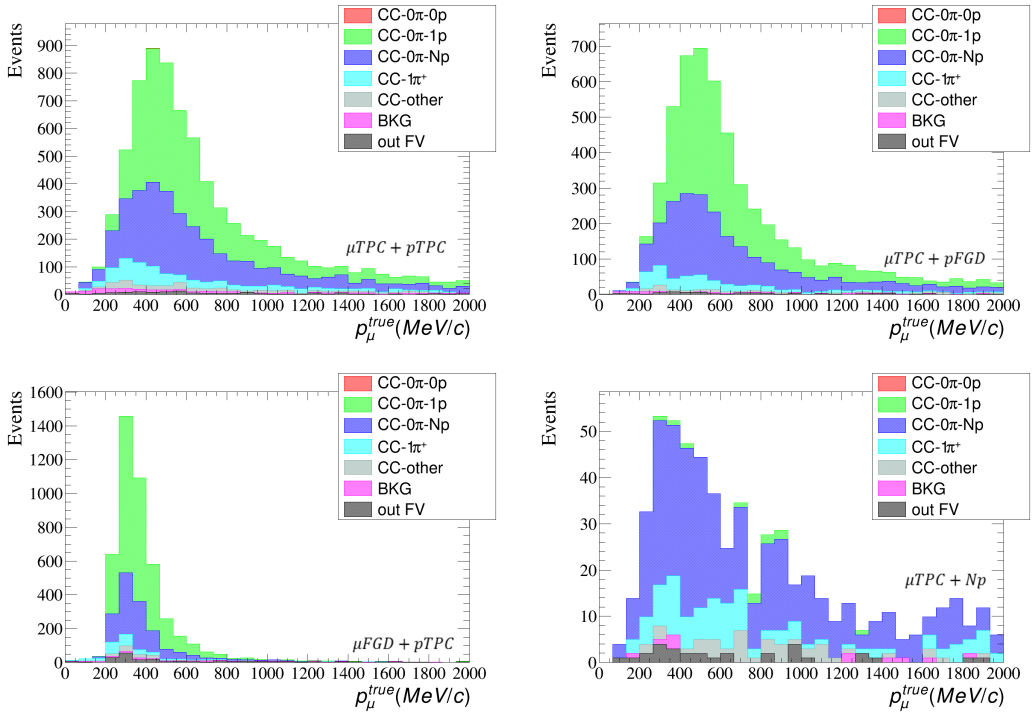


Figure 6.3: The true muon momentum of events passing the cuts for each signal sample, stacked by true event topology (predicted by NEUT). Note that here Np refers to a multi-proton final state.

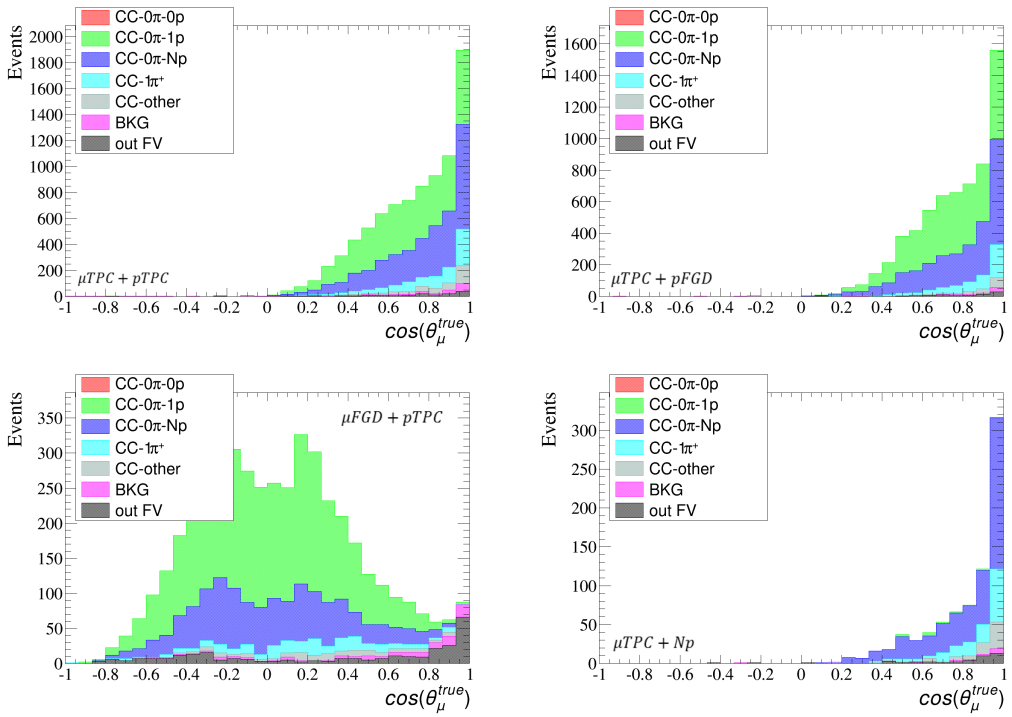


Figure 6.4: The true muon angle with respect to the incoming neutrino direction ($cos(\theta_\mu)$) of events passing the cuts for each signal sample, stacked by true event topology (predicted by NEUT). Note that here Np refers to a multi-proton final state.

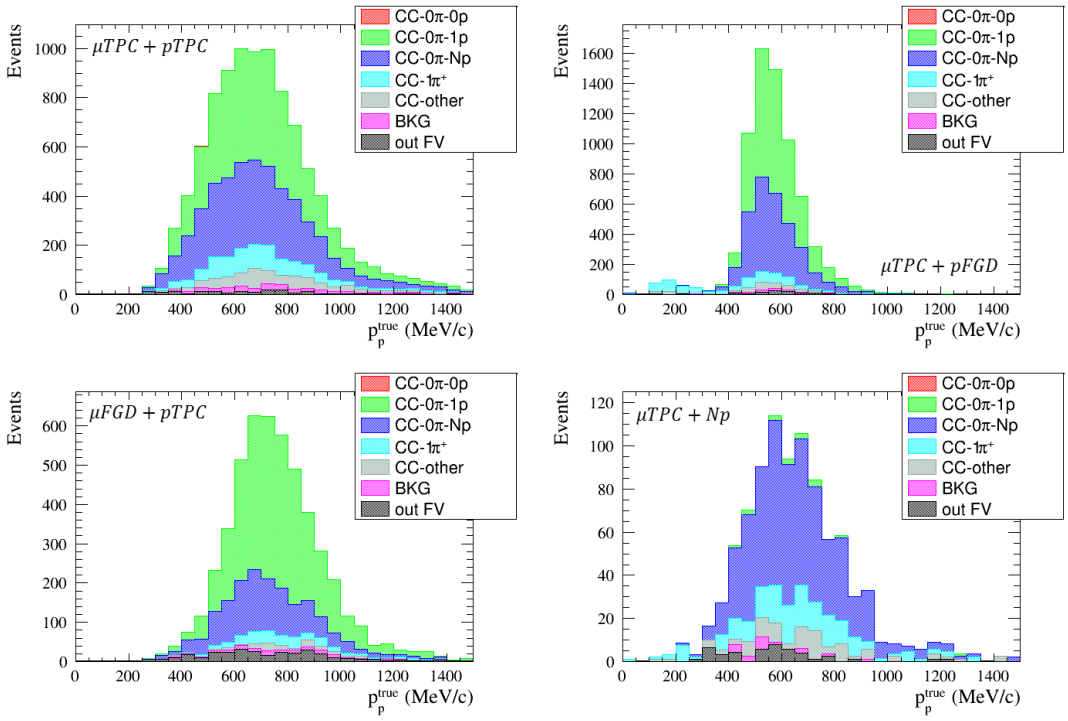


Figure 6.5: The true proton momentum of events passing the cuts for each signal sample, stacked by true event topology (predicted by NEUT). Note that here Np refers to a multi-proton final state.

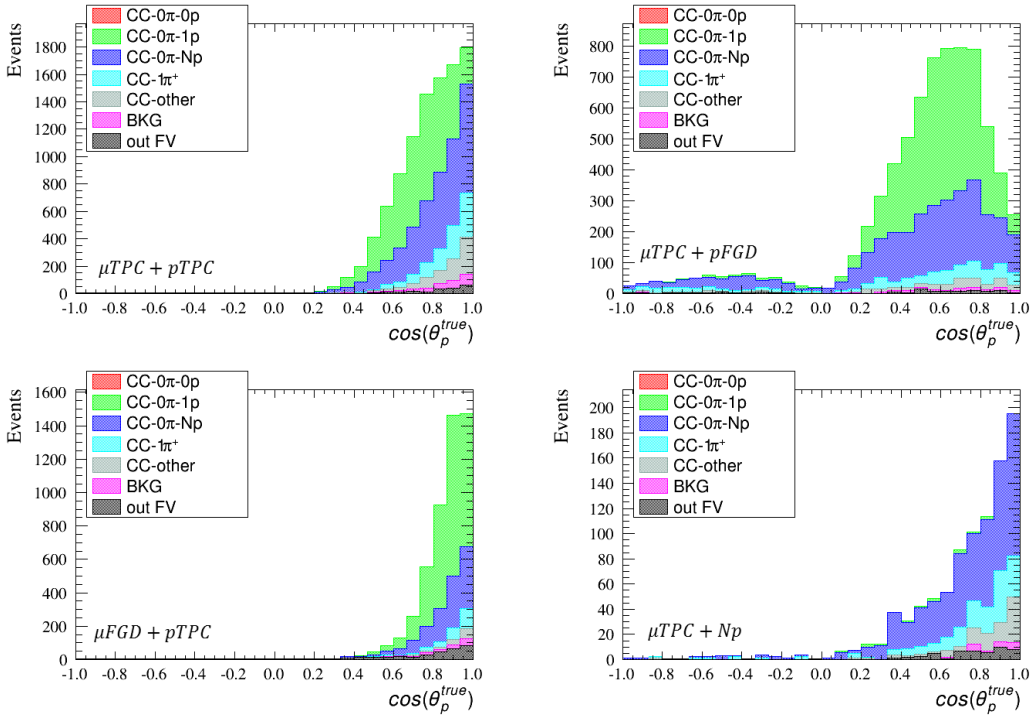


Figure 6.6: The true proton angle with respect to the incoming neutrino direction ($cos(\theta_p)$) of events passing the cuts for each signal sample, stacked by true event topology (predicted by NEUT). Note that here Np refers to a multi-proton final state.

6.2.1 Backgrounds

To understand the limitations of the selection it is useful to consider backgrounds by selection failure mode. In this categorisation, there are three main backgrounds: π^0 tag failures, π^+ tag failures and pion mis-PIDs.

π^0 Events When a π^0 is produced, it may not be identified by the tracker. In this case a CC1π^0 event may be identified as having only a muon and a proton in the final state, and therefore looking like a CC0π+Np event.

π^+ Tag Failure If a pion is too low momentum to be seen in the tracker and its Michel electron is missed, it is possible to identify a CC1π event as CC0π+Np.

π^+ Mis-identification Section 4.2 demonstrates that the FGD may struggle to distinguish low energy pions and protons, whilst the TPC may struggle to distinguish high momentum pions and protons. Therefore such pions may be mis-reconstructed as a proton by the PID algorithms. In this case, a CC1π event may be tagged as CC0π+Np.

The breakdown of these backgrounds are shown in table 6.4 for both the NEUT 6B and GENIE 2.8.0 ND280 simulations discussed in section 4.2.3 (more details of the simulations can be found in appendix A). From this it can be seen that over two thirds of the backgrounds come from failing to tag a π^+. It is therefore important to focus on characterising this background when constructing sidebands. These will be discussed in section 6.5.

		NEUT	GENIE
CC0π+Np selection purity		87.7%	84.8%
Relative contribution of backgrounds	π^0	18.3%	18.2%
	π^+ Tag Failure	71.6%	74.6%
	π^+ Mis-identification	3.2%	3.4%
	other	6.9%	3.8%

Table 6.4: The breakdown of various backgrounds to the CC0π+Np selection, separated by selection failure mode, for the NEUT 6B and GENIE 2.8.0 simulations.

6.3 Efficiency and phase-space constraints

When extracting a cross section, all backgrounds will need to be corrected for using the input simulation constrained by the sideband samples. Therefore a low purity can lead to excessive dependence on the background models in the input simulation. Fortunately,

section 6.1 has demonstrated that the event selection can identify CC0π+Np events with a high purity, so there is limited scope for input simulation background model dependence. However, this is not the only consideration: the signal events that were not identified by the selection must also be accounted for in an efficiency correction. It is therefore also important to ensure that ND280 has a reasonable efficiency for identifying signal events. Moreover, as discussed in section 2.5.2, it is important not to attempt to efficiency correct into regions of kinematic phase-space that ND280 is not sensitive to. For example, it is clearly unrealistic to try and measure protons with momentum below the FGD detection threshold, since there will be very few events and would require a large (and very model dependent) efficiency correction. To ensure this does not happen, it is necessary to set a signal definition as events exhibiting a true CC0π+Np topology, alongside appropriate limitations on the muon and proton momentum and angle to ensure a good selection efficiency and therefore reduce reliance on the input simulation. This section aims to identify such constraints and review the scope for bias.

6.3.1 Choosing the phase-space restrictions

Choosing the exact phase-space limitations to apply is not trivial. Ideally the efficiency correction should be able to be determined from GEANT4 and elecsim and be entirely independent from the neutrino interaction model. To first order, if single muons are being measured, an event with a particular true momentum and angle should either be seen or not, so the efficiency should be a 2D equivalent of a step function in muon p, $cos(\theta)$[1]. However, if the efficiency curve of just momentum or angle in 1D is considered, then the input simulation interaction model is relied upon to fold in the effects of the other dimension (i.e. to inform which p distribution falls into each $cos(\theta)$ bin or vice-versa).

In this analysis both muons and protons are identified and therefore it is necessary to identify a region of 4D-kinematic phase-space (muon and proton momentum and angle) where both are consistently reconstructed. Since 4D spaces are difficult to visualise, this can been split into 6 2D phase-spaces which fully encompass the correlations between the kinematic variables. The distribution of events and of the efficiency are shown in these spaces in figure 6.7. The 1D projected efficiency for each variable is also shown for both the ND280 NEUT 6B and GENIE 2.8.0 simulations in figure 6.8. The difference in generator predictions on the efficiency gives a rough indication of model bias. Ideally, the selected phase-space restrictions should leave a flat efficiency in all kinematic variables, since this ensures that the integrated/projected efficiency over any variable is independent of the distribution of

[1]In reality this is not quite the case since factors such as vertex position and secondary interaction rates also play a role in whether an event can be reconstructed, but this is still independent of the neutrino-nucleus interaction model.

events within it. Indeed, during the cross-section extraction for this analysis the efficiency is projected onto the STV, so any deviation from flatness in the 4D efficiency gives potential for model bias. The efficiency correction procedure is described in chapter 8.

Phase-space limits are chosen by examining efficiency and selected event distributions in these spaces when the signal is just CC0π+Np without phase-space constraints. It's important to only consider efficiency in light of selected event distributions since low simulation statistics can cause unrealistic efficiencies and a balance must be found between the number of events rejected and efficiency flatness. It is also important to try to retain interesting signal regions that are likely to be particularly sensitive to nuclear effects (e.g. low proton momentum).

The distributions in figure 6.7 clearly show that the variables driving the efficiency hardest are proton momentum and angle (with the exception of regions of very low muon momentum). The bottom two plots of figure 6.8 show that events with low momentum and backward-going protons have very low efficiencies. These low efficiencies at large angles, and also low momentum, stems from ND280's limited acceptance. For example the FGD is not capable of reliably reconstructing muons below 250 MeV/c, whilst high angle particles (which are therefore unlikely to enter the TPC) will only be well reconstructed if they have low enough momentum to be ranged out in the FGD. Additionally there is an efficiency dip for protons above about 1 GeV, this is due to high momentum protons exhibiting similar rates of energy loss with distance to other charged particles thereby making consistent PID difficult in the TPC, as was shown in figure 4.11.

From this evaluation of the efficiency the following constraints are to be placed on proton kinematics in the signal definition for the cross-section extraction:

- True proton momentum $>$ 450 MeV/c,
- True proton momentum $<$ 1000 MeV/c,
- True proton $cos(\theta) > 0.4$.

Figure 6.9 shows the efficiencies as function of true muon kinematics after the above phase-space constraints on proton kinematics are applied (for both the NEUT 6B and GENIE 2.8.0 simulations). From an evaluation of these, alongside consideration of the correlations between muon momentum and angle, the following constraints on muon kinematics are chosen:

- True muon momentum $>$ 250 MeV/c,
- True muon $cos(\theta) >$ -0.6.

Following the application of all constraints the efficiency is shown to be flat in proton kinematics in figure 6.10 and that the NEUT 6B and GENIE 2.8.0 efficiency predictions are similar. It is particularly important to note that the efficiency is mostly flat in proton kinematics, such that any efficiency formed from integrating over any portion of the proton kinematics does not depend on the poorly understood distribution of events in the proton kinematic phase-space.

6.3.2 Total efficiency and purity

The total efficiency and purity of events adhering to the phase-space restrictions on the signal definition are 29% and 88% respectively (in NEUT). The contribution from each topology and the change in efficiency and purity as a function of applied cut is shown in figure 6.11. Note that implementing the restricted phase-space does not involve any cuts, just an alteration of what is regarded as signal and background.

In reality the events that are selected but fall out of phase-space should be included as a new type of background. Including these causes the purity of the selection to drop to 65% (i.e. although the vast majority of selected events exhibit a CC0π+Np topology, only 65% are CC0π+Np and fall within the phase-space constraints). This shows that there is a substantial contribution from CC0π+Np events that fall out of the signal phase-space requirements. This background will be accounted for in the cross-section extraction.

6.4 Reconstructing the single-transverse variables

The previous sections in this chapter have outlined a selection of CC0π+Np events with a high purity, whilst the application of phase-space constraints ensures a relatively flat efficiency. This section will review the post-selection distributions of the STV (defined in chapter 3).

In order to reconstruct the STV the initial neutrino direction is required. In other analyses that require an outgoing particle angle this has previously simply assumed to be the z-axis of the detector (as is defined in figure 4.8), which is approximately aligned with the incoming neutrino beam, but it is relatively easy to get a much better estimate. The vertex position of each event and the mean point in which a neutrino's parent decays, in the decay tunnel of the secondary neutrino beamline (PDP), is well known [111]. Therefore, the vector connecting them provides a good estimate of the initial neutrino direction, which can then be used to define outgoing particle angles. This is shown schematically in figure 6.12.

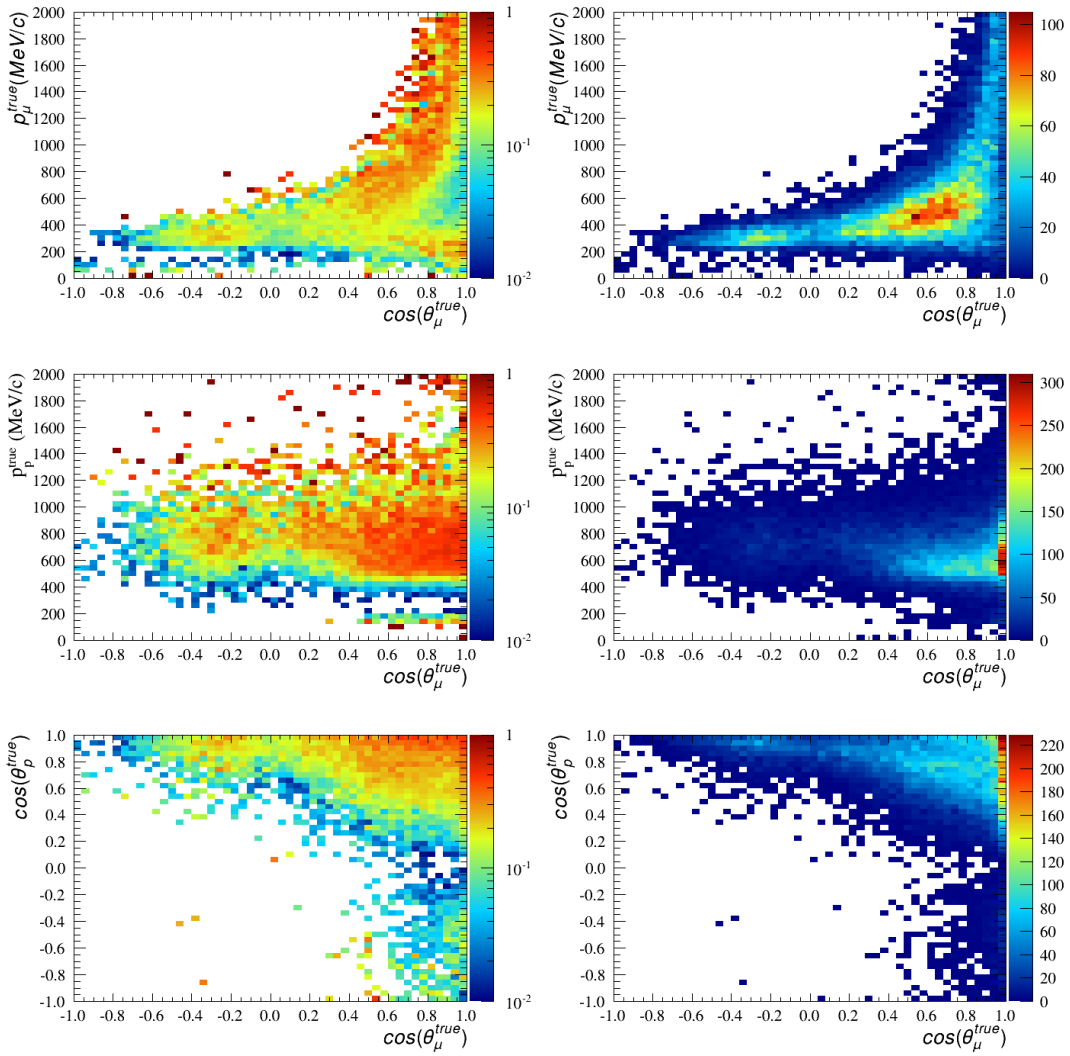


Figure 6.7: The efficiency under a CC0π+Np signal definition (left), and the true distribution (right), of proton and muon momentum and angle. These plots are used to choose phase-space limits to create a more realistic and less model dependent signal definition. This is done by looking for regions of consistently high efficiency across the 4D phase-space.

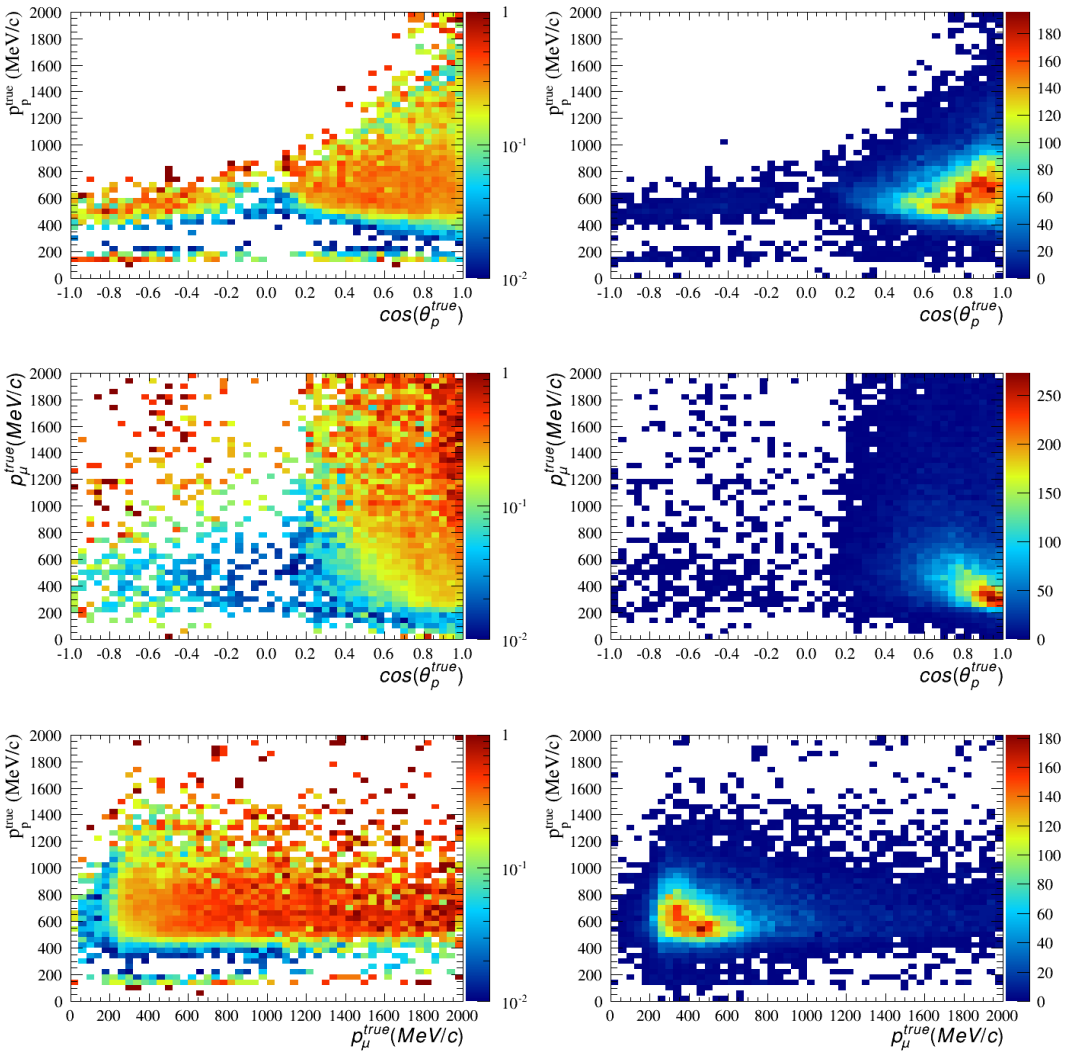


Figure 6.7: Cont.

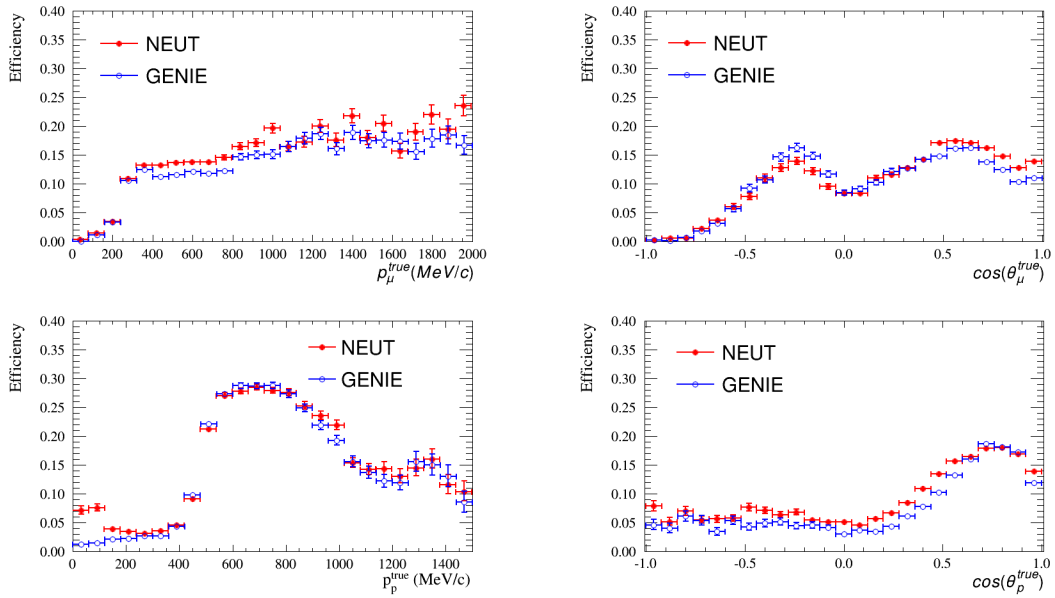


Figure 6.8: Selection efficiencies as function of true muon (upper plots) and proton (lower plots) kinematics.

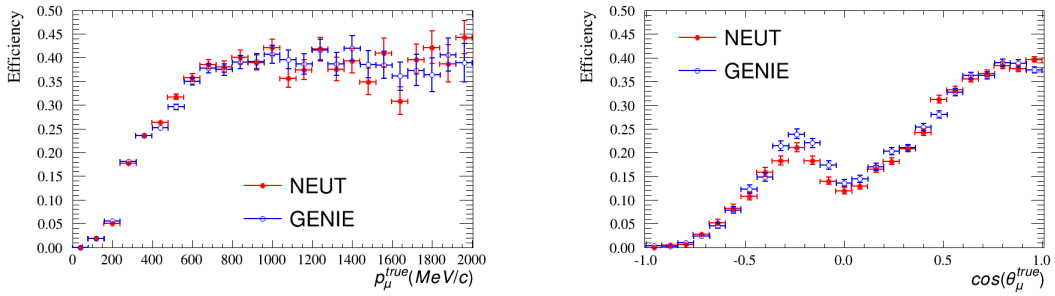


Figure 6.9: Selection efficiencies after proton phase-space constraints are applied, as a function of true muon kinematics.

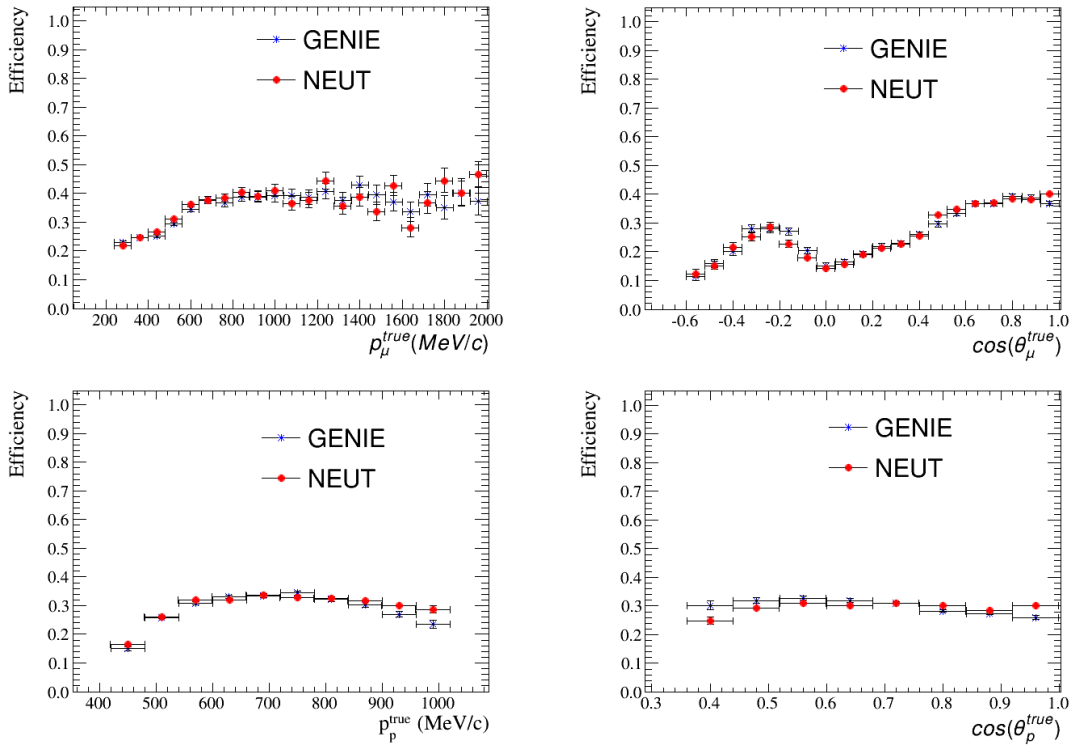


Figure 6.10: Selection efficiencies after proton and muon phase-space constraints are applied, as function of true muon (upper) and proton kinematics. The difference in the predictions is indicative of the interaction model dependence that will feed into the efficiency correction during cross section extraction.

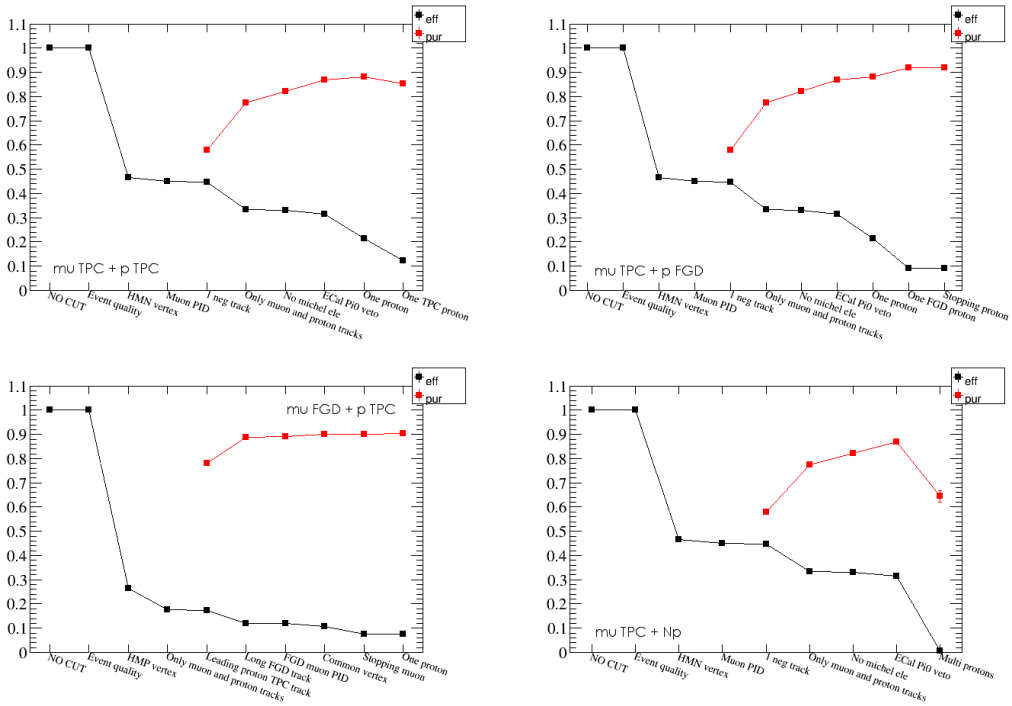


Figure 6.11: Selection efficiency and purity after p phase-space constraints (excluding out of phase-space events) as a function of cut in each sample. To ensure manageable analysis trees, sufficient information to calculate the purity is only stored if the event passes a pre-section of cuts (which identify CC-inclusive events).

.

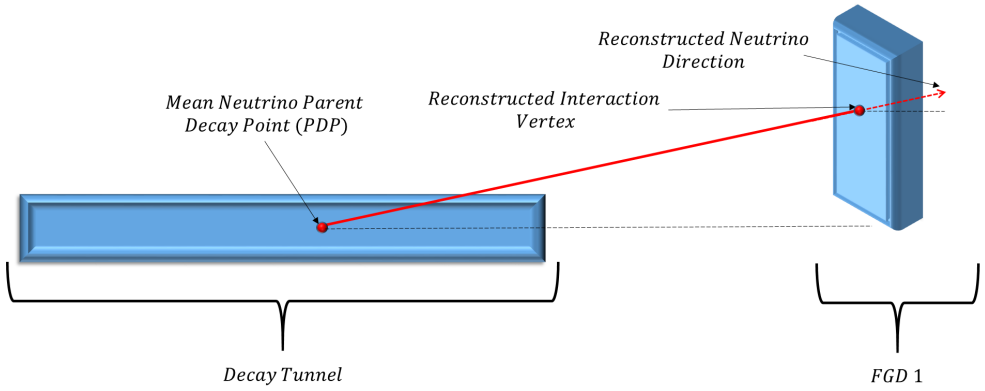


Figure 6.12: The schematic of the procedure to reconstruct the incoming neutrino direction from the neutrino PDP and the FGD1 vertex position.

6.4.1 Expected STV distributions

Once an estimate of incoming neutrino direction is acquired, the muon and proton kinematics can be used to form the STV distributions of selected events, as described in chapter 3. The STV formed from the true kinematics of selected events (using the true simulated neutrino direction rather than that derived from the neutrino PDP), in the binning that will be used in cross section extraction, is shown in figure 6.13 stacked by topology and in figure 6.14, stacked by interaction mode. The same is shown for reconstructed events in figure 6.15 and 6.16 whilst the detector smearing matrices are shown in figure 6.17 and the efficiency as a function of the true STV is shown in figure 6.18 from both the GENIE 2.8.0 and NEUT B input simulations. Finally the signal component of the true STV are shown in figure 6.19, split by interaction mode to show the expected composition of events following cross section extraction (according to the input NEUT 6B simulation).

Overall these distributions highlight that the selection is fairly pure for all values of the STV but that the truth to reconstructed-level smearing between adjacent bins is substantial. The background is typically most prominent at the high end of all three STV but for δp_T and $\delta\phi_T$ the smearing between true bins is typically larger at lower values.

Figure 6.18 primarily shows that the efficiency distribution in the STV is quite flat, particularly in regions most susceptible to detector smearing. This ensures that the misplacement of events into adjacent bins will not be blown up by the efficiency correction. There is mostly excellent agreement for the predicted efficiency as a function of the STV between NEUT 6B and GENIE 2.8.0, indicating that the deviations from a flat efficiency distribution are more likely to be well modelled. The exception to this is a deviation between the GENIE and NEUT efficiencies in figure 6.18 at high $\delta\alpha_T$ and $\delta\phi_T$, which is problematic, but somewhat expected. Even if phase-space constraints were chosen to ensure a 100% selection efficiency for individual outgoing protons or muons, a lower efficiency would be seen once events include the protons and muons together. This is because there are some cases where the outgoing particles are so close to being co-linear, that they cannot be individually resolved. The probability of finding a co-linear proton and muon in a selection is dependent on the interaction model and so unfortunately introduces a practically irreducible model dependence in any analysis which considers events with more than one final state particle. High values of $\delta\phi_T$ (and, albeit more indirectly, $\delta\alpha_T$) correspond to events in which the outgoing selected proton and muon candidates are more likely to be co-linear in the transverse plane. These events are then also much more likely to be co-linear in 3D, and are therefore more likely to suffer from this inefficiency effect. The fake data studies in chapter 8 will allow assessment of any subsequent bias in the cross-section result.

Figure 6.19 reiterates some of the conclusions from chapter 3 regarding sensitivity to

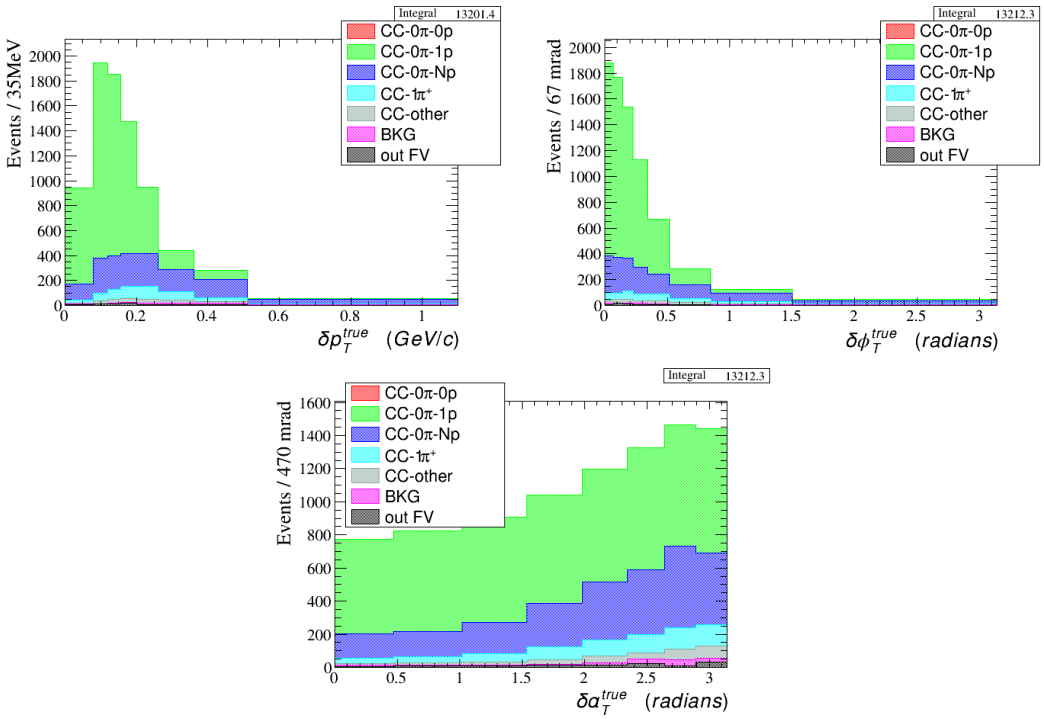


Figure 6.13: The STV formed from the true kinematics of selected events broken down by true signal topology. The binning is that which will be used for cross section extraction.

2p-2h contributions. In particular, the shape of the higher end of the signal portion of the STV distributions is largely impacted by 2p-2h contributions, indicating potential scope for extracting the 2p-2h component of a CC0π+Np cross-section. However, the proton kinematic phase-space constraints exclude 43% of the selected 2p-2h contribution in NEUT, making the relative 2p-2h contribution fairly small in any analysis bin. A substantially improved reconstruction of low momentum protons would improve the sensitivity to 2p-2h (as implemented in NEUT).

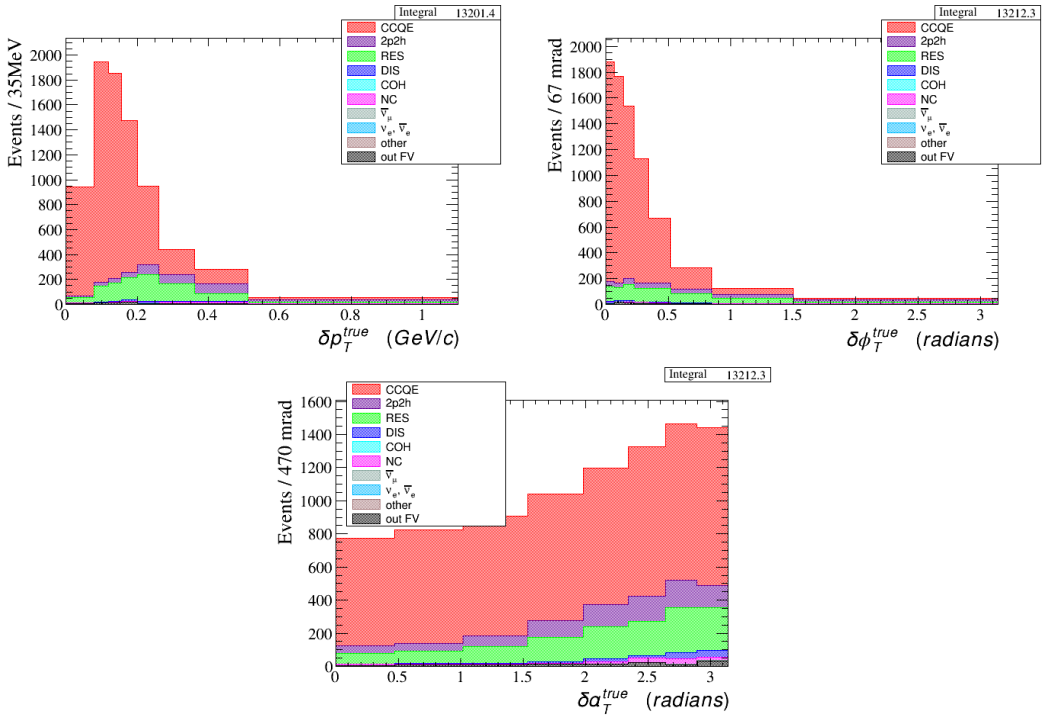


Figure 6.14: The STV formed from the true kinematics of selected events broken down by true signal interaction mode. The binning is that which will be used for cross section extraction.

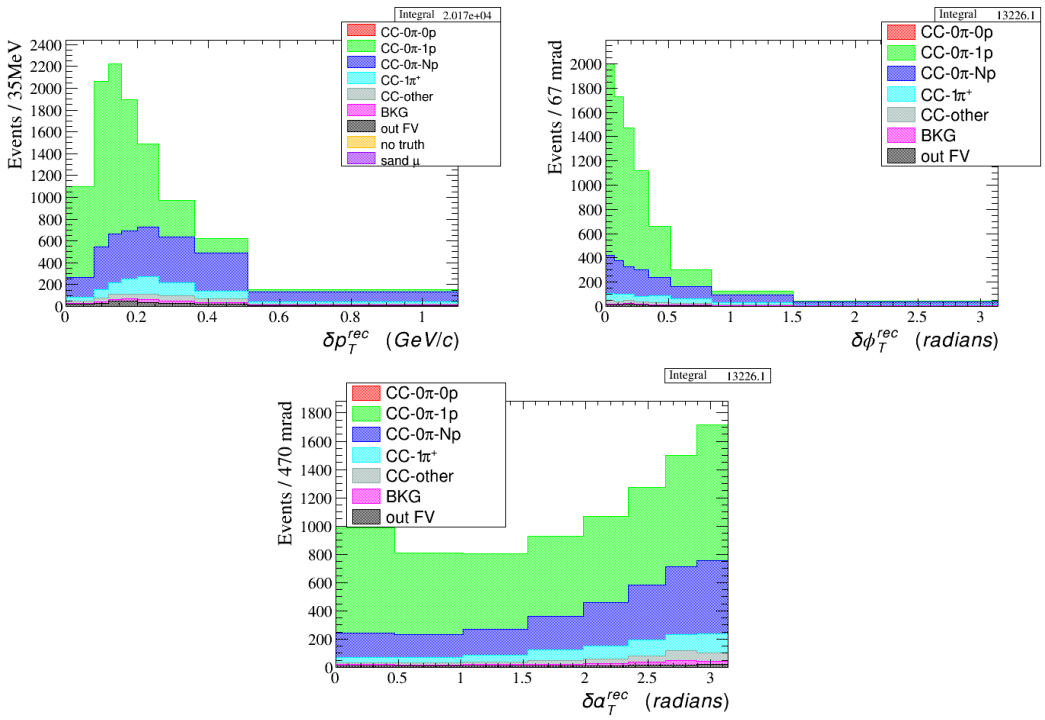


Figure 6.15: The STV formed from the reconstructed kinematics of selected events broken down by true signal topology. The binning is that which will be used for cross section extraction.

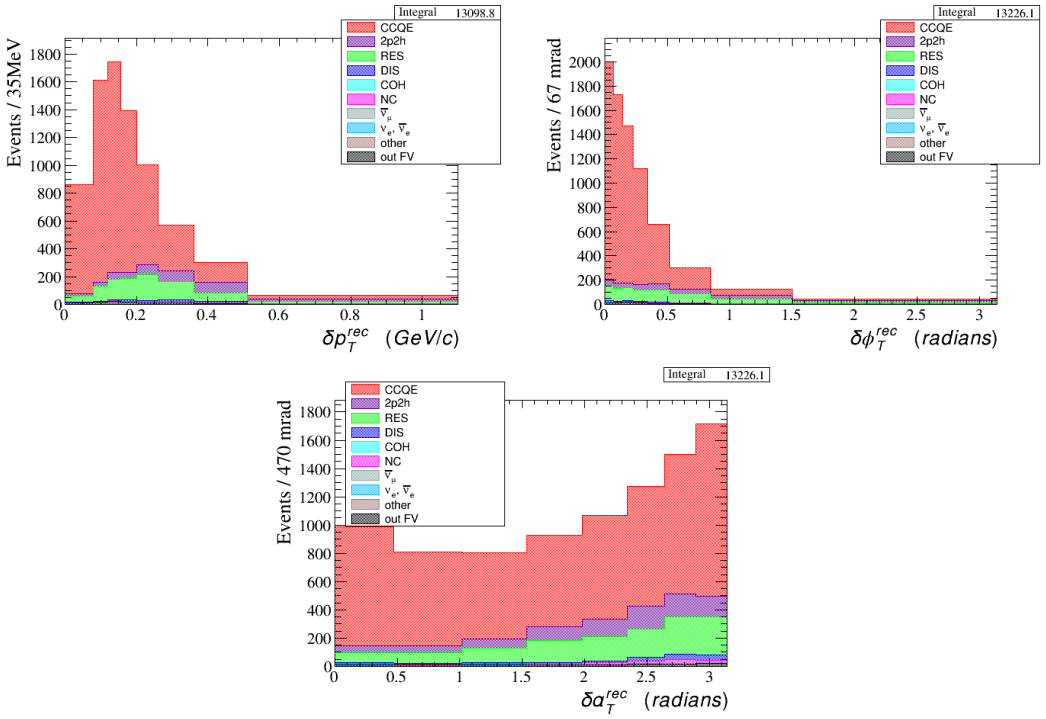


Figure 6.16: The STV formed from the reconstructed kinematics of selected events broken down by true signal interaction mode. The binning is that which will be used for cross section extraction.

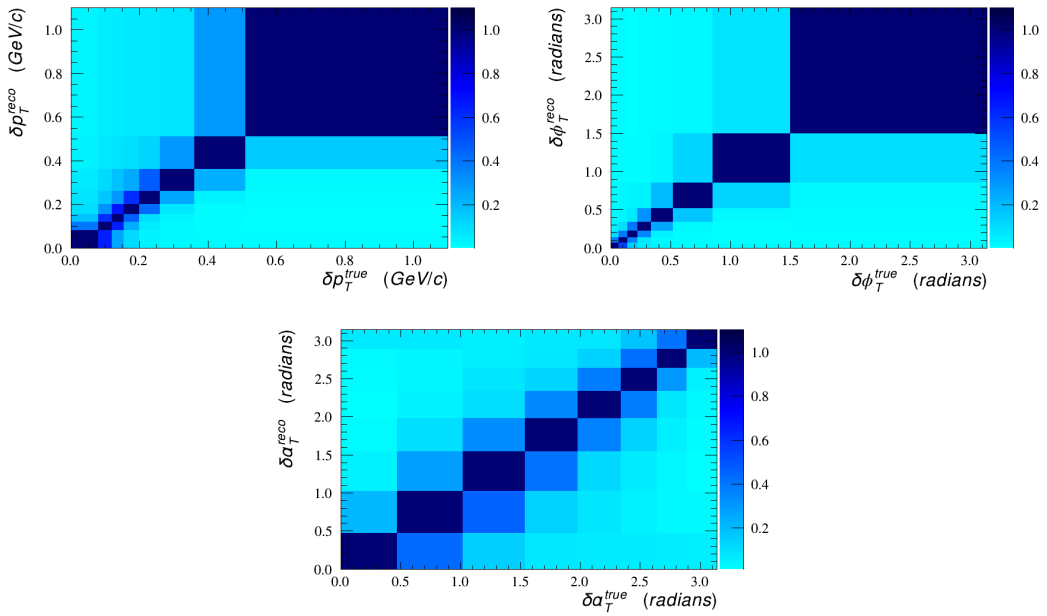


Figure 6.17: The row normalised smearing matrices for the STV. These indicate ND280's ability to accurately reconstruct the STV. The binning is that which will be used for cross section extraction.

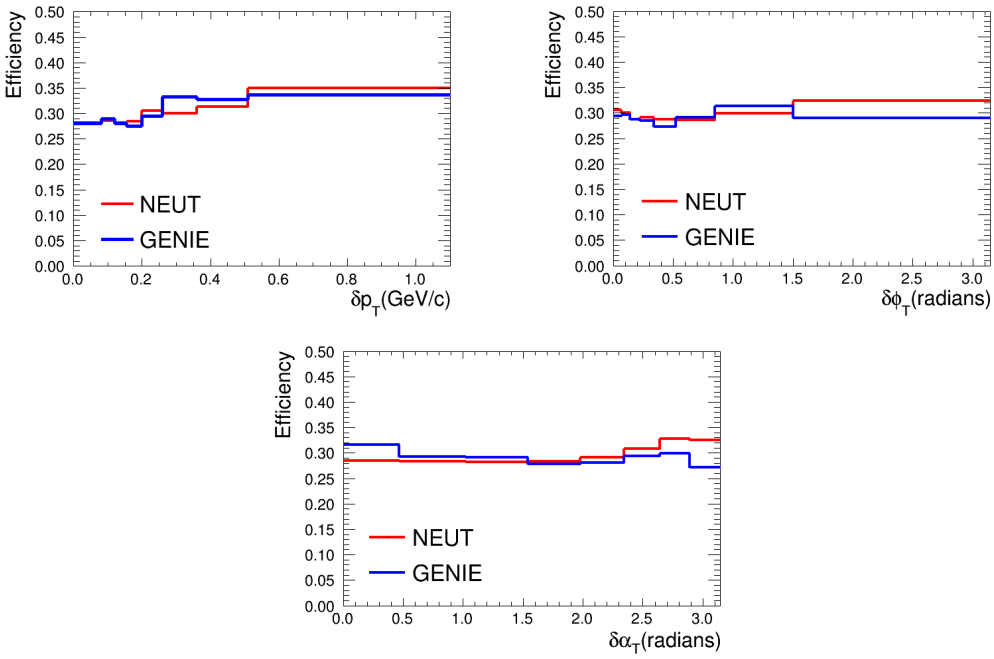


Figure 6.18: The selection efficiency as a function of STV bin is shown for both GENIE and NEUT. The difference in the predictions is indicative of the interaction model dependence that will feed into the efficiency correction during cross section extraction. The binning is that which will be used for cross section extraction.

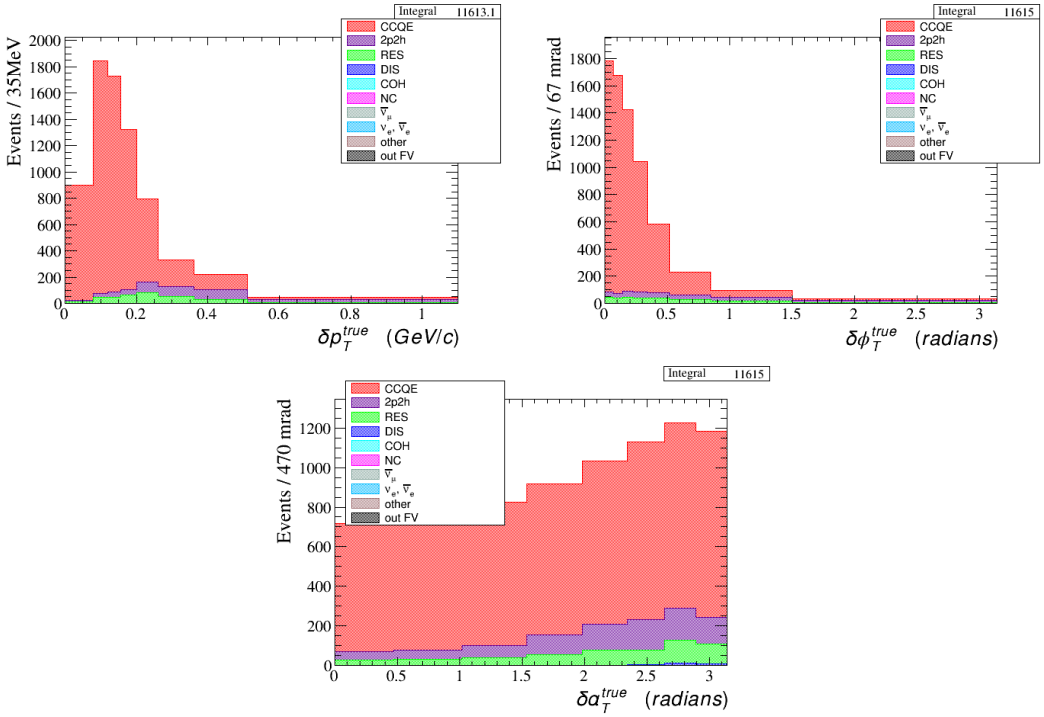


Figure 6.19: The signal portion of selected events broken down by interaction mode. This indicates the expected interaction mode contribution in each STV bin following cross section extraction. Although it is clear that the 2p-2h contribution has an impact on the shape of the STV distributions, it's contribution is small due to the rejection of all events without a proton with momentum greater than 450 MeV/c. The binning is that which will be used for cross section extraction.

6.5 Sideband samples

Although the selection presented in this chapter selects a high purity of CC0π+Np events, there are still non-negligible backgrounds. Section 6.2 demonstrates that the majority of these come from CC1π^+ events where the pion (and associated Michel electron) are missed, but there is also notable contribution from other (multi-pion) CC events.

These backgrounds can be constrained through an analysis of data-simulation agreement in 'sideband' selections which characterise them (also known as control regions). The method of using sidebands to constrain the backgrounds varies depending on the cross-section extraction strategy. In the fitting methodology used in this thesis (detailed in chapter 7), background model parameters (such as M_A^{RES} and C_A^5, discussed in chapter 2) are fit in the sideband region simultaneously to the parameters describing the signal. It is important to note that, although the background is constrained by them, the sidebands are only free to alter the background within the parametrisation of a particular model, and so the analysis is still somewhat dependent on the background models. However, such a dependence is relatively safe because of the high purity ($\sim$85%) nature of the analysis and because the dominant background (CC1π^+ events, corresponding to approximately half the background) are reasonably well parametrised by existing models (which were discussed in chapter 2). Furthermore, if the background model is insufficient to describe the sidebands, it will be clear in the goodness of fit to them. This will be discussed further in chapter 7, suffice to say that the sidebands should contain a high concentration of the dominant background.

To achieve these constraints on the analysis backgrounds, two sidebands are employed. One to specifically fit CC1π^+ events, and another to include CC1π^++Other events (called the 'CCOther' sideband). An illustration of the topologies these aim to identify is shown in figure 6.20. The cuts to form these samples are shown in figure 6.21 and are chosen to select high purities of their respective signal, whilst ensuring no event appears in more than one sample. These sideband selections are heavily based on those in [133] and used in the first analysis discussed in [106]. The cuts have mostly already been detailed in section 6.1.1, but the sideband specific cuts are briefly described below. Following this the selected events in the sideband regions are summarised and the formation of the STV in the sideband regions is discussed.

6.5.1 Summary of sideband cuts

HMP track is a π^+ This cut is passed if the HMP track is determined to look π^+-like based on the TPC PID method outlined in section 4.2.1. Note that π^+-like means the track is likely to be a π^+ or an antimuon, since the TPC cannot reliably distinguish these.

2 tracks This cut is passed if there is exactly two tracks within the FGD and TPCs.

Multi TPC Tracks This cut is passed if there are more than two tracks in the TPCs.

6.5.2 Summary of selected sideband events

To verify the sideband regions are selecting high purities of their respective backgrounds, table 6.5 and 6.6 show the breakdown of events selected in each sample by topology and interaction mode respectively, whilst Figures 6.22 and 6.23 show the true proton and muon kinematics of events selected within the CC1π^+ and CCOther sidebands respectively. Overall these show that each sideband well characterises that background it is built to constrain and that the most substantial impurity in each is constrained by the other.

	Topology						
Sample	CC0π0p	CC0π1p	CC0πNp	CC1π^+	CCOther	Other	Events
CC1π^+ Control	0.00%	3.28%	1.53%	72.5%	13.7%	9.10%	877
CCOther Control	0.00%	0.31%	0.37%	11.3%	70.9%	17.2%	1307

Table 6.5: The proportion of each true signal topology in each reconstructed sideband sample alongside the expected number of events in available data (predicted by NEUT).

	Interaction Mode					
Sample	CCQE	2p2h	CCRES	CCOther	Other	Events
CC1π^+ Control	4.34%	0.22%	51.7%	17.8%	25.8%	877
CCOther Control	0.45%	0.01%	9.86%	71.95%	17.7%	1307

Table 6.6: The proportion of each true interaction mode in each reconstructed sideband sample alongside the expected number of events in available data (predicted by NEUT).

6.5.3 Formation of the STV in the sideband samples

In principle, by employing the sidebands to constrain background model parameters, the sidebands could be binned in different variables than the signal samples. However, in this analysis both the signal region and sidebands are binned identically in the STV. This ensures that the background model constraints are established from a sideband sample that well-characterises the background in the signal region. Furthermore the STV are formed in the sideband using the selected muon and the HMP track, such that are as representative as possible of the backgrounds in the signal region. The STV formed from the CC1π^+ and CCOther sideband regions are shown in figure 6.24 and 6.25 respectively.

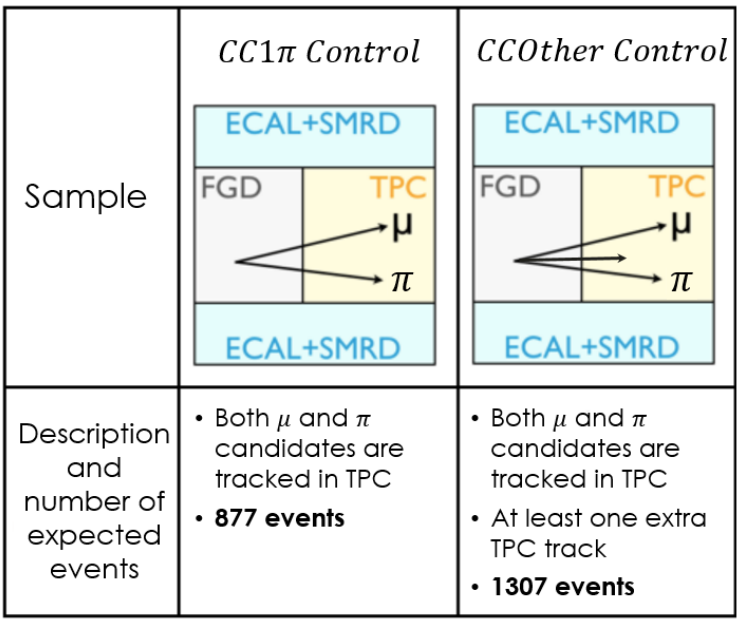


Figure 6.20: A simple breakdown of the different sideband samples used in this analysis. The number of events expected in each sample within available ND280 data is indicated.

CC1π Control	*CCOther Control*
Event Quality	
HMN Vertex	
Muon PID	
1 *Negative Track*	*HMP track is a* π^+
HMP track is a π^+	*Multi TPC Tracks*
2 *tracks*	

Figure 6.21: An overview of the selection cuts for each of the two sideband samples.

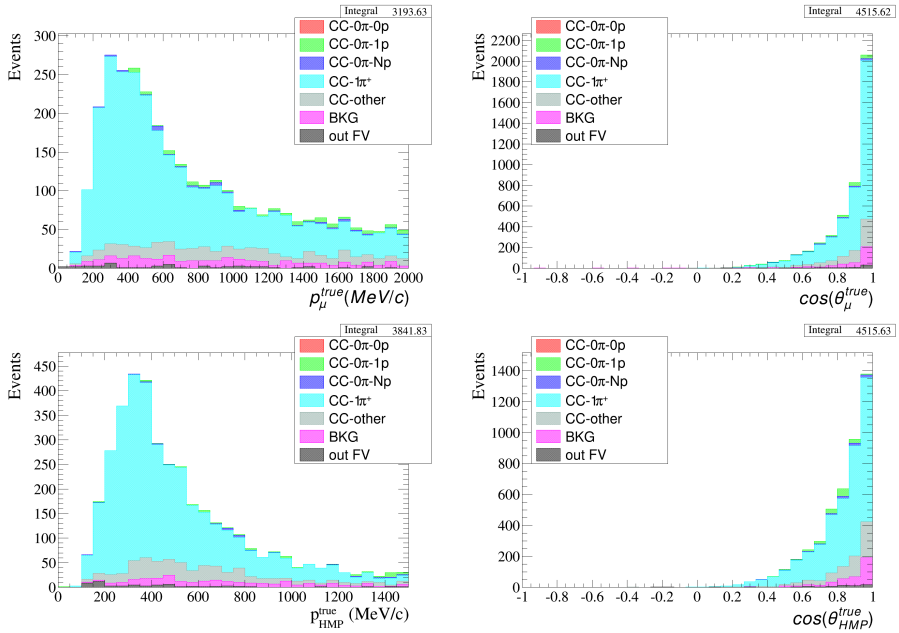


Figure 6.22: The true proton and muon kinematics of events passing the cuts for the CC1π^+ sideband sample, stacked by true event topology.

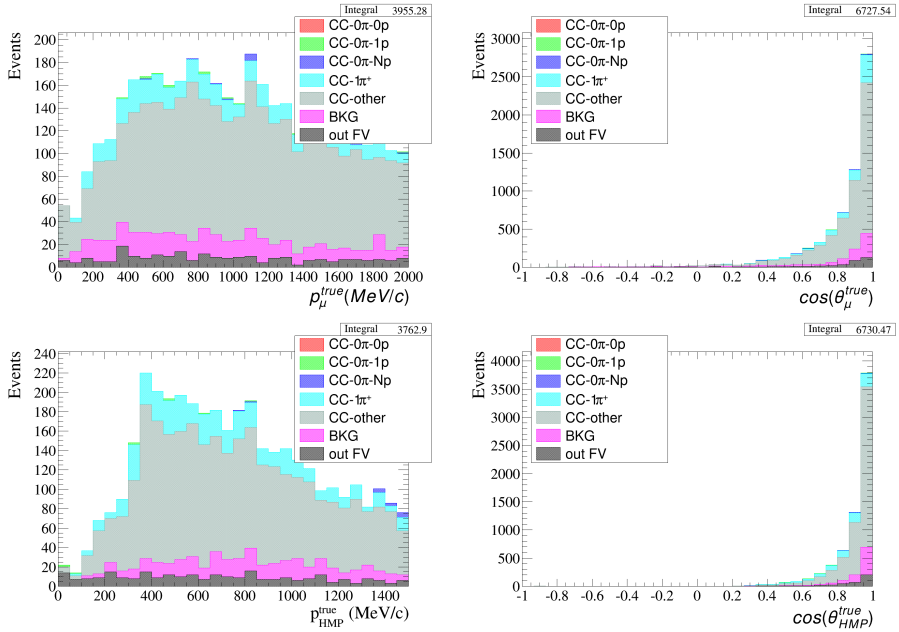


Figure 6.23: The true proton and muon kinematics of events passing the cuts for the CCOther sideband samples, stacked by true event topology.

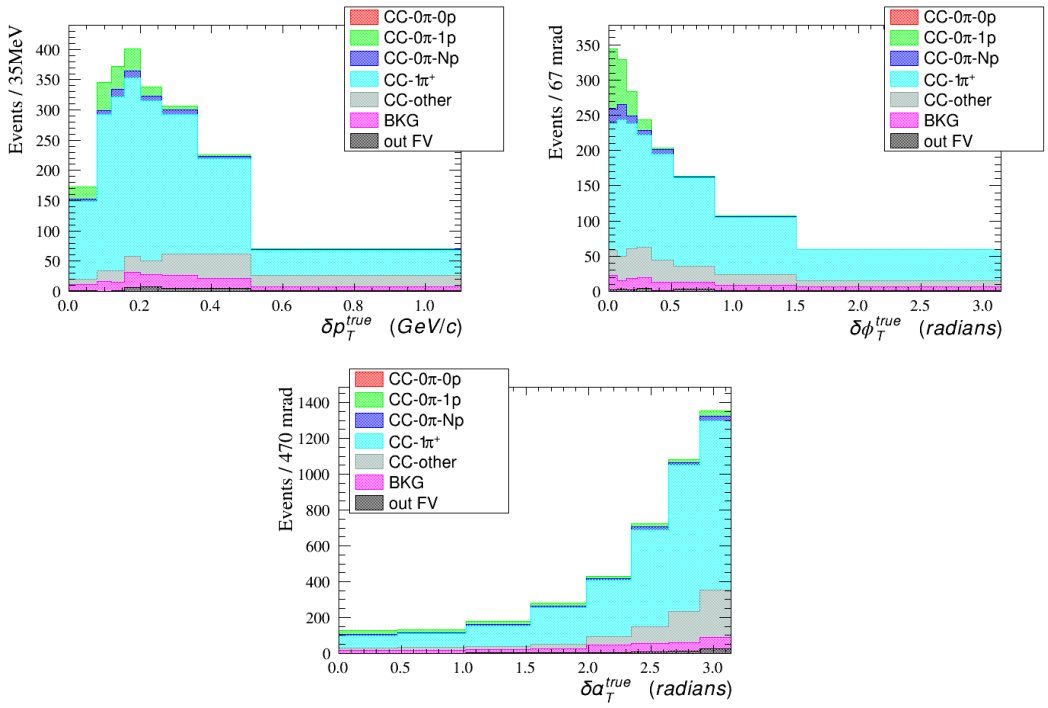


Figure 6.24: The STV formed in the CC1π^+ sideband sample from the true kinematics of selected events, stacked by true event topology. The binning is that which will be used for cross section extraction.

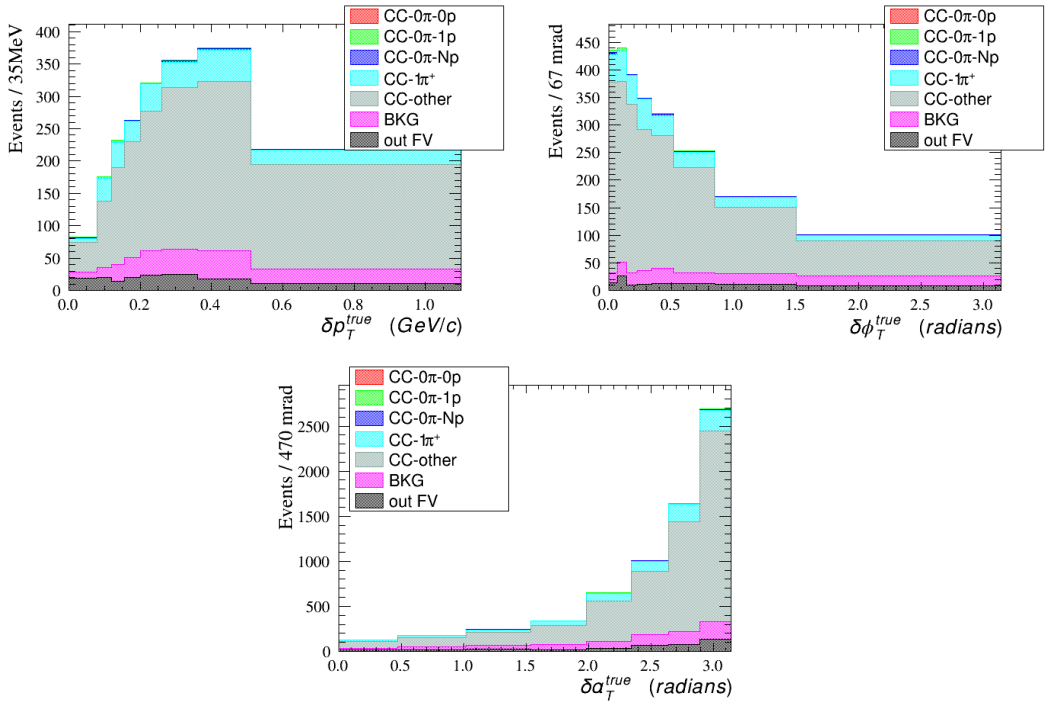


Figure 6.25: The STV formed in the CCOther sideband sample from the true kinematics of selected events, stacked by true event topology. The binning is that which will be used for cross section extraction.

6.6 Systematic uncertainties in the selection

Neither the propagation of particles through ND280 (via GEANT4), nor ND280's response to the ensuing tracks (via elecsim), are perfectly modelled in the input simulation. Uncertainties on these models must be accounted for in cross-section extraction. Thankfully, such systematic uncertainties have been extensively studied for both previous cross-section analyses, and so that ND280 can be used to supply useful constraints for oscillation analyses. A review of most of these uncertainties can be found in section IV.B.4 of [78] and throughout [133]. This section will review the role of the different detector systematic uncertainties associated with the selection, with particular focus on the new uncertainty required to account for the possible mis-modelling of the mean neutrino parent decay point in the secondary neutrino beamline (which is used to determine the incoming neutrino direction as described in section 6.4). Following this, the approximate effect on the uncertainty in the event rate in bins of the STV will be discussed.

6.6.1 Sources of detector systematic uncertainties

There are many potential sources for uncertainty in both GEANT4 and elecsim. This section reviews the uncertainties that are accounted for in this analysis, which are summarised in figure 6.26. The scale of each of most of these have been constrained using in-situ control samples (such as cosmic ray or sand muons), test beam data or, in the case of the GEANT4 uncertainties, through using data from other experiments. Systematics that have a substantial impace (greater than 1% on a single proton signal sample) are quantified in table 6.7.

Matching efficiencies The algorithm for matching TPC-FGD, TPC-ECal and FGD-ECal tracks has the potential to miss some hits or to fail entirely. These systematics account for the plausible variation in the failure rate for the corresponding matching [78, 133].

Outgoing particle kinematics There are several ways for the mis-modelling of the detector to directly influence the outgoing particle kinematics. Firstly, both the TPC and FGD methods of momentum reconstruction (described in section 4.2) have a limited resolution. This resolution has been shown to appear to be marginally different in data and simulation [78] and has a corresponding uncertainty to account for this (this is done separately for the TPC and FGD momentum resolutions). Furthermore the absolute scale of the momentum measurements are also subject to uncertainty. In the TPC this is dominated by variations from nominal magnetic field strength, whilst in the FGD this stems from imprecise track start and end positions alongside imperfect readings of the deposited energy [78, 133].

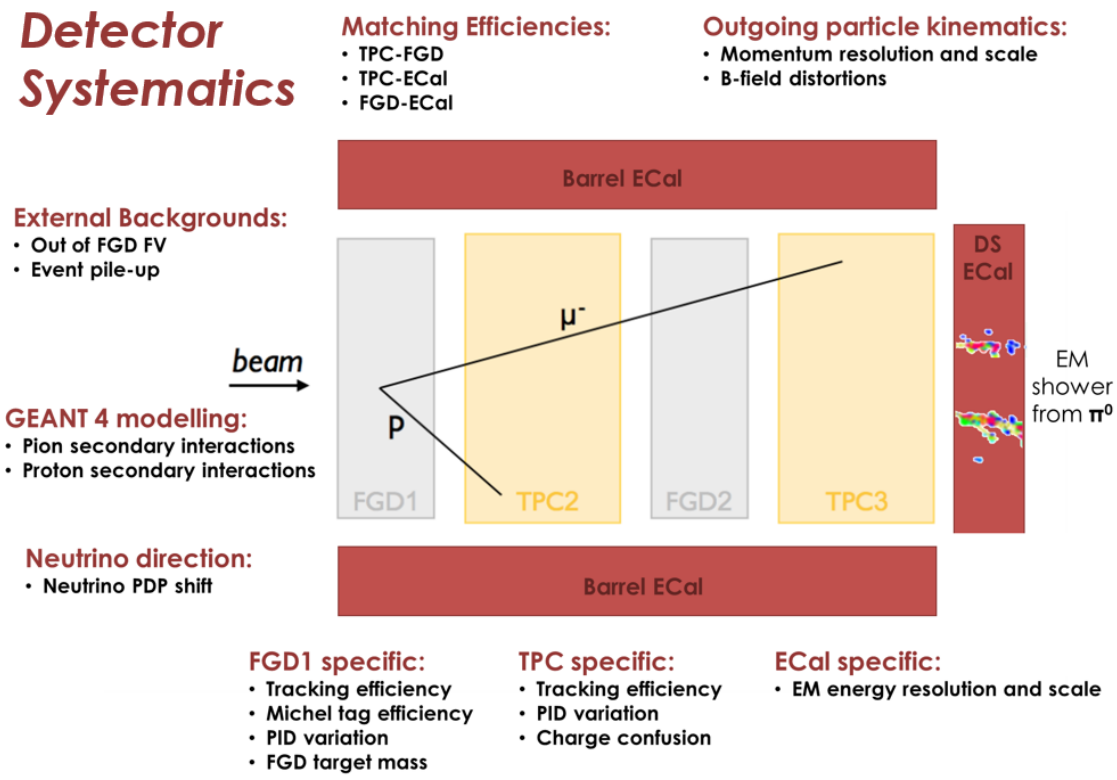


Figure 6.26: A schematic of the detectors relevant to this analysis alongside a list of the corresponding systematics. Image of the ND280 TPCs and FGDs taken from [133].

ECal specific uncertainties The ECal's efficiency at identifying the photons from π^0 events is altered by the precision to which it can measure the deposited energy. There is an associated systematic on the resolution and scale of this energy deposit [134].

TPC specific uncertainties A plausible variation in the number of events missed because of a failure in the TPC reconstruction to identify tracks or include hits in tracks causes an uncertainty on the selection efficiency. Furthermore, the potential for more or less failures in the TPC PID algorithm (described in section 4.2.1) due to a variation in measurement of energy resolution is also covered by a TPC PID systematic [78, 133].

FGD specific uncertainties The FGD is subject to similar tracking efficiency and PID systematics as the TPC. The FGD has an additional systematic to account for the plausible variation in the Michel electron tagging efficiency and another to account for a possible change in interaction rate due to a small uncertainty on the FGDs mass [78, 133].

Neutrino direction uncertainty This systematic accounts for possible variations in the mean neutrino parent decay point inside the secondary beamline, which is used to estimate the incoming neutrino direction. It is discussed in more detail below in section 6.6.1.

GEANT4 modelling uncertainties Pion and proton secondary interactions (SI) within the detector (but after leaving the nuclear medium of the interaction) have the potential to significantly alter the selection efficiency and the measured kinematics of events. The extent to which this happens can be constrained with dedicated nucleon and pion scattering experiments, but there is still significant uncertainty. This uncertainty is treated with the application of a systematic that allows the re-weighting of events that undergo secondary interactions [78, 133].

External backgrounds The rate at which entering backgrounds or event pile-up causes the mis-selection of events impacts the selection efficiency. There is a systematic to account for the plausible variation of this rate [78, 133].

Neutrino parent decay point systematic

As described in section 6.4, the mean neutrino parent decay point is required to calculate the incoming neutrino direction. However, there is a systematic uncertainty associated with the reconstruction of the neutrino direction in this way, since there is significant variation in the true PDP around its mean. This is accounted for by adding a systematic in which the

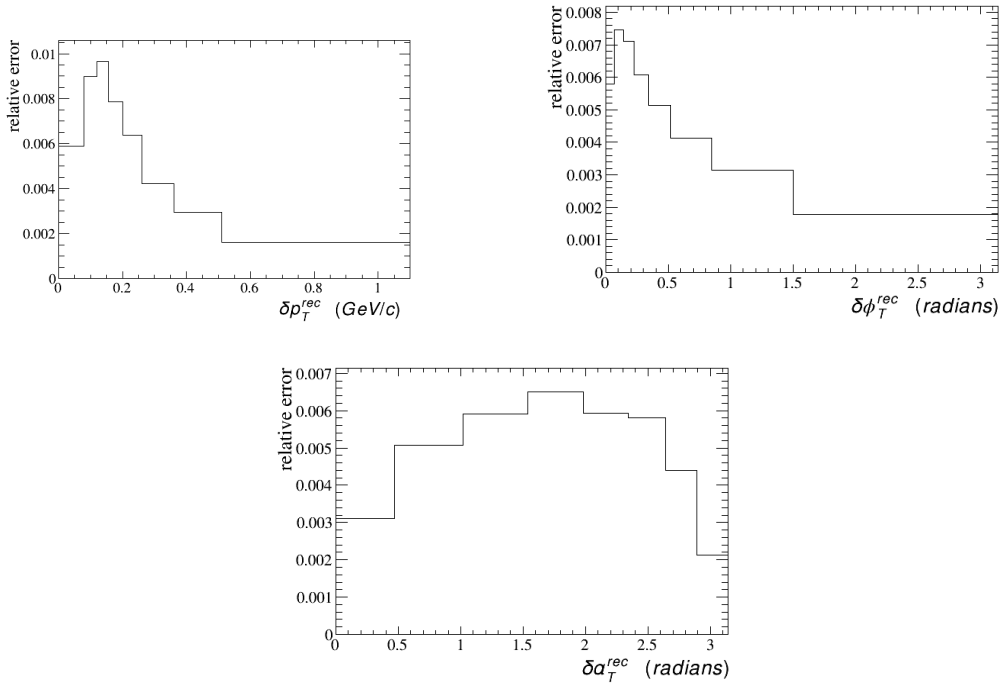


Figure 6.27: The relative standard deviation of the variation in the selected event rate in reconstructed STV bins after 500 variations of the neutrino mean PDP, as described in section 6.6.1. The binning is that which will be used for cross-section extraction.

effect of varying the mean PDP around a nominal suggested by the beam group is assessed. This is approximately 1-2 m in the the plane transverse to the beam, and around 50 m in the longitudinal direction [135]. 'Toy' experiments of different mean PDP co-ordinates are produced using a 3D sampling method based on the distribution of actual PDPs in the beamline simulation. More details can be found in [135]. Since this analysis is one of the first on T2K to use this method of neutrino direction reconstruction, the size of this systematic is considered separately to verify a reasonable response. The standard deviation in the STV for an ensemble neutrino PDP toys are shown in figure 6.27. This is found to be small compared to other detector systematics.

6.6.2 Total detector systematic uncertainty

Since there are a large amount of detector uncertainties that will all have potential correlations in the way they alter observables, it is difficult to analytically find a systematic error from the detector response on observables of interest. Instead, a systematic error is found numerically. To do this, the way in which a systematic allows the detectors response to vary, to reflect the associated uncertainty, is parametrised such that this response can be altered in a smooth way. The variation of these parameters is then described by a multi-variate Gaussian with a mean corresponding to the detector response and covariances that reflect the scale of

uncertainty, normally determined from comparison of data and simulation in control samples and test beam/external data. This multi-variate Gaussian can be described by a covariance matrix, of which the Cholesky decomposition is found and used to sample an ensemble of plausible variations in the detector. For each 'toy' variation the selected events are weighted or otherwise modified to reflect the change in detector response. Many such toys are created and used to build covariances between bins of the observables of interest, which describe the propagated detector systematic uncertainty. The covariance matrices themselves are shown and discussed further, as an input to the cross-section extraction framework, in section 7.2.3.

For illustration, the standard deviation in each bin of the observables can be calculated to form a relative error. Such relative errors are shown in figure 6.28, when 500 detector systematic toys are used to consider the variation of detector systematics on the STV. This demonstrates a typical variation of detector systematic parameters has a 3-5% effect on the event rate in each bin of the STV, but it should be noted that significant correlations exist between the bins. Since the parameters from each uncertainty source are sampled together it is not possible to determine the largest systematic uncertainty from this study. To attempt to determine the dominant sources of uncertainty, the parameters describing each systematic are sampled separately and the effect on the total event rate in each sample is considered. The results of doing this have shown that the dominant systematics are from uncertainty in secondary interactions, alongside a notable contribution from the TPC and FGD tracking efficiency. The systematics contributing a greater than 1% uncertainty individually in any single proton signal sample are shown in table 6.7 for each selected sample. Note that considering the total event rate in each sample disregards the effect of systematics shifting events between kinematic bins.

	Sample					
Source	μTPC+pTPC	μTPC+pFGD	μFGD+pTPC	μTPC+multi-p	CC1π^+	CCOther
Pion SI	3.5%	2.3%	0.6%	10.5%	3.8%	18.8%
Proton SI	0.6%	2.6%	0.8%	1.1%	1.6%	2.2%
FGD Tracking	1.4%	4.7%	5.6%	3.2%	2.2%	4.7%
TPC Tracking	1.4%	0.4%	0.8%	1.4%	1.1%	6.5%
TPC Clustering	1.2%	0.6%	0.7%	1.4%	1.4%	1.9%

Table 6.7: The relative detector systematic error on the total event rate in each sample from sources contributing greater than 1% to a single proton signal sample. SI stands for secondary interactions. Note that considering the total event rate in each sample disregards the effect of systematics shifting events between kinematic bins. TPC clustering refers to the variation in TPC reconstruction identifying the relevant hits in a track.

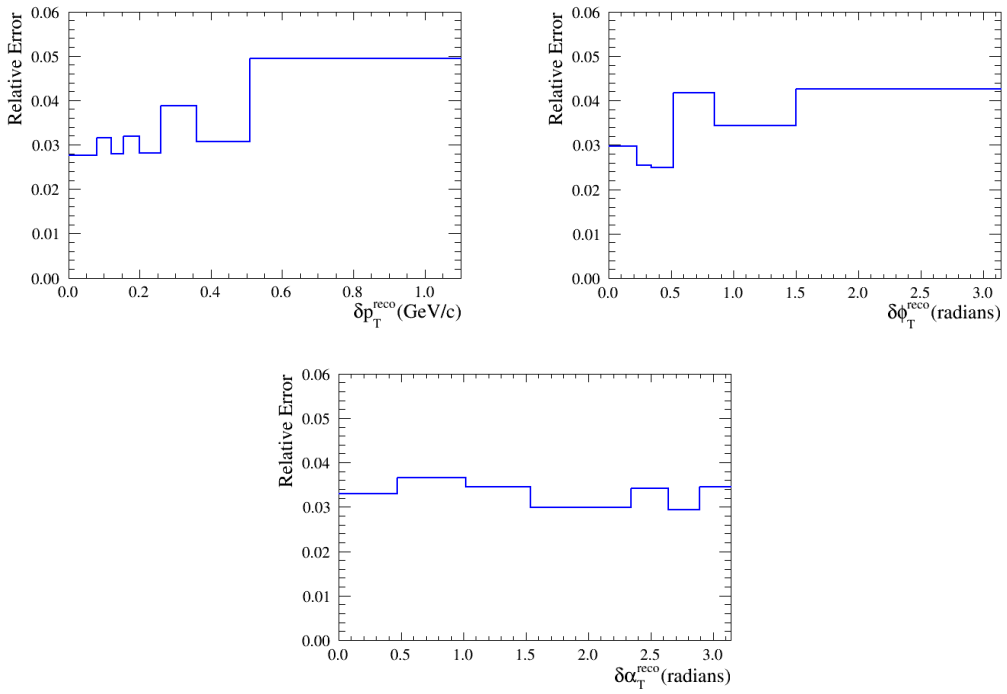


Figure 6.28: The relative standard deviation of the variation in the selected event rate in reconstructed STV bins after many toy variations of all detector systematics, as described in section 6.6.2. The binning is that which will be used for cross-section extraction.

Chapter 7

Binned likelihood fitting as a cross-section extraction technique

Binned maximum likelihood fitting has recently become the primary approach to cross-section extraction at T2K, due to the options it offers to allow the mitigation of model-dependence and an transparent interpretation of the results. It has thus far been used in two T2K publications: in [99] and [106]. The analysis presented in this thesis uses a regularised binned likelihood fit to go from the event selection presented in chapter 6, to a cross section. This chapter first, in section 7.1, gives a short general overview of binned likelihood fitting as a cross-section extraction technique with a particular focus on regularisation techniques, since neither of T2K's published results using this method employ these. Following this, the implementation of such a likelihood fit to extract a CC0π+Np differential cross section in the STV is described in section 7.2.

7.1 General overview of binned likelihood fitting

Binned likelihood fitting is a conceptually simple method of cross-section extraction, the idea is to allow the input simulation the freedom to vary via a set of parameters, such that a best fit set can be extracted once these parameters have been altered to the point where the simulation best describes data in some post event selection reconstructed bins. These parameters should describe the signal to extract and any background processes which may affect either the best set of fit parameters or their uncertainty. For example, in a very simple CCQE model fit, parameters that describe that signal (e.g. M_A) could be used as fit parameters but it would also be necessary to include nuisance parameters that describe background modes (e.g. M_A^{RES}, 2p-2h normalisation); detector effects (momentum resolution, PID ability e.t.c) and the flux. The best fit parameters are then those that maximise the following likelihood:

$$L = L_{stat(fit\,goodness)} \times L_{syst(penalty)} \times L_{reg}. \tag{7.1}$$

Because finding the point of maximum likelihood will involve taking the derivative of the likelihood function, it is actually much more convenient to work with the log-likelihood whose components can be split up into sums. Moreover, it is even more convenient to work with negative twice the log-likelihood since, in the limit of an infinite sample size, this is equivalent to a measure of the χ^2 (by Wilks' Theorem [137]). Because the logarithm is a monotonically increasing function, the log-likelihood achieves it's maximum value at the same place as the likelihood, so can simply be used in place of it. Therefore, equation 7.1 is equivalent to minimising the following negative log-likelihood:

$$-2\log(L) = -2\log(L)_{stat(fit\,goodness)} - 2\log(L)_{syst(penalty)} - 2\log(L)_{reg}. \tag{7.2}$$

Where:

$$-2\log(L)_{stat} = \sum_{j}^{reco\,bins} 2(N_j^{sim} - N_j^{obs} + N_j^{obs}\log\frac{N_j^{obs}}{N_j^{sim}}) \tag{7.3}$$

and

$$-2\log(L)_{syst} \approx \chi^2_{syst} = (\vec{a}^{syst} - \vec{a}^{syst}_{prior})(V^{syst}_{cov,prior})^{-1}(\vec{a}^{syst} - \vec{a}^{syst}_{prior}). \tag{7.4}$$

Here N_j^{sim} and N_j^{obs} are the number of events in each reconstructed bin, j. $\vec{a}^{syst}$ are systematic parameters which can describe the effect of parameters running over reconstructed bins. $\vec{a}^{syst}_{prior}$ are the prior values of these systematic parameters and $V^{syst}_{cov,prior}$ is a covariance matrix describing the confidence in the prior parameter values as well as the prior correlations between the parameters.

The term in equation 7.3 is negative twice the logarithm of the Poisson likelihood (where Stirling's approximation has been used to evaluate $\log(N_j^{obs}!)$ [138]). The term in equation 7.4 is the χ^2 goodness of fit between the parameters in $\vec{a}^{syst}$ and in $\vec{a}^{syst}_{prior}$, including the correlations between them described by $V^{syst}_{cov,prior}$, which corresponds to a multi-variate Gaussian likelihood.

Overall, $-2\log(L)_{stat}$ describes how well the data and the simulation agree in reconstructed bins: the smaller the discrepancy, the smaller this term is. $-2\log(L)_{syst}$ increases as parameters are moved far from their prior values (relative to the prior covariances set on them). This term allows the addition of prior knowledge (from theory or external data fits) into a fit. For example, if it was found that there is good agreement between data and a toy simulation with background model parameters far from what external data would imply, the $-2\log(L)_{syst}$ term will be large and therefore this will not likely be a best fit point. Another way of seeing this is that, without prior knowledge, there are often degeneracies in the fit solutions; it may be possible to fit the data by only changing the signal parameters or by only changing the background parameters. By adding prior knowledge regarding what some or all of these parameters should be, it becomes possible to resolve these degeneracies and extract useful information from the fit. Finally $-2\log(L)_{reg}$ is an optional regularisation term that will be discussed in section 7.1.2.

7.1.1 Template fitting as an unsmearing method

So far the fitting method has been discussed in the context of extracting model parameters, but this is by definition a maximally model dependent fit (there is only freedom for the simulation to move within the confines of the input models). As discussed in chapter 2, neutrino-nucleus interactions are not well understood and therefore should certainly not be confined to one model, so it is important to extract a result that is not dependent on the input signal model.

One method of facilitating such model independence is to allow 'template weights' on the nominal number of signal events in each *true* analysis bin to be fit (with no prior constraint), instead of fitting signal model parameters. The aim of a fit using this method is the same as the aim of other unfolding or unsmearing methods; to produce a data spectrum that is deconvoluted from detector smearing. The number of selected signal events in each true bin of the input simulation is varied and the effect is propagated through to the number of selected events in reconstructed bins of some observable (using the known detector smearing). This is then compared with the selected data such that a best fit number of signal events in each true bin can be found. In other words, this method constructs reconstructed 'templates' for the signal portion of each true analysis bin, as illustrated in figure 7.1, before fitting their normalisations (the template weights).

In the absence of any systematic fit parameters, these template weights can be written as c_i (where i is a true bin index) to describe the modified number of signal events in the i^{th} true bin:

$$N^{signal}_{i,\,true} = c_i N^{sim\ signal}_{i,\,true}, \tag{7.5}$$

where $N^{sim\ signal}_{i,\,true}$ is the number of simulated signal events in the i^{th} true bin.

The propagation of altered template weights to give the number of signal events in each reconstructed bin, j, can then be written as:

$$N^{signal}_{j,\,reco} = \sum_{i}^{true\ bins} c_i N^{sim\ signal}_{i,\,true} U_{ij}, \tag{7.6}$$

where U_{ij} is the detector smearing matrix from true bin i to reconstructed bin j. The *total* number of events in each reconstructed bin ($N_{j,\,reco}$), used to calculate the likelihood in equation 7.3, is then simply:

$$N_{j,\,reco} = \sum_{i}^{true\ bins} \left(c_i N^{sim\ signal}_{i,\,true} + N^{sim\ bkg}_{i,\,true} \right) U_{ij}, \tag{7.7}$$

where $N^{sim\,bkg}_{i,\,true}U_{ij}$ is the number of *reconstructed* background events in reconstructed bin j[1].

This can then be rewritten to allow for the variation of the neutrino flux, detector and interaction model systematic parameters:

$$N_{j,\,reco} = \sum_{i}^{true\,bins} \left[c_i \left(\prod_{a}^{model} w(a)_i^{signal} \right) N^{sim\,signal}_{i,\,true} + \left(\prod_{a}^{model} w(a)_i^{bkg} \right) N^{sim\,bkg}_{i,\,true} \right] U_{ij} r_j^{det} \sum_{n}^{E_\nu} w_n^i f_n. \quad (7.8)$$

- The product $\prod_a^{model}$ runs over the systematics related with the theoretical modelling of signal and background. Each $w(a)_{ij}$ acts as a weight to account for the effect of altering some model systematic parameter, a, (e.g. M_A^{RES}) on either signal events in a particular true bin or events of a particular background type, both characterised by the index i. Each model parameter has pre-fit prior values and covariances, which can be obtained by fits to external data.

- r_j^{det} are parameters describing the detector systematics encoded as uncertainty on the number of events in each reconstructed bin. The prior values are all unity (the nominal detector simulation) and a covariance matrix can be calculated using an ensemble of toys which sample detector systematic parameters, as described in section 6.6.2.

- The index n runs over bins of neutrino energy, the weights w_n^i describe the relative contribution of each neutrino energy bin in true analysis bin i. The parameters f_n describe the uncertainty on the flux and they are distributed around unity (which represent the nominal flux tuning) with some covariances. On T2K, these covariances are typically provided by the beam group from the flux prediction discussed in section 4.1.3.

The number of background events in each bin can be further constrained by fitting sideband regions simultaneously to constrain background model related systematic parameters. These sideband regions must be statistically independent from each other and all signal regions and should contain as few signal events as possible, such that they are minimally altered under variations of the template weights.

The final result of the fit is a set of template weight, model, flux and detector parameters, alongside covariances between them. The input simulation can be reweighted with this set of parameters to obtain an unsmeared (truth-level) distribution of selected signal events:

$$N_i^{signal} = \left[c_i \prod_{a}^{model} w(a)_i^{signal} \sum_{j}^{reco\,bins} r_j^{det} (U_{ij})^{-1} \sum_{n}^{E_\nu} w_n^i f_n \right] N^{sim\,signal}_{i,\,true}, \quad (7.9)$$

[1]Note that the background events do not really belong to a well-defined true bin in this framework. To make sense of this, consider i to run over the desired true signal bins but then to continue over some pseudo-truth bins which each contain the total background contribution of a particular background type, typically split by topology (CC1π, CCOther, OOFV etc.). U_{ij} then smears the contents of these effective bins into the reconstructed bins that are measured.

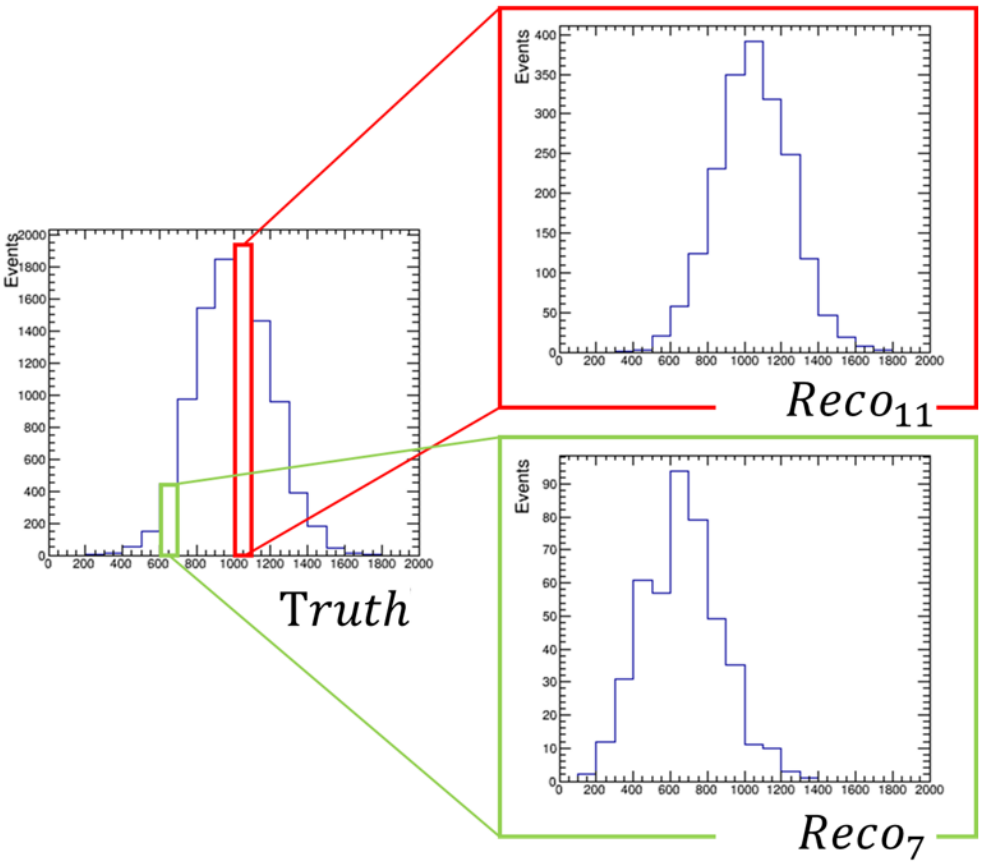


Figure 7.1: Toy simulations were produced using a Gaussian to first draw a "true" distribution, within which each event was smeared by a separate Gaussian to give a "reconstructed" distribution. The distribution on the left is the true distribution whilst the red and green boxes show the reconstructed distribution of events in the corresponding true bin. Here it can be seen how the reconstructed templates can be drawn from each true bin. Additionally the figure illustrates how true bins may have overlapping reconstructed templates, thereby motivating the need for regularisation.

whilst the covariances describe the uncertainty on this. One advantage of this method over a fit for model parameters, and many other types of unfolding, is that, with fine enough bins, the result is not inherently biased by some pre-conceived notion of what the input simulation suggests it should look like. The input simulation is still relied upon to model backgrounds, but the result is mostly independent of signal models (since the template weights have no prior constraint). Furthermore, it can easily be assessed whether the background models are suitable to describe the data by assessing the goodness of fit in sideband regions and monitoring whether the background model parameters converge to physically meaningful values.

Although this section has discussed the general concepts of maximum-likelihood fitting, the details of the likelihood maximisation method and the precise implementation of the fit parameters will be discussed in section 7.2.

7.1.2 Ill-posed problems and the role of regularisation

A potentially major issue with this method is that finding the best set of fit parameters is often an 'ill-posed problem', i.e. there may be multiple discrete sets of parameters which give a similarly good fit. The reason for this is shown with some toy data in figure 7.1. Here it can be seen that each true bin being fit in this example has a wide template in reconstructed bins, this means that small statistical fluctuations in the data can allow the fit to freely raise the template weight associated with some true bin so long as it reduces the template weights in adjacent bins. This leads to very strong anti-correlations between the fitted number of events in adjacent bins. This problem can be solved by reducing the overlap between the reconstructed templates for each true bin through phase-space constraints on the selected events; increasing the number of reconstructed bins; or decreasing the number of true bins, but all of these solutions damage either cross-section resolution or statistical power. An alternative method is to use an implementation of Tikhonov regularisation, which allows the insertion of knowledge that - between thin bins - cross sections (or event rates) should be relatively smooth, in order to loosely regulate adjacent bins [139].

In this analysis, regularisation enters as an additional penalty term in the fit, $-2\log(L)_{reg}$, where:

$$-2\log(L)_{reg} = p_{reg} \sum_{i}^{truebins-1} (c_i - c_{i+1})^2. \tag{7.10}$$

Here c_i is the template weight for the i^{th} true bin and p_{reg} controls the regularisation strength. This adds a penalty for having large differences to the shape of the input simulation. Different terms within the summation are possible, the quadratic term is chosen because it is the simplest stable option; the only simpler option ($|c_i - c_{i+1}|$) has a discontinuity in the first derivative for equal adjacent template weights.

It is clear that this implementation of regularisation adds a constraint which biases the fit toward the shape of the signal model in the input simulation. One may think that a way to avoid this model bias would be to regularise the difference between the absolute number of events in adjacent bins, rather than the relative change from nominal (equivalent to using a flat input simulation) but this would bias the result toward a flat result, which is less realistic than any simulation. The difficulty with implementing a regularisation term comes with choosing a sensible regularisation strength such that the impact of the bias is mitigated. If the strength is too low then the regularisation will not resolve the fit degeneracies but if the strength is too high then the extracted unsmeared distribution will be biased toward the shape of the input simulation. Indeed in the case of a very large regularisation parameter the template weights become degenerate and the fit effectively just determines a scaling factor to the signal component in the input simulation.

The impact of altering regularisation strength is demonstrated in figure 7.2, which shows the effect of increasing p_{reg} in a toy fit. In this toy, an input simulation is produced to be effectively flat in true-bins of some variable of interest, this is series 2 in figure 7.2 (the enumeration of the series is defined in the legend). Alongside this, the Gaussian distributed toy simulation from figure 7.1 is taken as fake data, this is series 3. Fake 'reconstructed' bins of the input simulation and the fake data are produced by applying the same smearing shown in figure 7.1 to series 2 and 3: these form series 5 and 6 respectively - the latter of which is what would be measured in a real experiment. Template weights of each true bin are produced from the input simulation and a likelihood fit to the fake data is used to minimise $-2\log(L)_{stat} - 2\log(L)_{reg}$. The best fit reconstructed-level result (i.e. N_j^{sim} in equation 7.3) is shown in series 4. The extrapolated truth-level result (the best fit template weights multiplied by each truth bin the input simulation) is shown in series 1. This fit is repeated for different values of p_{reg}. It can be seen that with little or no regularisation there are large degeneracies in the template weights, leading to a truth-level fit result with large anti-correlations between the fitted number of events in adjacent bins (as was discussed in the text above equation 7.10). This causes the result to 'zig-zag' around the fake data truth. By increasing the regularisation strength results with these zig-zags are penalised so the result is successfully smoothed. However, when p_{reg} is set too large the result is biased toward the flat input simulation such that the truth-level result becomes 'squashed' and there is substantially worse agreement between the fit result and the fake data at the reconstructed level. This demonstrates the importance of choosing p_{reg} carefully to find a balance between these extremes.

A simple method of choosing p_{reg} is the 'L-curve' technique presented in [140] and well described in TUnfold software documentation [141]. In this method, a compromise is found between the impact of the regularisation (defined by the normalised regularisation penalty: $-2\log(L)_{reg}/p_{reg}$) and the compromised goodness of fit (decreased $\log(L)_{reg}$). The regularisation parameter is found by plotting one against the other and extremising curvature of the resultant L-shaped curve. The application of this method to the toy fit is shown in figure 7.3, demonstrating it to be effective at extracting a reasonable regularisation parameter.

It is important to check that a fit is not too sensitive to the regularisation strength. An effective regularisation bias can be found by seeing how the fit result changes when p_{reg} is moved around the point of extremised curvature.

Whilst the fit result may appear to be more compatible with the fake data truth (i.e. series 1 better resembles series 3) after the introduction of regularisation, it is important to note that the least biased result is really the unregularised result ($p_{reg} = 0$). In general the unregularised result will also be in best statistical agreement with the truth so long as the

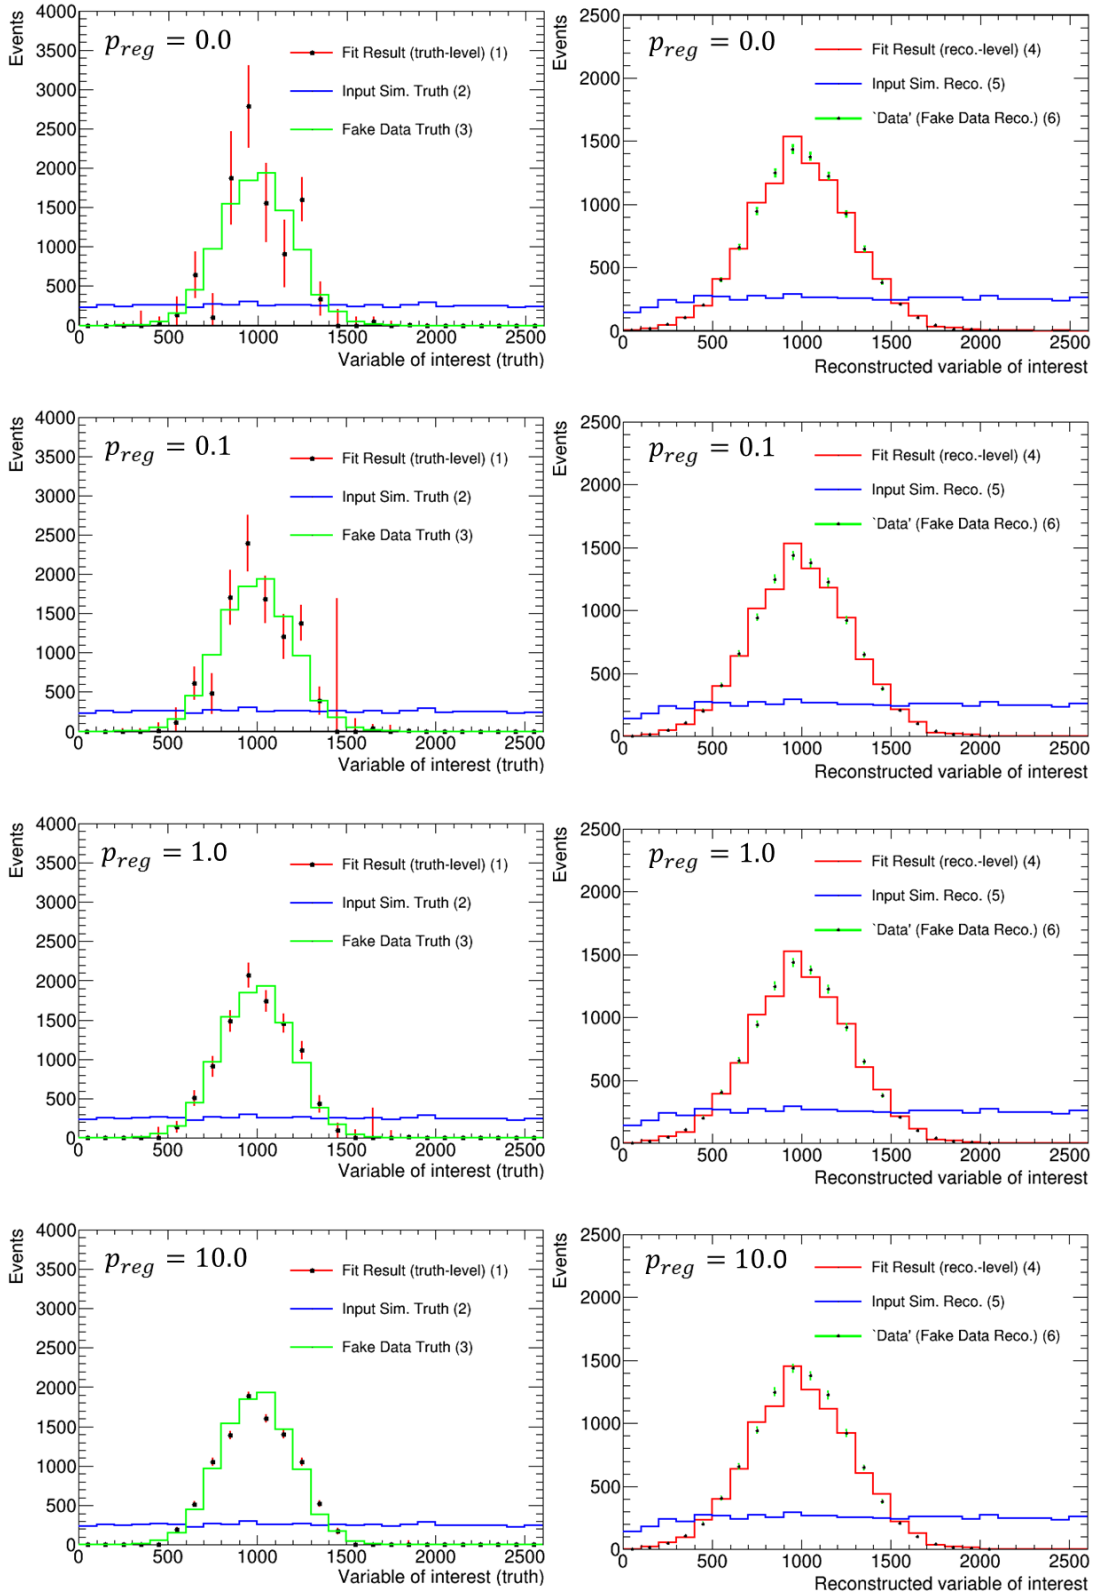


Figure 7.2: A toy data study to demonstrate the effect of changing the strength of regularisation on the fit result. The toy used in figure 7.1 is the fake data whilst the input simulation is a very wide Gaussian (such that it is effectively flat) in true bins smeared the same way as the fake data. A fit is performed with various regularisations strengths (indicated on the plot), as described in section 7.1.1, to obtain the truth and reconstructed level result - shown on the left and right respectively. This demonstrates how regularisation can resolve degeneracies in the fit, removing the apparent anti-correlations between adjacent bins in the fit result with no regularisation. More details are given in the text.

full covariance of the fit result is taken account of, where the aforementioned bin to bin anti-correlations will allow for the zig-zag shape. Overall regularisation allows the production of a result where the statistical compatibility can be judged by eye without the need for the full covariance. A regularised result is therefore particularly well suited for result-theory comparison plots but the unregularised result is likely more suitable for forming quantitative conclusions.

Regularisation as a penalty term

Examining the form of the $-2\log(L)_{reg}$ term, it may seem that some number is being added to the total log-likelihood which has no reason to be statistically meaningful. However, regularisation can be identically implemented as an extra component of the systematic penalty term:

$$-2\log(L)_{reg} \approx \chi^2_{reg} = p_{reg} \sum_{i}^{truebins-1} (c_i - c_{i+1})^2 = p_{reg}(\vec{c} - \vec{c}_{prior})(V^{reg}_{cov})^{-1}(\vec{c} - \vec{c}_{prior}). \quad (7.11)$$

For:

$$(V^{reg}_{cov})^{-1} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ -1 & 2 & -1 & \ddots & \vdots \\ 0 & -1 & 2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & -1 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}. \quad (7.12)$$

In order to make $(V^{reg}_{cov})^{-1}$ invertible a modification is required (the lower right element is modified)[2]:

$$(V^{reg}_{cov})^{-1} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ -1 & 2 & -1 & \ddots & \vdots \\ 0 & -1 & 2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & -1 \\ 0 & \dots & 0 & -1 & 2 \end{bmatrix}. \quad (7.13)$$

Now the $-2\log(L)_{reg}$ term is identical to the $-2\log(L)_{syst}$ term (see equation 7.4) and the regularisation can just be thought of as a constraint on the signal from prior knowledge that the cross sections should not deviate dramatically from the input simulation.

7.1.3 From a fit result to a flux-integrated cross section

The fit to the data gives the number of selected signal events in true bins. To get to a flux-integrated cross-section result, it is still necessary to correct for the efficiency, the

[2] in the analysis this modification isn't actually necessary as the regularisation is not applied to all template weights, this makes the regularisation covariance matrix invertible. See 7.2.2 for more details.

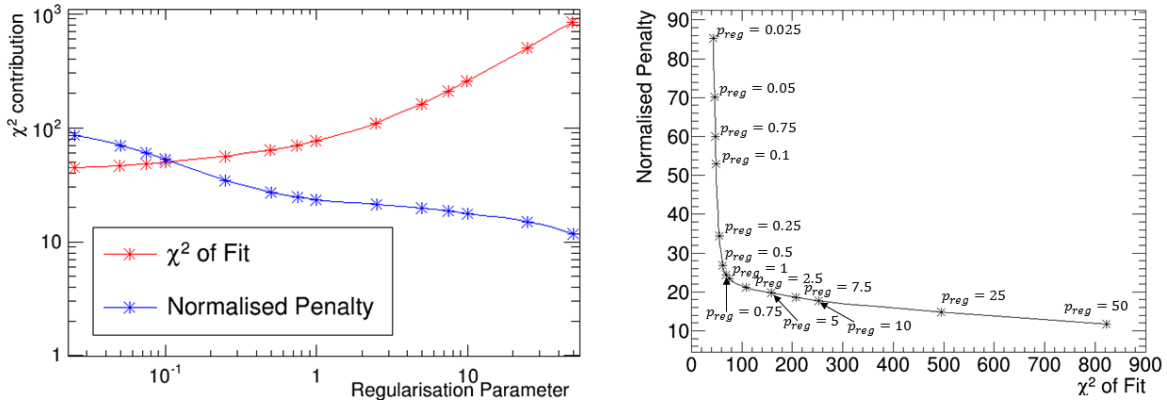


Figure 7.3: A demonstration of the L-curve method to extract an optimal regularisation strength from the toy study shown in figure 7.2. The left panel compares the variation of the normalised penalty term from regularisation and the overall χ^2 with the regularisation parameter. These two quantities are then plotted against one another in the "L-Curve" plot in the right panel. The chosen regularisation parameter is that which extremises the curvature of the L-Curve (approximately $p_{reg} = 1$).

integrated flux and the number of targets in the fiducial volume of the target detector, as shown in equation 2.20. Although the efficiency can only be estimated from the input simulation, the topology based and acceptance motivated signal definition should ensure that this correction is mostly coming from the (relatively well modelled) detector and particle propagation simulation (elecsim and GEANT4), rather than the neutrino scattering model (as discussed in section 6.3).

7.2 Implementation of a likelihood fitter

To extract the CC0π+Np differential cross section in the STV, a regularised likelihood template fit to the number of selected events in reconstructed STV bins, as described in the previous section, is used. This section will describe how such a likelihood fit is implemented and how its various inputs are obtained.

7.2.1 Analysis binning

The binning for each of the STV are shown in tables 7.1 to 7.3. Bins 0-7 are restricted such that they only contain events that are within the proton and muon kinematic phase-space that ND280 is most sensitive to (defined in section 6.3), these will be called 'within phase-space' (IPS) bins. In addition to the IPS bins, there is also a single out of phase-space (OOPS) bin (bin 8 for $\delta\phi_T$ and $\delta\alpha_T$, bin 9 for δp_T). δp_T additionally contains a 'buffer bin' (bin 8) discussed below. Although it is likely only the STV binned region that contains

interesting information about nuclear effects, considering this region in isolation leads to a substantial OOPS background that's largely immovable by the fit. By including a single OOPS bin, the fit gains freedom to rescale the true OOPS background (which contains some events reconstructed as IPS) with a single normalisation parameter.

The binning is chosen such that: there are a similar number of events in each IPS reconstructed bin (integrated over all samples); the statistical error on any bin in the primary signal samples (all signal samples apart from the multi proton sample) is not much larger than 10% (so that it is comparable to the minimum plausible systematic error); the change in expected cross section within a bin is small; and the fineness of the binning is not thinner than the detector resolution (as indicated by figure 6.17).

The reconstructed and true binning are identical but each signal and sideband sample from chapter 6 has its own set of reconstructed bins. Overall this means that there are 6 (the number of samples, including sidebands) $\times$ 9 or 10 (the number of true bins for $\delta\phi_T/\delta\alpha_T$ or δp_T) reconstructed bins (i.e. 60 or 54 entries in the j index of equation equation 7.8).

The δp_T buffer bin contains events with δp_T larger than the upper limit on what will be reported as a cross-section, it contains very few events and will not have a template weight that is allowed to move. It's main purpose is to facilitate the inclusion of events with a reconstructed δp_T above 1.1 GeV/c but a lower true value δp_T.

For the sideband samples, no events are placed in the OOPS bins, since OOPS in proton kinematics is not well defined for a sideband that often has no proton.

Bin Number	p_{proton} (GeV)	$cos(\theta_{proton})$	p_μ (GeV)	$cos(\theta_\mu)$	δp_T (GeV)
0	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.0-0.08
1	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.08-0.12
2	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.12-0.155
3	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.155-0.2
4	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.2-0.26
5	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.26-0.36
6	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.36-0.51
7	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.51-1.1
8	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	1.1-100
9	0.0-0.45 or 1.0-10.0	0.0-0.4	0.0-0.25	0.0-0.6	0.0-100

Table 7.1: The chosen binning for δp_T cross section extraction. The first column is an enumerator used to refer to the bins throughout the text and the others give the 5 dimensional bin boundaries in proton/muon kinematics and in δp_T.

Bin Number	p_{proton} (GeV)	$cos(\theta_{proton})$	p_{μ} (GeV)	$cos(\theta_{\mu})$	$\delta\phi_T$ (radians)
0	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.0-0.067
1	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.067-0.14
2	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.14-0.225
3	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.225-0.34
4	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.34-0.52
5	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.52-0.85
6	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.85-1.50
7	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	1.50-π
8	0.0-0.45 or 1.0-10.0	0.0-0.4	0.0-0.25	0.0-0.6	0-π

Table 7.2: The chosen binning for $\delta\phi_T$ cross section extraction. The first column is an enumerator used to refer to the bins throughout the text and the others give the 5 dimensional bin boundaries in proton/muon kinematics and in $\delta\phi_T$.

Bin Number	p_{proton} (GeV)	$cos(\theta_{proton})$	p_{μ} (GeV)	$cos(\theta_{\mu})$	$\delta\alpha_T$ (radians)
0	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.0-0.47
1	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	0.47-1.02
2	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	1.02-1.54
3	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	1.54-1.98
4	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	1.98-2.34
5	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	2.34-2.64
6	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	2.64-2.89
7	0.45-1.0	0.4-1.0	0.25-10.0	-0.6-1.0	2.89-π
8	0.0-0.45 or 1.0-10.0	0.0-0.4	0.0-0.25	0.0-0.6	0-π

Table 7.3: The chosen binning for $\delta\alpha_T$ cross section extraction. The first column is an enumerator used to refer to the bins throughout the text and the others give the 5 dimensional bin boundaries in proton/muon kinematics and in $\delta\alpha_T$.

7.2.2 Regularisation implementation

As discussed in section 7.1.1, each analysis bin has a template weight associated with it. For the IPS STV bins, these template weights are then subject to the regularisation described in section 7.1.2.

The OOPS contribution is regularised slightly differently, since the last STV bin and OOPS bin are not really adjacent. Instead the OOPS bin template weight incurs a penalty as it moves far from the weighted (by number of events) average of the template weight for the IPS bins. In other words deviations from the input simulation in the total cross section predicted in the restricted phase-space are loosely required to be similar in size to the rest of the phase-space. This OOPS-IPS bin regularisation is allowed to have a different regularisation strength than the IPS-IPS bin regularisation. This can also be determined by the L-Curve method described in section 7.1.2. The effect of this additional regularisation term is universally found to have a negligible impact on the fit results in all fake data studies

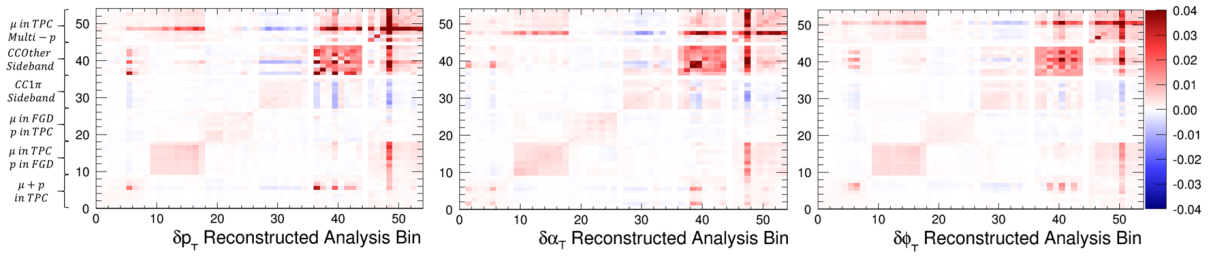


Figure 7.4: The covariance matrix used to construct the penalty term for detector systematic parameters, for each STV. Each sample used in the analysis corresponds to 9 detector systematic bins, as described in the text. The samples are in the following order: μTPC-pTPC, μTPC-pFGD, μFGD-pTPC , CC1π sideband, CCOther sideband, μTPC-MultiProton
.

performed (detailed in chapter 8) and so is left typically left to have $p_{reg}^{OOPS} = 1.0$.

7.2.3 Systematic parameters in the Fit

As described in section 7.1, the fit must include effects from plausible variation of detector, flux, and neutrino interaction models by fitting systematic parameters, which characterise these, simultaneously to the template weights. The fit is free to move these parameters in accordance to a covariance matrix defined alongside them, picking up a penalty in the log-likelihood as it does, as quantified by equation 7.4. The mechanism by which each of the parameters alter the fitted number of selected events in each bin is shown quantitatively in equation 7.8. The following paragraphs will briefly reiterate the role of each systematic parameter whilst discussing how the prior values of each of the parameters and their covariance matrices are chosen.

Detector systematic parameters

The detector systematics describe the uncertainty in ND280's reconstruction of events due to detector and particle propagation modelling. The systematic uncertainties considered are described in section 6.6 and are propagated to a set of covariances in the reconstructed STV bins, as described in section 6.6.2, using the reconstructed analysis binning discussed in section 7.2.1. The covariance matrix for each STV are shown in figure 7.4[3]. Row and column 35 and 44 are exactly 0, this is because these correspond to the OOPS bin for the sideband samples, which are empty. Particularly large event variances can be seen for the multi-proton and CCOther sideband sample, this is not surprising since ND280's ability to disentangle multiple tracks is limited and dependent on its modelling.

[3]The only deviation from this is that the buffer bin is merged with the final IPS bin in δp_T, so all three STV have 54 reconstructed bins.

Model systematic parameters

The model systematics parametrise the uncertainty in the neutrino-nucleus interaction models used during cross-section extraction (primarily the background models). These parameters are able to enter the fit as scaling factors on the associated physics parameter in the prior simulation, such that they have starting values of one. The prior covariances for each parameter are motivated from external data fits performed by the T2K neutrino interactions working group (NIWG) [101, 142, 143] with some modifications to give more conservative systematics. The covariance matrix used is shown in figure 7.5 and a brief description of the model systematic parameters used to calculate systematic uncertainties is given below:

E_BC**:** The binding energy of carbon nuclei.

M_A^{RES}**:** A relatively unconstrained Rein-Seghal RES cross-section model parameter described in section 2.3.1.

C_5^A**:** A relatively unconstrained Rein-Seghal RES cross-section model parameter described in section 2.3.1.

BG_{RES}**:** A scale factor on the ratio of non-resonant (isospin=$\frac{1}{2}$) to resonant (isospin=$\frac{3}{2}$) pion-production interactions within the Rein-Seghal model.

DIS_{NP}^{ISHP}**:** A parameter to allow a plausible neutrino energy dependent variation of the DIS cross-section, altering lower neutrino energy events more than larger neutrino energy events. This shape of variation reflects the uncertainties in underlying parton distribution functions, which are generally better understood at larger neutrino energy.

$CC_{COH}C$**:** A scale factor on the nominal CC coherent scattering cross section on carbon. The prior uncertainty on this parameter is inflated, relative to the NIWG external data fits, to allow the fit freedom to drastically reduce the number of CC coherent scattering events without much penalty. This is motivated by the MINERνA results showing a very poor agreement with the NEUT model predictions [144].

$CC\nu_e$**:** A scale factor on the total electron neutrino scattering cross section.

NC_{COH}**:** A scale factor on the nominal NC coherent scattering cross section.

CC_{OTHER}**:** A scale factor on all nominal NC non-coherent scattering cross sections.

M_A**:** The nucleon axial mass discussed in chapter 2. This parameter is omitted from the fit but is included when finding the systematic uncertainty on the efficiency correction in section 7.2.5.

$2p2h_{norm}$**:** A scale factor on the total 2p-2h contribution. This parameter is omitted from the fit but is included when finding the systematic uncertainty on the efficiency correction in section 7.2.5.

p_fC**:** The Fermi momentum of a carbon nucleus for CCQE events. This parameter is omitted from the fit but is included when finding the systematic uncertainty on the efficiency correction in section 7.2.5.

πFSI**:** These are scale factors on the different possible pion FSI processes.

Note that, as discussed in section 7.1.1, parameters that predominately describe the signal (M_A, $2p2h_{norm}$ and p_fC) are omitted from the fit. This is because the effects of these parameters are already included in the fit via the template weights. Including these parameters in addition to the template weights would therefore give the fit degenerate degrees of freedom that cause problems in cross-section uncertainty propagations. The effect of these parameters on the efficiency is accounted for and is evaluated separately (see section 7.2.5).

Since chapter 6 demonstrates that the background is predominantly from RES events, it is expected that the free parameters of the Rein-Seghal RES cross-section model described in section 2.3.1, M_A^{RES} and C_5^A, are likely to be the dominant sources of model systematics (although ideally these will be well-constrained by the CC1π^+ sideband sample). Although external data fits tend to anti-correlate M_A^{RES} and C_5^A (as they are both largely normalisation factors on RES events), the pre-fit parameters are left uncorrelated to give a more conservative uncertainty estimation. This same approach is followed by the most recently published T2K neutrino oscillation analysis [142].

To find the effect of altering a model parameter the input simulation can be re-weighted, as shown in equation 7.8. This involves calculating the effect of the altered model parameter and giving each event in the simulation a weight to reflect this using NEUT's and the NIWG's re-weighting tools, which are described in [145]. This avoids having to reproduce the entire input simulation with a tweaked parameter. However, as will become clear in 7.2.4, the method of likelihood fitting employed is necessarily single-threaded, and even this re-weighting is too slow to be used for every iteration toward the point of maximum likelihood.

	NC_{OTHER}	NC_{COH}	$CC\nu_e$	CC_{COH} C	DIS_{MP}^{ISHP}	BG_{RES}	M_A^{RES}	C_A^5	E_B C	M_A	p_f	$2p2h_{norm}$	πFSI_{CEX}^{HighE}	πFSI_{CEX}^{LowE}	πFSI_{ABS}	πFSI_{PROD}	πFSI_{INEL}^{HighE}	πFSI_{INEL}^{LowE}
NC_{OTHER}	0.09																	
NC_{COH}		0.09																
$CC\nu_e$			0.004															
CC_{COH} C				9														
DIS_{MP}^{ISHP}					1													
BG_{RES}						0.024												
M_A^{RES}							0.025											
C_A^5								0.014										
E_B C									0.13									
M_A										0.021		0.053						
p_f											0.020							
$2p2h_{norm}$										0.053		0.54						
πFSI_{CEX}^{HighE}													0.077			0.14		
πFSI_{CEX}^{LowE}														0.32				
πFSI_{ABS}															0.17			0.023
πFSI_{PROD}													0.14			0.25		
πFSI_{INEL}^{HighE}																	0.11	
πFSI_{INEL}^{LowE}															0.023			0.17

Figure 7.5: The covariance matrix used to construct the penalty term for model parameters, blank entries are zero.

Instead, each parameter is re-weighted up and down by three standard deviations (as defined from figure 7.5) and the effect on each true analysis bin in each sample of each interaction type (split by true final state topology) is calculated for reference during fitting and stored in a 'response function'. For example, one such 'response function' might see the effect of varying M_A^{RES} up and down by three standard deviations on true CC1π^+ events selected in the μTPC-pTPC falling within the first true δp_T bin. To allow the fit to evaluate the effect of any value of any model parameter, the response functions are turned into splines by interpolating between the discrete values of the standard deviations evaluated. Since cross-section parameters should vary event weights smoothly, this should adequately describe the parameter variations.

Flux systematic parameters

As discussed in section 4.1.3 and shown in figure 7.6, the T2K beam group consider a plethora of different error sources when producing a flux prediction. In the same way that all detector systemic error sources are combined into a covariance matrix in reconstructed analysis bins, the beam group project flux systematic error sources into covariances in true neutrino energy bins. The flux systematic parameters in the fit are then scale factors on the number of events in a particular true neutrino energy bin. The neutrino energy binning used, and the associated covariance matrix parametrising the scale of the flux uncertainty, is shown in figure 7.6. Since the beam contamination is so small, only the uncertainty of the ν_μ component of the flux is considered.

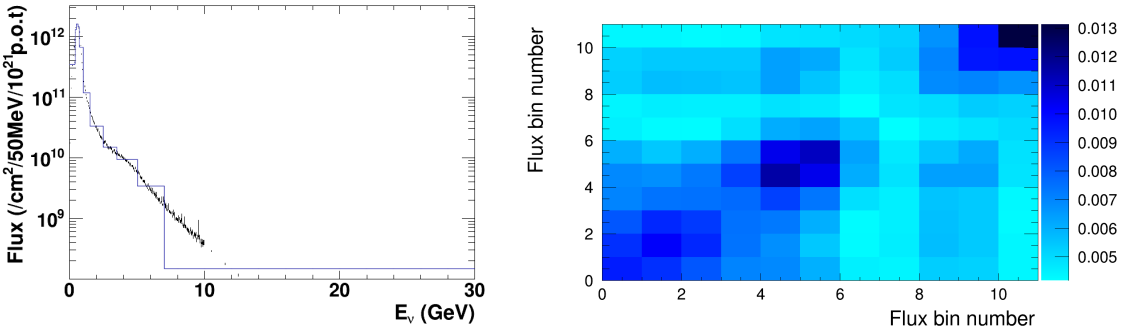


Figure 7.6: The left plot shows the ND280 flux prediction (black points) and binning for systematic variation (blue solid line). On the right is the covariance matrix used to construct the penalty term for flux parameters, as provided by the T2K beam group.

Summary of fit parameters

A summary of all parameters included in the fit and their index (used in later plots) is shown in table 7.4.

Index	Parameter	Nominal value	Pre-fit constraints
0-8	Template weight for STV in each analysis bin	1.0	regularisation
9-19	Weight for events in each flux bin	1.0	flux covariance matrix
20-28	Weight for events in μTPC, pTPC sample in each reconstructed bin	1.0	detector covariance matrix
29-37	Weight for events in μTPC, pFGD sample in each reconstructed bin	1.0	detector covariance matrix
38-46	Weight for events in μFGD, pTPC sample in each reconstructed bin	1.0	detector covariance matrix
47-55	Weight for events in CC1π sideband sample in each reconstructed bin	1.0	detector covariance matrix
56-64	Weight for events in CCOther sideband sample in each reconstructed bin	1.0	detector covariance matrix
65-73	Weight for events in μTPC, multi-p sample in each reconstructed bin	1.0	detector covariance matrix
74	C_A^5 scaling	1.0	model covariance matrix
75	M_A^{RES} scaling	1.0	model covariance matrix
76	Resonant Background scaling	1.0	model covariance matrix
77	CCν_e scaling	1.0	model covariance matrix
78	DIS (DISMPISHP) parameter scaling	1.0	model covariance matrix
79	CCCoh on C scaling	1.0	model covariance matrix
80	NCCoh scaling	1.0	model covariance matrix
81	NCOther scaling	1.0	model covariance matrix
82	Binding energy (E_B) parameter scaling	1.0	model covariance matrix
83	Pion FSI - Inelastic low energy scaling	1.0	model covariance matrix
84	Pion FSI - Inelastic high energy scaling	1.0	model covariance matrix
85	Pion FSI - pion absorption scaling	1.0	model covariance matrix
86	Pion FSI - pion production scaling	1.0	model covariance matrix
86	Pion FSI - Charge exchange low energy scaling	1.0	model covariance matrix
87	Pion FSI - Charge exchange high energy scaling	1.0	model covariance matrix

Table 7.4: A summary of all parameters included in the fit.

7.2.4 Fitting mechanism and interpretation of errors

To efficiently find the point of maximum likelihood (i.e. that which minimises the function in equation 7.2) in the many-dimensional fit parameter space, the MINUIT [146] numerical minimiser is used. MINUIT employs a maximum gradient decent method, called MIGRAD, to find the minimised negative log-likelihood. To do this, MIGRAD calculates the gradient of the log-likelihood surface to identify the steepest path toward the minimal negative log-likelihood. A step is then taken in that direction and the process is repeated.

Within each step it is necessary to calculate the new number of reconstructed events after applying appropriate weights, as in equation 7.8. This involves smearing the true, events which many of the weights are applied to, to reconstructed space. In equations 7.7 to 7.9, the smearing is handled by a detector smearing matrix built in analysis bins. This would not be suitable for this analysis since the bins do not well-characterise the detector response (since the binning is in the STV rather than the underlying muon and proton kinematics). For this reason the truth to reconstructed smearing is handled on an event-by-event basis using the information from the input simulation (thereby effectively using an unbinned version of the detector response function).

Once the fit identifies a point of maximum likelihood it proceeds to map the likelihood surface around it with a view to building an accurate covariance matrix to characterise the fit uncertainty, this is done using MINUIT's HESSE method. In HESSE, this covariance matrix is more accurately determined by calculating the inverse of the matrix of second derivatives of the negative log-likelihood, with respect to fit parameters, by the finite differences approximation. Essentially this assumes the likelihood surface can be described by a multi-variate Gaussian (including correlations) and uses the information very close to the point of maximum-likelihood to determine the parameters. This first "Gaussian errors approximation" is important as it can cause a misrepresentation of the uncertainty if the likelihood surface is not distributed as a Gaussian.

At this stage the fit has provided a maximum likelihood set of fit parameters and a covariance matrix describing the uncertainty on those parameters. Before a cross section can be calculated these must be converted to a best fit number of events (in each of the analysis bins) and covariances. To do this the Cholesky decomposition of the fit parameter covariance matrix is found and used, alongside the best fit set of parameters, to create around 1,000 toy sets of parameters which effectively sample from the likelihood space encoded in the covariance matrix. The toys then represent the spread of plausible parameters according to the statistical and systematic uncertainties from the fit. In each toy the best-fit simulation is re-weighted to the toy parameter set and the number of events in each analysis bin is found. The mean event rate in each bin is taken as the best fit and a new covariance matrix (binned with the analysis binning) is built from the toys. This whole process is really a numerical error propagation from a set of fit parameters to a set of event rates (and the corresponding covariance matrix). In doing this a second Gaussian errors approximation due to the requirement that the uncertainty in the number of events in each bin is be parametrised by a covariance matrix. This second approximation is clearly essential if a covariance matrix is to be provided along with the cross-section result. Avoiding either or both of these Gaussian error approximations would require the provision of a much more complicated data release and a more advanced fitting method, such as Markov chain Monte Carlo. This is discussed in more detail in appendix C.

It is important to consider how to interpret the final covariance matrix that is provided alongside a cross-section result. To calculate it, the likelihood space of the post-fit parameters, which itself incorporated information from the parameter's priors, is used. If the assumptions made by HESSE are valid, the resultant covariances essentially describe an N (number of analysis bins) dimensional contour with a constant log-likelihood around the best fit point, which contains some particular proportion of the probability of the resultant fit parameters. Therefore (assuming the second Gaussian errors approximation is also valid) the propagated covariance matrix describes a contour of constant log-likelihood around the best fit number of events in each analysis bin and contains the same particular probability of the possible event variation. Such a contour is an approximation of a Bayesian credible region, specified to have the best-fit point at the centre and is limited to being symmetric.

7.2.5 Additional systematics for cross-section extraction

Following the procedures in 7.2.4, the analysis framework provides the best fit number of selected signal events in each analysis bin along with corresponding covariances. As described in equation 2.20, to get to a cross section it is then necessary to account for the efficiency in each bin; the integrated flux; and the number of targets in FGD1's FV. To do this specifically for this analysis, equation 2.20 can be specified as:

$$\frac{d\sigma_{flux\,int.,\,fiducial}}{d\delta x_{T,i}} = \frac{N_i^{CC0\pi+Np,\,IPS}}{\epsilon_{i,sim}^{CC0\pi+Np,\,IPS}\Phi_{ND280}N_{targets}^{FGD1\,FV}} \times \frac{1}{\Delta\delta x_{T,i}}, \tag{7.14}$$

where i is the bin index; x is either p, ϕ or α such that the differential variable is one of the STV; $\Delta\delta x_{T,i}$ is the STV bin width; Φ_{ND280} is the integrated flux over ND280; $N_{targets}^{FGD1\,FV}$ is the number of targets in the fiducial volume of FGD; $N_i^{CC0\pi+Np,\,IPS}$ is the extracted number of signal events in the i^{th} IPS bin of the STV; and $\epsilon_{i,sim}^{CC0\pi+Np,\,IPS}$ is the efficiency from the input simulation, computed the same bins.

The nominal integrated flux provided by the T2K beam group is $1.927\nu_\mu/cm^2/10^{21}POT$ [122], whilst the nominal number of target nucleons in the FGD1 FV has been calculated to be 5.537×10^{29} [147]. When using equation 7.14 the integrated flux is modified based on the fitted flux parameters and the efficiency is modified from the nominal shown in chapter 6 based on the fitted values of all systematic parameters.

The use of equation 7.14 allows the propagation of a fit result to a cross section but there are uncertainties on the associated efficiency correction, the flux normalisation and the number of targets scaling. One method of accounting for these uncertainties would be to use the flux, model and detector response in each toy of the error propagation procedure, described in section 7.2.4, to calculate a cross-section (using equation 7.14), rather than

just the unsmeared number of selected signal events. This would also incorporate the constraints on the efficiency and flux normalisation in each toy, from having fit the systematic parameters, and the spread of these toy cross sections would then describe the uncertainty in the result. However, these constraints come almost entirely from the sideband samples and, when fitting these, it is possible that an incomplete model, flux or detector parametrisation causes one set of systematic parameters to alter the event rate in lieu of another, thereby leading to unrealistically constrained parameters. For example, limitations of the DIS model parametrisation may strongly constrain the Rein-Seghal cross-section parameters in the CCOther sideband, such that a data deficit/excess in the CC1π^+ sideband can only be accounted for by altering the flux or detector parameters (thereby allowing a constraint on the flux normalisation or efficiency - and therefore the cross section - in lieu of missing background model freedoms).

To avoid the potential for such unrealistically strong constraints, the pre-fit uncertainties will be used when considering plausible variations in the efficiency and flux normalisation. The sections below describe how covariances to account for these uncertainties are formulated.

Systematic error on the efficiency

From model and flux systematic parameters As discussed in section 6.3, despite detector acceptance motivated phase-space constraints in the signal definition, the efficiency with which the signal events are reconstructed does have some dependence on cross section and nuclear models. As previously discussed, a particular problem is that some combination of model parameters may predict that the proton and muon in signal events are more often co-linear (and therefore not individually resolvable) leading to a lower efficiency. Additionally it is possible that alterations to the flux shape could exaggerate regions of kinematic phase-space with large/small efficiencies (although the flatness of the efficiency in the underlying kinematic phase-space, discussed in section 6.3, should largely negate this effect).

To assess the effect of varying flux and model parameters on the efficiency in a conservative way, 250 toys are produced by re-weighting the input simulation using both the flux covariance matrix from the beam group (shown in figure 7.6) and full model systematics covariance matrix (shown in figure 7.5). Unlike in the fit, this includes the variation of the predominant signal model parameters M_A, $2p2h_{norm}$ and p_f. The covariance on the efficiency is found from the toys in each of the STV, to illustrate the scale of this systematic, the relative error is shown in figure 7.7. The small errors highlight the very limited neutrino interaction/flux model dependence of this analysis. Since the efficiency is just a scale factor on the cross section, the calculated relative covariances can be added to those from the fit.

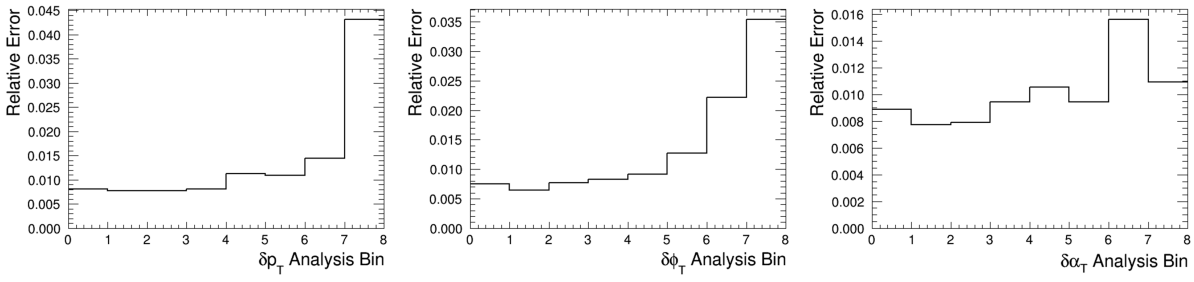


Figure 7.7: The error in each STV truth bin due to variation of the efficiency with the cross-section model parameters.

From nucleon FSI The uncertainty due to the plausible variation of the nucleon FSI model is addressed separately. Unfortunately the NEUT 6B input simulation does not contain the necessary information to facilitate the re-weighting of events to account for nucleon FSI. However, the special NEUT 6D simulation contains enough to use the only currently available nucleon FSI re-weighting for NEUT: a rescaling of the nucleon FSI mean free path (MFP). As discussed in section 4.2.3, NEUT 6D cannot be easily used for general analysis since it generates events only within the ND280 detectors and not in the magnet, but it can be used to extract a relative error on the efficiency correction due to plausible variation of the nucleon FSI MFP. To validate that the systematic 'dial' that controls the scaling of the MFP behaves as expected, the effect of altering the dial on δp_T is shown to be consistent with the effect of turning FSI on/off at the generator level in in figure 7.8. The corresponding uncertainty on the efficiency is calculated using a dial uncertainty of 10% to ensure that any plausible variation of the nucleon FSI model is covered. The resultant uncertainties are shown in figure 7.9. For a conservative treatment of nucleon FSI, this is added to the cross section covariance as an uncorrelated uncertainty (i.e. the uncertainties from figure 7.9 are squared and added to the diagonal of the cross-section covariance).

The 10% variation can be justified as conservative by the very large change in the signal event rate it causes. The change of the integrated signal cross section with a -10% dial variation is approximately 28%, the largest from any 1 standard deviation change of a model parameter. This covers the variation between FSI on and FSI off at the generator level and is enough to cover any of the other FSI model predictions, within in the FSI sensitive region of the STV, considered in the summary generator comparison (in section 9.2).

From detector systematic parameters The variation of detector systematic parameters alters the rate at which events are reconstructed, but has no impact on the true distribution of events. Therefore, the relative variation of the number of signal events under representative sampling of the detector systematic parameters is just an uncertainty on the efficiency

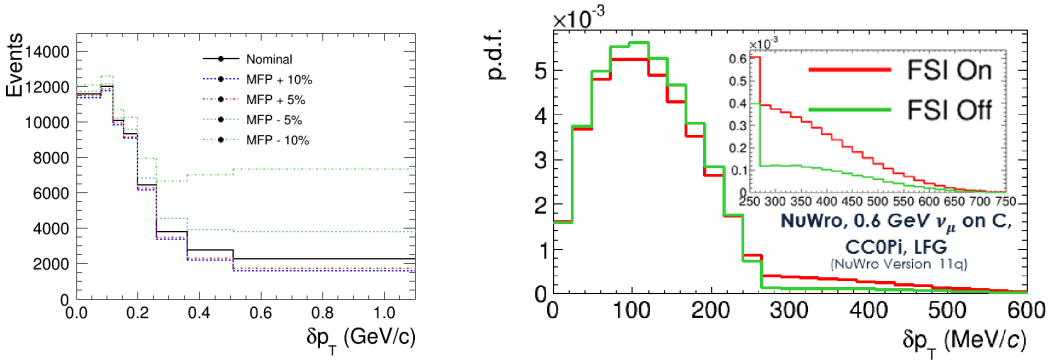


Figure 7.8: The plot on the left shows the effect of altering the mean free path dial on δp_T formed from all (full phase-space) events in signal topologies from the selection in chapter 6. The plot on the right shows the effect of turning FSI on/off on δp_T in the NuWro simulation (previously shown in chapter 3).

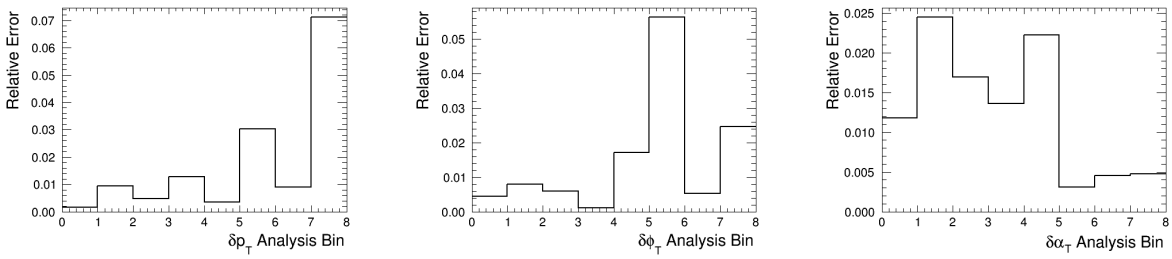
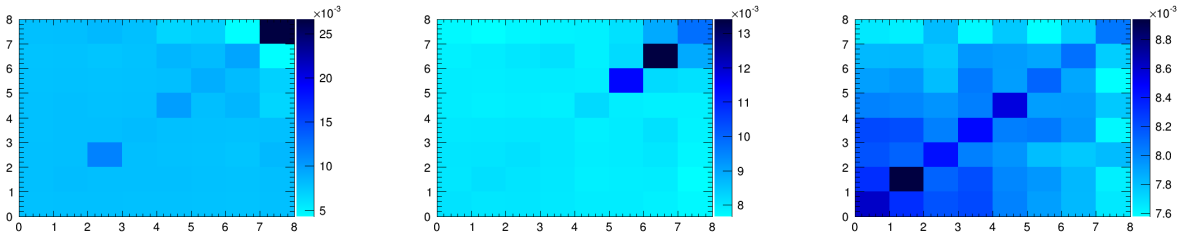


Figure 7.9: The error in each STV truth bin due to variation of the efficiency with the nucleon FSI MFP model parameter.



(a) δp_T (b) $\delta \phi_T$ (c) $\delta \alpha_T$

Figure 7.10: Fractional covariances to describe the systematic uncertainty on the efficiency, flux normalisation and number of targets, formulated in each analysis bin of the STV (the binning is described in tables 7.1 to 7.3). These must be applied when moving from a fit result to a cross-section (using equation 7.14).

correction. Covariances to describe such an uncertainty, in each analysis bin of the STV, are found from the ensemble of 500 detector systematic toys described in section 6.6.2.

Systematic error on the flux normalisation and number of targets

Flux normalisation systematic In addition to causing migration between bins in the fit and potentially altering the efficiency, the integrated flux is required to calculate the cross section. There is a significant error on this total integrated flux from the beam simulation which can been calculated from the flux covariance matrix from the T2K beam group to be 8.5% (see, for example, [106, 148]).

This systematic is considered to be fully correlated across all bins (since the effect is a total normalisation shift which does not change with bin).

Number of targets systematic The flux integrated cross section will be quoted per nucleon in the FGD1 FV. In order to extract this cross section it is necessary to scale by the number of nucleons. This has some very small uncertainty associated with it which has previously been calculated to be 0.67% (see, for example, [148, 106]).

Like the flux normalisation, this systematic is fully correlated across all bins.

Summary of additional systematic uncertainty

The (co)variances from each source of systematic described in this section are combined into a single relative covariance matrix, shown for each STV in figure 7.10. This must then be added to the relative covariance matrix from the fit when applying equation 7.14 to go from a fit result to a cross section.

7.2.6 Input simulation for fitting and associated model dependence

As described in section 7.1, likelihood fitting requires an input simulation from which the STV templates, the effect of systematic parameters and the constraints from regularisation can be derived. It is clearly essential that the input simulation contains realistic smearing from reconstructed to truth level observables but this information should be almost entirely encoded in the ND280 detector simulation and the GEANT4 particle propagation simulation. Providing that the fit is only attempting to extract what ND280 can realistically measure, this smearing should have a minimal dependence on the neutrino scattering simulation[4].

It is also important that the input simulation contains a background model with sufficient parametrisation to allow its plausible variation. Since the fit allows only a constrained variation of the backgrounds within priors set by external fits to a specific model, the result is dependent that the background model can describe the real background within said priors. However, since the purity of the selection in this analysis is greater than 80%, mis-modelling of the background would have to be quite dramatic in order for it to have a significant impact on the fit results. Moreover, if such mis-modelling existed, it would be clear in the goodness of fit to the sideband samples.

Since there is no constraint on the normalisation of the template weights, there is no dependence on the normalisation of the signal model in the fit. However, the regularisation adds a constraint which depends on the shape of the input simulation and therefore biases the fit towards this. The implemented L-curve method of choosing a regularisation strength helps to mitigate such bias. So long as the regularisation strength is low, such that the regularisation term does not dominate the likelihood at the best fit point, the effect of regularisation should be a coarse smoothing that has only a vague dependence on the input simulation. This can be assessed directly by considering the relative contribution of the regularisation to the total likelihood. Furthermore, the fit can be repeated for various regularisation strengths to verify a minimal bias.

There is an additional potential bias to the input simulation shape in the construction of the STV templates. Each template contains a particularly shaped distribution of muon and proton kinematics which depends on the input simulation (although this dependence is somewhat alleviated by the phase-space constraints on the STV analysis bins). Since the different selection samples have acceptances which depend on these underlying kinematics, the relative impact of altering a template weight on each sample depends on the input IPS signal model. For example: if the true cross section is similar to the input simulation but is more concentrated into higher angle muon phase-space, there may be a deficit of events

[4] As discussed in section 6.4.1, the exception to this is the effect on the smearing matrix from the rate of co-linear outgoing particles causing one or more not be reconstructed.

in the samples which require a muon in the TPC and an excess in the FGD muon sample. However, since the acceptances of the different samples widely overlap, in order to have a substantial effect on the fit result the mis-modelling of the shape of the proton or muon kinematics would have to be quite large. Moreover, this issue is strongly mitigated by the muon and proton phase-space constraints and the resultant flat overall acceptance in the underlying kinematics (discussed in section 6.3). These ensure that a mis-modelling of these kinematics can only migrate IPS events between the IPS bins in each signal sample. This would lead to an excess in some sample(s) and a deficit in other(s), which is irresolvable by the template weights in the fit[5]. Therefore, this source of bias is diagnosable from a poor goodness of fit to the signal samples, an aspect of the fit result that will be carefully scrutinised in chapter 8's validation studies. Although it is beyond the scope of application in this thesis, potential methods of avoiding this bias completely are discussed in appendix C.

As the only available simulation with events generated within the magnet (essential to understand the impact of OOFV background on the efficiency) for which comprehensive model systematics can easily be assessed, the input simulation used for this analysis is the T2K NEUT 6B simulation which, as discussed in section 4.2.3 (and further detailed in appendix A), contains a bug in the nucleon FSI simulation. However, the arguments made above demonstrate that, so long as the truth to reconstructed smearing is well described and the background models are plausible, the results dependence on the input simulation should be minimal. Moreover, chapter 8 will provide extensive tests of whether there is any bias to the input simulation.

[5] Although it may be possible to resolve such a discrepancy by altering the detector systematic parameters, their tight priors means that doing so would pick up a large penalty term

Chapter 8

Cross-section extraction validations and sensitivity

The previous chapter outlined the methodology for using a binned likelihood fit as a cross-section extraction method. This chapter will apply this to various fake data sets in order to validate that the cross-section extraction method is functioning as expected, that there is not significant bias to the input simulation and to qualitatively asses the sensitivity of the proposed measurement to different interaction models.

This chapter will begin, in section 8.1, with an assessment of the relative size of each systematic error source in an 'Asimov' fit (a fit in which the fake data is the same as the input simulation) to determine the dominant sources of uncertainty. Following this, section 8.2 critically analyses the analysis' performance with different fake data inputs.

8.1 Systematic uncertainties in an Asimov fit

In this section an Asimov fit is performed using the NEUT 6B simulation (defined in section 4.2.3 and further detailed in appendix A) and a typical optimal regularisation strength of 3.0 (from the fake data studies in section 8.2), to obtain an estimation of the uncertainty contribution from each class of systematics. To account for the expected statistical error, the input simulation and the identical fake data are modified to only contain the expected data POT. In an Asimov fit any movement of any parameter cannot raise the likelihood. It is therefore expected that the fit returns the nominal parameters, this is the most basic validation of the fitting framework. This is shown to be true in the Asimov fit result, including uncertainties from the fit and those described in section 7.2.5, shown in figure 8.1.

To better understand the predominant source of the error in each cross-section bin, an Asimov fit is first performed fitting only the template weights (c_i) before adding a particular class of systematic parameters and refitting. The extrapolated error from each class of systematic parameter is calculated as the difference in quadrature in the error on the template

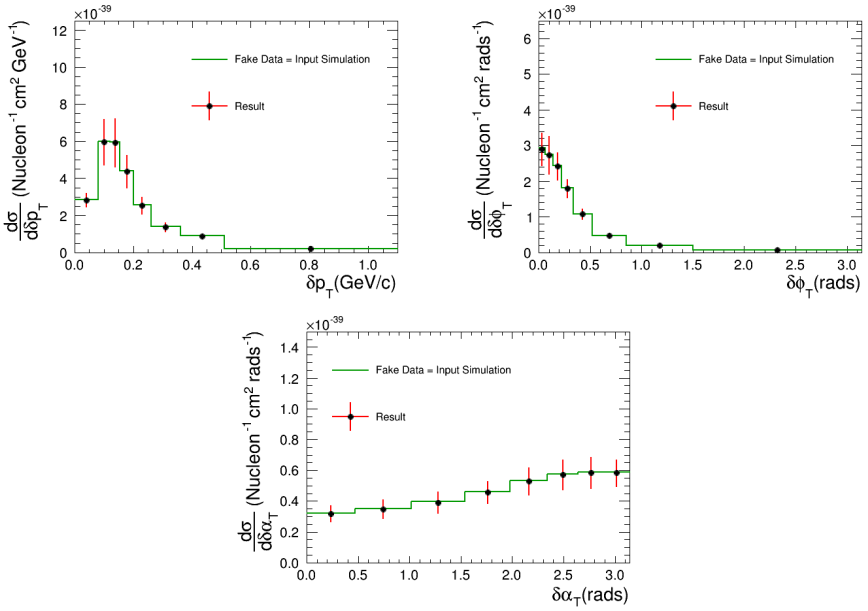


Figure 8.1: The differential cross-section in the three STV obtained using an Asimov fit where the input simulation is NEUT 6B.

weights returned from fit. This is repeated for each class of systematic separately and also for all systematic parameters together. These extrapolated error contributions for each parameter source are shown in figure 8.2. Note that these are not the actual uncertainty on the cross-section result (this is found using the methods described in section 7.2), but still gives a useful indication of the error contribution.

Figure 8.2 indicates that the error is generally dominated by the uncertainty in the template weights, this is partially from limited statistics but also due to the detectors resolution making the separation of adjacent template weights difficult. This exercise is repeated without including the sideband regions in the fit to highlight the control they offer over the combination of the flux and background model systematics, this is shown in figure 8.3. Finally this is repeated once further with the full available simulation statistics ($\sim 5\times$data POT) to demonstrate the scope for larger statistics to reduce the uncertainty, this is shown in figure 8.4[1]

[1]With the larger number of events, it is found that the matrix of second derivatives calculated by HESSE (see section 7.2.4) was often non-invertible due to degeneracies between the Rein-Seghal model parameters which scale the resonant pion production cross section. This is addressed by merging the effect of the offending parameters but it is possible that the process of doing this slightly misrepresents the model systematics in figure 8.4.

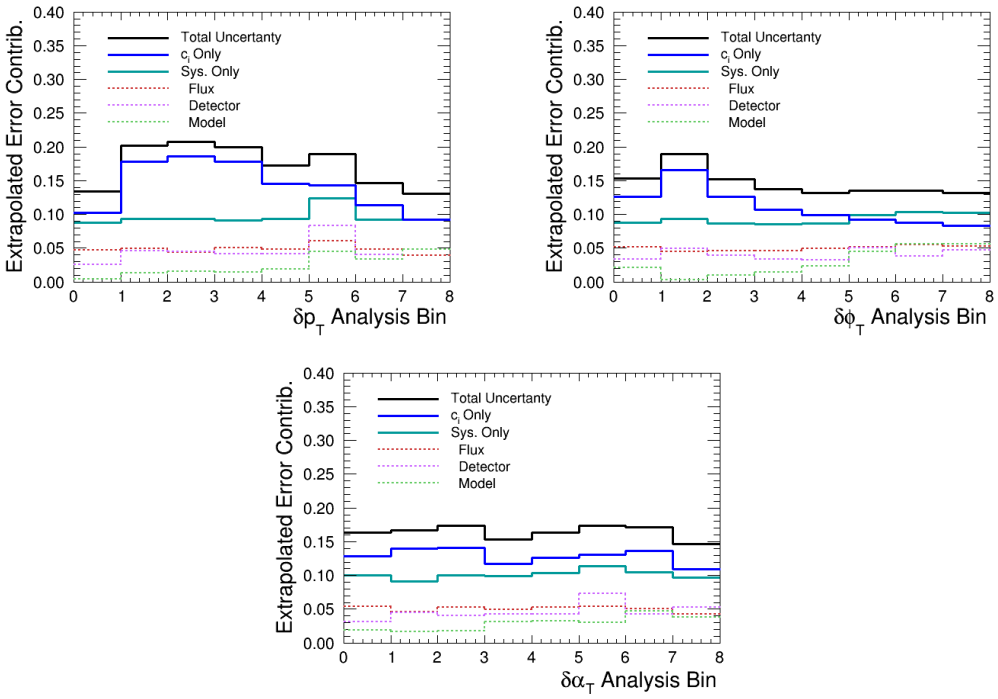


Figure 8.2: The contribution of each class of uncertainty on the CC0π+Np template weights in true STV analysis bins following an Asimov fit representative of real data POT.

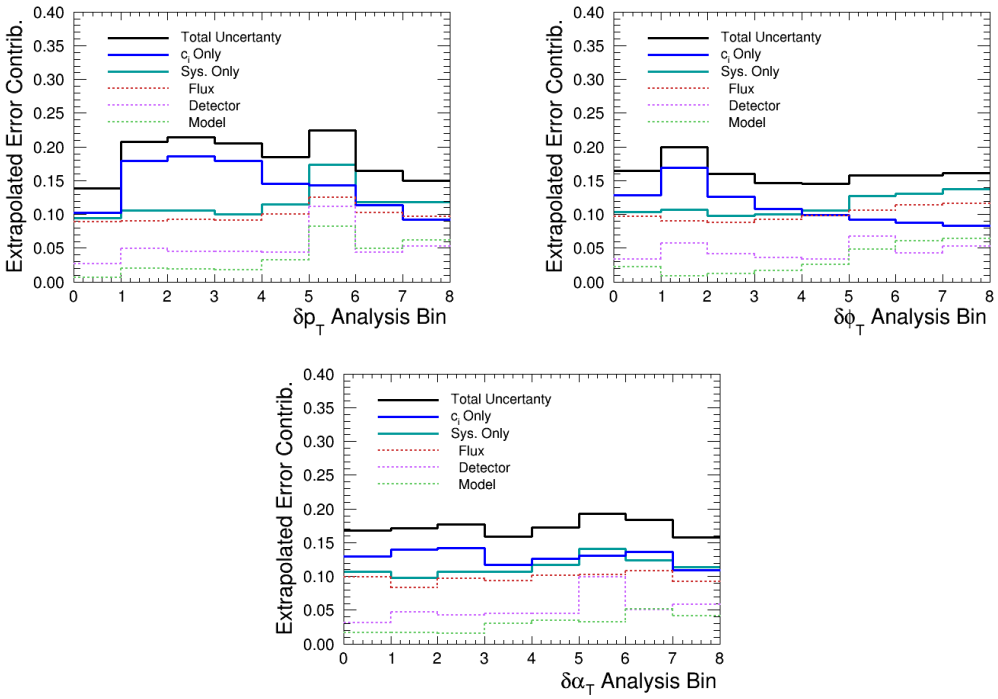


Figure 8.3: The contribution of each class of uncertainty on the CC0π+Np template weights in true STV analysis bins following an Asimov fit representative of real data POT in which the sideband regions are excluded.

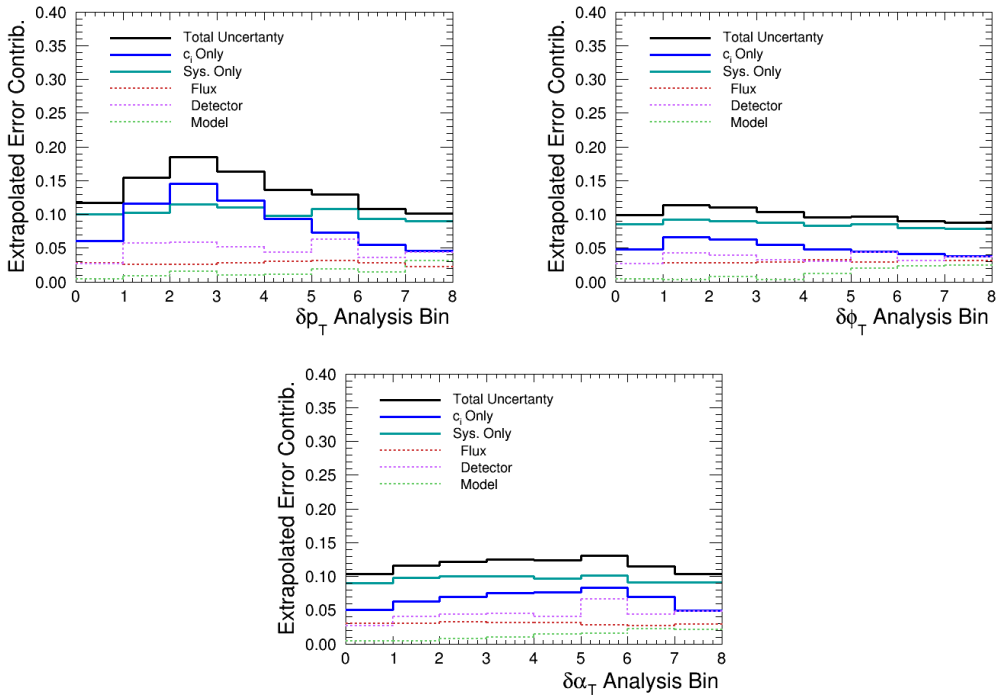


Figure 8.4: The contribution of each class of uncertainty on the CC0π+Np template weights in true STV analysis bins following an Asimov fit using the full simulation statistics ($\sim 5\times$ data POT).

8.2 Validating the analysis framework

In this section several fake data tests are used to further validate that the analysis framework is functioning correctly and that there are no significant biases in the cross-section results, particularly from: mis-modelled backgrounds; mis-modelled underlying IPS proton/muon kinematics; or regularisation. All fake data are scaled to have a number of events that is representative of the expected statistics in the real data and the cross sections from each are found using the full cross-section extraction method, described in chapter 7, that will be applied to data.

First, in section 8.2.1, a fit result and cross-section extraction is demonstrated using NEUT 6B as an input simulation and a statistically independent NEUT 6B as the fake data. The goodness of fit at the reconstructed-level, the post-fit parameters and the relative contribution of the different components of the fit are critically analysed. The effects of varying the regularisation strength and removing the sidebands are then also assessed to investigate potential bias in the cross-section extraction.

This procedure is then repeated for a plethora of other fake data sets, starting with the ND280 GENIE 2.8.0 simulation. Following this, NuWro 11q, NEUT 6D (NEUT 6B with the nucleon FSI bug fixed) and NEUT 6D with the 2p-2h contribution removed are used as fake data. The latter two of these tests explicitly evaluates the bias due to the nucleon FSI bug in the input simulation, whilst the use of NuWro and GENIE as fake data ensures the cross-section extraction framework is validated over a wide range of the plausible model phase-space. The simulations used are defined in section 4.2.3 and further detailed in appendix A.

As discussed in section 4.2.3, the NuWro 11q and NEUT 6D simulations only generate events in selective regions of ND280, so therefore have an unrealistically large selection efficiency. For this reason the efficiency in the input simulation is used to correct the reconstructed level predictions of the simulations to give 'effective truth' distributions which will be shown throughout this section[2]. Therefore the cross-sections do not necessarily fully describe the predictions of the simulation used to make the fake data, but do allow them to serve as very useful test data sets.

The truth (or effective truth) distributions for each of the fake data used are shown alongside the input simulation prediction in figure 8.5. This further emphasises the wide range of fake data models used for validation.

[2]It should be noted that the way in which the NuWro simulation was produced makes it difficult to ascertain the exact effective POT, which may lead to a small normalisation shift from the true NuWro cross-section predictions. This does not stop the simulation serving as useful fake data.

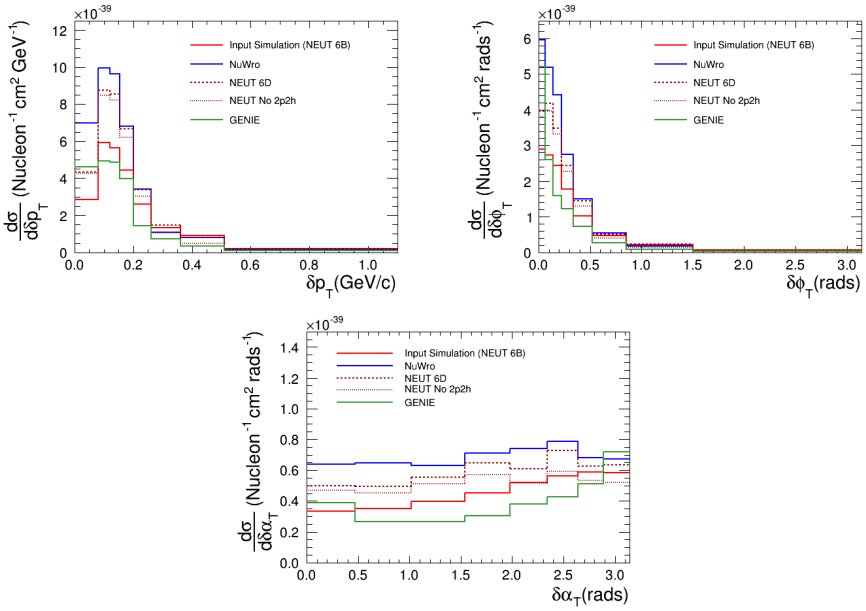


Figure 8.5: The truth (or effective truth) distributions for each of the fake data used.

8.2.1 NEUT fake data

A statistically independent NEUT 6B simulation, scaled to data POT, is treated identically to fake data to serve as an example of cross-section extraction and to validate the analysis framework.

First a regularisation study is made to determine the best choice of regularisation parameters (this procedure is described in section 7.1.2). The resultant L-curve for each fit is shown in figure 8.6 which justify a regularisation strength of 5 for all STV. The IPS-OOPS regularisation strength was found to make very little difference to the result and is fixed at 1.

Following the regularisation optimisation, the fake data is fit to extract a maximum likelihood set of fit parameters and an associated covariance matrix to characterise their uncertainty and correlations, this is done separately for the three STV. The resultant $\chi^2(\approx -2\log(L))$ contribution of each term in the maximum likelihood fit (see equation 7.1) is shown in table 8.1.

The extracted parameters are shown in figure 8.7, whilst comparisons of the pre and post-fit input simulation are shown alongside the 'measured' fake-data in reconstructed bins in figure 8.8. In these figures, correlation matrices, rather than covariance matrices, are shown since these are easier to interpret but do not show the magnitude of the variances (this is instead shown in the parameter error bars).

The best fit set of parameters and covariances are then propagated to a best fit number

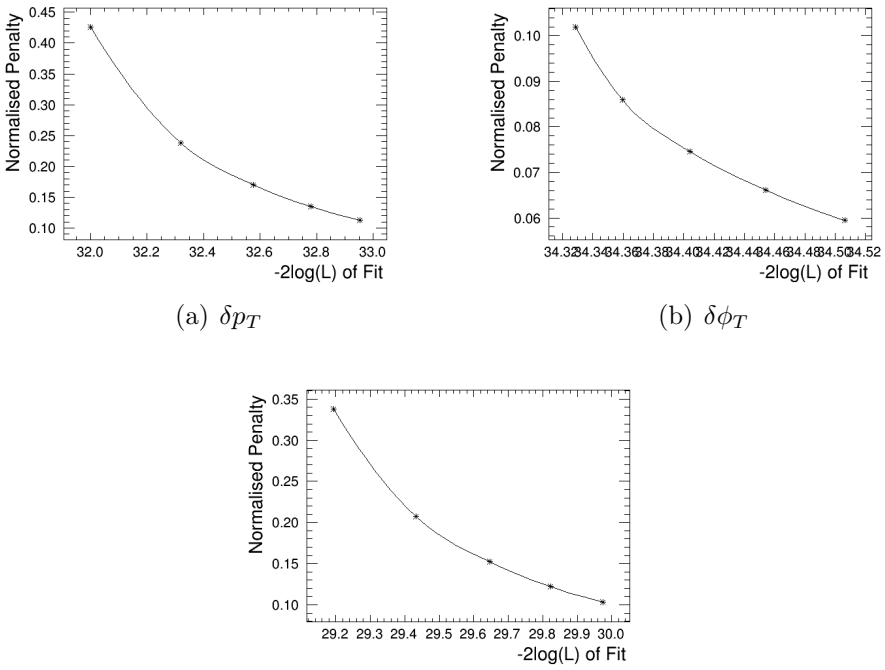
of selected signal events in each STV truth bin, as described in section 8.1. Following this, equation 7.14 is used to obtain a cross section and the associated systematic uncertainties in doing this are applied. The results, alongside their correlation matrix, are shown in figure 8.9.

To verify the fitter is in good agreement with - and has good sensitivity to - the fake data truth, a χ^2 goodness of fit can be calculated using the full covariance matrix for several test simulations including the fake data truth. This is calculated as follows:

$$(\chi^2)^{testsim.} = (\vec{\sigma}^{result} - \vec{\sigma}^{testsim.})(V_{cov}^{result})^{-1}(\vec{\sigma}^{result} - \vec{\sigma}^{testsim.}). \tag{8.1}$$

Here $\vec{\sigma}^{result}$ is a vector containing the extracted cross-section in each bin, $\vec{\sigma}^{testsim.}$ is the equivalent for a test simulation and V_{cov}^{result} is the covariance matrix characterising the extracted cross-section uncertainty. This method of calculating the goodness of fit includes the full correlated statistical and systematic uncertainty on the cross-section result, but it is important to note that it treats the test simulations as a single nominal prediction with no error band. This makes it a good diagnostic test but not suitable for calculating further test statistics.

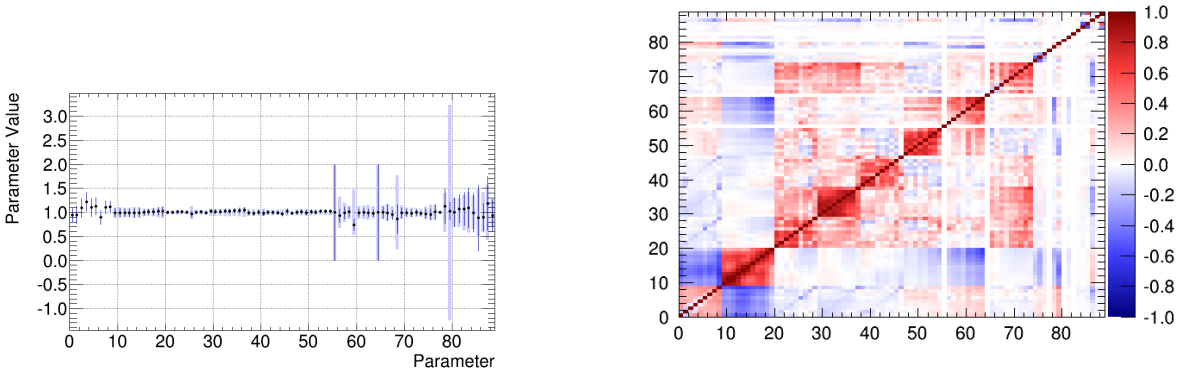
The results of comparing the goodness of fit to several test simulations is shown in figure 8.10.



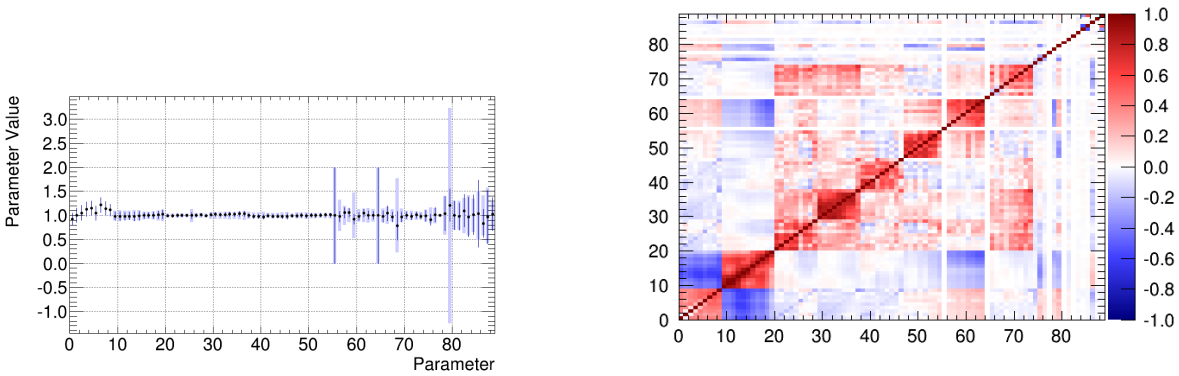
(a) δp_T (b) $\delta\phi_T$

(c) $\delta\alpha_T$

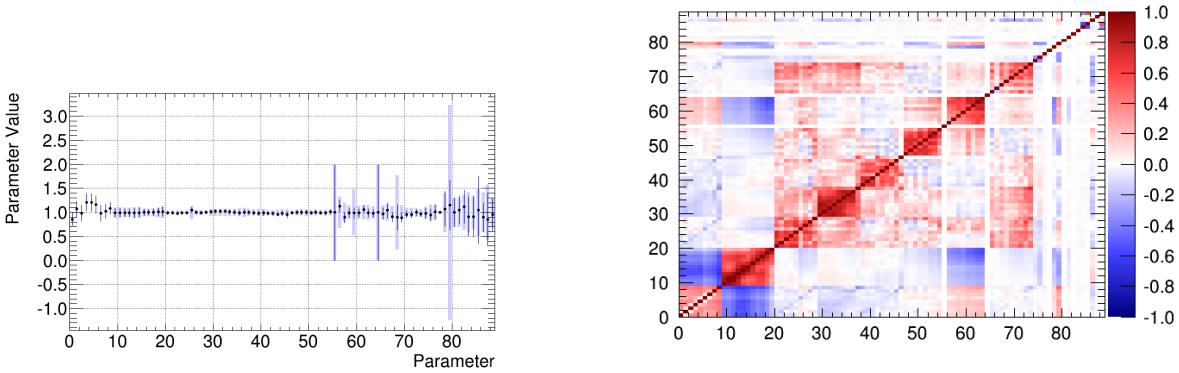
Figure 8.6: The L-curve used to determine the optimal regularisation parameters for each fit to NEUT fake data, which is statistically independent from the input NEUT simulation. The regularisation strengths run in sequence down the curves (from left to right) and are the following: 1, 3, 5, 7 and 9. For this fake data study, a regularisation strength of 5 is chosen for all three STV.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 8.7: The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NEUT fake data, which is statistically independent from the input NEUT simulation. The data points in the left plot show the extracted parameters and their standard deviations alongside the light blue lines which show each parameters prior standard deviation. The parameters are labelled by the index in table 7.4.

Variable fit	χ^2_{stat}	χ^2_{syst}	χ^2_{IPSreg}	$\chi^2_{OOPSreg}$
δp_T	25.2	3.7	0.8	0.0
$\delta \phi_T$	30.9	3.7	0.4	0.0
$\delta \alpha_T$	25.0	4.6	0.8	0.0

Table 8.1: The size of each term in the maximum likelihood fit to NEUT fake data.

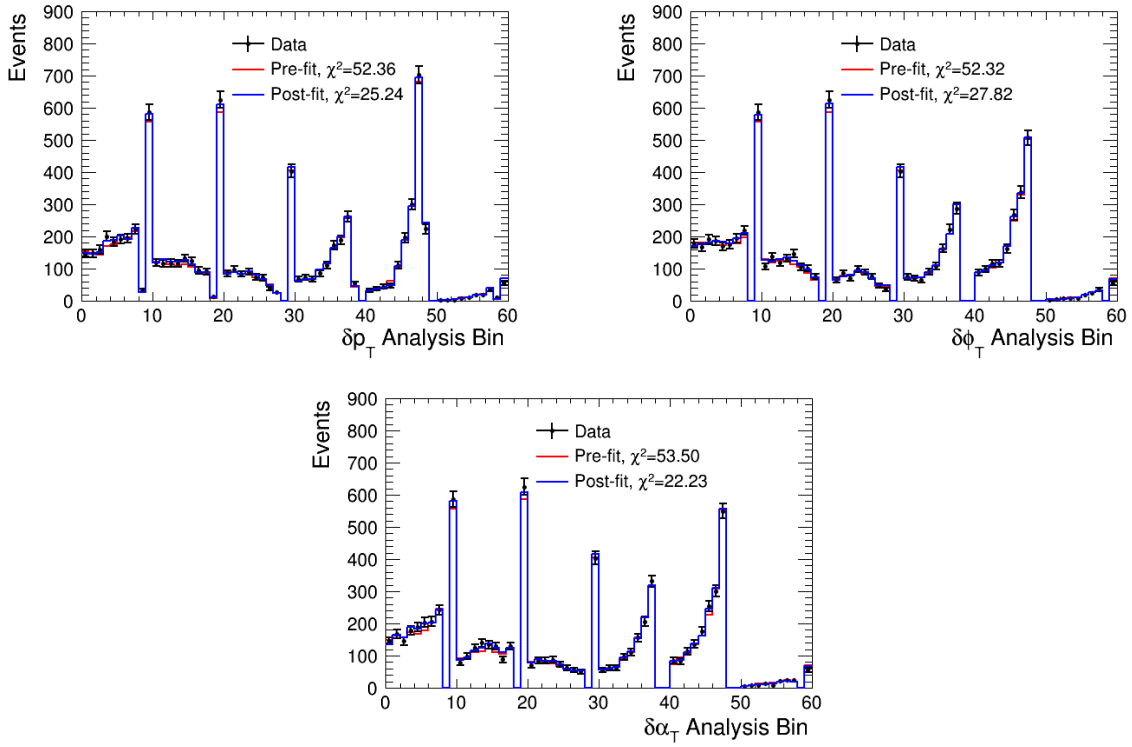


Figure 8.8: A comparison of the pre and post-fit simulation in reconstructed bins from a fit in each of the STV (separately) to NEUT fake data, which is statistically independent from the input NEUT simulation. The χ^2 goodness of fit is shown. The reconstructed bins for each variable are listed in tables 7.1 to 7.3. These are repeated for each sample in the following order: μTPC pTPC, μTPC pFGD, μFGD pTPC, CC1π sideband, CCOther sideband, μTPC multi-p. Note that the buffer bin (δp_T bin 8) is included for all three STV but is always empty for $\delta \phi_T$ and $\delta \alpha_T$ (i.e. the 9th bin in each sample is only relevant for δp_T).

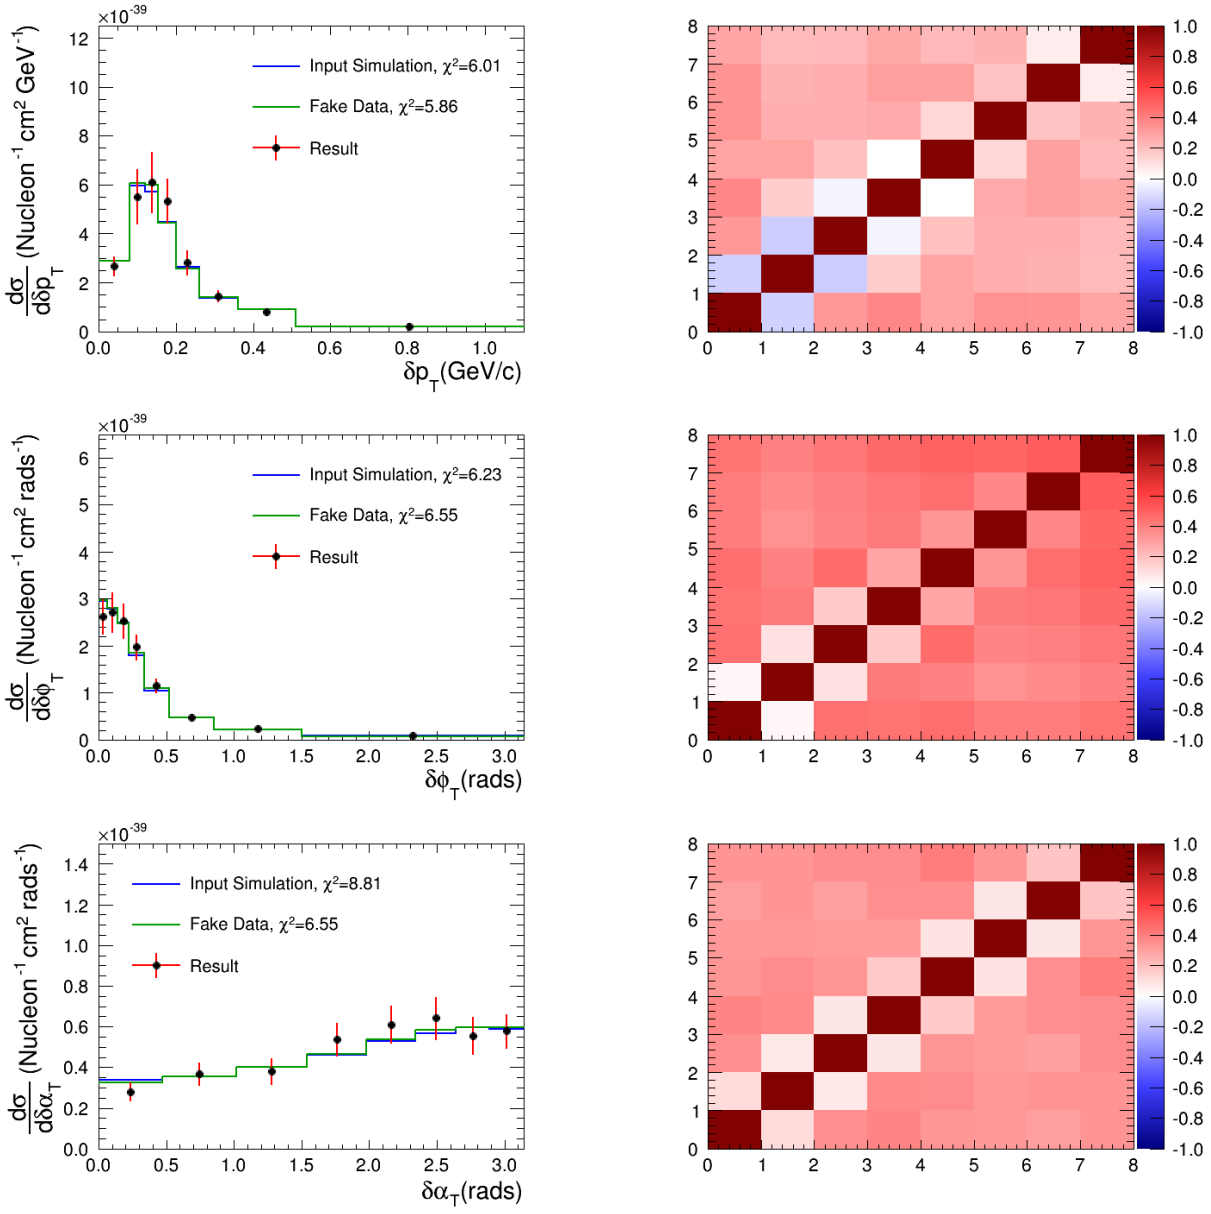
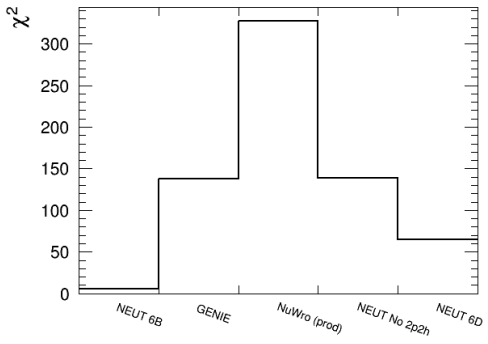
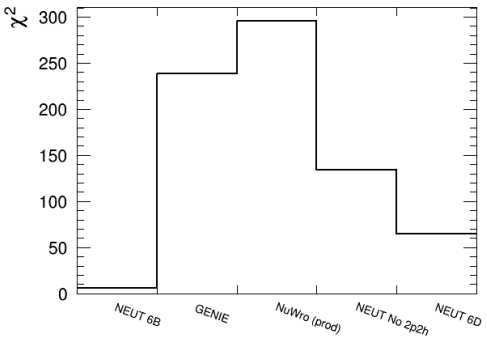


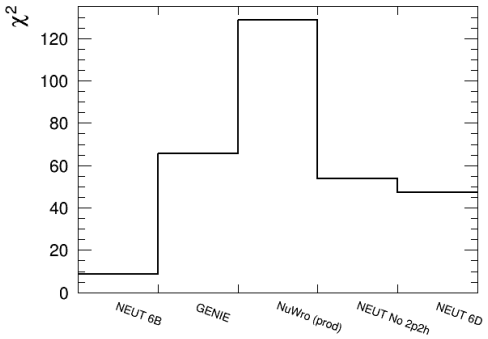
Figure 8.9: The differential cross section in the three STV and the associated correlation matrices obtained from separate fits to NEUT fake data, which is statistically independent from the input NEUT simulation.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 8.10: The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to a statistical fluctuation of the NEUT input simulation (labelled 'NEUT 6B'), to all test simulations, calculated as in equation 8.1. The simulation predictions are extracted from their ND280 productions. 'NEUT 6D' is the bug-fixed NEUT and 'NEUT No 2p2h' is NEUT 6D without any 2p-2h events.

NEUT fake data discussion Figure 8.7 demonstrates that the fit result is clearly compatible with the fake data truth. A significant anti-correlation between the flux and template weight parameters can be seen, which is expected as both sets of parameters directly scale the signal events (although the flux parameters scales the background events too): the only separation offered between them is via the sidebands. Figure 8.8 demonstrates an excellent goodness of fit (to be expected when the fake data is a statistical fluctuation of the input simulation) in all three STV.

Table 8.1 then illustrates that χ^2_{stat} is substantially larger than any other term in the fit, thereby showing that the dominant factor in finding the best fit point is the data and not the systematic parameter priors or regularisation. This is good evidence that the fit result is not biased by these.

Figure 8.9 shows that the cross-section result is in excellent agreement with the fake data truth and is confirmed by figure 8.10, produced using equation 8.1. Additionally, figure 8.10 shows that the fit finds the NEUT input simulation to have the lowest χ^2 by a substantial margin, thereby demonstrating that the three fits are all able to identify NEUT input as the best fitting test simulation.

Overall these tests validate the stability of the cross-section extraction framework under a statistical fluctuation of the input simulation and show that the result is capable of offering interesting model separation.

Effect of regularisation and sidebands on the NEUT fake data As previously discussed in section 7.1.2, 7.2.5 and 7.2.6, it is possible that regularisation or the sideband can bias the result toward the input simulation. To further assess the extent to which this happens, the NEUT fake data is fit with and without the sidebands and with various strengths of regularisation.

Figure 8.11 compares the differential cross section and χ^2 goodness of fit relative to all fake data, whilst figure 8.12 shows the same with the most favoured regularisation strength with and without sidebands.

Overall it can be seen that both the regularisation strength and inclusion of sidebands makes very little difference to the nominal fit result but their presence modestly increases the fits power to separate between fake data sets. For the case of regularisation this extra sensitivity stems from the extra constraint that cross sections should be smooth lowering uncertainties, whilst for the case of the sidebands this comes from the constraints they offer on the background event rate.

Generally speaking, the results display very little potential bias to regularisation. This is shown by the very small regularisation contribution to the best fit point likelihood (see table 8.1) and the negligible change on in model separation power with regularisation strength.

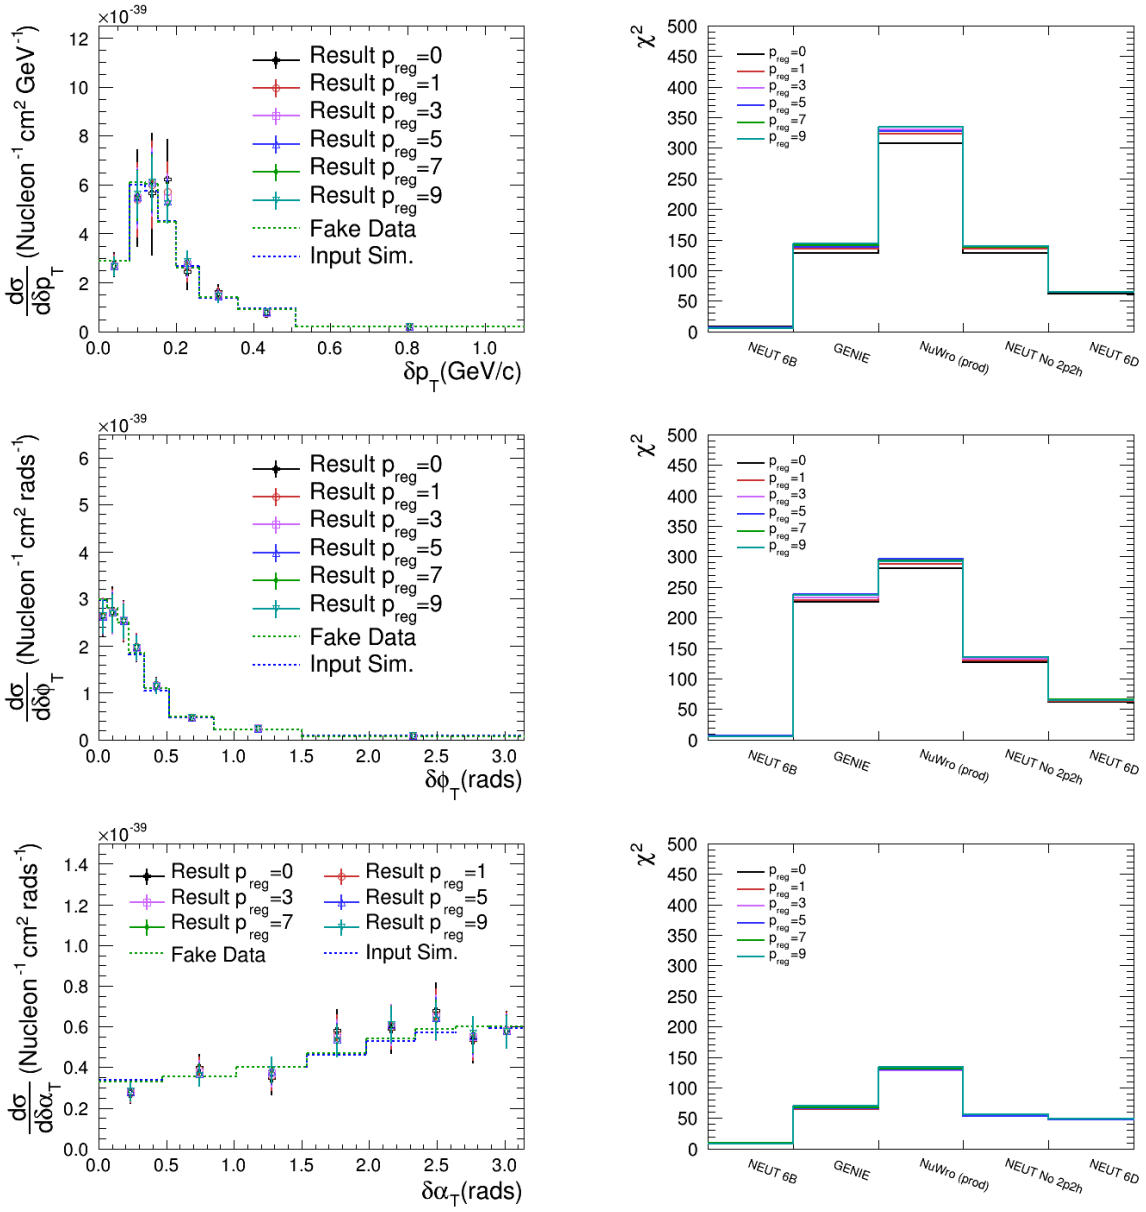


Figure 8.11: The differential cross section in the three STV obtained from a fit to NEUT fake data with various regularisation strengths, alongside the χ^2 of the results relative to each fake data set.

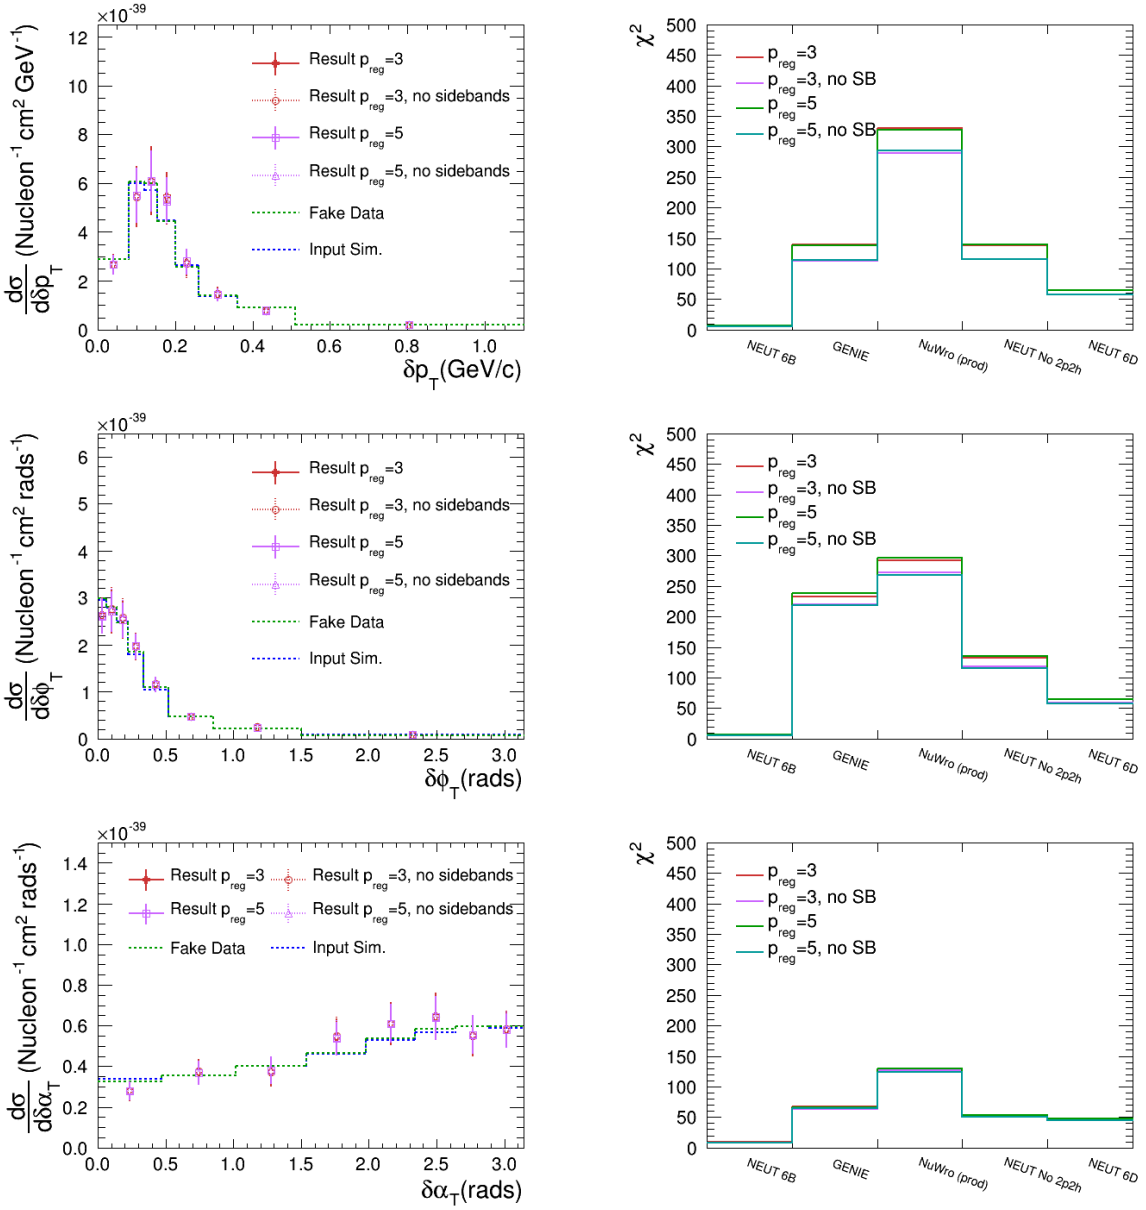
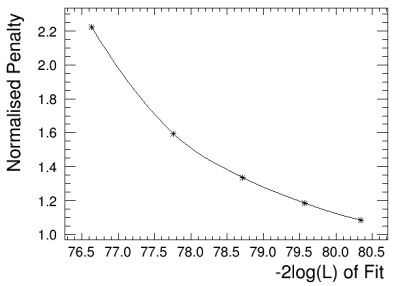


Figure 8.12: The differential cross section in the three STV obtained from a fit to NEUT fake data using, regularisation strengths deemed reasonable by the L-curve method, with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set.

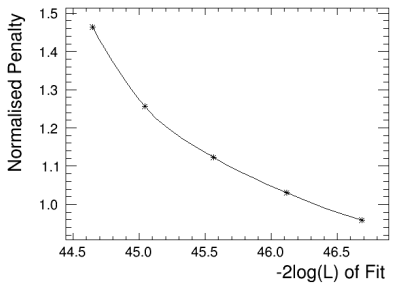
8.2.2 GENIE fake data

The GENIE 2.8.0 neutrino interaction simulation is used to produce a fake data set scaled to real data POT. GENIE 2.8.0 is quite different from NEUT in that is has no 2p-2h contribution, a very different FSI implementation and an RFG spectral function (see section 4.2.3). It therefore serves as a test of whether the fit is in any way prevented from matching fake data with different signal and background models though bias to the input simulation.

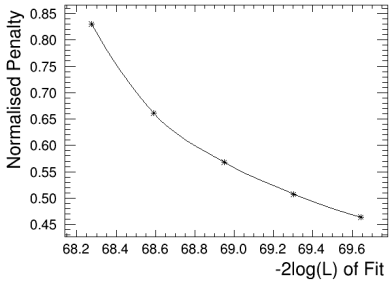
Following the same procedure as in section 8.2.1, a regularisation study is made to determine the best choice of regularisation parameters. The resultant L-curve for each fit is shown in figure 8.13 which justify a regularisation strength of 3 for all three STV. Once again, the IPS-OOPS regularisation strength was again found to make very little difference to the result and is fixed at 1. The resultant χ^2 contribution of each term in the maximum likelihood fit (see equation 7.1) is shown in table 8.2. Following this, the extracted fit parameters and correlations are shown in figure 8.14, whilst comparisons of the pre and post-fit simulation with the fake data are shown in reconstructed bins in figure 8.15. The cross-sections, alongside their correlation matrix, are shown in figure 8.16 and the χ^2 goodness of fit results are shown in figure 8.17.



(a) δp_T

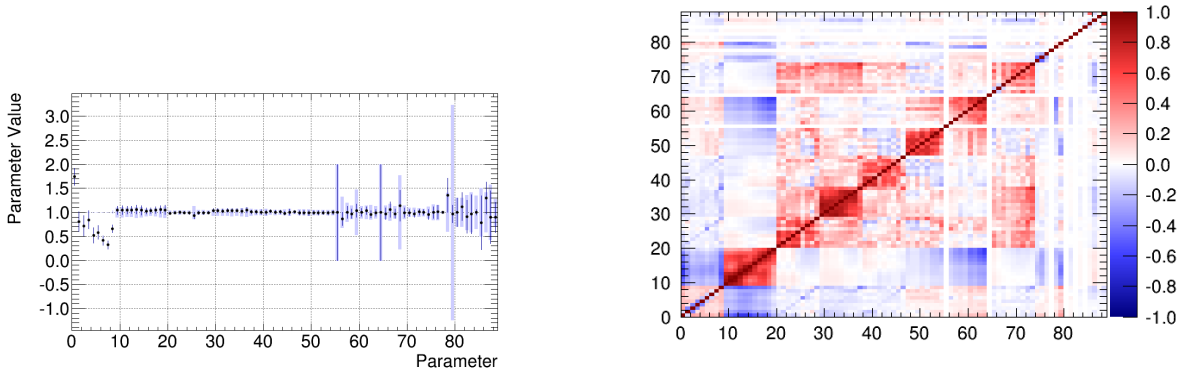


(b) $\delta \phi_T$

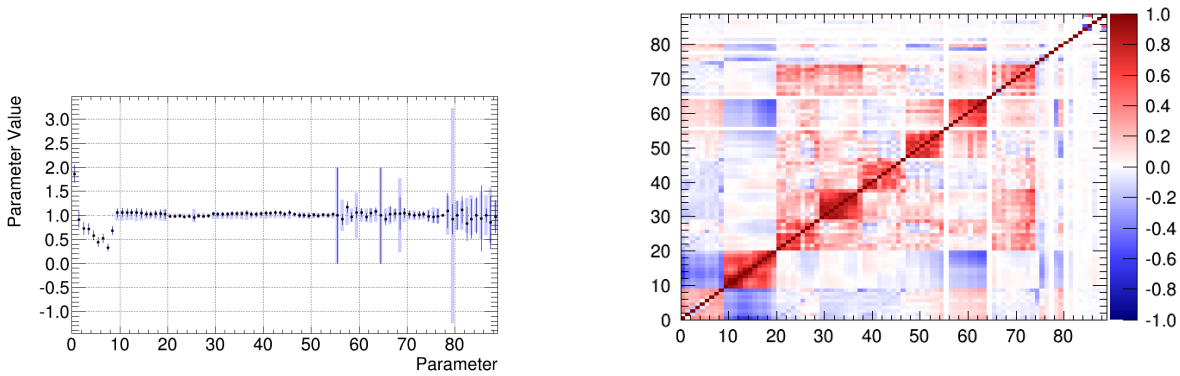


(c) $\delta \alpha_T$

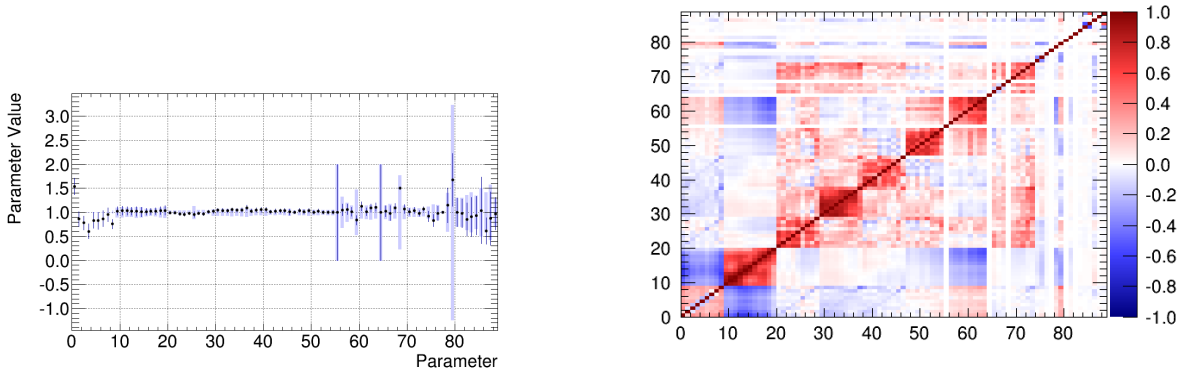
Figure 8.13: The L-curve used to determine the optimal regularisation parameters for each fit to GENIE 2.8.0 fake data. The regularisation strengths run in sequence down the curves (from left to right) and are the following: 1, 3, 5, 7 and 9. For this fake data study, a regularisation strength of 3 is chosen for all three STV.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 8.14: The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to GENIE 2.8.0 fake data. The data points in the left plot show the extracted parameters and their standard deviations alongside the light blue lines which show each parameters prior standard deviation. The parameters are labelled by the index in table 7.4.

Variable fit	χ^2_{stat}	χ^2_{syst}	χ^2_{IPSreg}	$\chi^2_{OOPSreg}$
δp_T	70.6	7.1	4.6	0.1
$\delta \phi_T$	38.4	6.5	3.7	0.0
$\delta \alpha_T$	57.9	10.7	1.9	0.0

Table 8.2: The size of each term in the maximum likelihood fit to GENIE 2.8.0 fake data.

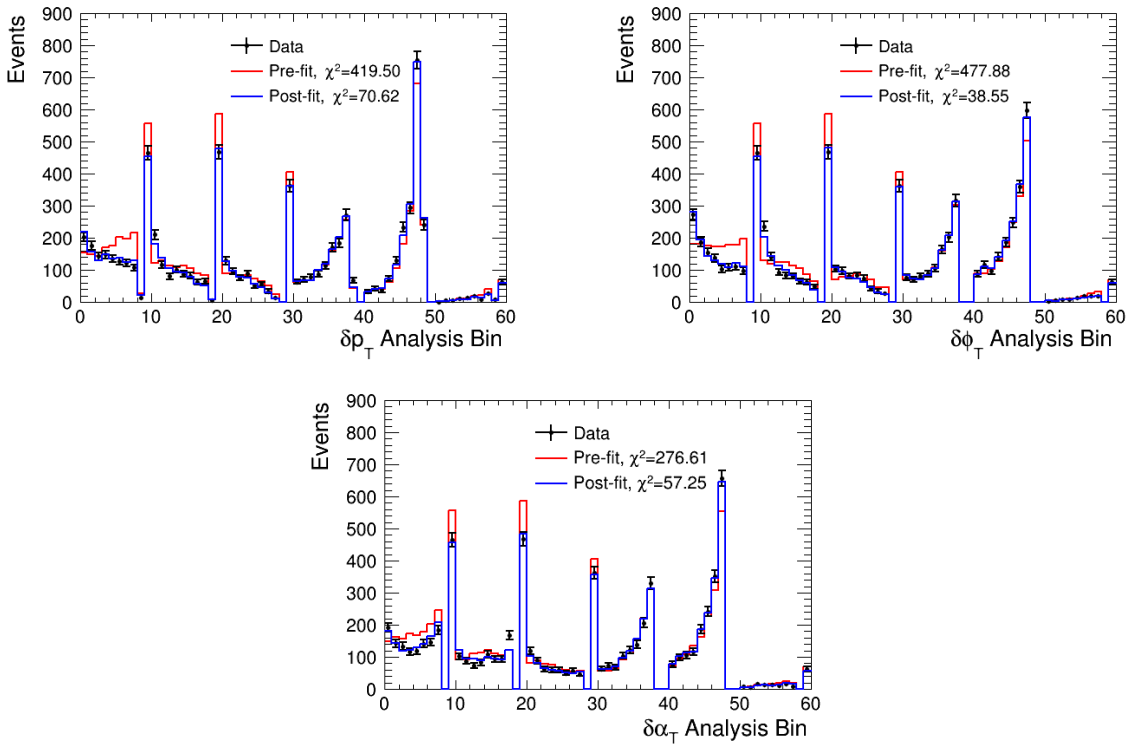


Figure 8.15: A comparison of the pre and post-fit simulation in reconstructed bins from a fit in each of the STV (separately) bins to GENIE 2.8.0 fake data. The χ^2 goodness of fit is shown. The reconstructed bins for each variable are listed in tables 7.1 to 7.3. These are repeated for each sample in the following order: μTPC pTPC, μTPC pFGD, μFGD pTPC, CC1π sideband, CCOther sideband, μTPC multi-p. Note that the buffer bin (δp_T bin 8) is included for all three STV but is always empty for $\delta \phi_T$ and $\delta \alpha_T$ (i.e. the 9th bin in each sample is only relevant for δp_T).

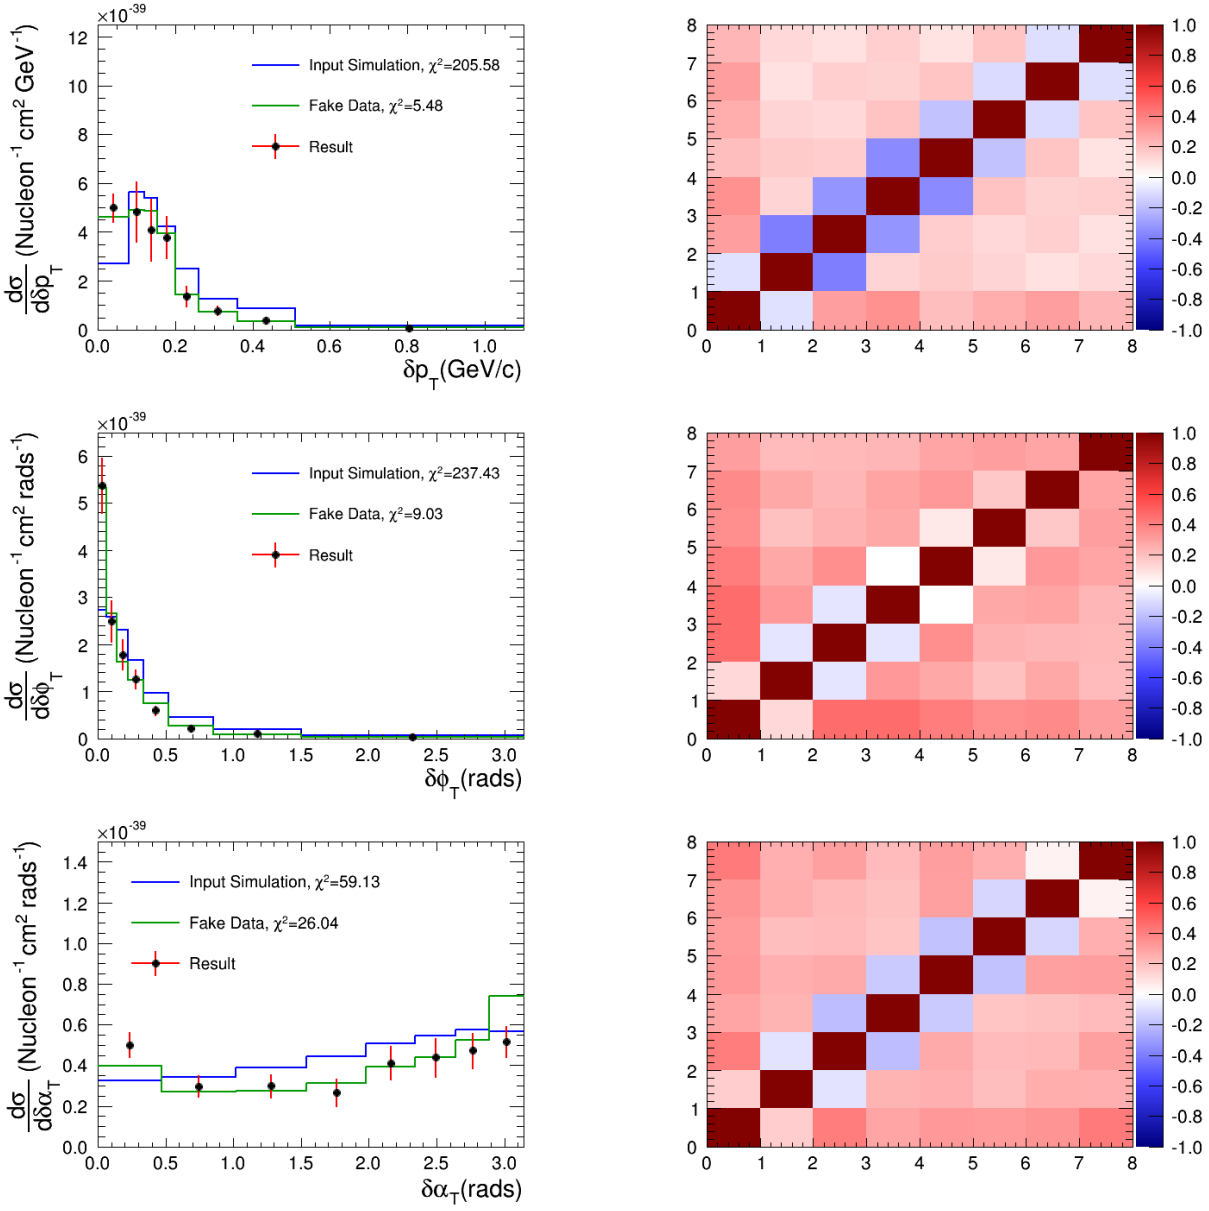
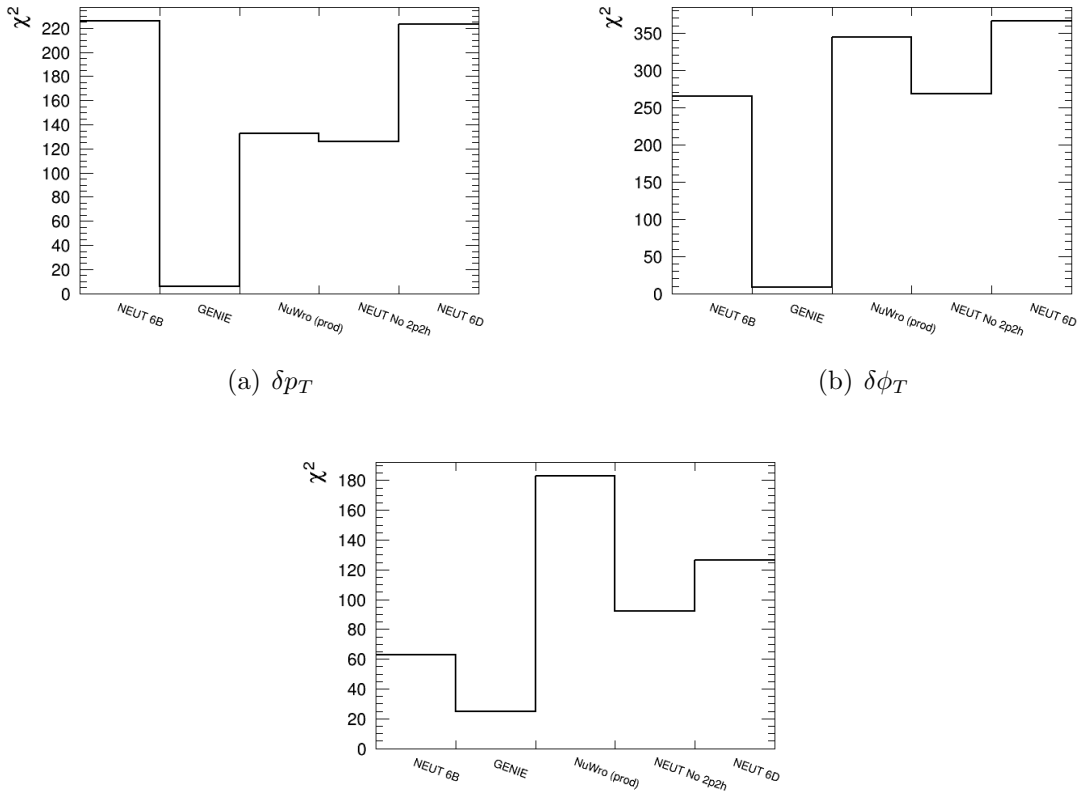


Figure 8.16: The differential cross section in the three STV and the associated correlation matrix obtained from separate fits to GENIE 2.8.0 fake data.



(a) δp_T

(b) $\delta \phi_T$

(c) $\delta \alpha_T$

Figure 8.17: The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to GENIE 2.8.0 fake data (labelled 'GENIE'), to all test simulations, calculated as in equation 8.1. NEUT 6B is the input simulation. The simulation predictions are extracted from their ND280 productions. 'NEUT 6D' is the bug-fixed NEUT and 'NEUT No 2p2h' is NEUT 6D without any 2p-2h events.

Genie fake data discussion The comparison of the pre and post-fit results in figure 8.15, alongside the parameter fit results in figure 8.14, demonstrates that the fits parametrisation is mostly flexible enough to allow a good fit to the GENIE 2.8.0 fake data across all three STV fits. However, it should be noted that, whilst good agreement is seen for the $\delta\alpha_T$ final IPS bin in the TPC proton samples (bins 7 and 27), there is a substantial discrepancy in the FGD proton sample (bin 17). This is indicative of different IPS underlying proton kinematic distributions between the fake data and the signal sample, discussed in section 7.2.6.

This issue is propagated through the $\delta\alpha_T$ cross section in figure 8.16, where there is a notable bias in the final bin and subsequent fairly poor agreement with the fake data truth. This demonstrates the importance of looking for poor agreement data and post-fit reconstructed-level results to diagnose potential bias in the extracted cross section. If a similar bias is manifest in a real data fit it would be identified from the discrepancy in the fit at the reconstructed level (the equivalent of figure 8.15).

Similarly to the NEUT fake data case, table 8.2 demonstrates that χ^2_{stat} is substantially larger than any other term in the fit, thereby showing that the dominant factor in finding the best fit point is the data. This is good evidence that the fit result is not biased by the systematic parameter priors or the regularisation.

The δp_T and $\delta\phi_T$ cross sections and χ^2 comparisons in figure 8.16 show that, despite the very large differences to the NEUT predictions in the early bins in each STV (especially $\delta\phi_T$), the fit is able to produce a result that is fully compatible with the GENIE 2.8.0 cross-section predictions and can clearly distinguish this from the other test simulations.

Overall this fake data test demonstrates that the analysis framework is clearly able to fit and distinguish GENIE 2.8.0 fake data in δp_T and $\delta\phi_T$ without significant bias to the input simulation, despite GENIE's signal and background models being quite different from the input's. The notable bias in the final bin of the $\delta\alpha_T$ cross section is understood and has highlighted the importance of post-fit to data comparisons.

Effect of regularisation and sidebands on the GENIE fake data fit The bias from regularisation and sidebands are assessed in the same way as in section 8.2.1. Figure 8.18 compares the differential cross section and χ^2 goodness of fit relative to all fake data, whilst figure 8.19 shows the same with the most favoured regularisation strengths with and without sidebands.

With the exception of the case with no regularisation in δp_T, it can be seen that both the regularisation strength and inclusion of sidebands makes very little difference to the nominal fit result but their presence modestly increases the fits power to separate between fake data sets, just as seen in section 8.2.1. For the no regularisation case in δp_T there is a

minimal effect on the χ^2 goodness of fit, but a substantial effect on the nominal result in the second and third bin. The zig-zag shape of the unregularised result in these bins is exactly what is expected when the fit encounters a highly ill-posed problem, it can be seen that similar behaviour is predicted in the toy study shown in figure 7.2. The fact that increasing the regularisation strength moves the best fit point but does not greatly alter the results sensitivity indicates that the regularisation is acting mostly as an aesthetic smoothing and is not causing any significant bias: the regularisation is effectively selecting a smooth result out of many otherwise degenerate possibilities. This is exactly what it is intended to do (as discussed in section 7.1.2).

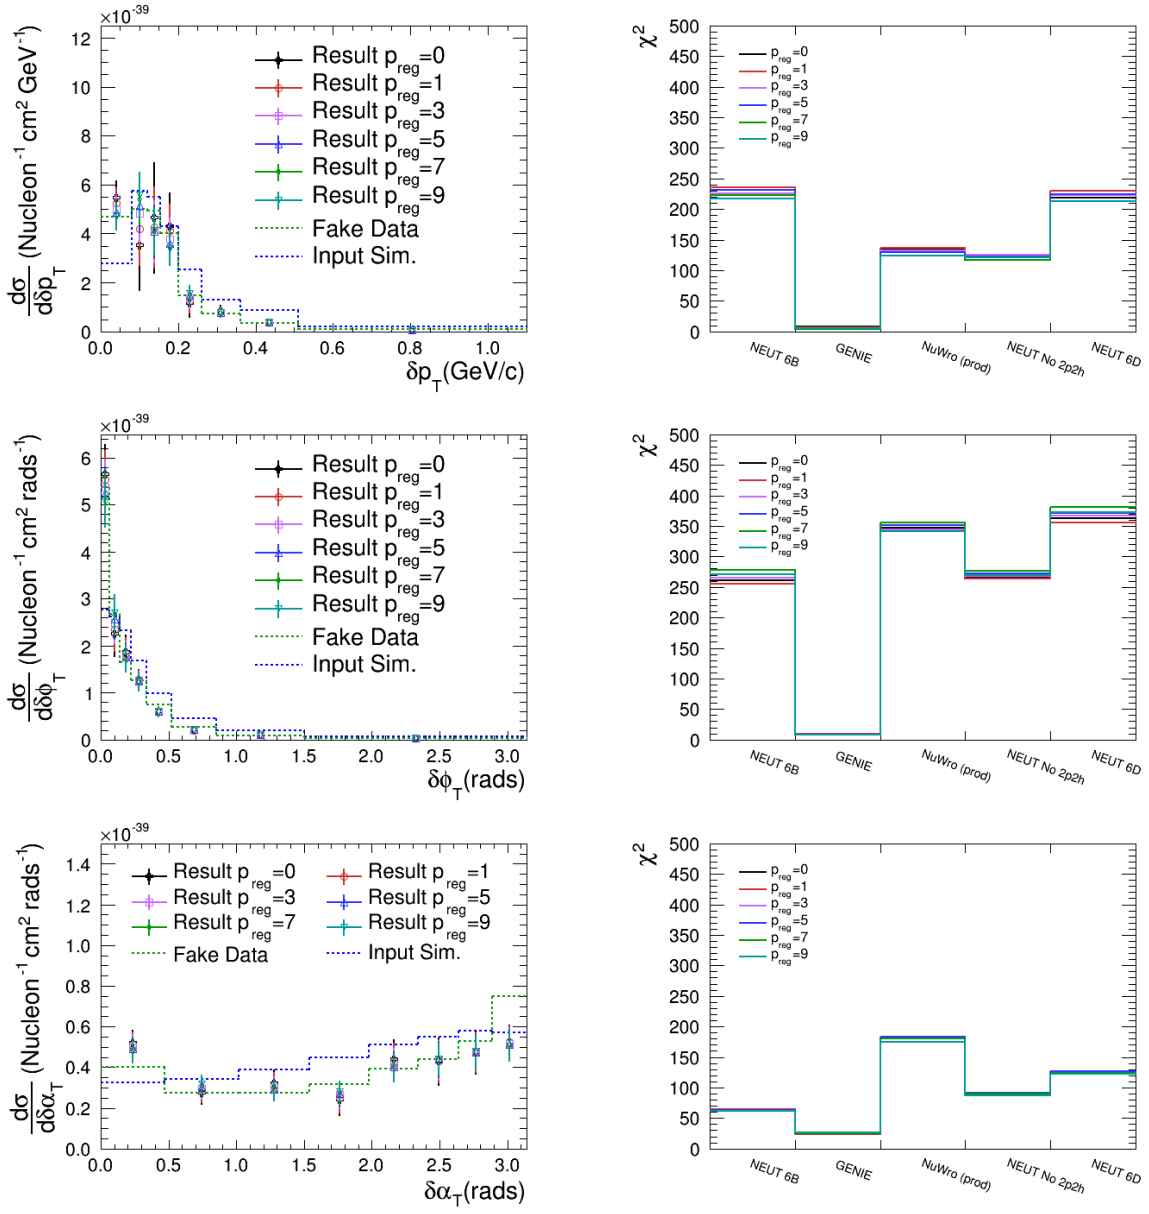


Figure 8.18: The differential cross section in the three STV obtained from a fit to GENIE 2.8.0 fake data for various regularisation strengths, alongside the χ^2 of the results relative to each fake data set.

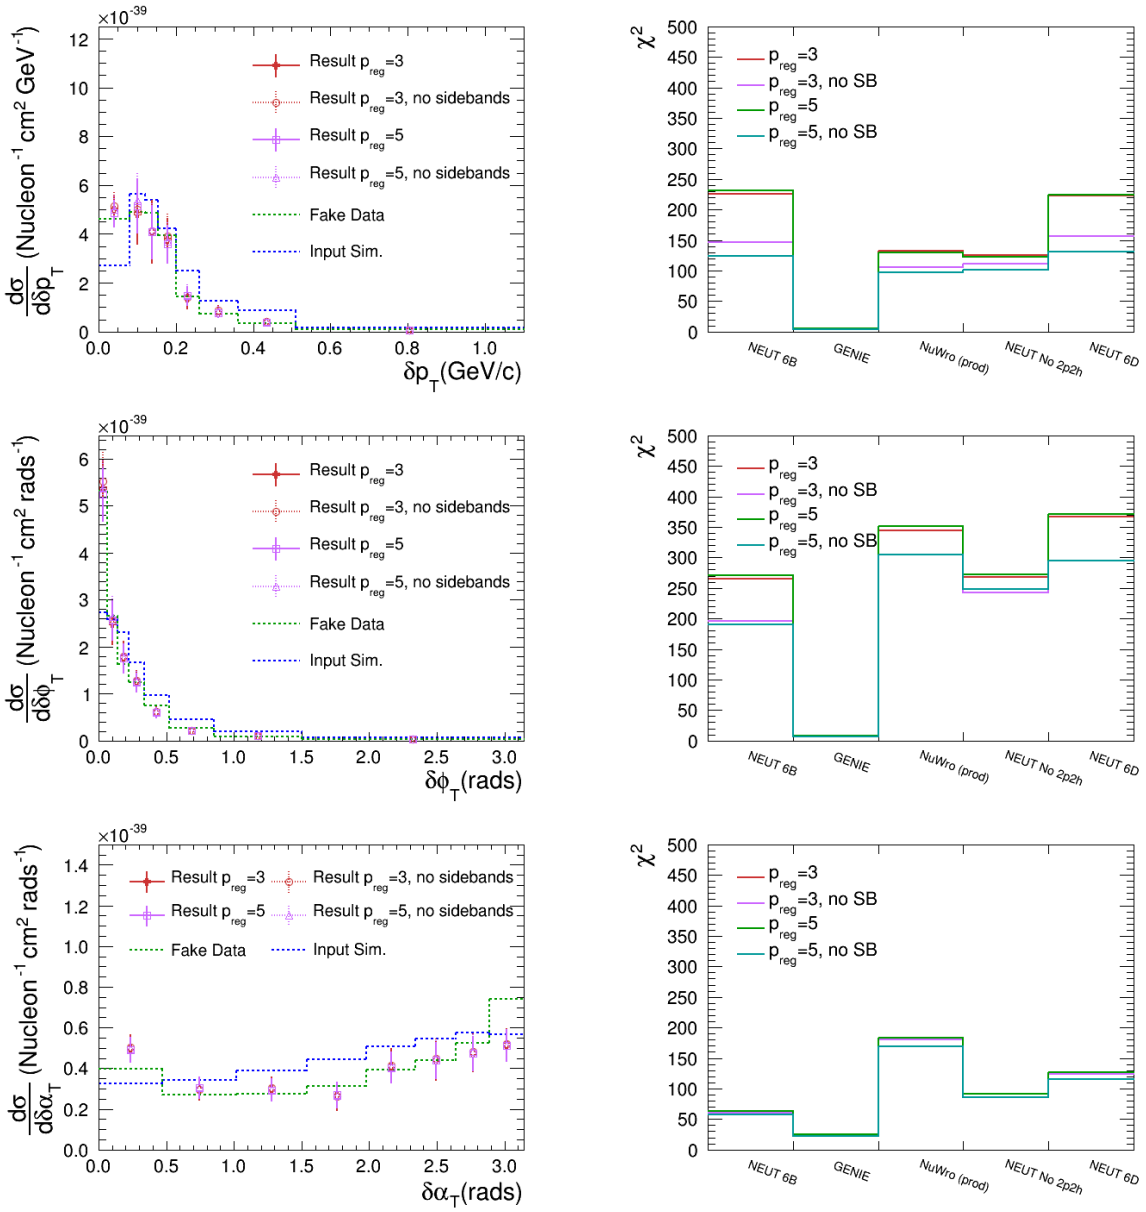


Figure 8.19: The differential cross section in the three STV obtained from a fit to GENIE 2.8.0 fake data using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set.

8.2.3 NuWro fake data

The NuWro 11q simulation is used to produce a fake data set scaled to real data POT. This simulation was used to produce events only in FGD1. This gives the prediction an unrealistic efficiency, which is remedied using the use of the 'effective truth' in a procedure described at the beginning of this section (8.2). This NuWro 11q simulation has a different spectral function (LFG) and a different implementation of signal and background models, as described in section 4.2.3. This allows a NuWro fake data test to complement the GENIE test to firmly establish the scope for bias to the input simulation when fitting quite different models.

Following the same procedure as in section 8.2.1, a regularisation study is made to determine the best choice of regularisation parameters. The resultant L-curve for each fit is shown in figure 8.20, from which a regularisation strength of 3 is chosen for all three STV. Once again, the IPS-OOPS regularisation strength was found to make very little difference to the result and is fixed at 1. The resultant χ^2 contribution of each term in the maximum likelihood fit (see equation 7.1) is shown in table 8.3. Following this, the extracted fit parameters and correlations are shown in figure 8.21, whilst comparisons of the pre and post-fit simulation with the fake data are shown in reconstructed bins in figure 8.22. The cross-sections, alongside their correlation matrix, are shown in figure 8.23 and the χ^2 goodness of fit results are shown in figure 8.24.

NuWro fake data discussion The comparison of the pre and post-fit results in figure 8.22, alongside the parameter fit results in figure 8.21, shows that the fits parametrisation is flexible enough to allow a good fit to the NuWro 11q fake data, suggesting that overall there is enough freedom in the fit to account for the model difference. There is no evidence of a similarly poor goodness of fit (in the signal samples or otherwise) that was inherent in the GENIE fake data result.

However, it is interesting to note that the fit reduces the flux down by ∼5-10%, picking up a notable penalty term in doing so (as shown in 8.3)[3]. Considering that the flux used to generate NuWro was the same as all the other fake data, this movement is potentially accounting for changes in background models that cannot quite be reached by varying the cross-section systematic parameters within their priors (the potential for this to happen was discussed in section 7.2.5 and is the reason for using pre-fit covariances when calculating uncertainties on the efficiency, number of targets and flux normalisation corrections). It is also possible that this flux shift stems a mis-evaluation of the effective POT the NuWro simulation was generated for (this possibility was highlighted in the introduction to section 8.2).

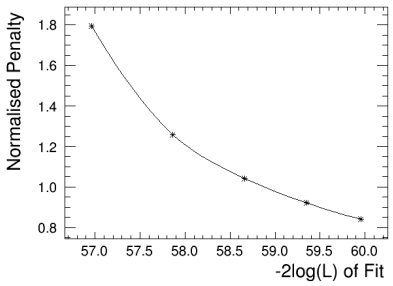
[3]It is worth noting that this penalty term is still much smaller than may be inferred from figure 8.21 alone, since the flux parameters are very highly correlated (and therefore are more free to move than is suggested by the prior variance, providing the parameters all move together).

Regardless of the source of the small shift, table 8.3 demonstrates that χ^2_{stat} is substantially larger than any other term in the fit, thereby showing that the dominant factor in finding the best fit point is the data. This is good evidence that any bias due to the systematic parameter priors is minimal.

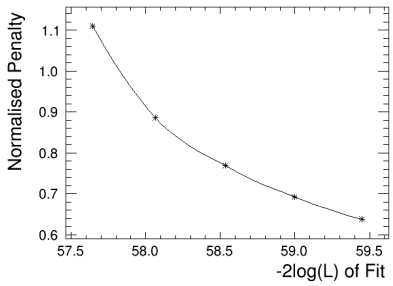
The STV cross-sections and χ^2 comparisons in figure 8.23 and 8.24 show that, despite the very large differences to the input NEUT 6B predictions, the fit is able to produce a result that is very compatible with the NuWro 11q cross-section predictions and can clearly distinguish this from the other test simulations. This highlights the lack of bias in the input simulation and the sensitivity of the fit to differentiate NuWro 11q from other models.

Effect of regularisation and sidebands on the NuWro fake data fit The bias from regularisation and sidebands are assessed as described in section 8.2.1. Figure 8.25 compares the differential cross section and χ^2 goodness of fit relative to all fake data whilst figure 8.26 shows the same with the most favoured regularisation strength with and without sidebands.

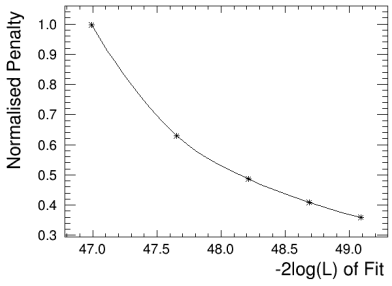
Figure 8.12 demonstrates that there is a small bias in the result if the sidebands are not used. This is due to the subsequent lack of constraint on the background model predictions



(a) δp_T

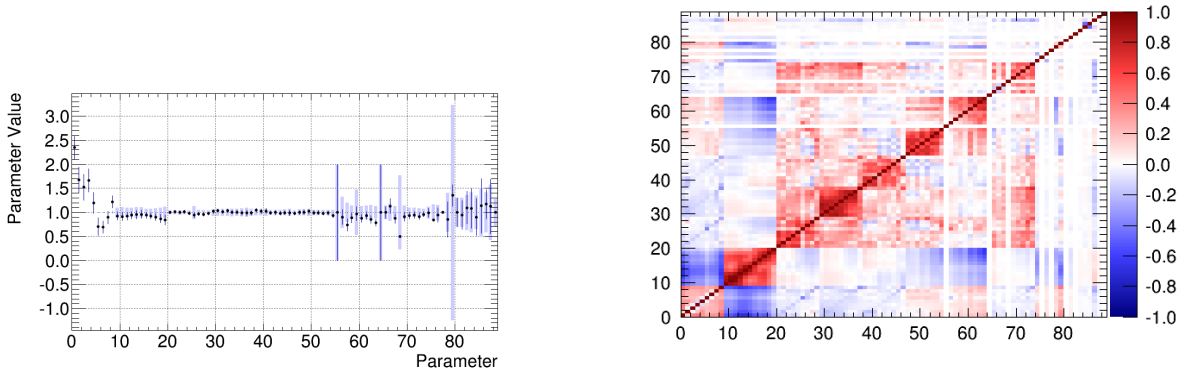


(b) $\delta \phi_T$

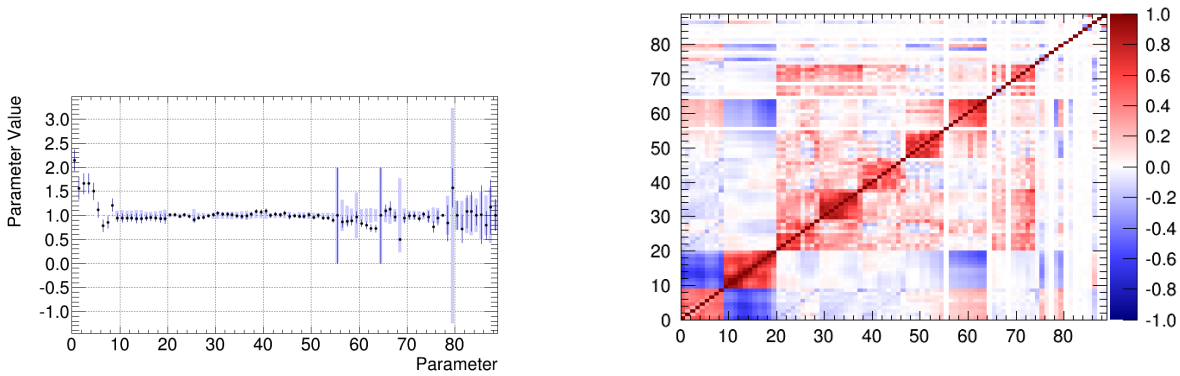


(c) $\delta \alpha_T$

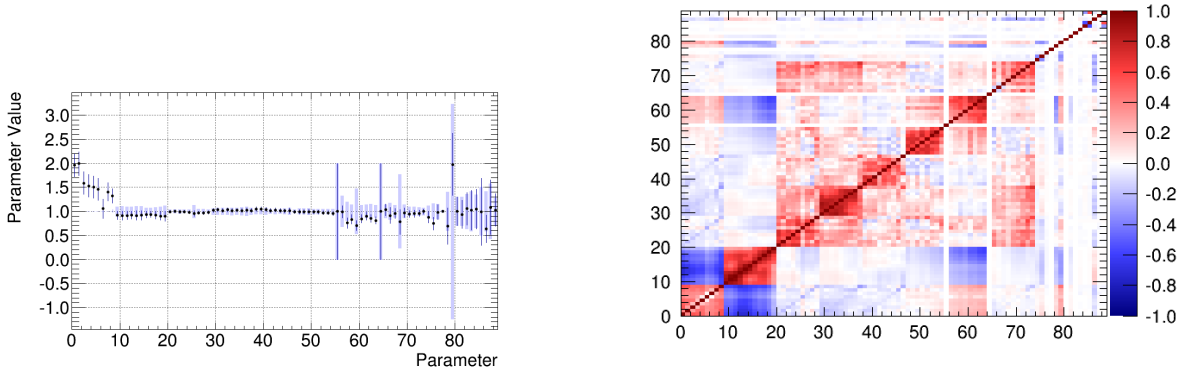
Figure 8.20: The L-curve used to determine the optimal regularisation parameters for each fit to NuWro 11q fake data. The regularisation strengths run in sequence down the curves (from left to right) and are the following: 1, 3, 5, 7 and 9. For this fake data study, a regularisation strength of 3 is chosen for all three STV.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 8.21: The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NuWro 11q fake data. The data points in the left plot show the extracted parameters and their standard deviations alongside the light blue lines which show each parameters prior standard deviation. The parameters are labelled by the index in table 7.4.

Variable fit	χ^2_{stat}	χ^2_{syst}	χ^2_{IPSreg}	$\chi^2_{OOPSreg}$
δp_T	45.0	12.3	3.1	0.2
$\delta \phi_T$	41.3	16.4	2.5	0.1
$\delta \alpha_T$	34.8	12.9	1.7	0.1

Table 8.3: The size of each term in the maximum likelihood fit to NuWro 11q fake data.

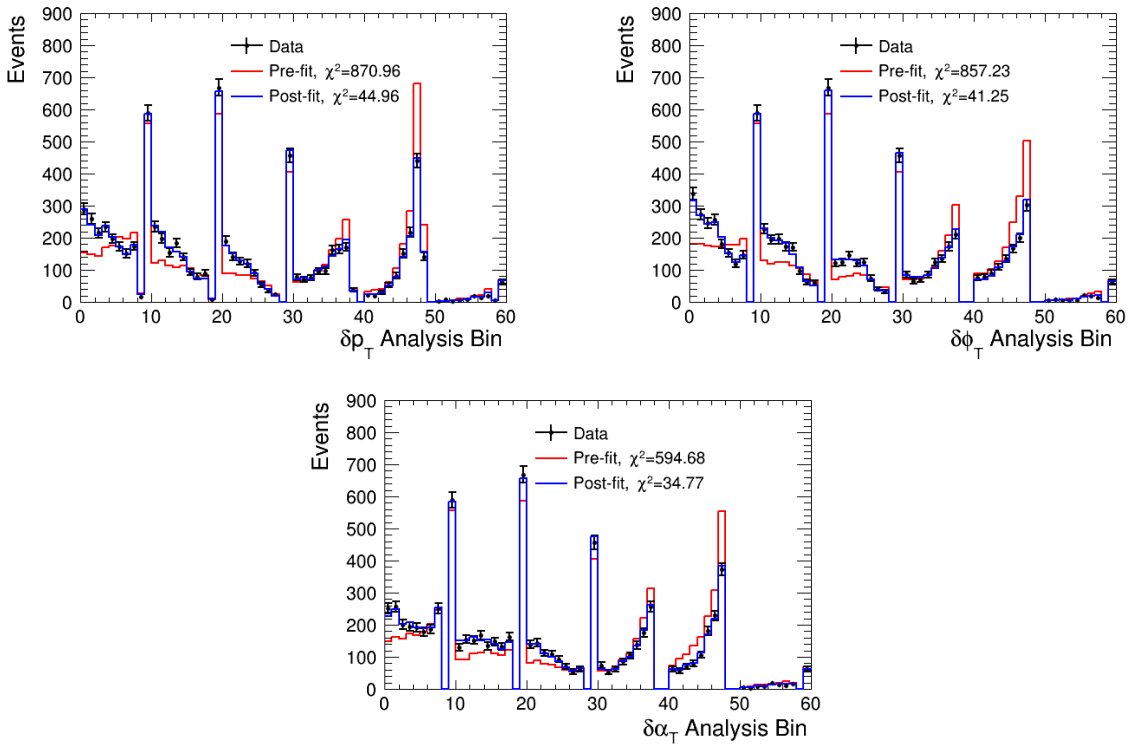


Figure 8.22: A comparison of the pre and post-fit input simulation in reconstructed bins from a fit in each of the STV (separately) bins to NuWro 11q fake data. The χ^2 goodness of fit is shown. The reconstructed bins for each variable are listed in tables 7.1 to 7.3. These are repeated for each sample in the following order: μTPC pTPC, μTPC pFGD, μFGD pTPC, CC1π sideband, CCOther sideband, μTPC multi-p. Note that the buffer bin (δp_T bin 8) is included for all three STV but is always empty for $\delta \phi_T$ and $\delta \alpha_T$ (i.e. the 9th bin in each sample is only relevant for δp_T).

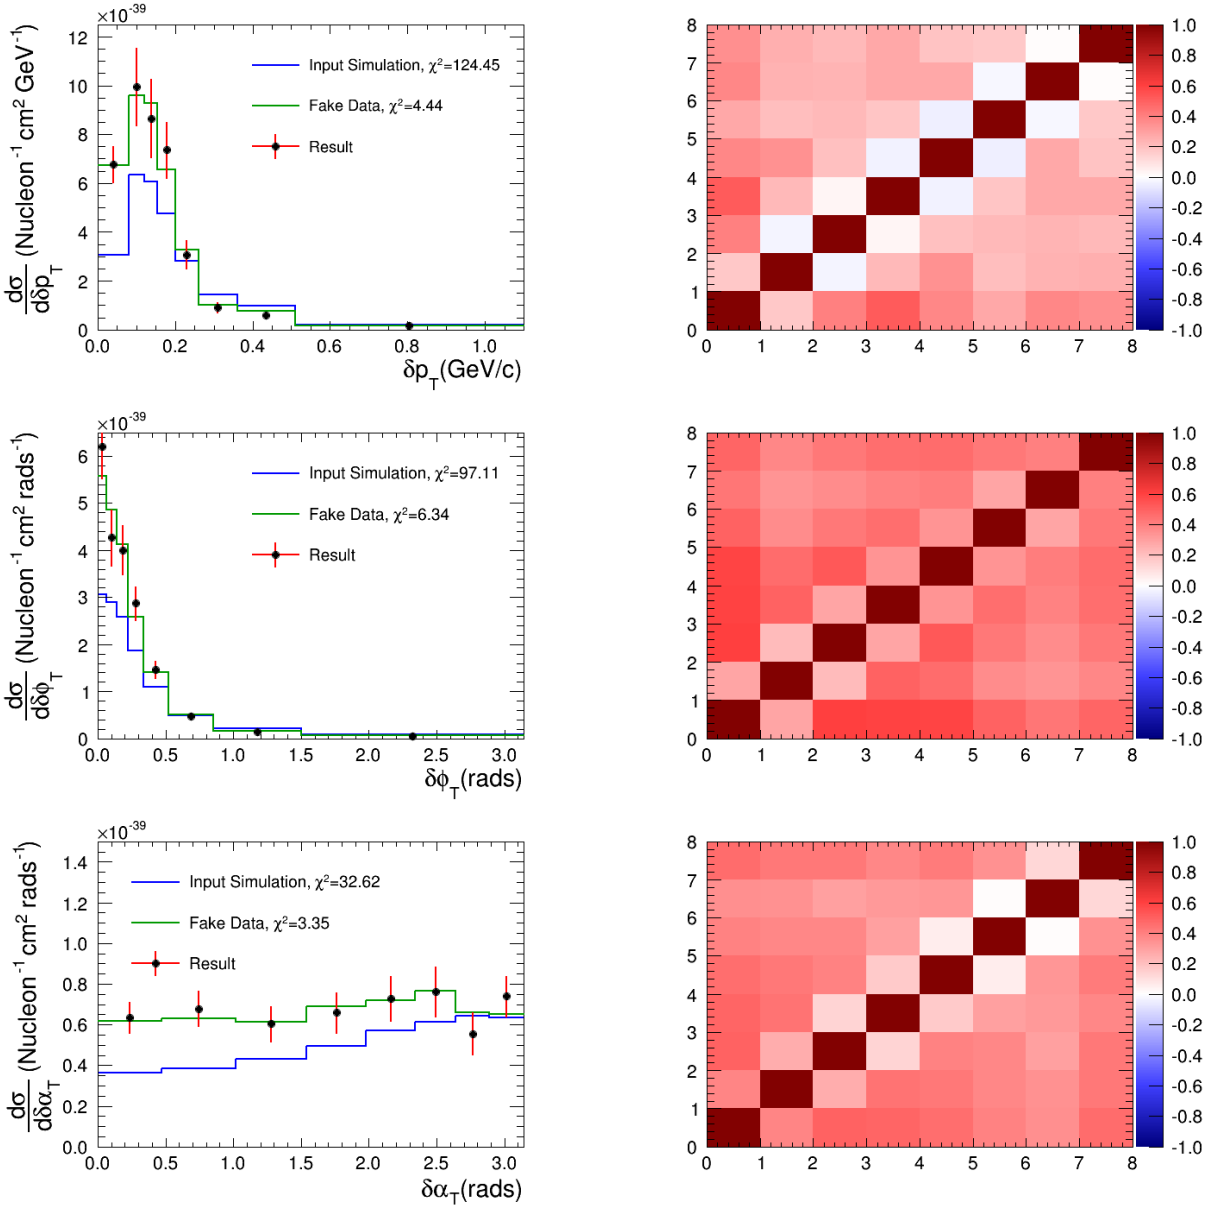
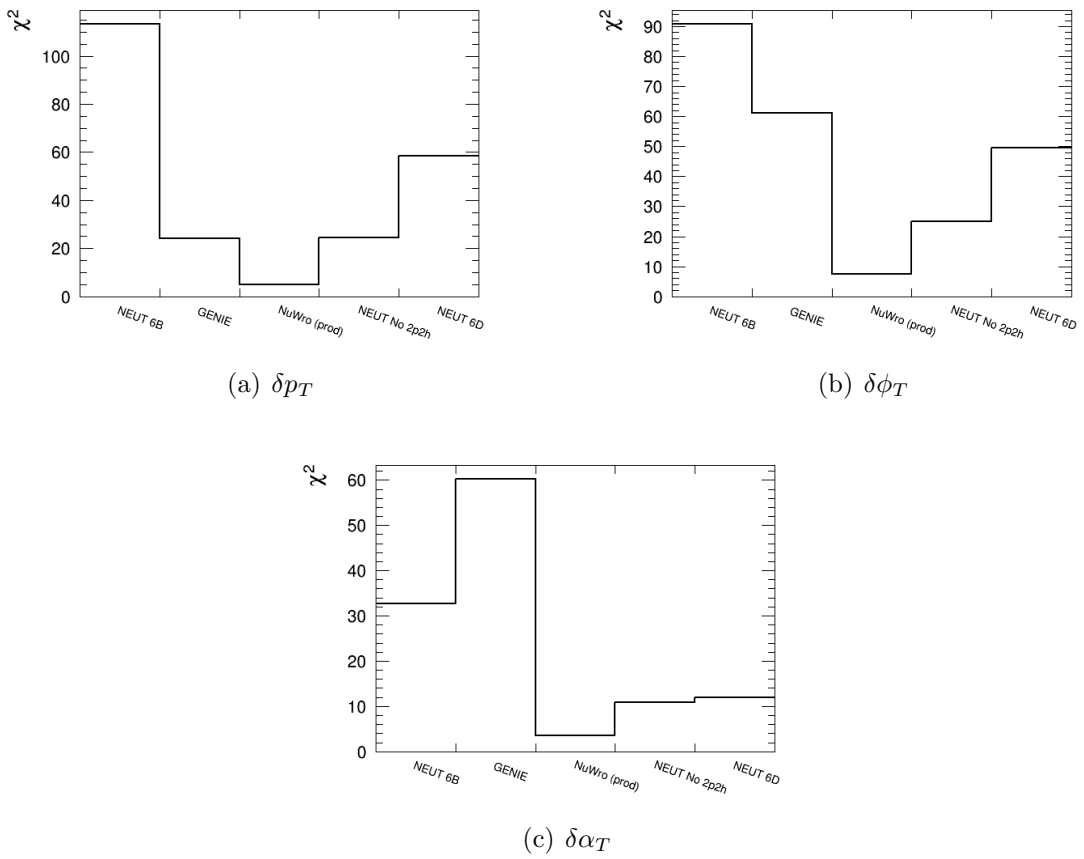


Figure 8.23: The differential cross section in the STV and the associated correlation matrices obtained from fits to NuWro 11q fake data.



(a) δp_T

(b) $\delta \phi_T$

(c) $\delta \alpha_T$

Figure 8.24: The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to NuWro 11q fake data (labelled 'NuWro prod'), to all test simulations, calculated as in equation 8.1. NEUT 6B is the input simulation. The simulation predictions are extracted from their ND280 productions. 'NEUT 6D' is the bug-fixed NEUT and 'NEUT No 2p2h' is NEUT 6D without any 2p-2h events.

(which are different in NuWro 11q and the NEUT input simulation). This indicates that, as far as extracting the cross sections in the STV is concerned, the parametrisation of the background model systematic parameters is sufficient to allow the sideband samples to successfully constrain the backgrounds to look similar to in NuWro 11q.

It can be seen that the regularisation strength makes very little difference to the nominal fit result or its sensitivity. Overall there is no suggestion of a bias from regularisation and that inclusion of sidebands are usefully avoiding bias.

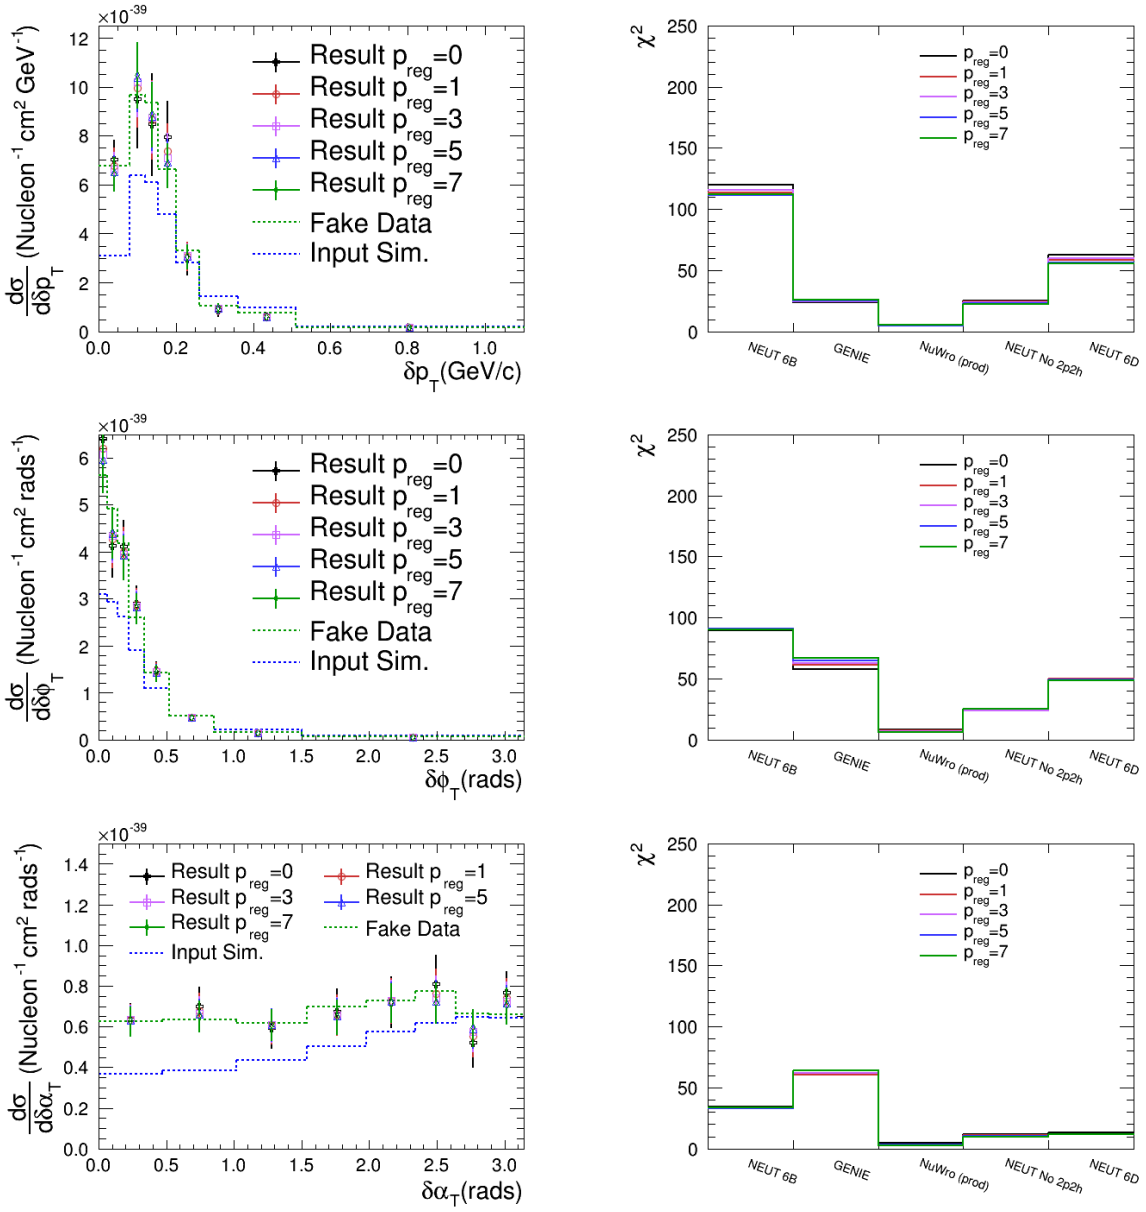


Figure 8.25: The differential cross section in the three STV obtained from a fit to NuWro 11q fake data for various regularisation strengths, alongside the χ^2 of the results relative to each fake data set.

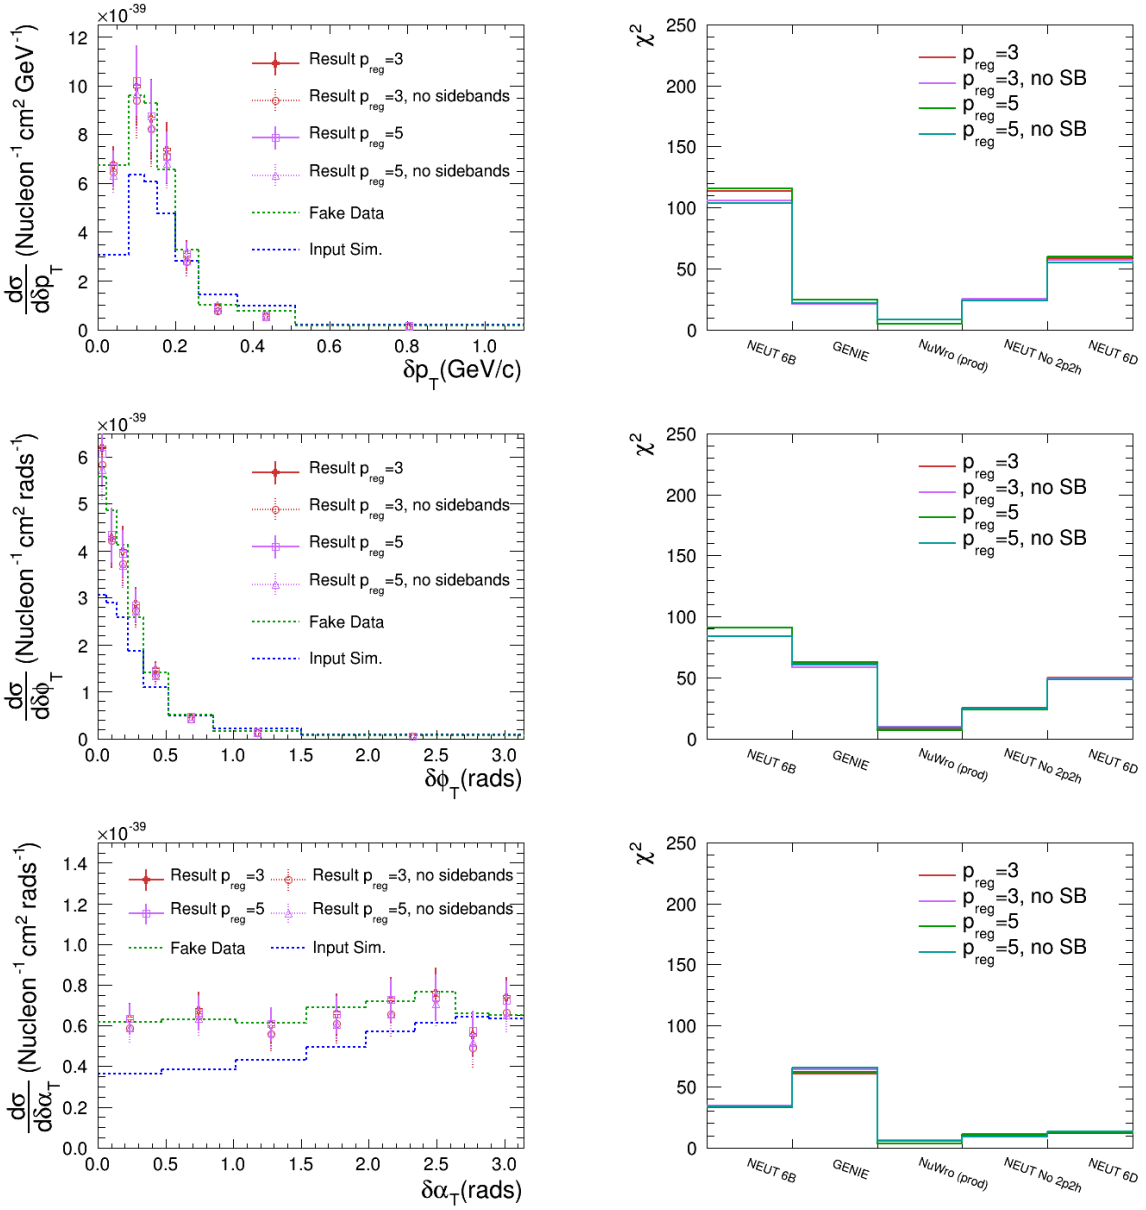


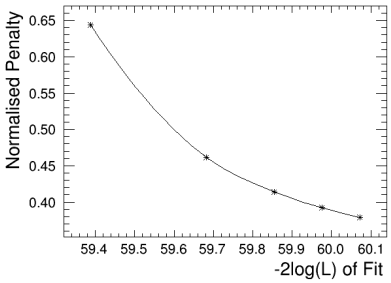
Figure 8.26: The differential cross section in the three STV obtained from a fit to NuWro 11q fake data using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set.

8.2.4 NEUT 6D fake data

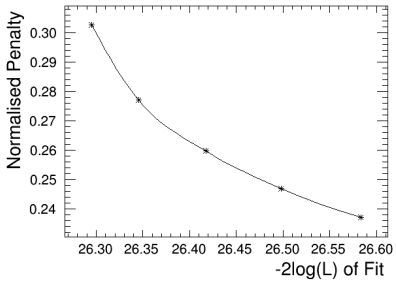
NEUT 6D, a version of the NEUT simulation without the nucleon FSI bug, is used to produce a fake data set scaled to real data POT. Since both the observables and the phase-space constraints are quite sensitive to FSI, the fake data NEUT differs from the input simulation considerably. NEUT 6D was used to produce events only inside the ND280 magnet, which (like with the NuWro fake data) gives it an unrealistic efficiency that's accounted for using the procedure described in the beginning of this section. The tests so far have demonstrated that the fit does not appear substantially biased to the input simulation, but showing that it can successfully extract a realistically modified version of its input simulation is essential to demonstrate that the bug in the input simulation is ineffectual.

Following the same procedure as in section 8.2.1, a regularisation study is made to determine the best choice of regularisation parameters. The resultant L-curve for each fit is shown in figure 8.27 which favours a regularisation strength of 3 for all three STV. Once again, the IPS-OOPS regularisation strength was found to make very little difference to the result and is fixed at 1. The resultant χ^2 contribution of each term in the maximum likelihood fit (see equation 7.1) is shown in table 8.4. Following this, the extracted fit parameters and correlations are shown in figure 8.28, whilst comparisons of the pre and post-fit simulation to fake data are shown in reconstructed bins in figure 8.29. The cross-sections, alongside their correlation matrix, are shown in figure 8.30 and the χ^2 goodness of fit results are shown in figure 8.31.

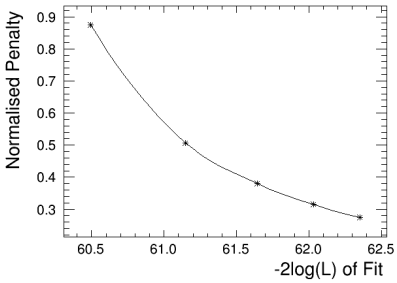
Note that, section 8.2.5 will use NEUT 6D with its 2p-2h contribution removed as fake data, to assess a model dependent sensitivity to 2p-2h. It is therefore particularly interesting to note the degree of separation between NEUT 6D with and without the 2p-2h contribution removed in figure 8.31.



(a) δp_T

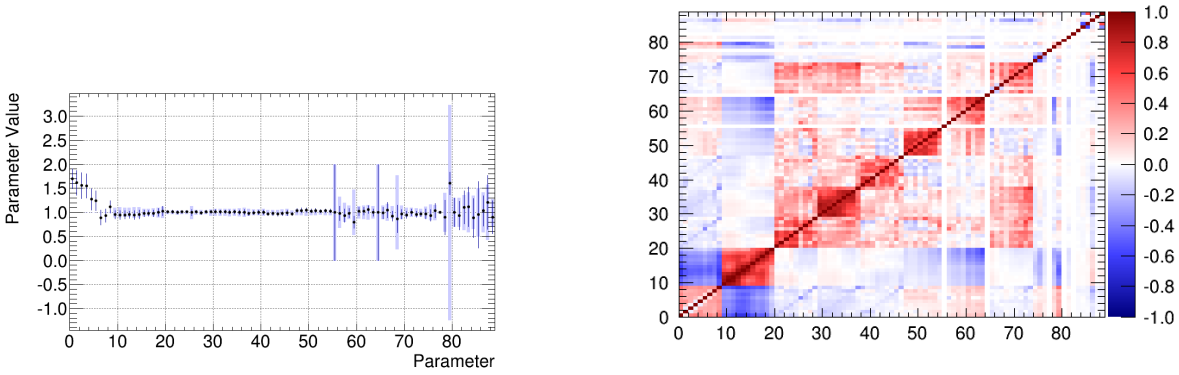


(b) $\delta \phi_T$

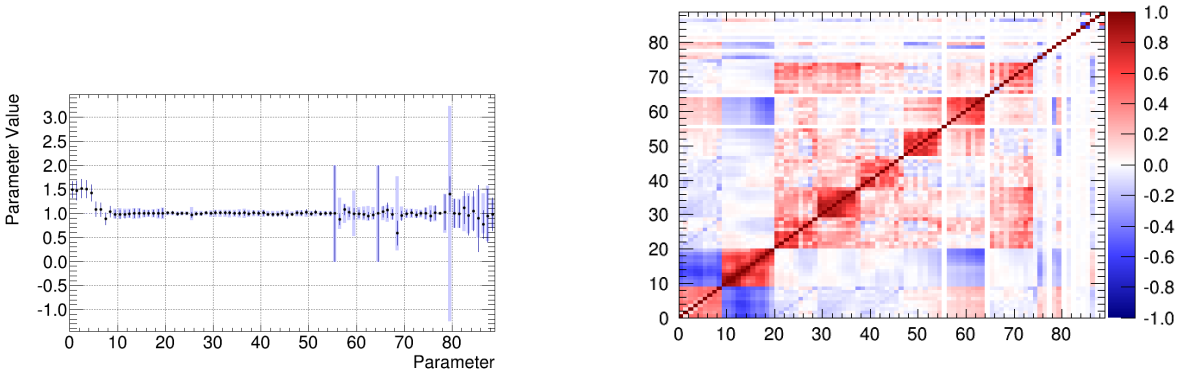


(c) $\delta \alpha_T$

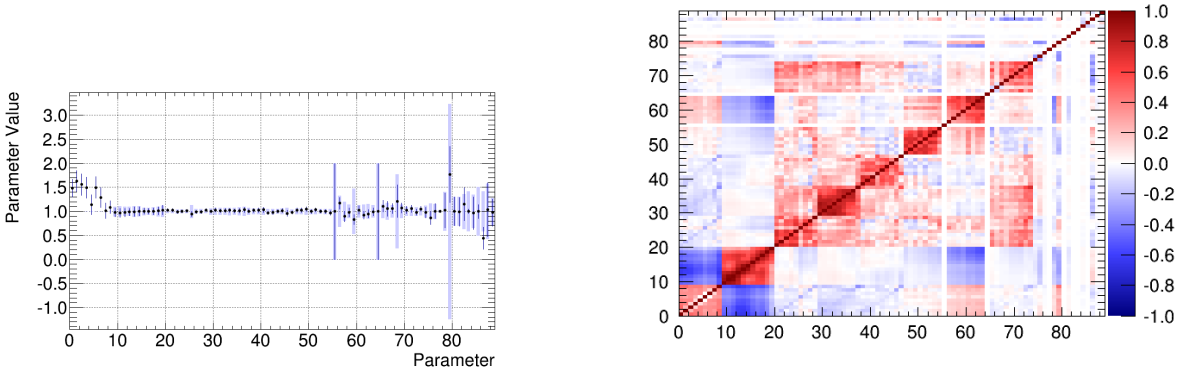
Figure 8.27: The L-curve used to determine the optimal regularisation parameters for each fit to NEUT 6D (bug-fixed) fake data. The regularisation strengths run in sequence down the curves (from left to right) and are the following: 1, 3, 5, 7 and 9. For this fake data study, a regularisation strength of 3 is chosen for all three STV.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 8.28: The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NEUT 6D (bug-fixed) fake data. The data points in the left plot show the extracted parameters and their standard deviations alongside the light blue lines which show each parameters prior standard deviation. The parameters are labelled by the index in table 7.4.

Variable fit	χ^2_{stat}	χ^2_{syst}	χ^2_{IPSreg}	$\chi^2_{OOPSreg}$
δp_T	52.1	6.7	0.7	0.2
$\delta \phi_T$	22.3	4.0	0.6	0.1
$\delta \alpha_T$	50.2	11.0	1.2	0.1

Table 8.4: The size of each term in the maximum likelihood fit to NEUT 6D (bug-fixed) fake data.

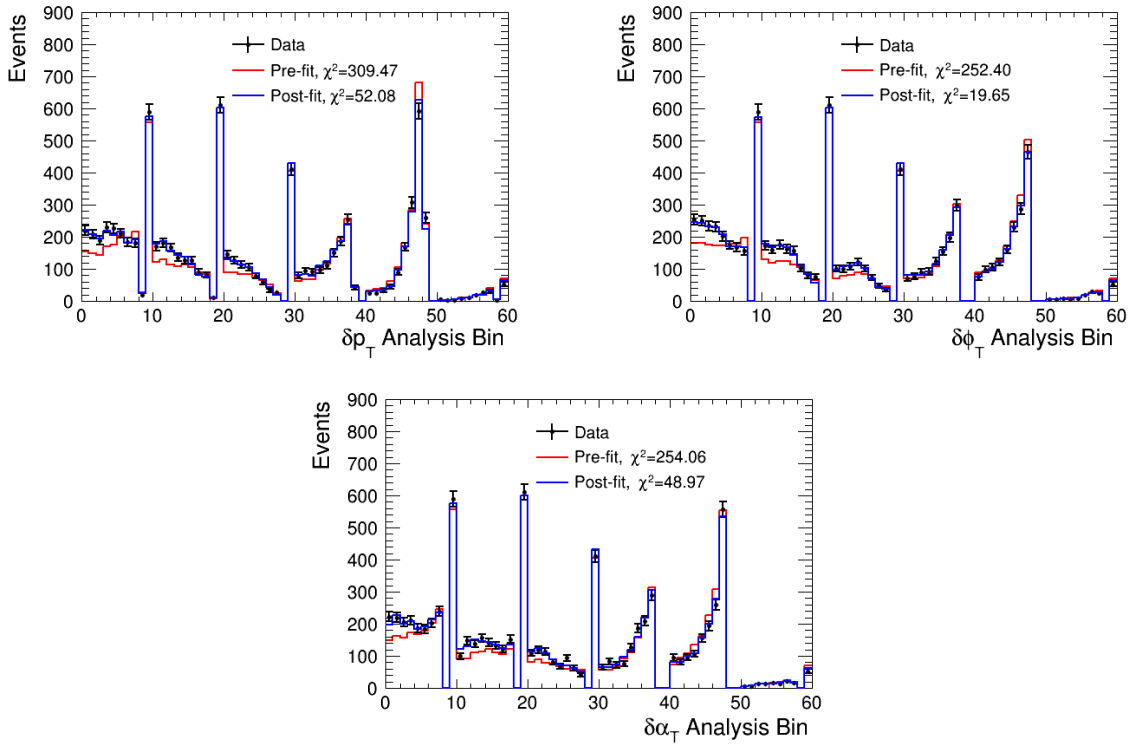


Figure 8.29: A comparison of the pre and post-fit input simulation in reconstructed bins from a fit in each of the STV (separately) bins to NEUT 6D (bug-fixed) fake data. The χ^2 goodness of fit is shown. The reconstructed bins for each variable are listed in tables 7.1 to 7.3. These are repeated for each sample in the following order: μTPC pTPC, μTPC pFGD, μFGD pTPC, CC1π sideband, CCOther sideband, μTPC multi-p. Note that the buffer bin (δp_T bin 8) is included for all three STV but is always empty for $\delta \phi_T$ and $\delta \alpha_T$ (i.e. the 9th bin in each sample is only relevant for δp_T).

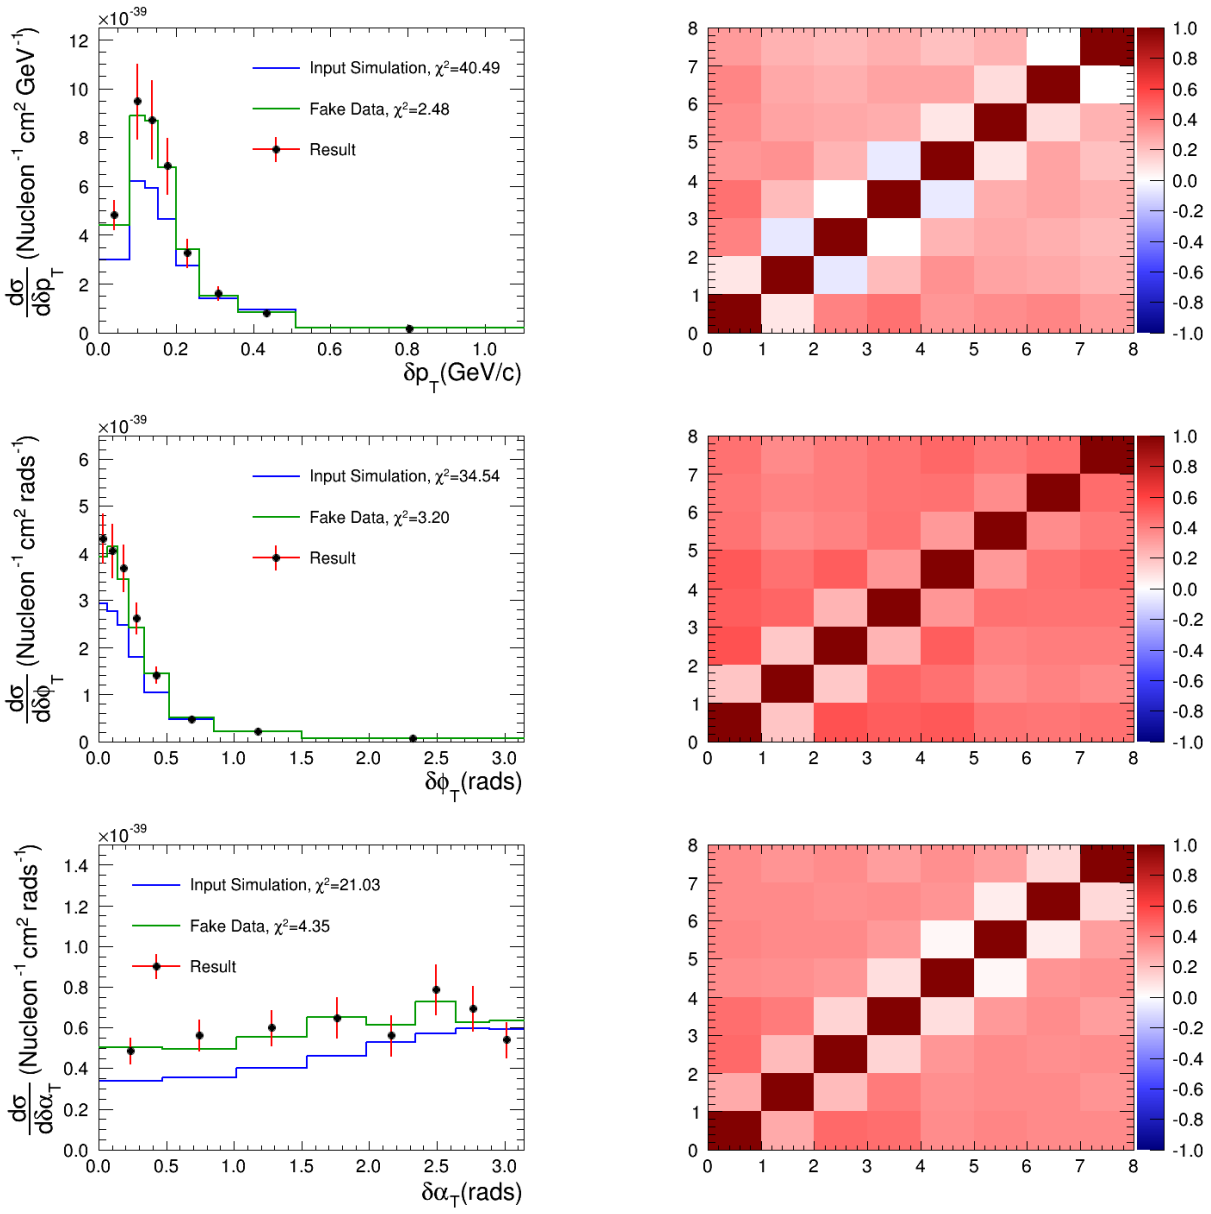
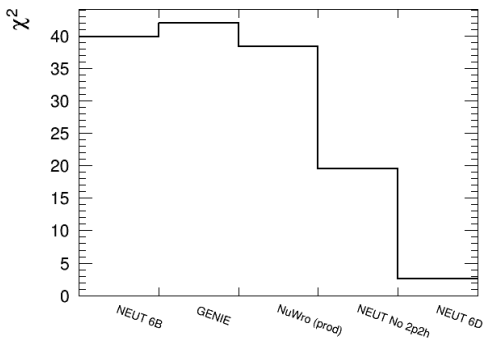
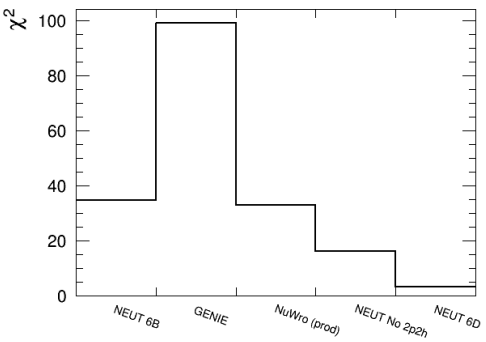


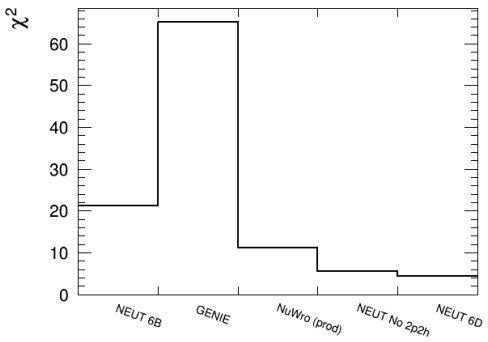
Figure 8.30: The differential cross section in the STV and the associated correlation matrix obtained from a fit to NEUT 6D (bug-fixed) fake data.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 8.31: The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to NEUT 6D (bug fixed) fake data (labelled 'NEUT 6D'), to all test simulations, calculated as in equation 8.1. The plots are ordered from the fit as δp_T in the top left, $\delta \phi_T$ in the top right and $\delta \alpha_T$ in the centre. NEUT 6B is the input simulation (with the nucleon FSI bug). The simulation predictions are extracted from their ND280 productions. 'NEUT No 2p2h' is NEUT 6D without any 2p-2h events.

NEUT 6D fake data discussion The comparison of the pre and post-fit results in figure 8.29 and the parameter fit results to the fake data in figure 8.28 demonstrates that the fits parametrisation is flexible enough to allow a good fit to the NEUT bug-fixed fake data. This is preliminarily indicates of a minimal bias toward the input simulation.

Similarly to the other data tests, table 8.4 demonstrates that χ^2_{stat} is substantially larger than any other term in the fit, thereby showing that the dominant factor in finding the best fit point is the data. This is good evidence that the fit result is not biased by the systematic parameter priors or the regularisation.

The STV cross-sections and χ^2 comparisons in figures 8.30 and 8.31 show that, despite the substantial differences to the input NEUT predictions, the fit is able to produce a result that is compatible with the bug-fixed NEUT cross-section predictions and can clearly distinguish this from the other test simulation. This includes, at least in δp_T and $\delta\phi_T$, a fairly clear distinction from the NEUT 6D prediction with its 2p-2h contribution removed, suggesting some sensitivity to the 2p-2h contribution in the result (this will be further analysed in the next section - 8.2.5).

Overall this fake data study indicates the fit is clearly able to identify a result that resembles the NEUT prediction without the nucleon FSI bug whilst using an input simulation that suffers from it. This suggests that the bug is unlikely to substantially constrain the fit.

Effect of regularisation and sidebands on the NEUT 6D fake data fit The bias from regularisation and sidebands are assessed as described in section 8.2.1. Figure 8.32 compares the differential cross section and χ^2 goodness of fit relative to all fake data whilst figure 8.33 shows the same with the most favoured regularisation strength with and without sidebands.

It can be seen that both the regularisation strength and inclusion of sidebands makes very little difference to the nominal fit result and that the sidebands offer the expected modest increase in sensitivity. The same zig-zag in the unregularised δp_T result is observed that was seen and understood in the GENIE fake data (section 8.2.2). Overall there is no suggestion of a bias from regularisation or the inclusion of sidebands.

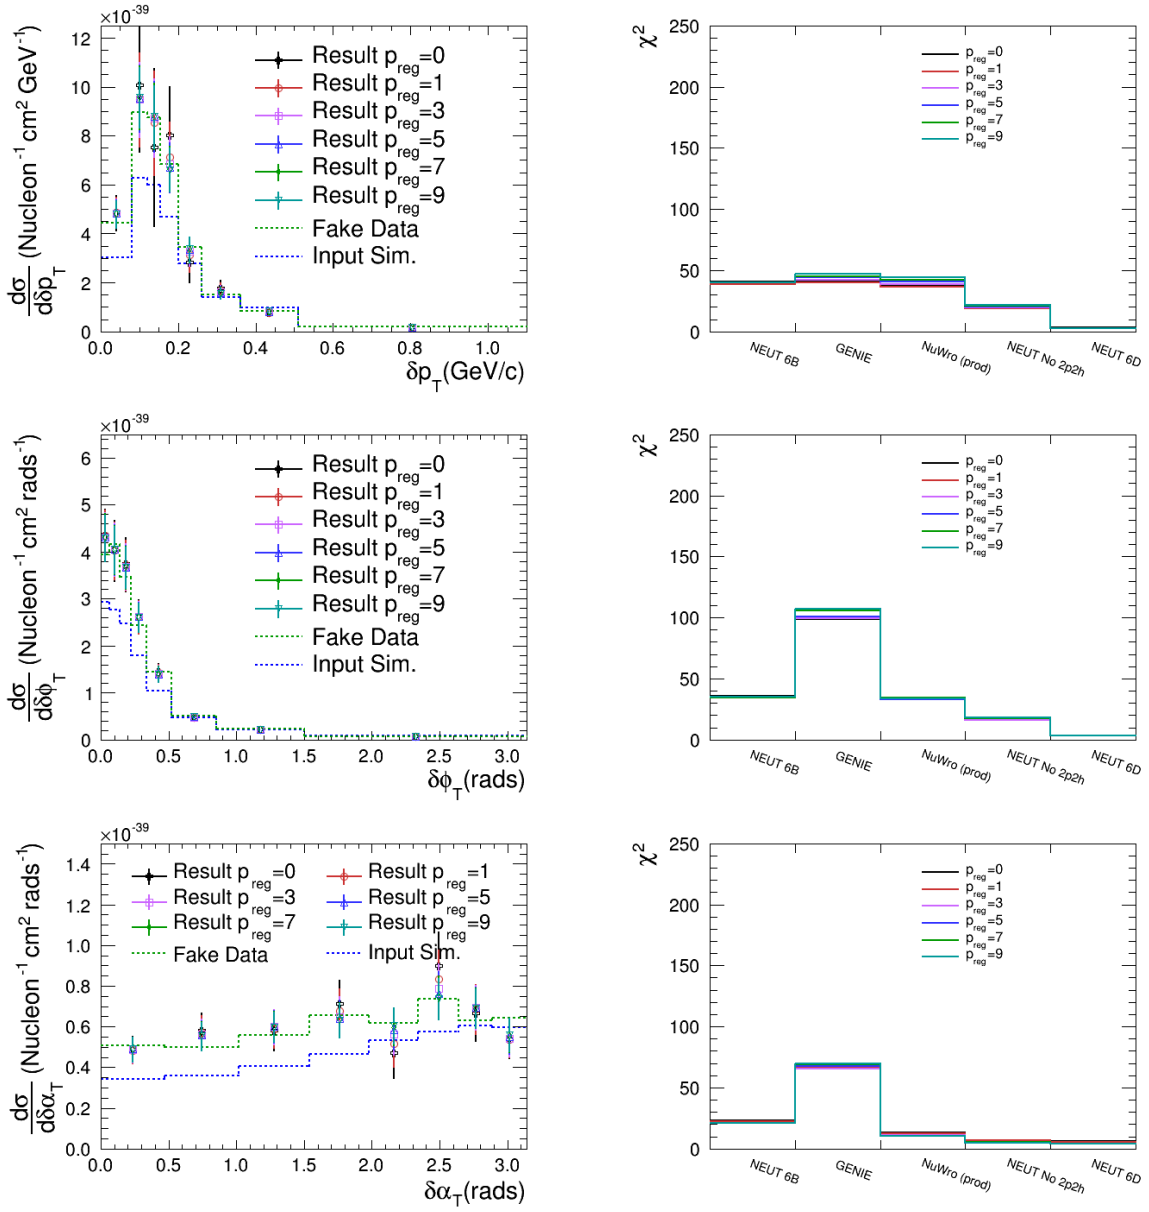


Figure 8.32: The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data for various regularisation strengths, alongside the χ^2 of the results relative to each fake data set.

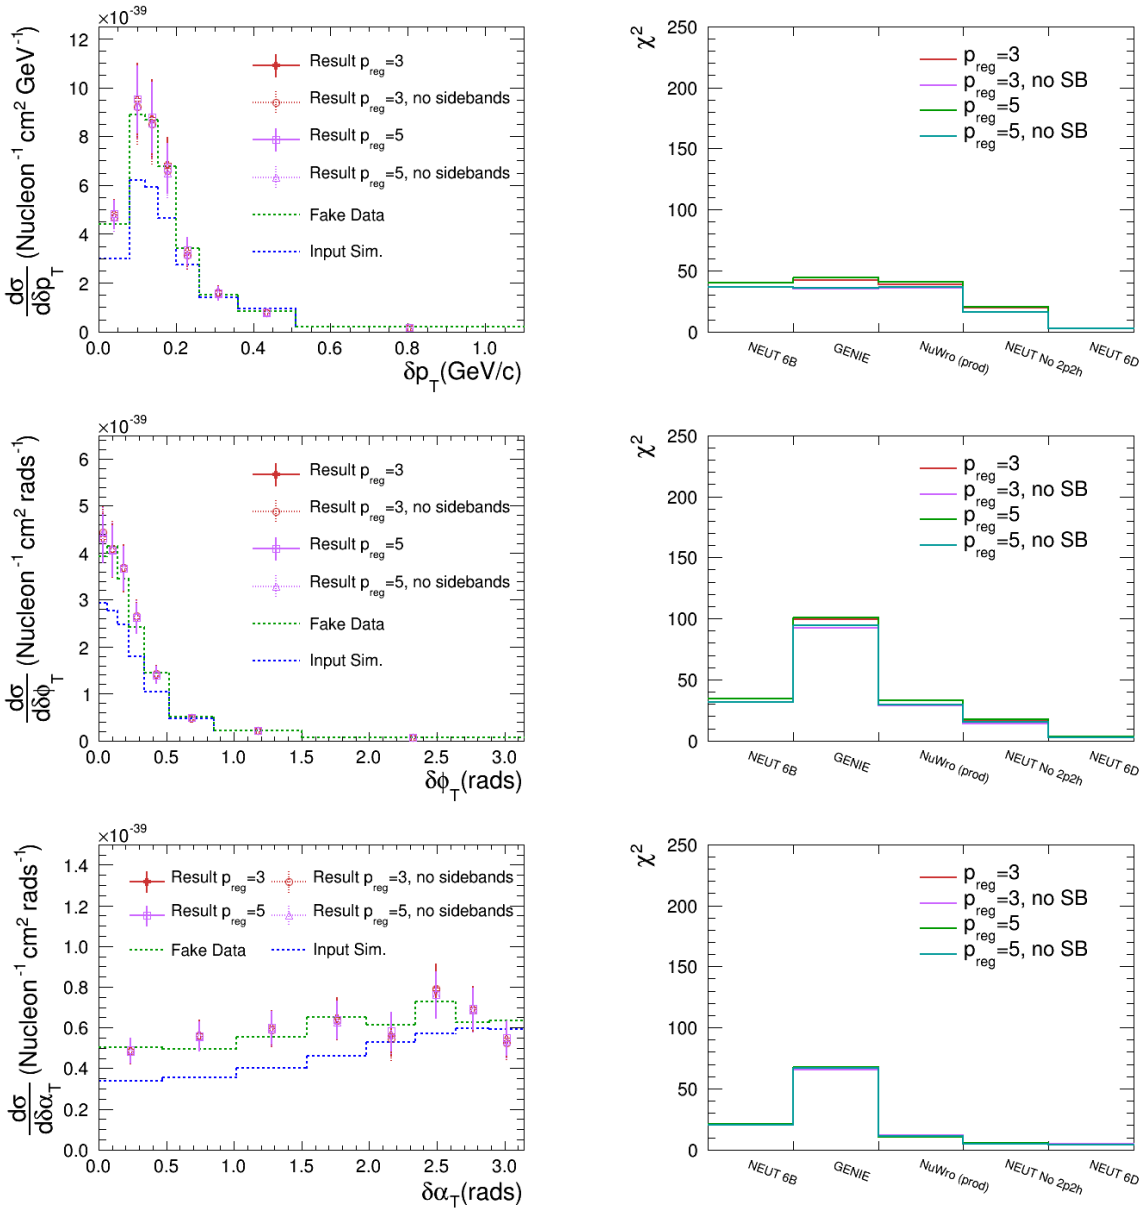


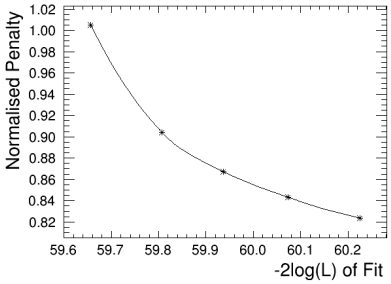
Figure 8.33: The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set.

8.2.5 NEUT 6D without 2p-2h fake data

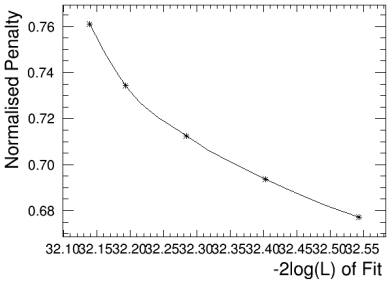
The NEUT 6D simulation used in section 8.2.4 is re-weighted to have its 2p-2h contribution removed. This simulation is then used to produce a fake data set scaled to real data POT. By assessing the analysis' ability to identify a 2p-2h contribution within the context of a particular model, it is possible to draw quantitative (but model dependent) conclusions regarding sensitivity to 2p-2h, whilst further validating the analysis framework.

The L-curve used to choose a regularisation strength of 3 for each fit is shown in figure 8.34. The IPS-OOPS regularisation strength was found to make very little difference to the result and is fixed at 1. The resultant χ^2 contribution of each term in the maximum likelihood fit (see equation 7.1) is shown in table 8.5.

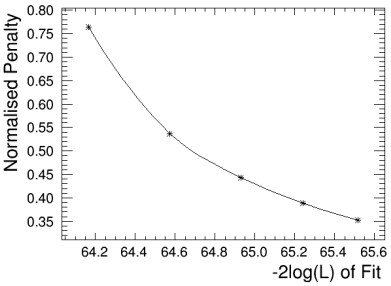
The extracted fit parameters and correlations are shown in figures 8.35, whilst comparisons of the pre and post-fit input simulation with the fake data are shown in reconstructed bins in figure 8.36. The cross-sections, alongside their correlation matrix, are shown in figure 8.37 and the χ^2 goodness of fit results are shown in figure 8.38.



(a) δp_T

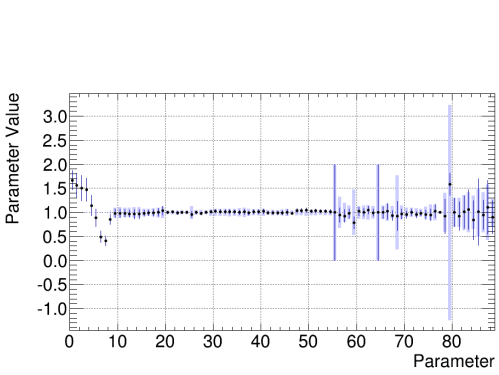


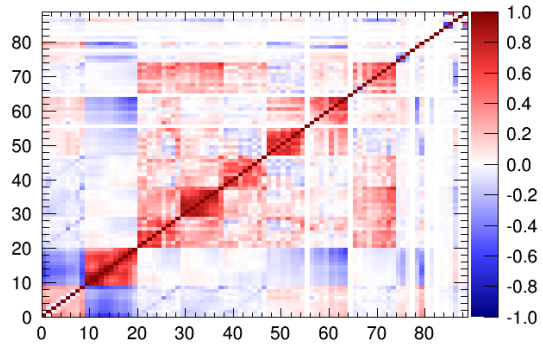
(b) $\delta \phi_T$



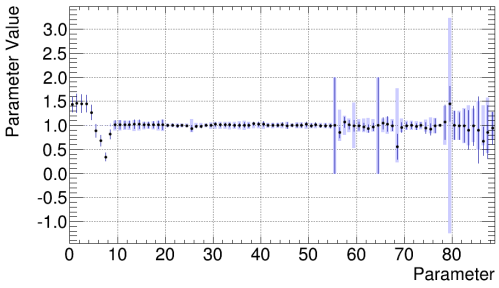
(c) $\delta \alpha_T$

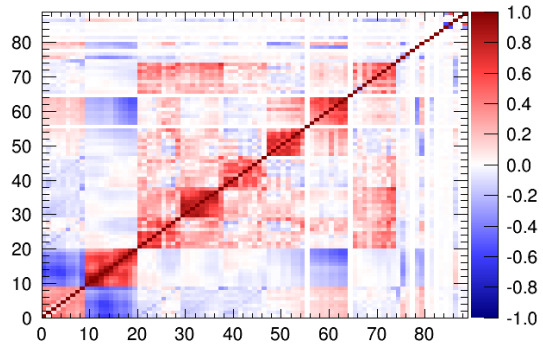
Figure 8.34: The L-curve used to determine the optimal regularisation parameters for each fit to NEUT (bug-fixed) fake data with the 2p-2h contribution removed. The regularisation strengths run in sequence down the curves (from left to right) and are the following: 1, 3, 5, 7 and 9. For this fake data study, a regularisation strength of 3 is chosen for all three STV.



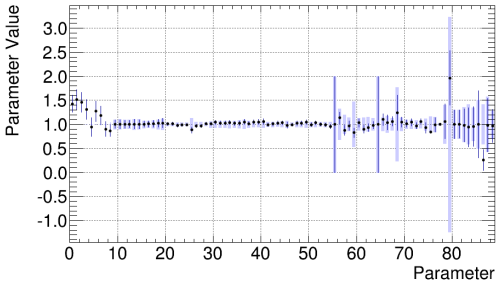


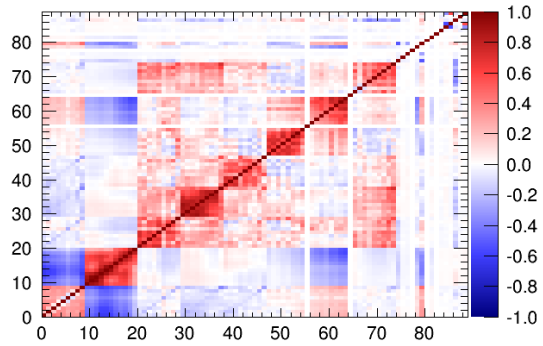
(a) δp_T





(b) $\delta \phi_T$





(c) $\delta \alpha_T$

Figure 8.35: The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to NEUT 6D (bug-fixed) fake data with the 2p-2h contribution removed. The data points in the left plot show the extracted parameters and their standard deviations alongside the light blue lines which show each parameters prior standard deviation. The parameters are labelled by the index in table 7.4.

Variable fit	χ^2_{stat}	χ^2_{syst}	χ^2_{IPSreg}	$\chi^2_{OOPSreg}$
δp_T	53.2	5.6	2.1	0.2
$\delta \phi_T$	25.5	4.2	1.8	0.2
$\delta \alpha_T$	49.5	14.0	1.1	0.2

Table 8.5: The size of each term in the maximum likelihood fit to NEUT fake data with the 2p-2h contribution removed.

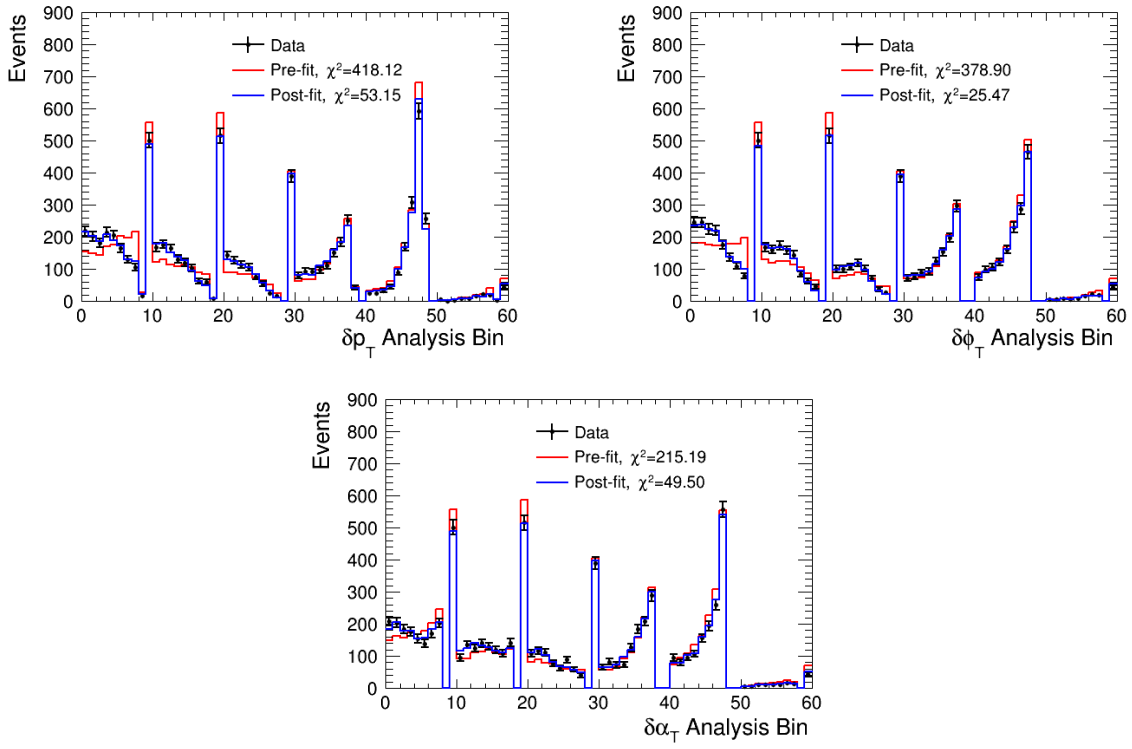


Figure 8.36: A comparison of the pre and post-fit input simulation in reconstructed bins from a fit in each of the STV (separately) to NEUT (bug-fixed) fake data with the 2p-2h contribution removed. The χ^2 goodness of fit is shown. The reconstructed bins for each variable are listed in tables 7.1 to 7.3. These are repeated for each sample in the following order: μTPC pTPC, μTPC pFGD, μFGD pTPC, CC1π sideband, CCOther sideband, μTPC multi-p. Note that the buffer bin (δp_T bin 8) is included for all three STV but is always empty for $\delta \phi_T$ and $\delta \alpha_T$ (i.e. the 9th bin in each sample is only relevant for δp_T).

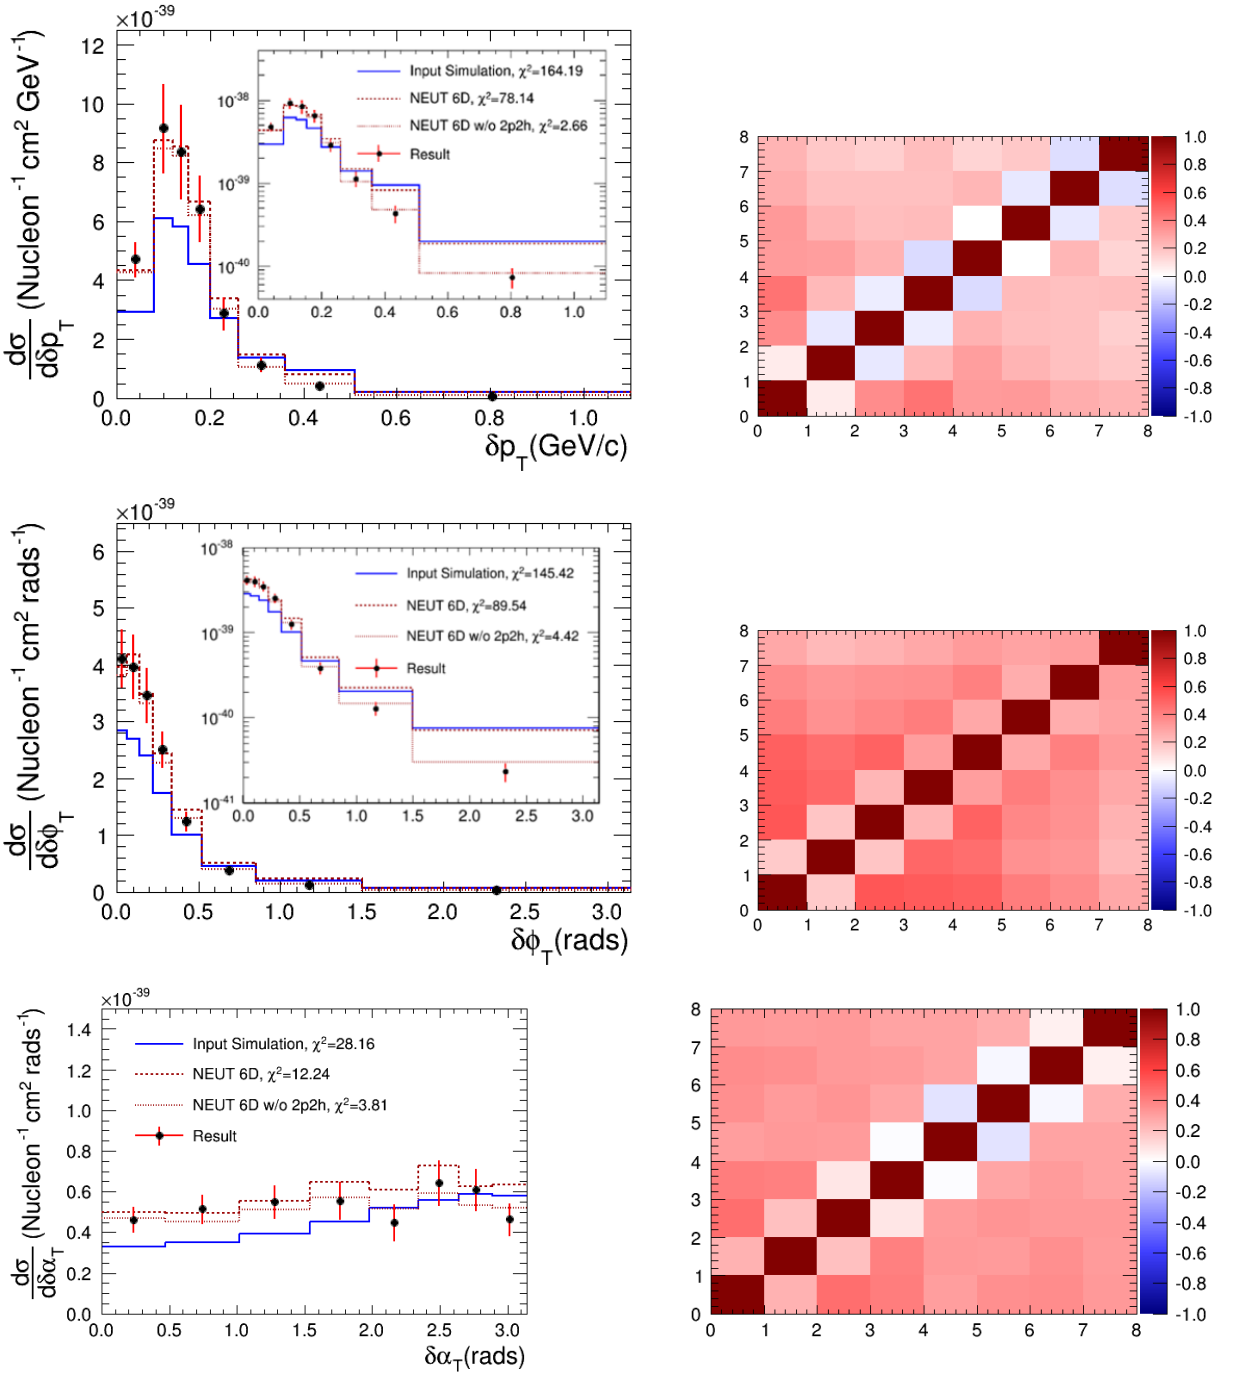
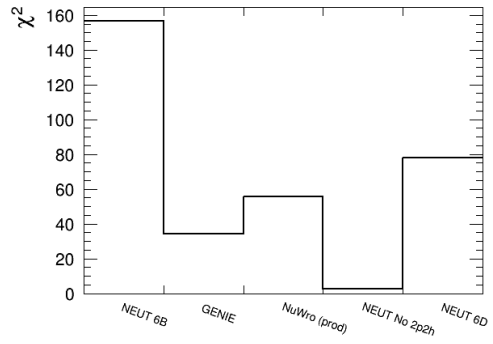
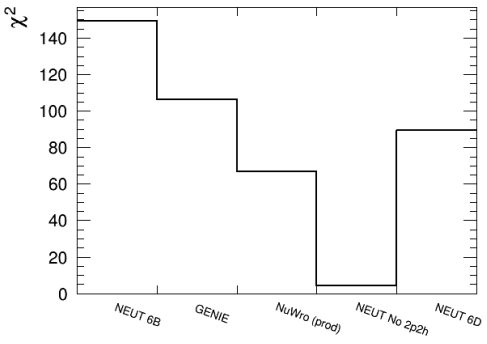


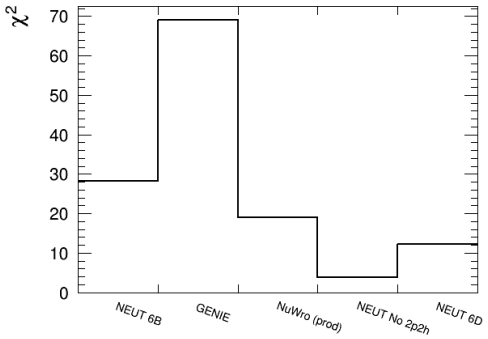
Figure 8.37: The differential cross section in the STV and the associated correlation matrix obtained from a fits to NEUT 6D fake data with the 2p-2h contribution removed. The same simulation before the removal of the 2p-2h contribution is shown alongside the fake data and the input simulation to allow an assessment of sensitivity to 2p-2h within the NEUT 6D simulation. The inlays show the same plot but on a logarithmic scale to highlight the difference in the tails in δp_T and $\delta\phi_T$. Although the inlays may make the χ^2 goodness of fit, shown in the legend, difficult to read, this infomation is also available in figure 8.38.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 8.38: The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to NEUT 6D fake data with the 2p-2h contribution removed (labelled 'NEUT No 2p2h'), to all test simulations, calculated as in equation 8.1. NEUT 6B is the input simulation (with the nucleon FSI bug). The simulation predictions are extracted from their ND280 productions.

2p2h-less fake data discussion The comparison of the extracted result to both the fake data truth and the fake data prior to the 2p-2h removal (figure 8.37), alongside the goodness of fit to the various fake data in figure 8.38, demonstrates that - within the context of the NEUT 6D simulation - the fit is able to clearly distinguish the lack of a 2p-2h contribution using δp_T or $\delta\phi_T$. Similarly, the previous section (8.2.4) demonstrates that the fit can also do the opposite (distinguish the NEUT 6D distribution with a 2p-2h contribution from the same without one). Overall this shows that if other nuclear effects (i.e. the impact of nucleon FSI and the nucleon initial state momentum distribution) are reasonably well understood, this analysis could be able to accurately characterise a 2p-2h contribution.

It's important to note that, as discussed in section 6.4.1, the phase-space constraints on the analysis remove 43% of the 2p-2h contribution from the signal in the input simulation. Therefore improved reconstruction of low momentum protons may allow access to even more of the 2p-2h contribution.

Effect of regularisation and sidebands on the 2p2h-less NEUT 6D fake data fit The bias from regularisation and sidebands are assessed as described in section 8.2.1.

Figure 8.39 compares the differential cross section and χ^2 goodness of fit relative to all fake data whilst figure 8.40 shows the same with the most favoured regularisation strength with and without sidebands.

It can be seen that both the regularisation strength and inclusion of sidebands makes very little difference to the nominal fit result or its sensitivity. Overall there is no suggestion of a bias from regularisation or the inclusion of sidebands.

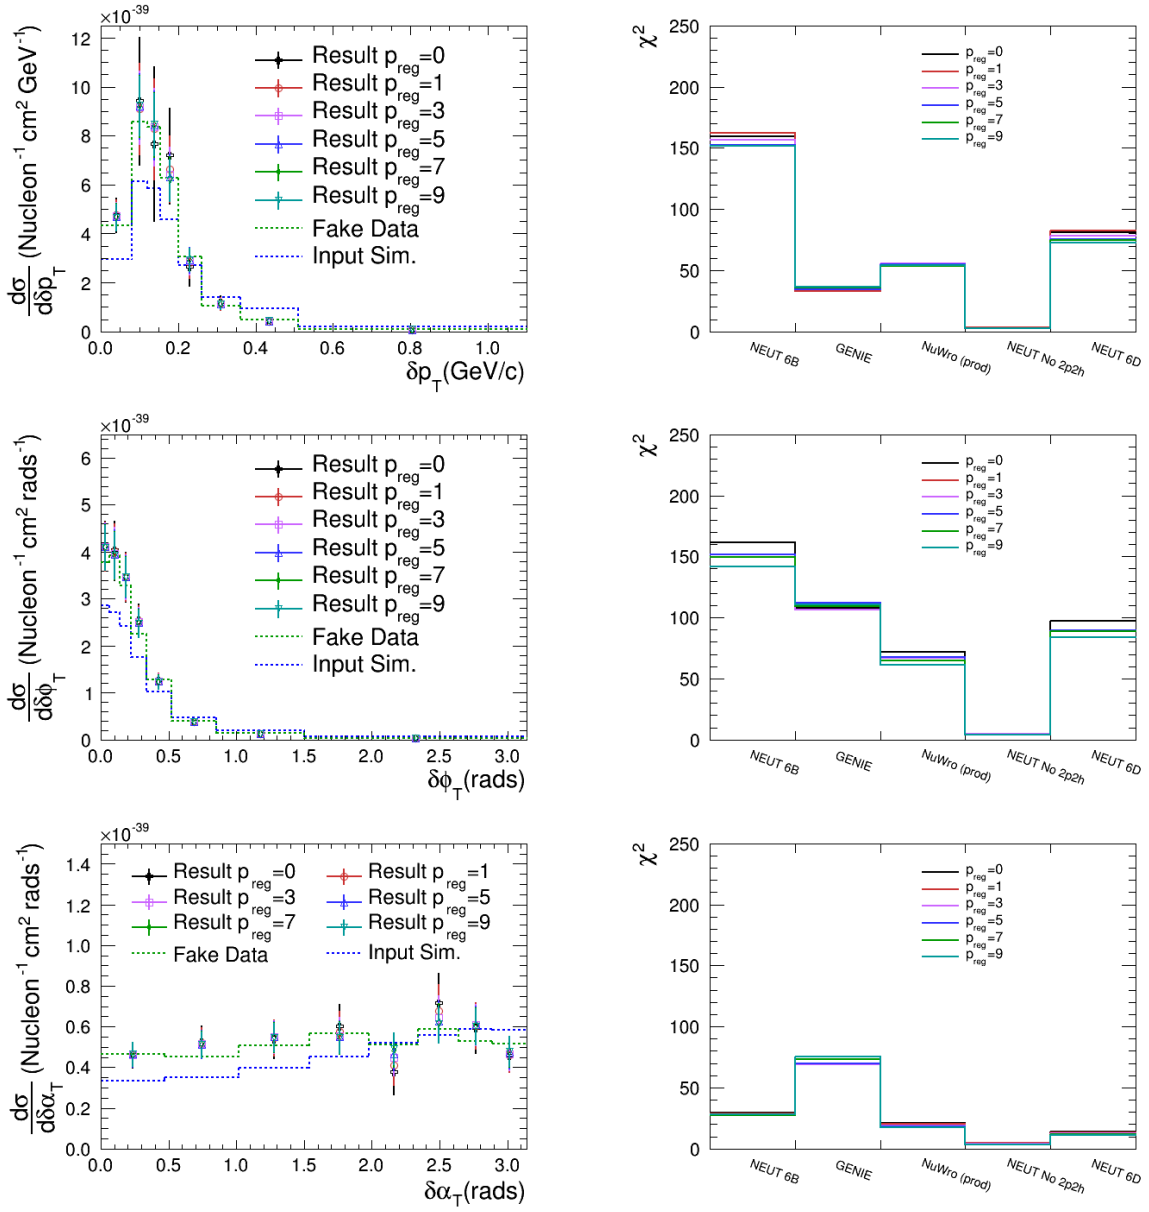


Figure 8.39: The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data with its 2p-2h contribution removed with various regularisation strengths, alongside the χ^2 of the results relative to each fake data set.

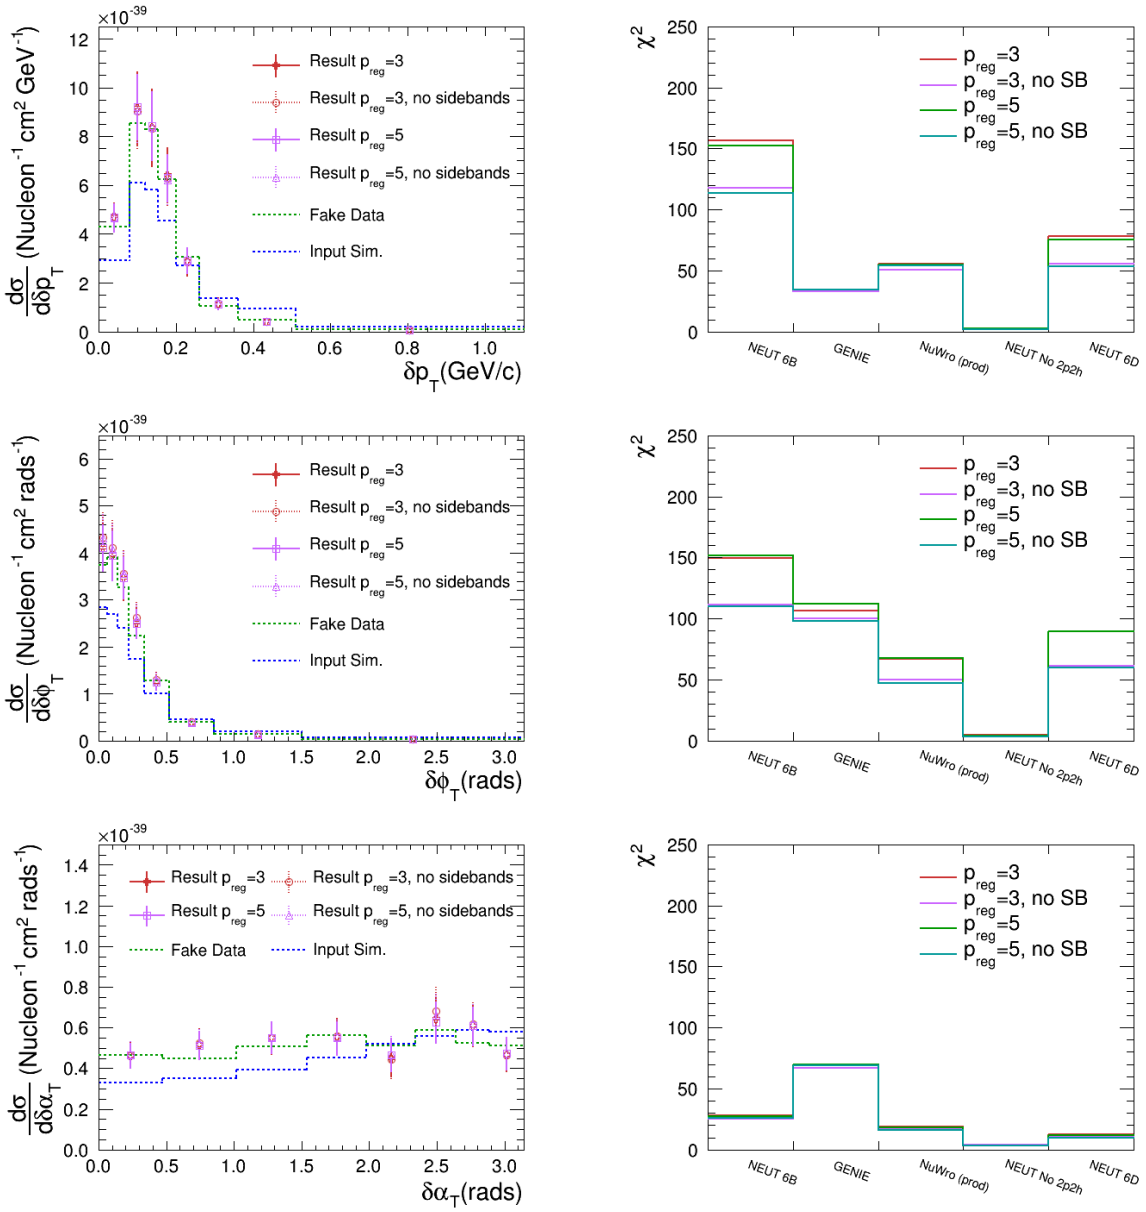


Figure 8.40: The differential cross section in the three STV obtained from a fit to NEUT 6D (bug-fixed) fake data with its 2p-2h contribution removed and using regularisation strengths deemed reasonable by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set.

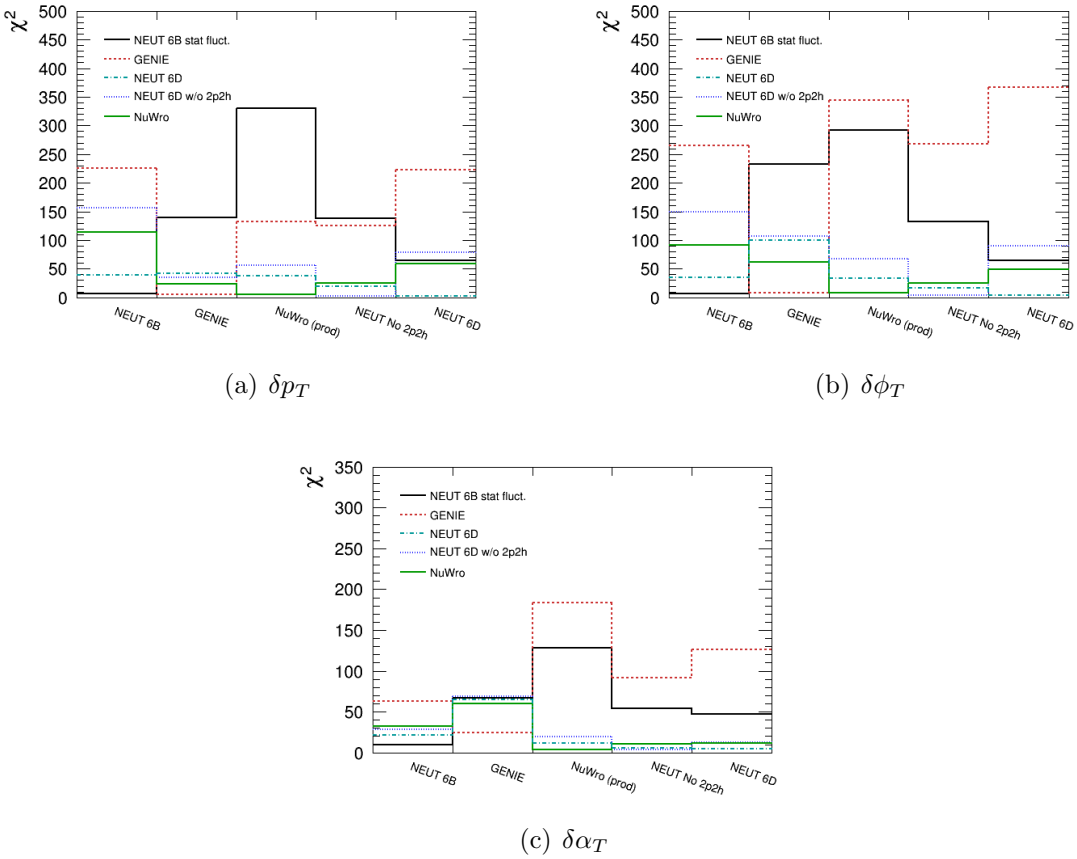
8.2.6 Fake data studies summary

Section 8.2 has demonstrated both that the cross-section extraction method is robust and that the analysis offers interesting model separation potential. To summarise the fake data studies, figure 8.41 shows the χ^2 goodness of fit of the result for a particular fake data set relative to all fake data sets. It can be seen that the χ^2 is always minimal and small for the fake data truth, thereby indicating that there is little bias. In addition, it can also be seen that it is almost always possible to distinguish the fake data truth from all other fake data sets, therefore demonstrating the analysis' interesting model separation potential. In particular figure 8.41 shows that the analysis can clearly distinguish the NEUT, GENIE and NuWro predictions in all three observables whilst also offering clear separation between NEUT (6D) with and without a 2p-2h contribution in δp_T and $\delta\phi_T$.

Throughout all of the fake data studies, the only sign of substantial bias in the fake data studies was in a single bin of $\delta\alpha_T$ when fitting GENIE fake data. The source of this bias is well-understood and the planned method of identifying it, through a data to post-fit simulation comparison at the reconstructed level, was shown to be effective.

Furthermore it has been shown that the analysis is not biased by relying on the input simulation model to describe the sidebands in any of the fake data tests (indeed it was found to be essential to have the sidebands to avoid bias in the case of the NuWro fake data). It has also been shown that regularisation successfully smooths the result whilst having an almost negligible impact on the best fit likelihood and the results sensitivity, therefore showing no subsequent bias to the input simulation.

That the fit can reliably extract NEUT 6D fake data (with the FSI bug fixed) from a NEUT 6B input (with the FSI bug), indicates that the bug in the input simulation is not an issue. This is expected since, as discussed in section 7.2.6, so long as the smearing is correct and the background models are plausible, the results dependence on the input simulation should be small. The primary way in which model bias can enter the cross-section extraction is through the regularisation and it has been clearly shown that this is not happening to an appreciable extent.



(a) δp_T

(b) $\delta \phi_T$

(c) $\delta \alpha_T$

Figure 8.41: The χ^2 goodness of fit, calculated as in equation 8.1, of the extracted differential cross-section from a fit in each of the STV (separately) to each fake data set (which one is indicated by the legend) relative to every fake data set (all fits using a regularisation strength of 3).

Chapter 9

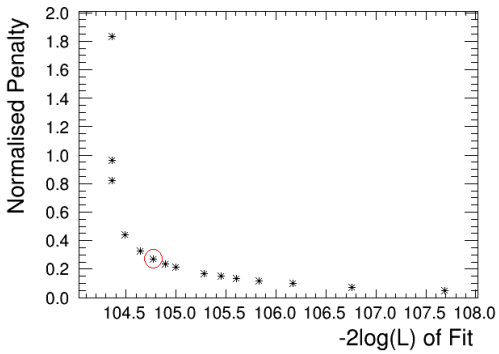
Cross-section results and interpretations

Chapter 8 validates that the cross-section extraction framework is fully functional and shows minimal bias across a variety of fake data studies. To begin this chapter, section 9.1 exercises the same procedure as has been applied to fake data now to real data to extract the principle result from this analysis and suggest preliminary interpretations. Following this, section 9.2 provides more concrete physics conclusions from an extensive set of data-simulation cross-section comparisons.

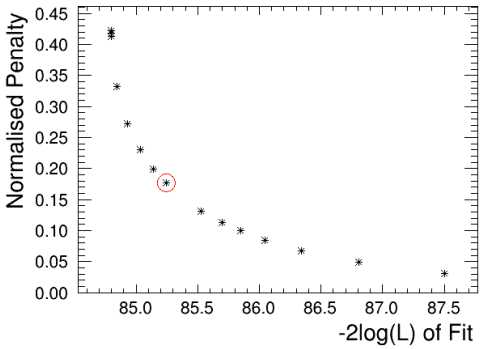
9.1 Real data results

As with the fake data studies, firstly an L-curve is produced, as described in section 7.1.2, to find the best choice of regularisation parameters. The resultant L-curves for the IPS regularisation parameter are shown in figure 9.1 to justify a regularisation strength of 3 for δp_T and 5 for $\delta\phi_T$ and $\delta\alpha_T$, whilst it was found that the OOPS regularisation parameter made only a negligible difference to the fit result and so is left at 1. Detailed results without regularisation are also presented in appendix B.

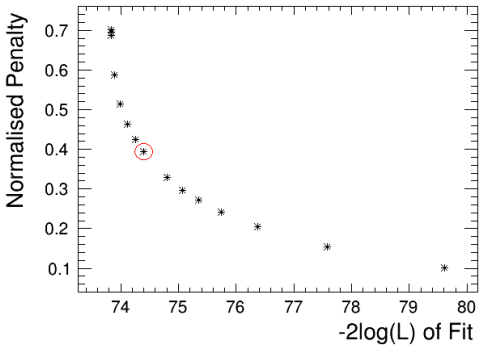
The approximated $-2\log(L)(\approx \chi^2)$ of each term in the maximum likelihood fit (see equation 7.1) is shown in table 9.1. The extracted fit parameters and correlations with the chosen regularisation strengthens are shown in figure 9.2, whilst comparisons of the pre and post-fit simulation with data are shown in reconstructed bins in figure 9.3. The final cross-section results from real ND280 data, alongside their correlation matrix, are shown in figure 9.4 and the χ^2 goodness of fit to the simulations used in the fake data studies are shown in figure 9.5.



(a) δp_T

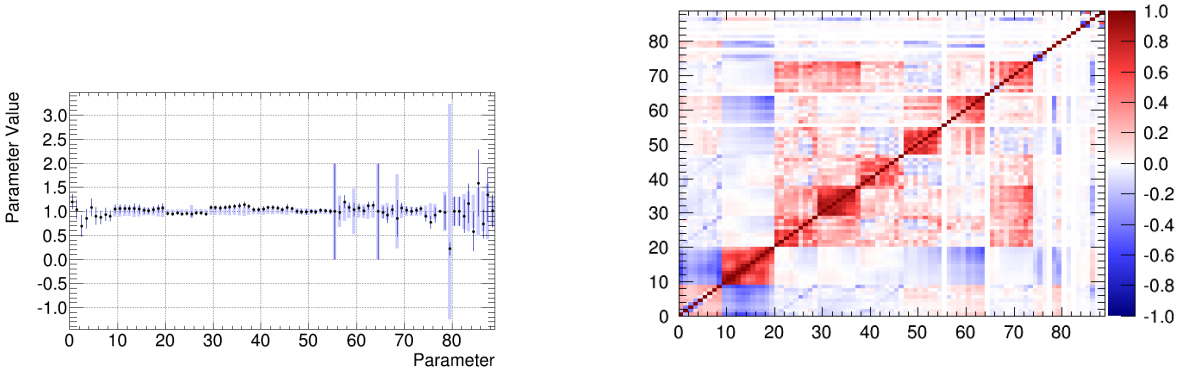


(b) $\delta \phi_T$

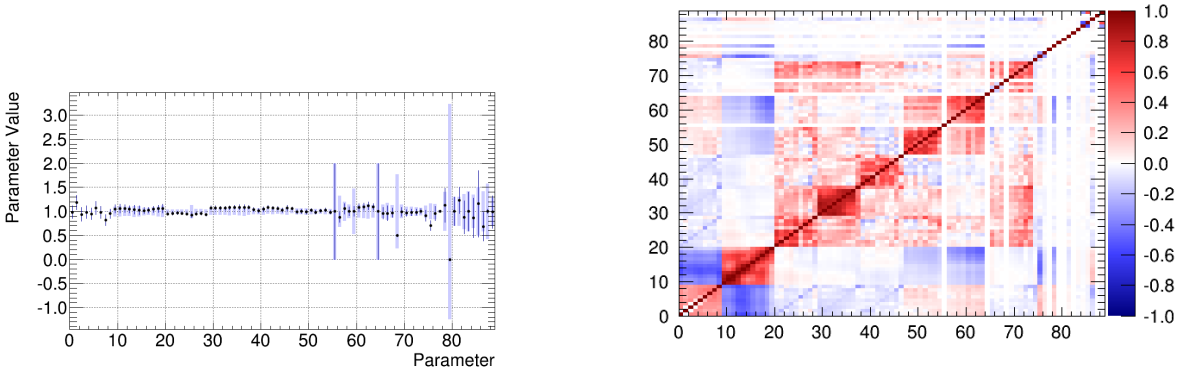


(c) $\delta \alpha_T$

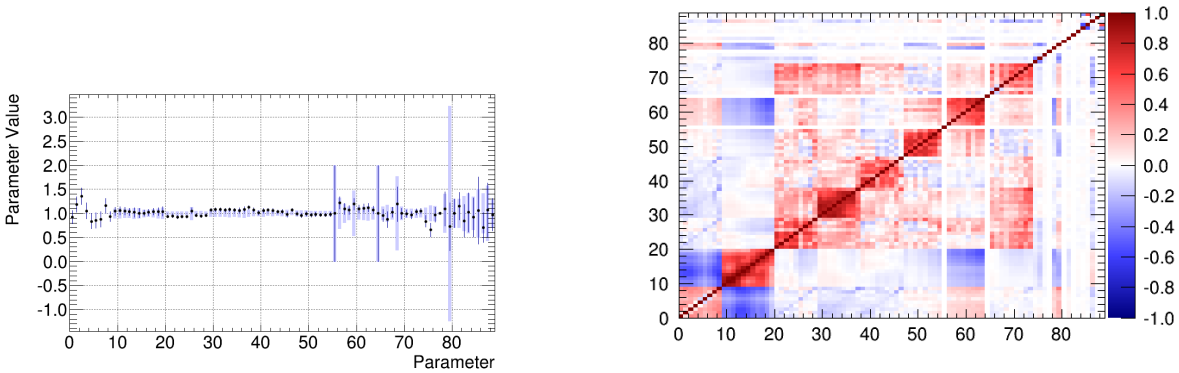
Figure 9.1: The L-curves, described in section 7.1.2, obtained from many fits to the real data. The regularisation strengths run in sequence down the curves (from left to right) and are the following: 0.01, 0.05, 0.1, 1, 2, 3, 4, 5, 8, 10, 12, 15, 20, 30 and 50. The chosen regularisation strengths are circled in red: 3 for δp_T and 5 for $\delta \phi_T$ and $\delta \alpha_T$.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 9.2: The parameters and associated correlation matrix extracted from a fit in bins of all three STV (separately) to real data. The data points in the left plot show the extracted parameters and their standard deviations alongside the light blue lines which show each parameters prior standard deviation. The parameters are labelled by the index in table 7.4.

Variable fit	χ^2_{stat}	χ^2_{syst}	χ^2_{IPSreg}	$\chi^2_{OOPSreg}$
δp_T	81.2	22.6	0.8	0.0
$\delta\phi_T$	59.7	24.0	0.9	0.0
$\delta\alpha_T$	51.4	22.7	1.9	0.0

Table 9.1: The size of each term in the maximum likelihood fit to real data.

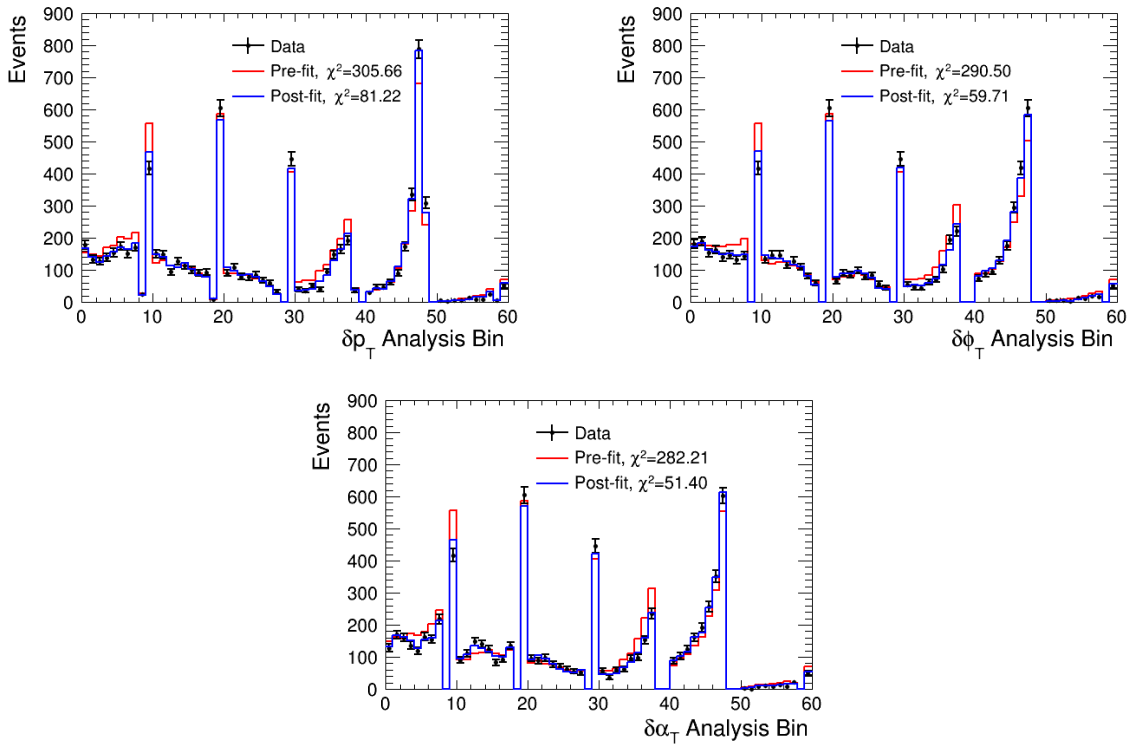


Figure 9.3: A comparison of the pre and post-fit simulation in reconstructed bins from a fit in each of the STV (separately) to real data. The χ^2 goodness of fit is shown. The reconstructed bins for each variable are listed in tables 7.1 to 7.3. These are repeated for each sample in the following order: μTPC pTPC, μTPC pFGD, μFGD pTPC, CC1π sideband, CCOther sideband, μTPC multi-p. Note that the buffer bin (δp_T bin 8) is included for all three STV but is always empty for $\delta\phi_T$ and $\delta\alpha_T$ (i.e. the 9th bin in each sample is only relevant for δp_T).

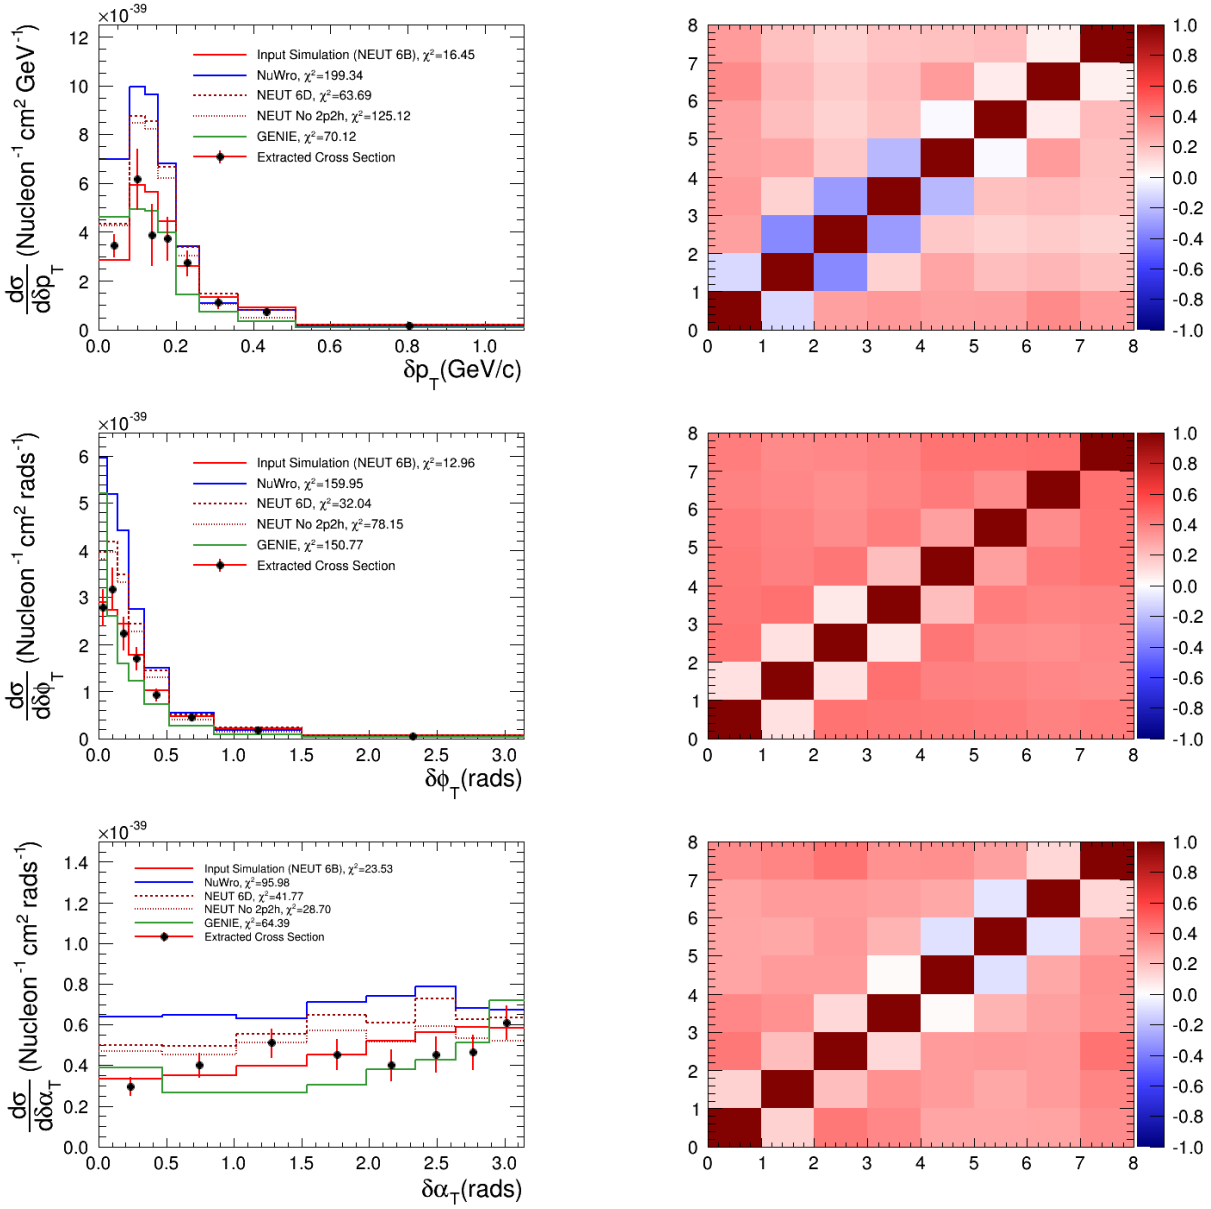
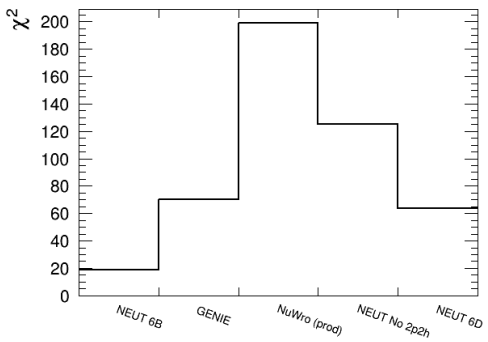
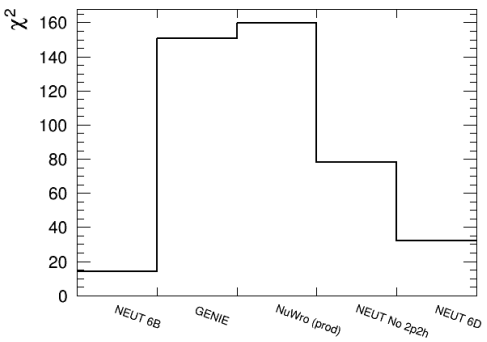


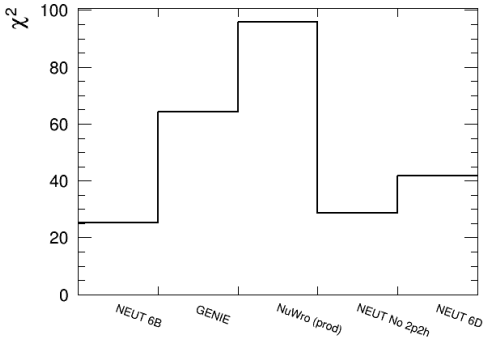
Figure 9.4: The differential cross section in the three STV and the associated correlation matrix obtained from (separate) fits to real data. The extracted cross-section is the primary result of this thesis. NEUT 6B in the input simulation. The simulation predictions are extracted from their ND280 productions. 'NEUT 6D' is the bug-fixed NEUT and 'NEUT No 2p2h' is NEUT 6D without any 2p-2h events. All the simulations are discussed in section 4.2.3 and further detailed in appendix A.



(a) δp_T



(b) $\delta \phi_T$



(c) $\delta \alpha_T$

Figure 9.5: The χ^2 goodness of fits, of the extracted differential cross-section from a fit in each of the STV (separately) to real data, to all test simulations, calculated as in equation 8.1. The NEUT 6B (the input simulation), GENIE and NuWro predictions are taken from their ND280 productions. 'NEUT 6D' is the bug-fixed NEUT and 'NEUT No 2p-2h' is NEUT 6D without any 2p-2h events. All the simulations are discussed in section 4.2.3 and further detailed in appendix A.

Real data discussion Figure 9.2 demonstrates the fit parameters shared between fits (the flux and model parameters) are consistent.

The fit prefers an inflation of the flux by ∼5-8%, similar to what is found by the T2K ND280 fits for oscillation analyses [142], but due to the strong correlations in the flux pre-fit covariance matrix, this is quite compatible with the pre-fit flux predictions. In addition, the flux is - as expected - quite anti-correlated with the template weights and impactful model parameters.

The fit parameters also prefer a large down-scaling of the CC_{COH} contribution (compatible with the Minerva results [144]) and a reduction of the Rein-Seghal cross-section parameters (M_A^{RES} and C_A^5) (also compatible with the oscillation analysis ND280 fit results [142]).

Figure 9.3 generally demonstrates a good fit to the real data at the reconstructed level. There are no stand out IPS signal bins that are not well matched, thereby indicating reliable cross-section results. The overall χ^2 in δp_T is fairly large but it can be seen this is largely not from the IPS signal bins.

The cross-section results themselves (shown in figure 9.4) and the corresponding generator comparisons (figure 9.5) are interesting but difficult to interpret from these plots alone. It should be noted that, although there is an apparent peak in the second bin of $\delta\phi_T$, the result is entirely consistent with a monotonically decreasing cross section (due to the strong anti-correlations between the bins). The results primarily indicate a good agreement with the input NEUT 6B simulation (that contains the proton FSI bug). That the real data results agree with an input simulation which is known to contain a significant bug could be taken to suggest bias in the cross-section extraction method, however extensive fake data studies - shown in chapter 8 - should clearly demonstrate that this is not the case. Furthermore figure 9.3 shows that this good agreement is also present at the reconstructed level. This may suggest stronger nucleon FSI (the effective result of the bug in the input simulation) but will be discussed in more detail in section 9.2.

It can be seen that the GENIE 2.8.0 prediction is in fairly good agreement with the total cross section, but the relatively large prediction at low δp_T and $\delta\phi_T$ is in poor agreement with the shape. This particular curiosity in the GENIE prediction is discussed further in section 9.2.

In general the shape of the data is compatible with the predictions of the other generators tested but not the normalisation. This suggests that the generators' total cross-section predictions in the restricted phase space of this analysis is likely not correct. Since most of the simulations considered so far represent quite different arrays of models, it is difficult to

draw further conclusions. Section 9.2 will compare the results to a wider set of simulations to allow further investigation.

Overall figure 9.4 contains the primary result of this analysis: the first cross-section measurement of single transverse kinematic imbalance in neutrino scattering.

Effect of regularisation and sidebands on the real data result As previously discussed in the previous chapters (but thus far found not to be an issue) it is possible that regularisation or the use of sidebands can bias the result toward the input simulation. To further assess the extent to which this happens, the real data is fit with and without the sidebands and with various strengths of regularisation.

Figure 9.6 compares the differential cross section and χ^2 goodness of fit relative to all fake data whilst figure 9.7 shows the same with the most favoured regularisation strength with and without sidebands.

As in the fake data studies, table 9.1 demonstrates that χ^2_{stat} is substantially larger than any other term in the fit, thereby showing that the dominant factor in finding the best fit point is the data. This is good evidence that the fit result is not biased by the systematic parameter priors or the regularisation. The χ^2_{reg} terms are extremely small relative to the total χ^2, showing the regularisation is doing very little to influence the best fit point. It is worth noting that the χ^2_{syst} term in the fit is slightly larger than seen in fake data studies, but this is not so surprising as it is certainly expected that the input models of the flux, detector and backgrounds are better representative of other simulations than of real data in previously unmeasured variables. The source of this larger χ^2_{syst} term is the slightly increased contention between the fit parameters and their pre-fit priors, as shown by figure 9.2, which has already been discussed.

For the no regularisation case in δp_T we see minimal effect on the χ^2 goodness of fit and in the best fit likelihood (see table 9.1) but a substantial effect on the nominal result. This zig-zag shape of the unregularised result is exactly what is expected when the fit encounters an ill posed problem and was seen in the GENIE and NEUT 6D fake data tests (in section 8.2.2 and 8.2.4). As was seen in the fake data, the regularisation is effectively selecting a smooth result out of many degenerate possibilities.

Overall it can be seen that, with the exception of the no regularisation case in δp_T (discussed above), both the regularisation strength and inclusion of sidebands makes very little difference to the nominal fit result. Additionally, there is negligible regularisation contribution to the best fit point likelihood (see table 9.1) and only a small change in model separation power with regularisation strength. This shows that the fits display very little potential bias to regularisation.

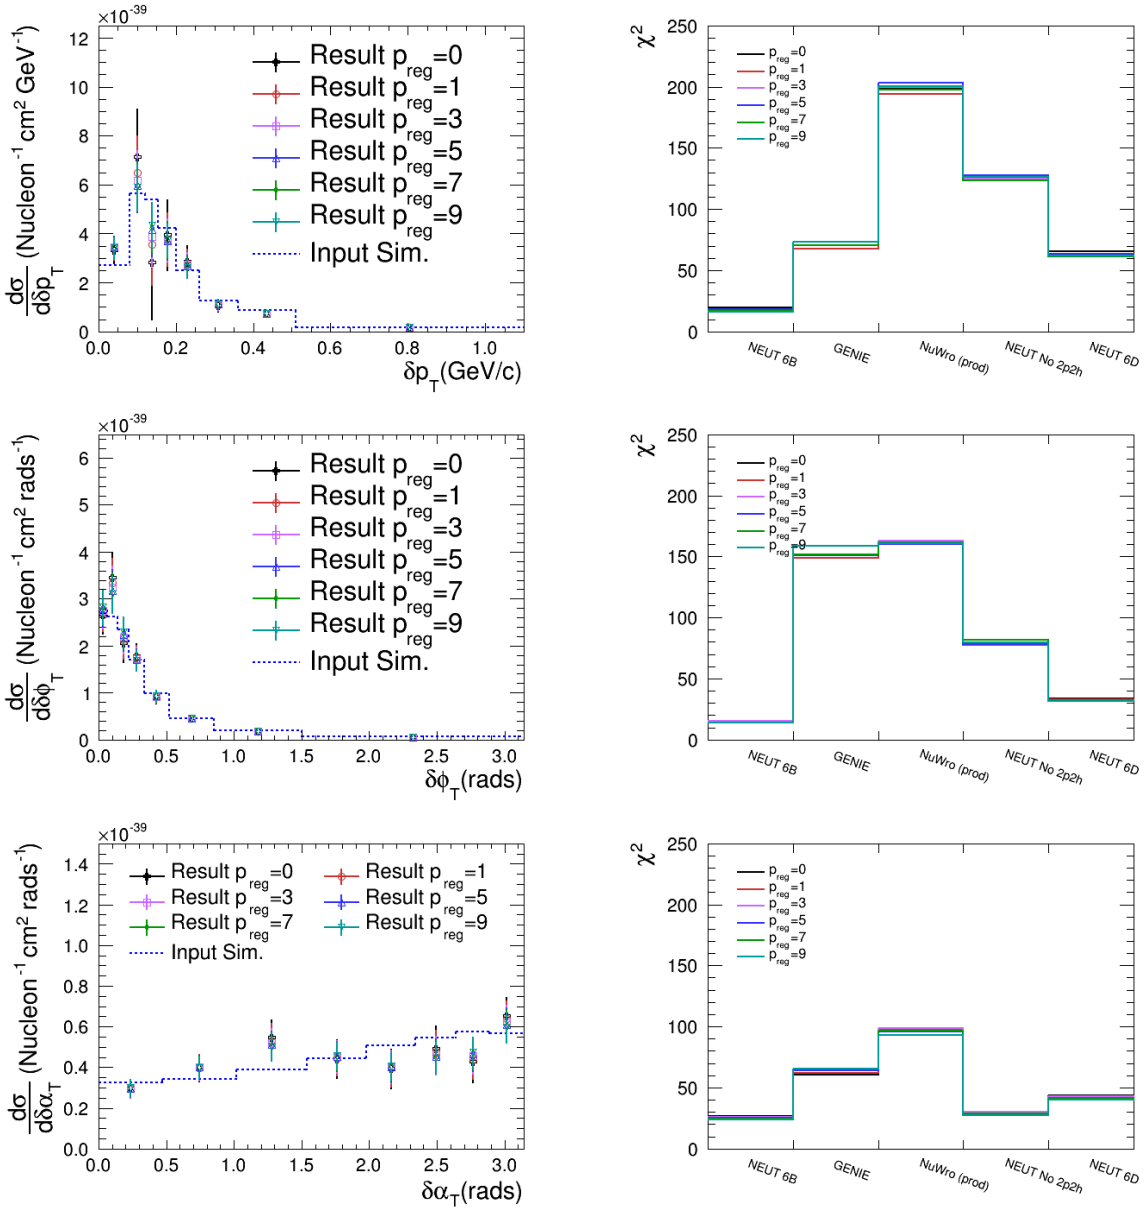


Figure 9.6: The differential cross section in the three STV obtained from a fit to real data with various regularisation strengths, alongside the χ^2 of the results relative to each fake data set.

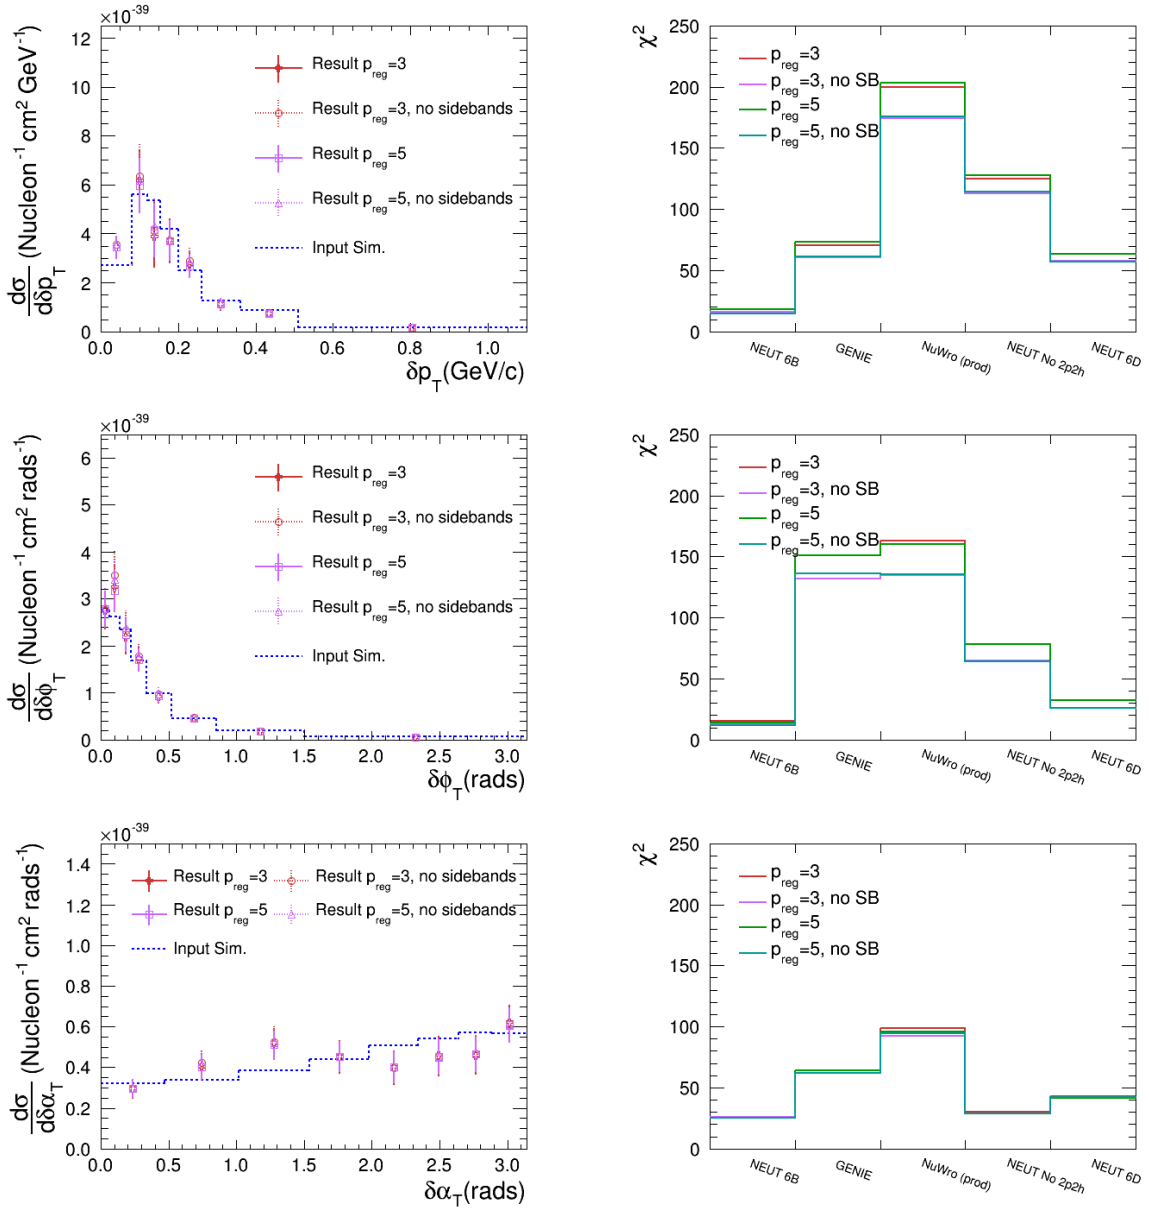


Figure 9.7: The differential cross section in the three STV obtained from a fit to real data with the regularisation strengths deemed plausible by the L-curve method with and without the sideband samples, alongside the χ^2 of the results relative to each fake data set.

9.2 Simulation comparisons and further interpretations

When the single transverse variables were introduced in chapter 3, it was shown that they had the potential to offer a novel characterisation of nuclear-effects and significant model separation power. The primary result of this analysis is a differential cross section measured in each of the STV, as presented in section 9.1, which offers a largely model independent measurement that can now be used to offer such a characterisation. One method of achieving this would be to fit cross section and nuclear model parameters in neutrino-nucleus interaction simulations to the extracted cross sections. However, currently available simulations do not yet offer offer enough freedom to allow for plausible variations in hadronic kinematics, but even if the results cannot be fit, interesting model comparisons can still be made.

This section first compares the results to a wide range of existing models. Although some simulation comparisons have already been offered in section 9.1, these were limited to the generator predictions that have been run through the ND280 detector simulation. This section uses the latest predictions from the current state-of-the-art models from the NEUT 5.3.2.2, GENIE 2.12.4, NuWro 11q and GiBUU 2016 neutrino interaction simulations in a variety of configurations (the available configurations are summarised in appendix A). By comparing the analysis results to variations of models within a simulation and also to the differences in models between simulations, it is then possible to draw some qualitative physics conclusions from the analysis results. Comparisons are made to both the full cross-section results shown in section 9.1, and also to area-normalised (shape only) results. The covariances are numerically propagated through the area normalisation using the cross-section correlation matrices shown in figure 9.4. To facilitate the easy comparison of generators with entirely different output formats, the NUISANCE software package has been used [149].

Generator nominal predictions

This section shows comparisons of the data to the predictions of the generators ran in both their default configurations and configurations closely resembling those used as the input simulation for T2K and NOνA oscillation analyses. These are shown in figures 9.8 to 9.14, each comparison is split by interaction mode to allow conclusions to be drawn regarding the data's agreement with each contribution in each model. A χ^2 goodness of fit is calculated using equation 8.1 and the full cross-section or shape-only covariance matrix. Figure 9.15 shows a comparison to the data from most of the generators in this section together.

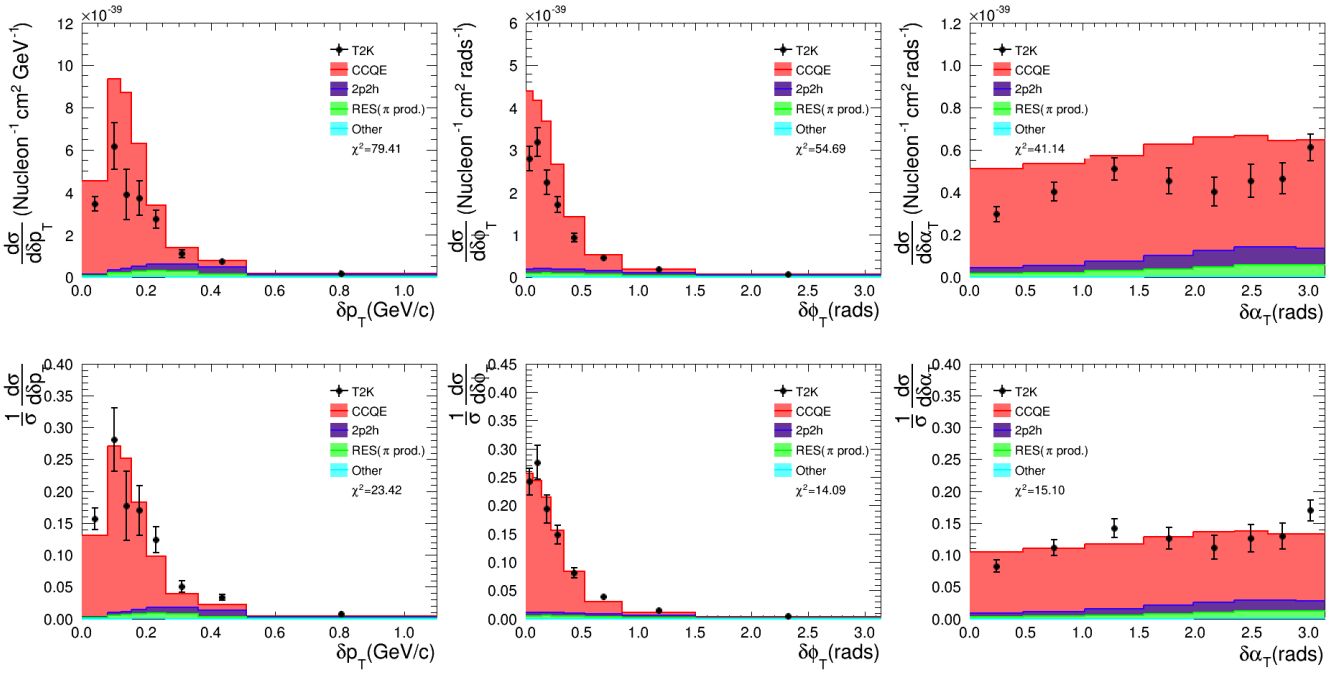


Figure 9.8: A comparison of analysis results to predictions from NEUT 5.3.2.2 using $M_A^{QE} = 1.03$ and an SF nuclear model, split by interaction mode.

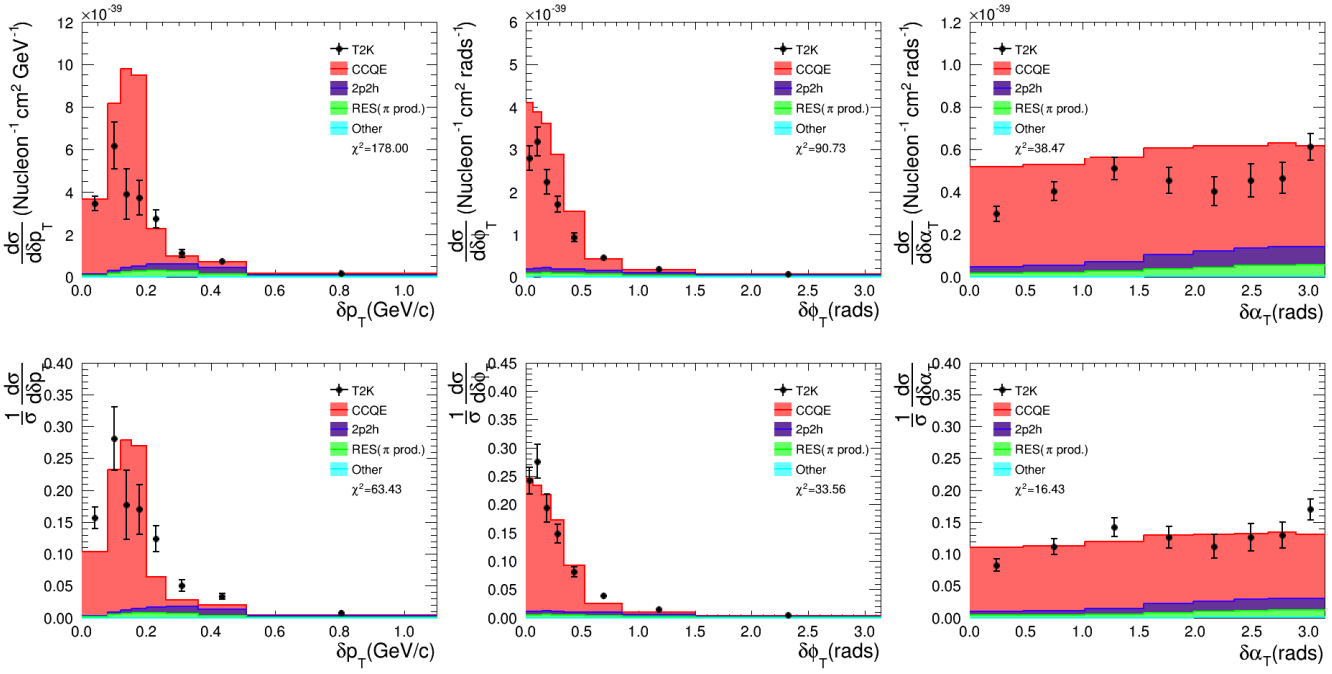


Figure 9.9: A comparison of analysis results to predictions from NEUT 5.3.2.2 using $M_A^{QE} = 1.03$ and an RFG+RPA nuclear model, split by interaction mode. This is similar to the starting point for T2K oscillation analyses.

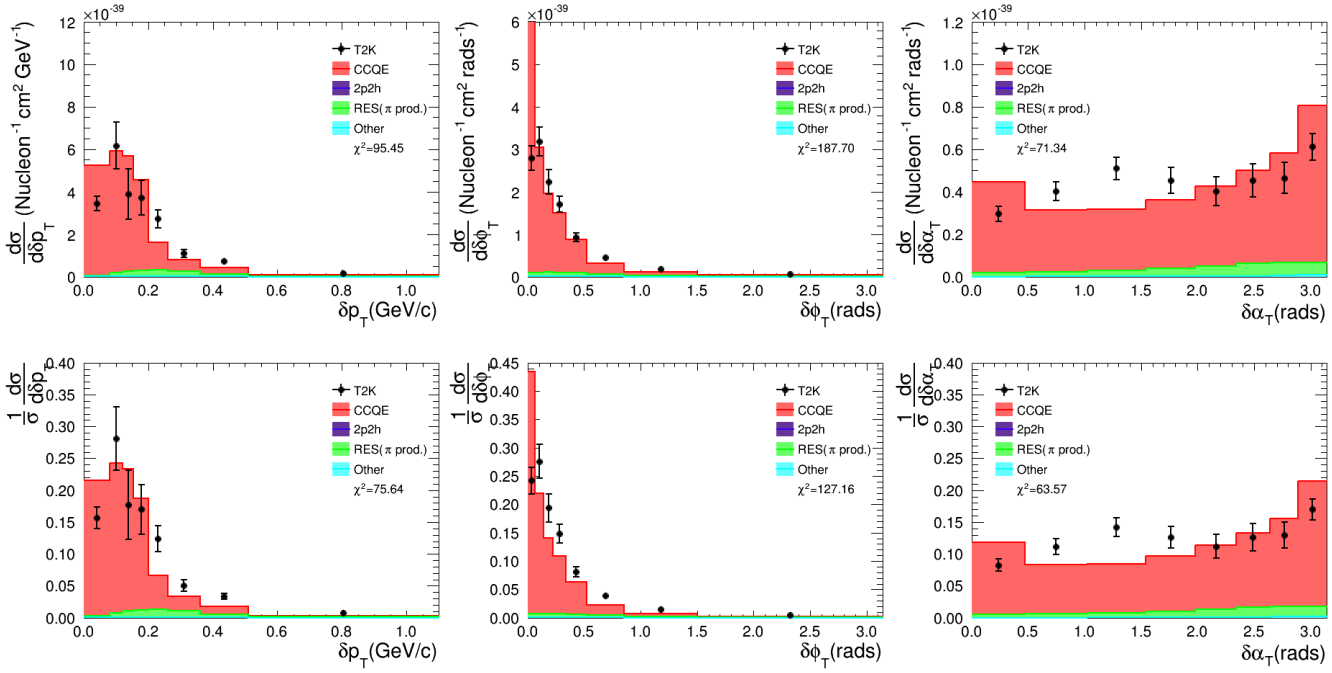


Figure 9.10: A comparison of analysis results to predictions from GENIE 2.12.4 using $M_A^{QE} = 0.99$ and an RFG nuclear model, split by interaction mode (note that the default GENIE implementation contains no 2p-2h).

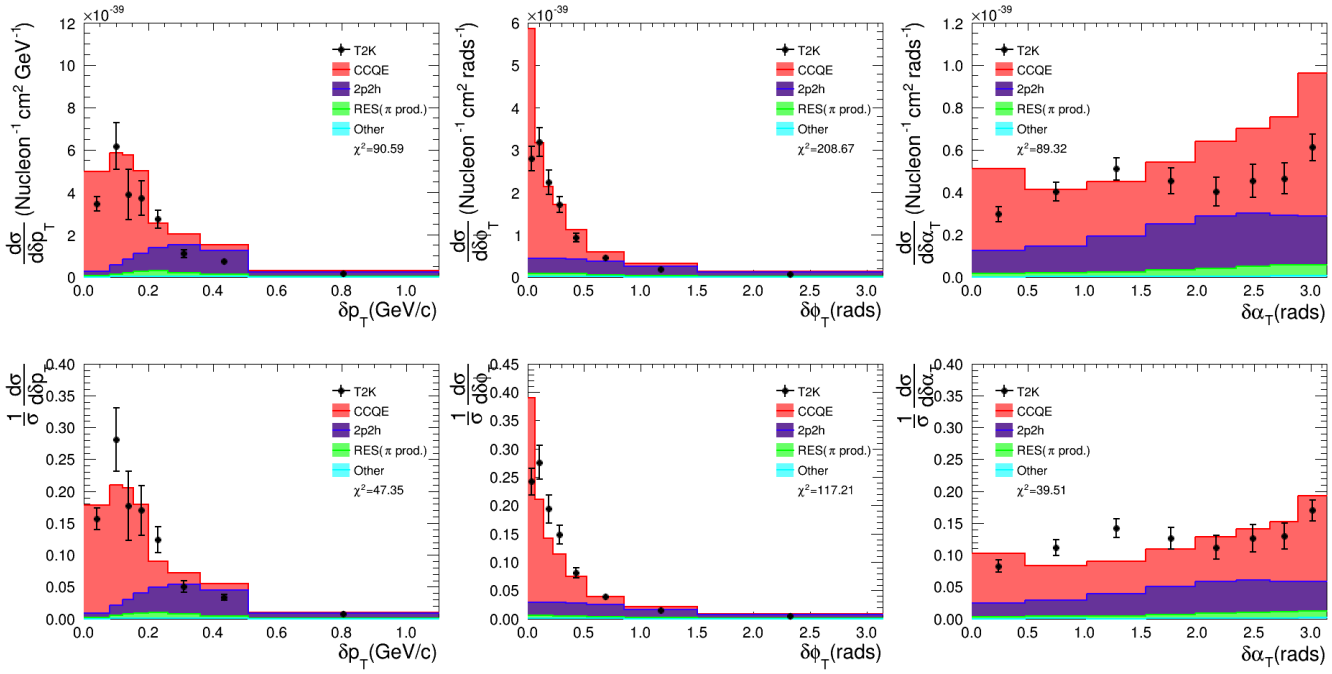


Figure 9.11: A comparison of analysis results to predictions from GENIE 2.12.4 using $M_A^{QE} = 0.99$ and an RFG nuclear model with the GENIE empirical 2p-2h model enabled, split by interaction mode [41]. This is similar to the starting point for NOνA oscillation analyses.

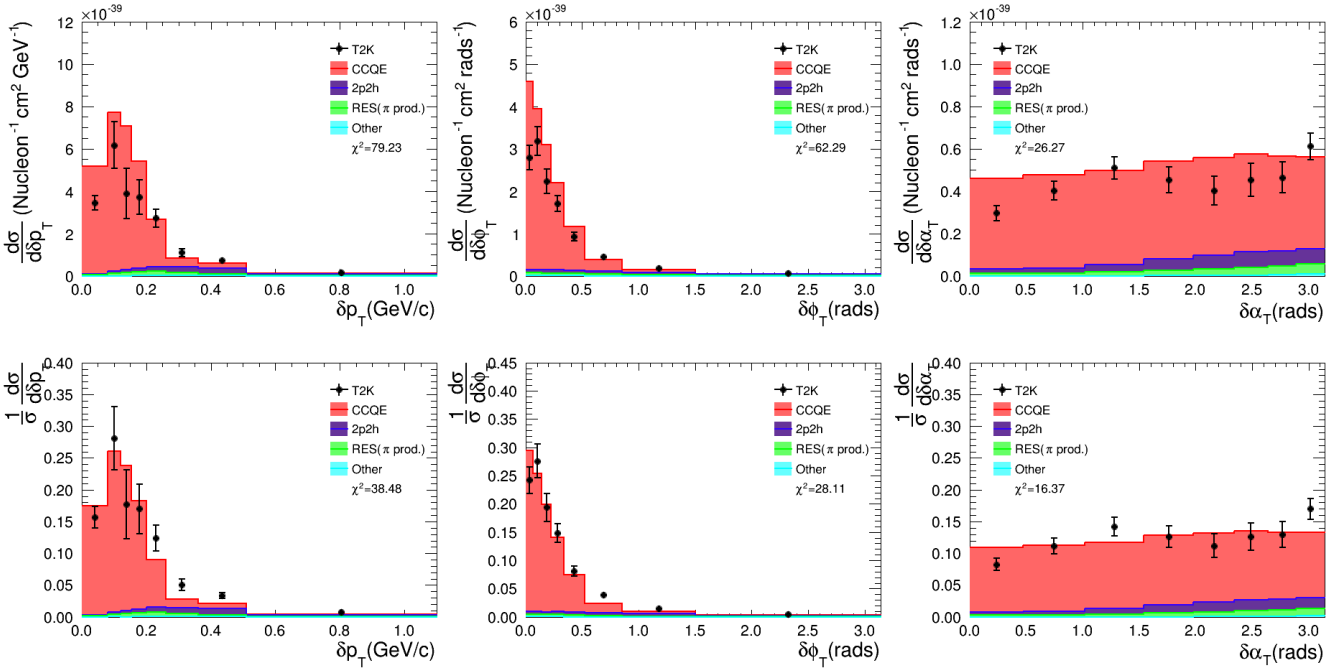


Figure 9.12: A comparison of analysis results to predictions from NuWro11 using $M_A^{QE} = 1.00$ and an LFG nuclear model, split by interaction mode.

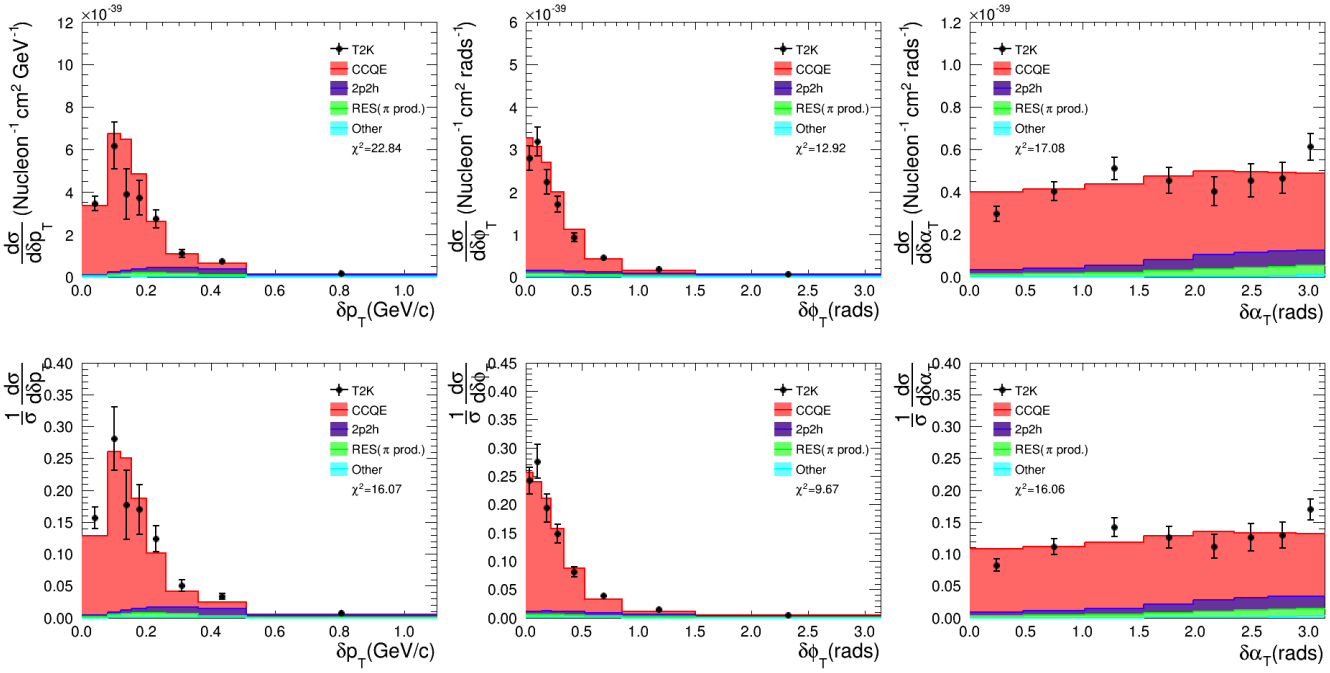


Figure 9.13: A comparison of analysis results to predictions from NuWro11 using $M_A^{QE} = 1.00$ and an SF nuclear model, split by interaction mode.

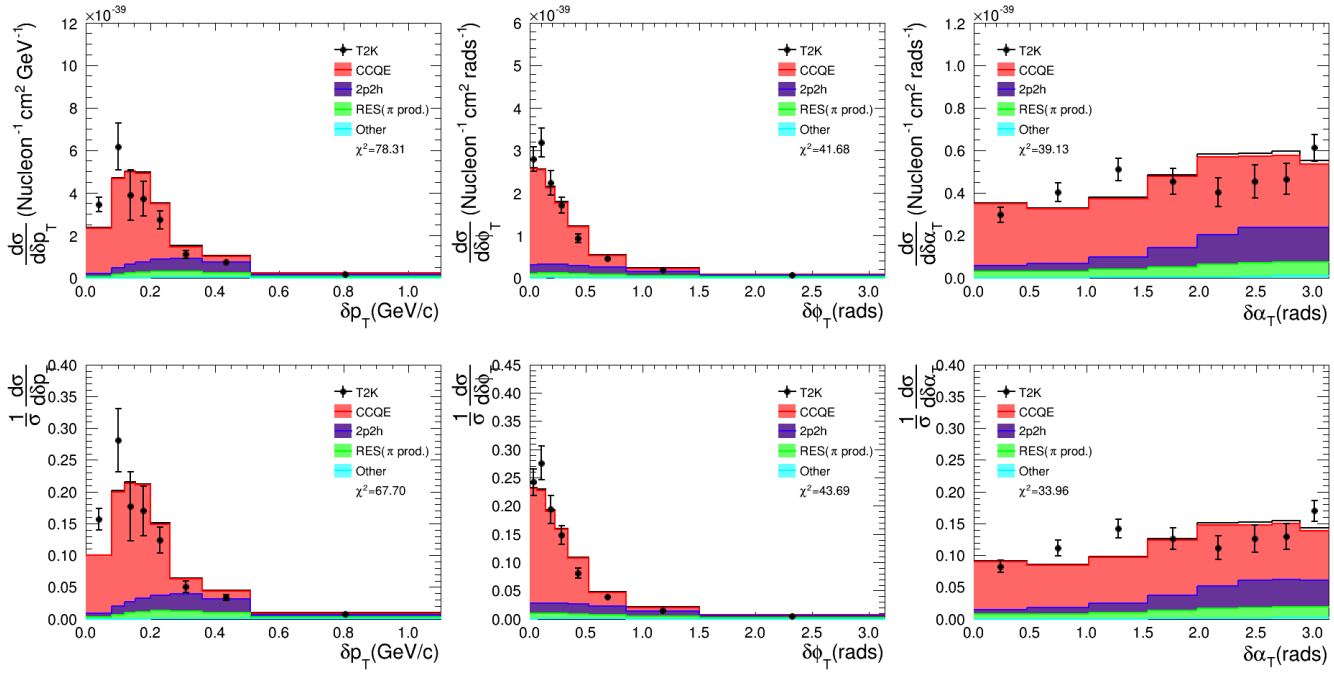


Figure 9.14: A comparison of analysis results to predictions from GiBUU 2016 using $M_A^{QE} = 1.00$ and an LFG nuclear model, split by interaction mode.

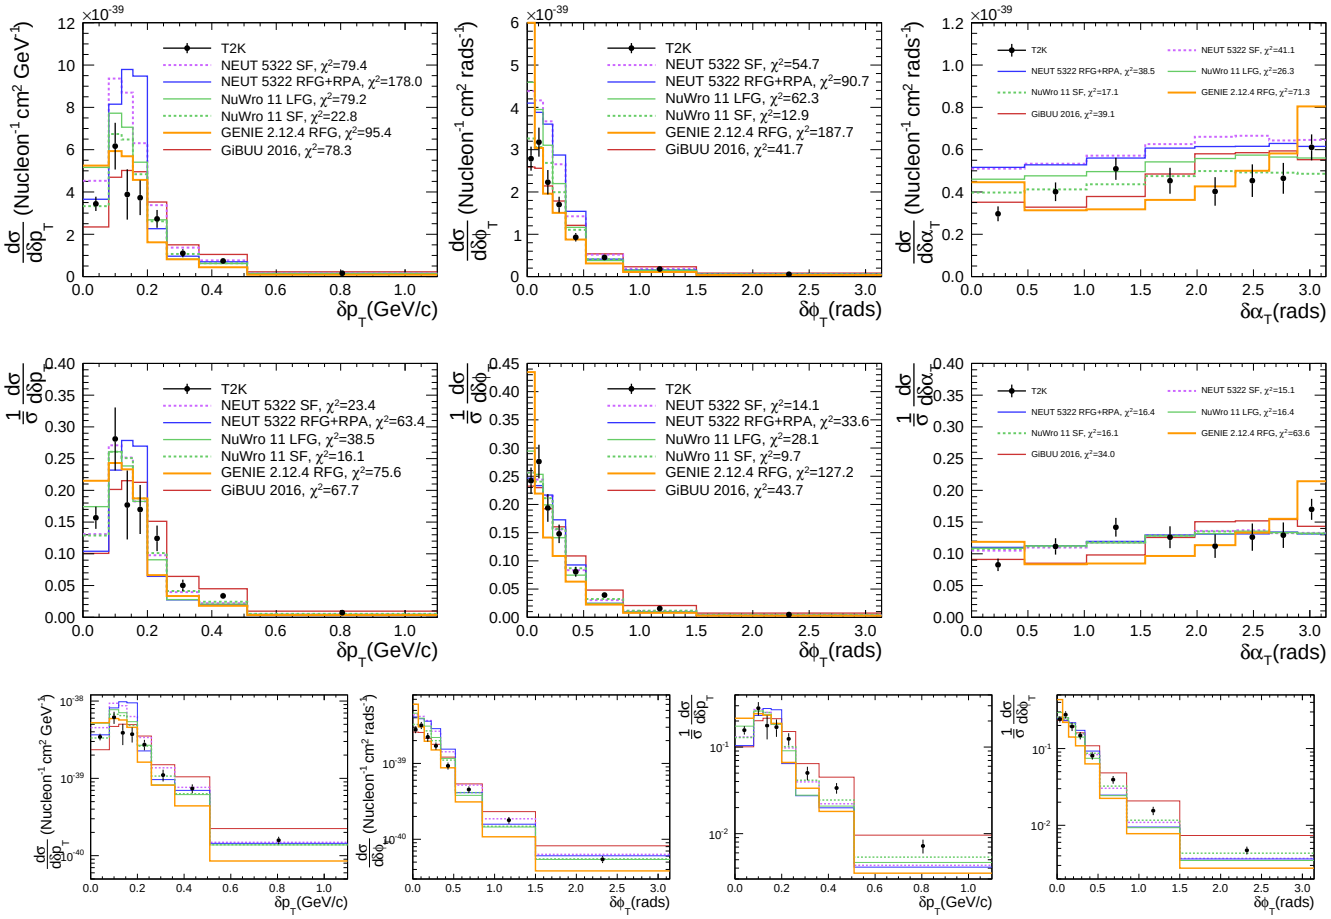


Figure 9.15: A comparison of analysis results to predictions from a selection of different generator predictions from this section. For δp_T and $\delta \phi_T$ the same plots are also shown on a logarithmic scale to allow easier model comparison to the result in the region characterising large transverse imbalance.

Generator nuclear model comparisons

This section shows comparisons of the data to the predictions of individual generators ran with different spectral functions. A comparison of those available within NEUT is shown in figure 9.16, and for NuWro is shown in figure 9.17. A χ^2 goodness of fit is calculated using equation 8.1 and the full cross-section or shape-only covariance matrix.

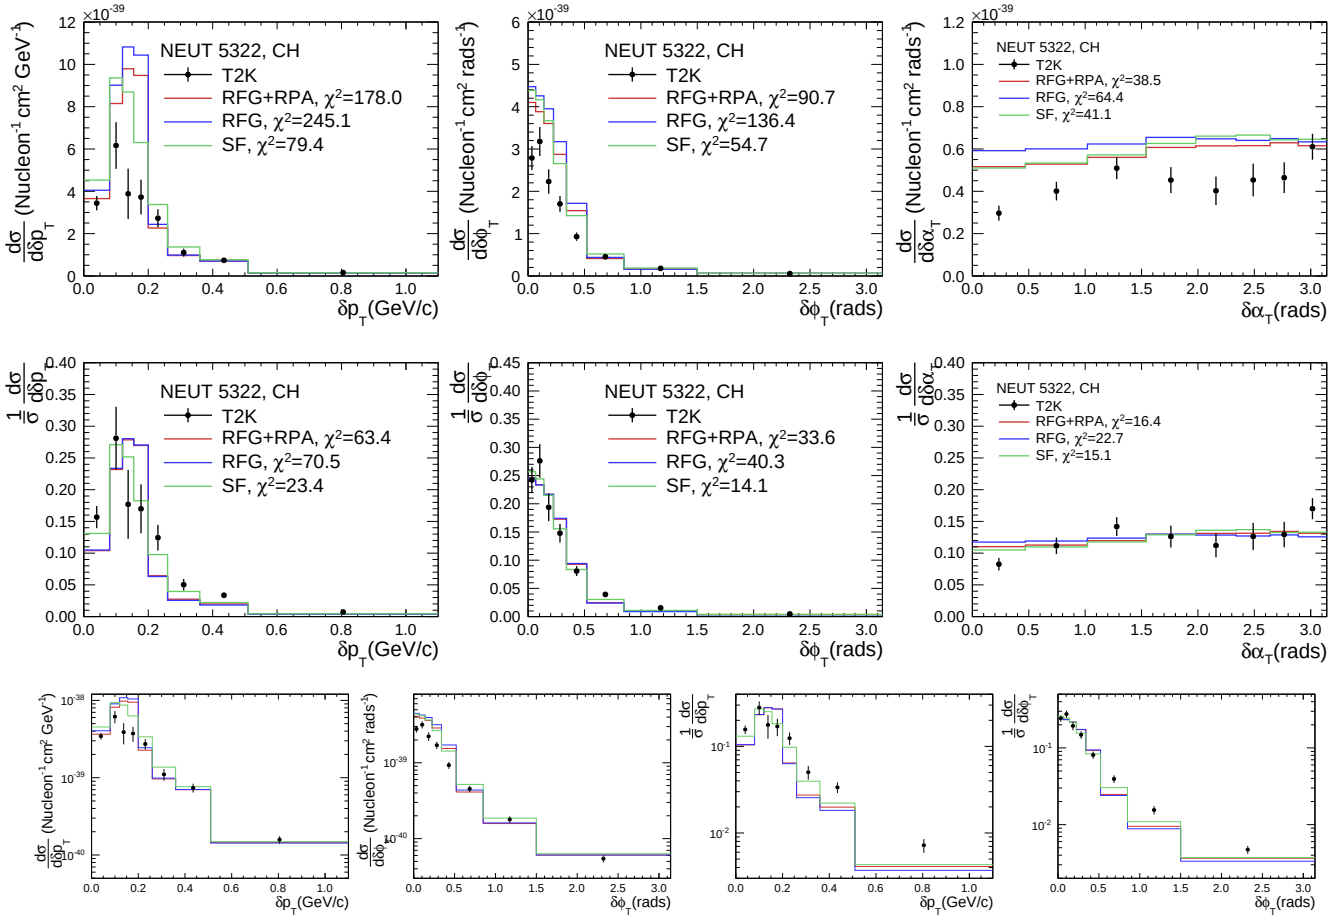


Figure 9.16: A comparison of analysis results to predictions from NEUT 5.3.2.2 using various nuclear models. For δp_T and $\delta\phi_T$ the same plots are also shown on a logarithmic scale to allow easier model comparison to the result in the region characterising large transverse imbalance.

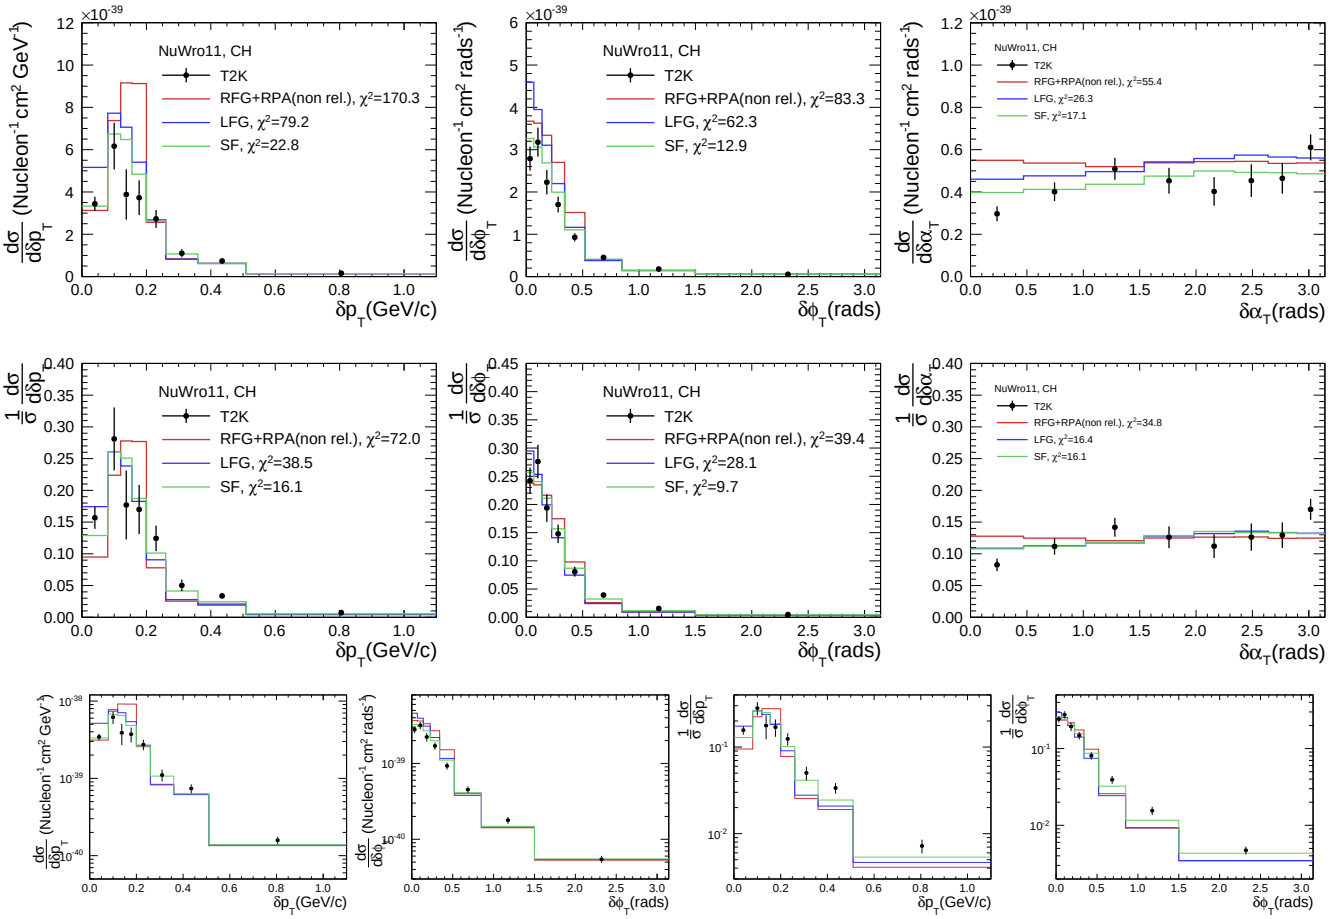


Figure 9.17: A comparison of analysis results to predictions from NuWro 11 using various nuclear models. For δp_T and $\delta\phi_T$ the same plots are also shown on a logarithmic scale to allow easier model comparison to the result in the region characterising large transverse imbalance.

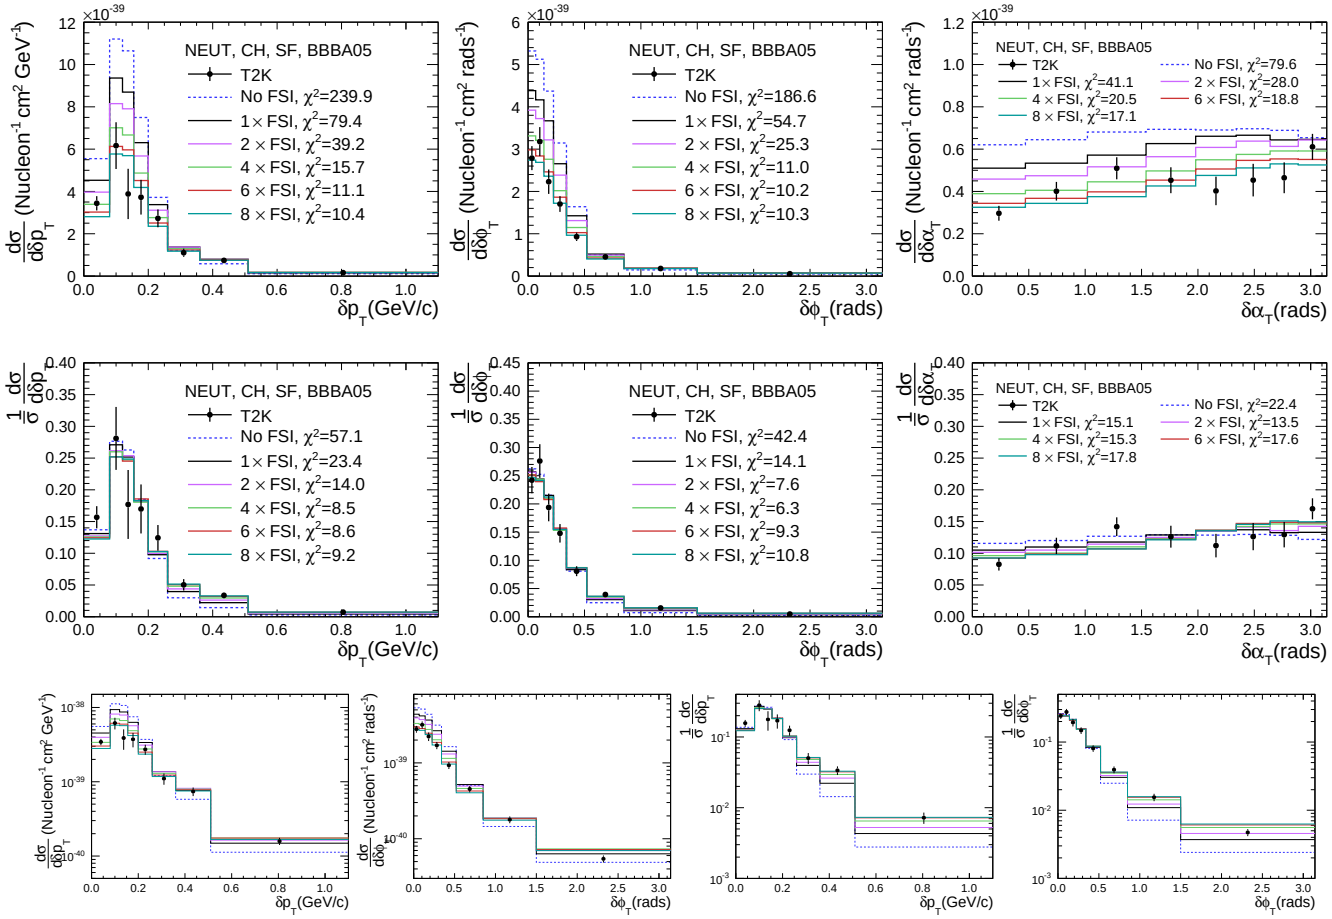


Figure 9.18: A comparison of analysis results to predictions from NEUT 5.3.2.2 using an SF nuclear model with various different FSI strengths. To achieve this the nucleon FSI mean free path is scaled by the factor shown in the legend, as described in section 9.2. For δp_T and $\delta\phi_T$ the same plots are also shown on a logarithmic scale to allow easier model comparison to the result in the region characterising large transverse imbalance.

NEUT FSI strength comparison

This section shows comparisons of the data to the predictions of NEUT with various nucleon FSI strengths. In each figure the simulated mean free path of a nucleon exiting a nucleus is decreased by the factor shown in the figure's legend. This is shown for the NEUT SF prediction in figure 9.18, and for the NEUT RFG+RPA prediction in figure 9.19. A χ^2 goodness of fit is calculated using equation 8.1 and the full cross-section or shape-only covariance matrix.

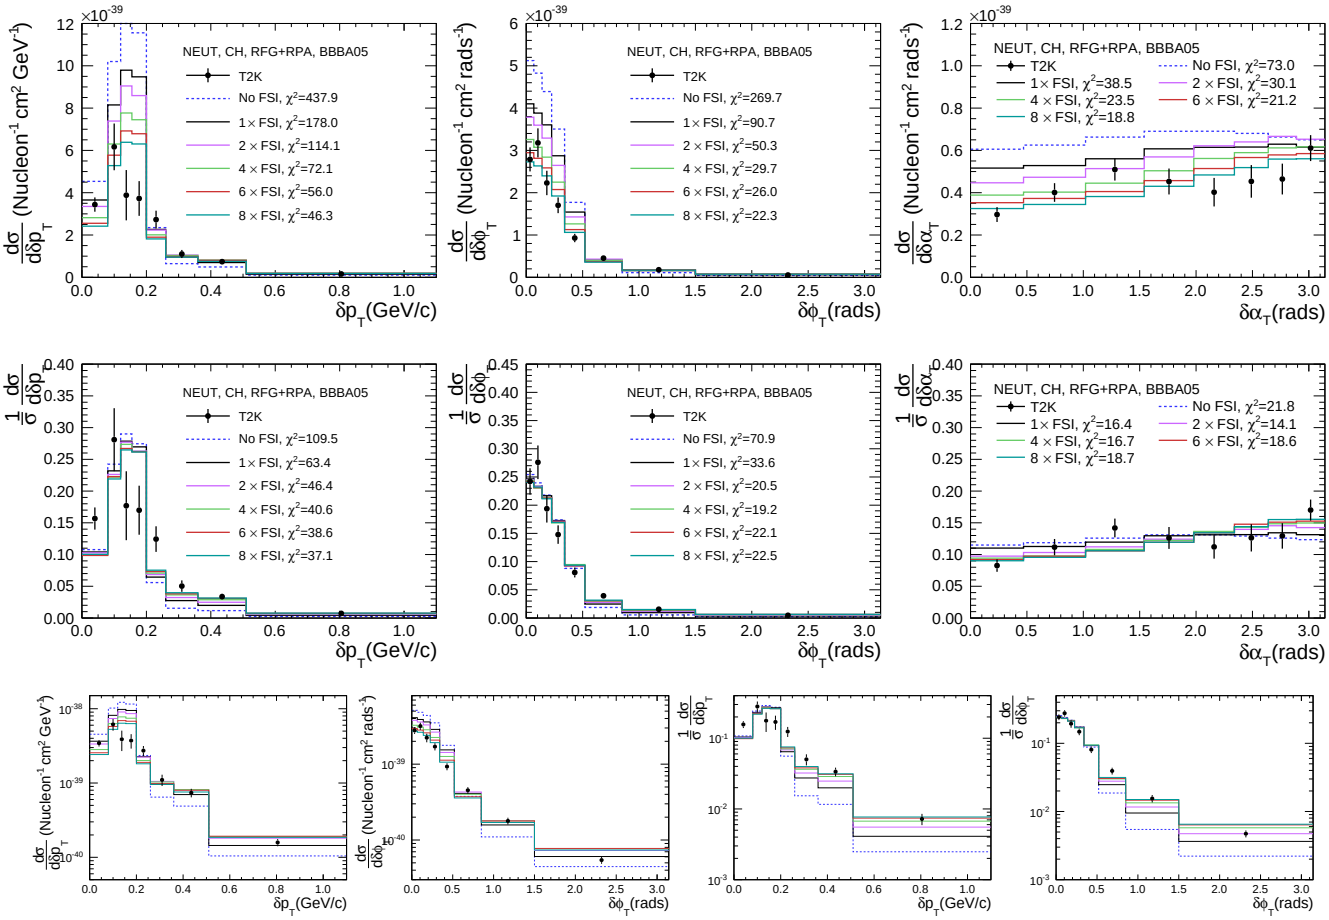


Figure 9.19: A comparison of analysis results to predictions from NEUT 5.3.2.2 using an RFG+RPA nuclear model with various different FSI strengths. To achieve this the nucleon FSI mean free path is scaled by the factor shown in the legend, as described in section 9.2. For δp_T and $\delta\phi_T$ the same plots are also shown on a logarithmic scale to allow easier model comparison to the result in the region characterising large transverse imbalance.

Conclusions

Since, in the absence of any other nuclear effects, δp_T is the transverse projection of the Fermi motion, it is not surprising to see that it offers some separation between initial state nuclear models. Figure 9.16 shows that, although the NEUT normalisation is consistently high, the data does not favour the Fermi cliff steep drop off at around 200MeV/c in the δp_T RFG+RPA (or standalone RFG) prediction, and instead prefers the smoother drop off as predicted by the SF initial state nuclear model. The data does, however, favour suppression of the CCQE contribution, as is offered by relativistic RPA correction (within the context of the NEUT simulation). This disfavouring of a Fermi cliff extends also to the NuWro predictions, as shown in figure 9.17. The NuWro predictions using a SF nuclear model are in particularly good agreement with the result, in both normalisation and shape, relative to the other generator comparisons. The substantial improvement in the agreement with the total measured cross section may be from the effect of short range correlations in a SF nuclear model causing a larger population of lower momentum protons, thereby causing more events to move outside of the restricted phase-space of the measurement and lowering the prediction of the total cross section. However, it should be noted that a model that includes the Nieves 2p-2h prediction on top of the SF nuclear model is not entirely theoretically sound. This is because the 2p-2h implementation is specifically designed to be applied over an LFG nuclear model and already includes some of the short-range correlations that are included in the SF model.

Figure 9.10 and 9.11 show that, although the total cross-section prediction is comparable, the shape of the GENIE prediction is not well favoured by the data. In particular the first bin of all three observables shows considerable contention. These features have previously been shown to stem from the elastic component of the GENIE 'hA' nucleon FSI model and, motivated by the large differences observed compared to other generators in the STV, has been under investigation by the GENIE collaboration [3].

Figure 9.14 shows that the GiBUU predictions achieve both a reasonable shape and normalisation using an LFG nuclear model. Figure 9.15 shows the GiBUU prediction is quite different from those of the other generators, possibly stemming from it's advanced FSI implementation (briefly discussed in section 2.2.2). If so, this demonstrates how a plausible change in nucleon FSI within the other generators may alter the normalisation of the predicted cross section to be in better agreement with the results.

The effect of altering nucleon FSI is further demonstrated in figures 9.18 and 9.19. This shows that, within the context of the NEUT simulation, altering the strength of nucleon FSI is able to achieve excellent agreement with the data using the SF nuclear model, but is unable to distort the Fermi cliff enough to see as good agreement with an RFG+RPA

nuclear model. This is in accordance with the observations in section 9.1, showing that the data was in good agreement with the NEUT simulation containing a bug which causes FSI to be generated as if they are on a heavier target (and are therefore stronger). It should be noted that increasing the FSI to the extent shown is probably not physical given analysis of proton FSI in electron-nucleus scattering experiments similar to those that the NEUT FSI model is tuned to (see, for example, [150]) .

Figures 9.8 to 9.14 allow some assessment of the necessity and strength of a 2p-2h contribution in the context of each generator. As demonstrated in chapter 3, the shape of the STV is almost independent of the alteration of M_A - the nucleon axial mass, such that comparing the shape of the simulated predictions to the result we are afforded a probe of 2p-2h which is independent of M_A. By assessing the tails of the shape only results in δp_T and $\delta\phi_T$ it can be seen that all NEUT and NuWro predictions have a small relative deficit of events in the region of the observables that characterise the 2p-2h contribution. Furthermore the large 2p-2h purity in the tails of the observables show that the removal of the 2p-2h contribution would cause a very large deficit, thereby indicating the need for at least some 2p-2h-like contribution. Although it would be possible to correct such a deficit by increasing the RES contribution or adding more proton FSI, figures 9.18 and 9.19 demonstrate that a dramatic increase in nucleon FSI would be required, whilst the small RES contribution in this region suggests that the necessary increase in the RES prediction would need a similarly remarkable change to the model. A similar deficit is also seen in the GENIE default generator prediction (that do not have a 2p-2h contribution) but figure 9.11 shows that GENIE's empirical 2p-2h is too large to allow good agreement with the result (when ran in its nominal configuration as it was here). It can be seen that GiBUU prediction slightly over-estimates the relative number of events in the tail, but that the 2p-2h contribution is certainly required to allow good agreement with the nominal prediction.

Overall it can be seen that the results offer a powerful new probe of nuclear effects. The shape of the observables are independent of M_A and RPA, but are sensitive to nucleon FSI, 2p-2h and Fermi motion, the latter of which is separated from the two former within the observables phase-space. Moreover, the cross section normalisation provides additional sensitivity to nucleon FSI, M_A and RPA. This unique sensitivity affords the result the ability to lift degeneracies which previous measurements have suffered from (particularly between M_A and 2p-2h), thereby partially isolating the different nuclear effects. However, since only very limited constraints on nucleon FSI, 2p-2h and Fermi motion exist, it remains very challenging to draw firm quantitative conclusions. Despite this an extensive comparison with generator predictions has allowed interesting qualitative conclusions to be drawn which are summarised below:

- The results disfavour the 'Fermi cliff' feature, characteristic of an RFG spectral function (shown most clearly in the δp_T plots in figure 9.16 and figure 9.17).
- The results disfavour the features seen in GENIE's 'hA' elastic nucleon FSI predictions (shown most clearly in the $\delta\phi_T$ plots in figure 9.15).
- The results favour a 2p-2h contribution of a stronger strength to what is implemented in NEUT and NuWro (shown most clearly by the excess of events over the prediction in the tails of the δp_T and $\delta\phi_T$ plots in figure 9.15).
- The results suggest GENIE's 'empirical MEC' 2p-2h model predicts too much transverse imbalance (shown in figure 9.11).
- The results favour an increase in nucleon FSI and an RPA-like suppression of CCQE in NEUT (shown in figure 9.18 and figure 9.16).
- The results best favour the normalisations predicted by GiBUU, GENIE and NuWro with an SF nuclear model (shown in figure 9.15).

If it is possible to further constrain nuclear effects through complementary measurements, whilst also improving the sensitivity of single transverse analyses through gathering more data, it may become reasonable to fit the results from this (or similar) analyses as part of a global neutrino scattering fit to allow more quantitative constraints on nuclear effects. However, currently it is very difficult to parametrise nucleon FSI and the difference in predictions of spectral functions in such a way that a fit could be given the freedom to explore all plausible model space. If such a parametrisation can be developed it may also become reasonable to include single transverse observables in the near detector fit of T2K oscillation analyses, thereby allowing the fit the sensitivity to directly place significant constraints on the principle systematic uncertainties.

Chapter 10

Evaluation

In order to determine the neutrino mass hierarchy and the extent of CP violation generated in the neutrino sector using long-baseline neutrino oscillation experiments, the systematic uncertainty on future measurements must be constrained to within a few percent (discussed in chapter 1). One of the primary difficulties in achieving such precision stems from our naivety of few-GeV neutrino-nucleus interactions. As demonstrated in chapter 2, such interaction uncertainties are already the dominant systematic on current measurements of the oscillation parameters and will soon become the principle limitation if an improved understanding cannot be achieved. In particular, it is essential to better understand the nuclear effects that obfuscate any attempt to measure the cross-section of interactions between neutrinos and free nucleons.

This thesis has presented a way of probing these nuclear effects in neutrino-nucleus scattering through a measurement of muon neutrino charged-current interactions with carbon nuclei that leave no pions but at least one proton in the final state – CC0π+Np interactions. The analysis utilises three novel observables, known as the single-transverse variables, that characterise the kinematic imbalance in the plane transverse to the incoming neutrino and are formed from the kinematics of both the outgoing muon and the highest momentum proton. The primary output of this thesis is the first CC0π+Np fiducial differential cross section in bins of the single-transverse variables, using T2K's ND280 off-axis near detector.

The use of single-transverse kinematic imbalance is motivated by a series of phenomenological studies which culminated in the publication in [3] and the extended studies presented in chapter 3. In particular, these demonstrate that the various features of the single-transverse variables' distributions correspond unambiguously to particular combinations of nuclear effects, whilst remaining shape-independent to plausible variations of the free-nucleon cross-section. This facilitates a much more clear distinction of the effects that dominate the systematic uncertainties in oscillation analyses than can be achieved using single-particle kinematics and is manifest in the strong model-separation potential shown in these studies. Moreover, the single-transverse variables are directly accessible in particle detectors and therefore do not strictly require any input from a neutrino interaction simulation to use as

the observable in a differential cross section, thereby allowing a model-independent measurement. The demonstrated capability of transverse kinematic imbalances to offer a unique handle on principle uncertainties in the oscillation analyses illustrates the relevance of this thesis.

To be able to measure a CC0π+Np cross section to probe single-transverse kinematic imbalance, it is necessary to identify an event selection that well characterises these interactions. This is presented, as a development of well established existing work, in chapter 6. Furthermore, to define the scope of the fiducial cross section to be extracted, the efficiency of the selection is carefully studied to identify the regions of kinematic phase-space that ND280 is sensitive to. Applying this fiducial limitation to the result of this thesis is instrumental in avoiding an excessive reliance on the input simulation in cross-section extraction.

In order to produce a differential cross section that can be compared to the predictions of theorists it is necessary to remove the effects of detector smearing. This is especially difficult when making measurements of observables that correlate kinematics from multiple outgoing particles, as is the case for the single-transverse variables, since these tend to suffer more severe smearing. To address this, chapter 7 exhibited a newly-developed cross-section extraction technique that employs a regularised maximum likelihood template fitter with a rigorous handling of systematic parameters. This method offers significant advantages over more commonly used unfolding techniques, particularly through its focus on minimising dependence on neutrino scattering simulations. This is achieved with a native data-driven regularisation strength optimisation and a rigorous propagation of uncertainties through the procedure. The application of this method to measure differential cross sections that utilise single-transverse kinematic imbalance has been extensively tested through its application to fake data in chapter 8. These tests verified the functionality of the cross-section extraction techniques; exhibited an assessment of sensitivity to some nuclear effects; and illustrated methods of diagnosing model bias in a fit to data.

By applying a rigorous cross-section extraction method to measure observables with unique sensitivity, the results of this thesis, presented in chapter 9, enable a novel characterisation of nuclear effects, including the partial isolation of nucleon initial state correlations and final state interactions. These results demonstrate deficiencies in some existing nuclear models and will facilitate the construction and tuning of improvements. Overall this allows a better understanding of the impact of nuclear effects on the observables used in neutrino oscillation measurements, thereby providing a valuable step to achieving the few percent precision required to precisely study neutrino CP violation and the neutrino mass hierarchy.

Improvements to the cross-section extraction framework; the ongoing development of better pre-fit constraints on the systematic uncertainties; and the imminent arrival of more

data all conspire to afford this analysis considerable scope for improved precision in the near-term future. Such improvements are further discussed in appendix C. Furthermore the results of this thesis are complimentary to other analyses that exploit longitudinal information and, although further phenomenological work is required to realise it, may offer even further model constraints if included in a global neutrino scattering fit.

Appendix A

Event generator implementations

Throughout this thesis a variety of different neutrino-nucleus interaction simulations are used from the NEUT [40], GENIE [41], NuWro [42] and GiBUU [43, 44] neutrino-nucleus interaction event generators. This appendix provides pertinent details of the models used or available for use in each of the generators.

With the exception of the NuWro simulation in chapter 3, all of the event generator predictions shown prior to section 9.2 are those that have been ran through the ND280 production (data simulation) chain discussed in section 4.2.3. These can therefore make predictions of reconstructed level kinematics but have only been ran for one particular configuration of the input generator. The models in each of these are summarised in table A.1.

The generator predictions shown in section 9.2 and the NuWro predictions in chapter 3 have been specially ran for this thesis. In each simulation models have been particularly chosen to either assess the sensitivity of the analysis to different nuclear effects (in chapter 3), or to facilitate the primary physics conclusions of the anlysis (in section 9.2). The event generators used are summarised in table A.2.

The majority of the models listed in these tables are discussed in chapter 2 and further details can be found in the generator references given above. The 'bug' referred to in table A.1 is discussed in section 4.2.3. All generator parameters other than the ones listed in the table are left at their default values for the corresponding generator and version.

Simulation name	Generator version	Nuclear model	M_A	2p-2h	FSI
NEUT 6B	5.3.2	Benhar SF	1.21	Nieves	Cascade (with bug)
NEUT 6D	5.3.2.2	Benhar SF	1.21	Nieves	Cascade
NEUT No 2p2h	5.3.2.2	Benhar SF	1.21	Nieves	Cascade
GENIE	2.8.0	RFG	0.99	None	Empirical (hA)
NuWro	11q	LFG	1.20	Nieves	Cascade

Table A.1: A summary of the models used in the simulations used for event generation at ND280. Further details are discussed in section 4.2.3.

Generator	Version	Utilised nuclear models	M_A	Utilised 2p-2h models	Utilised FSI models
NEUT	5.3.2	Benhar SF, RFG(+RPA)	1.03	Nieves	Cascade (optional MFP variation)
GENIE	2.12.4	RFG	0.99	Empirical	Empirical (hA)
NuWro	11q	Benhar SF, RFG(+RPA), LFG	1.00	Nieves	Cascade
GiBUU	2016	LFG	1.00	e^- scattering based	Hadronic transport

Table A.2: A summary of the event generators and the models utilised within them for analysis sensitivity studies and interpretation of the analysis results. These are used to build the generator predictions shown in section 9.2 as well as the NuWro predictions in chapter 3 and have been specially ran for this thesis. All generator parameters other than the ones listed in the table are always left at their default values for the corresponding generator and version.

Appendix B

Unregularised results

Section 9.1 contained the primary results of the analysis probing single-transverse kinematic imbalance. The small regularisation contribution to the fit result's best-fit likelihood, alongside fake data bias studies in chapter 8, inferred that it is very unlikely that regularisation is significantly biasing the result toward the input simulation.

Despite this, the only truly unbiased result is that which is not regularised. This result is equivalent to finding the cross-section by simply inverting the detector smearing matrix and using the result to transform the reconstructed-level measurement (but still with the rigorous treatment of systematic uncertainties that comes with using a likelihood fitting method as described in chapter 7). When constructing a data-release both the nominal and unregularised result should be provided to allow a user of this result to ensure there is no bias coming from the regularisation. This unregularised result in δp_T, $\delta\phi_T$ and $\delta\alpha_T$ are shown in figure B.1, alongside their associated correlation matrices.

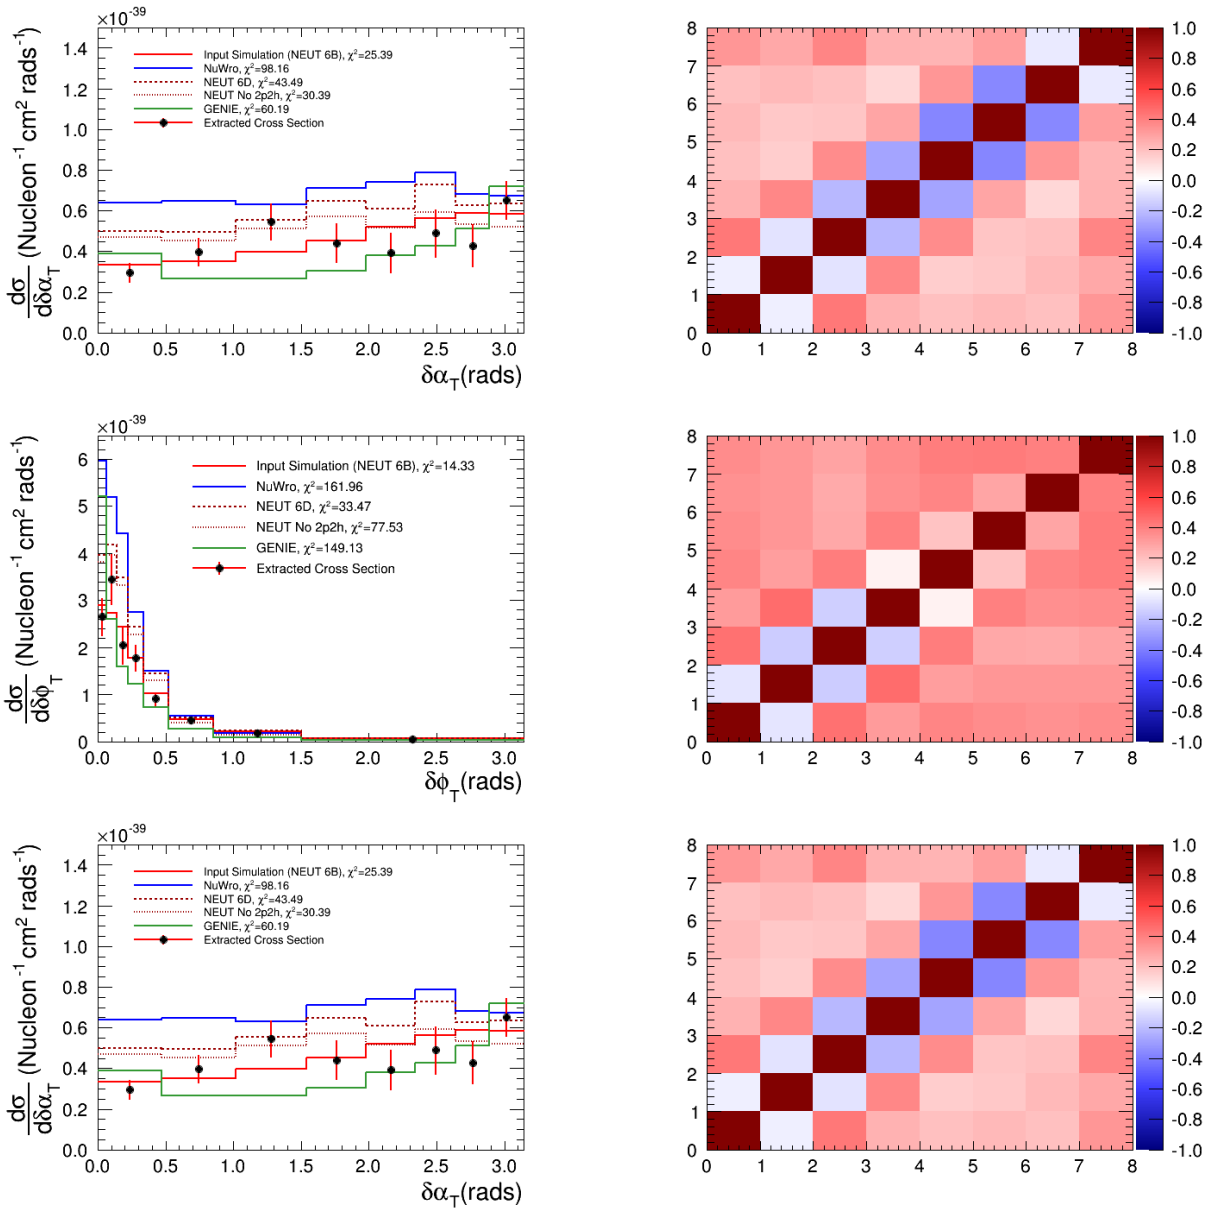


Figure B.1: The unregularised differential cross section in the STV and the associated correlation matrices obtained from fits to real data.

Appendix C

Outlook

The results of this thesis represent the first cross-section measurement using the STV. Achieving this has required the use of some non-standard analysis techniques, some of which have scope for improvement. Furthermore, the imminent availability of almost twice as much data (see section 4.3) will facilitate both improved sensitivity of the current result and the necessary statistics to investigate correlations between the different STV or between the STV and muon/proton kinematics. This section will discuss these potential improvements.

Improvements to the analysis techniques

Section 7.2.6 discussed the potential for model dependence in relying on the input simulation of the underlying muon and proton kinematics in each analysis bin[1]. This could be mitigated with an improved approach to fitting and efficiency correcting. Firstly, instead of fitting single dimensional templates binned in the STV, additional dimensions in proton and muon kinematics could be added. Ideally templates would be included in a STV, p_μ, $cos(\theta_\mu)$, p_p, $cos(\theta_p)$ and $cos(\theta_{\mu p})$ (the opening angle between the muon and the proton), as together these fully characterise the aspects of ND280's acceptance that would change with an alteration of signal model. Once a multi-dimensional cross section is formed, the underlying kinematic bins could then be marginalised to provide a single dimensional result in the STV. The issue with this approach is the immense amount of simulated events that would be needed to extract such a high dimensional cross section without becoming limited by the simulation statistics, but even a few coarse bins would help mitigate the model dependence. Moreover the efficiency correction could then be applied in the muon and proton kinematic phase-space that accurately characterises it. Overall this would afford the fit the freedom to vary the underlying kinematics in each STV analysis bin, whilst also facilitating a better motivated efficiency correction. The next ND280 data simulations (due to be produced within the next

[1] Although this model dependence is problematic, section 7.2.6 suggests it should be identifiable if it is substantial.

year) aim to have approximately 20 times available POT, which may facilitate aspects of this approach.

The error propagation method for the fit is rigorous so long as the two Gaussian error approximations, discussed in 7.2.4, are valid. These come from first characterising the fit uncertainties, and then later the cross-section uncertainties, with covariance matrices. The first assumption could be removed with a more advanced fitting method to fully map the likelihood surface around the best fit point; one option would be to use a Markov chain Monte Carlo to efficiently do this. The second assumption could then also be avoided by sampling from the posterior probability density surface to propagate the result to a cross section and an associated credible interval[2]. This method would allow a precise description of the likelihood surface, such that there would be no need to parametrise it as Gaussian.

More data

Section 8.1 demonstrated that a five times increase in statistics is likely to only modestly reduce the uncertainty on the fit result (from ∼20% to ∼15% for δp_T, and from ∼15% to ∼12% for $\delta\phi_T$ and $\delta\alpha_T$). Therefore the imminent arrival of about twice as much data (discussed in section 4.3) will not make large improvements to the sensitivity of the analysis in its current form. However, more statistics can also facilitate finer binning or the addition of extra dimensions. Although the detector resolution likely limits the potential of the former (see figure 6.17), it may be interesting to investigate a double-differential CC0π+Np cross section. Since the STV each offer their own characterisation of nuclear effects, looking at one in bins of another may allow further isolation of nuclear effects. For example measuring δp_T in bins of $\delta\alpha_T$ may allow a better separation of FSI and 2p-2h effects (enhanced at large δp_T and $\delta\alpha_T$) from the impact of the spectral function (best characterised at small δp_T and $\delta\alpha_T$).

It would also be interesting to investigate the STV in binning of the muon and proton kinematics, perhaps particularly in muon transverse momentum (p_T^μ). The use of p_T^μ is motivated by its potential as a proxy for transverse momentum transfer from the leptonic to the hadronic system in neutrino interactions, which itself may allow better isolation of the 2p-2h contribution to the result. Further details on the use of p_T^μ to probe nuclear effects can be found in [151].

[2]It should be noted that ,as discussed in 7.2.4, avoiding this second assumption would result in a much more complicated data release.

Reduced systematic uncertainties

It is seen in section 8.1 that the flux, detector and model systematics contribute similarly to the uncertainty in the cross section within the fit. However, with the addition of the flux normalisation uncertainty discussed in 7.2.5, the flux is the dominant systematic in the extracted cross section. Thankfully, there may soon be the potential to significantly reduce the flux covariances with the fitting of new T2K replica target data from NA61/SHINE (for more details see section 4.1.3).

Glossary

1p-1h A neutrino interaction mode in which the incoming neutrino scatters off a single nucleon, normally considered within a framework including 2p-2h events as an equivalent to CCQE interactions. 1p-1h stands for 1-particle 1-hole, referring to the single nucleon and the empty state it leaves following the interaction.

2p-2h A neutrino interaction mode in which the incoming neutrino scatters off a bound state of two nucleons. This leaves two empty states, hence the name: 2-particle 2-hole.

Asimov fit A fit in which the fake data is identical to the input simulation.

CCQE A neutrino interaction mode in which a W-boson is exchanged between a neutrino and nucleon, to produce a charged lepton and a nucleon with altered isospin. CCQE stands for charged-current quasi-elastic.

DIS A neutrino interaction mode in which a W- or Z-boson are exchanged with an individual quarks inside a nucleon, thereby fragmenting the nucleon and producing a jet of hadrons. DIS stands for deep inelastic scattering.

ECal The electromagnetic-calorimeter subdetector within ND280. The ECal surrounds the ND280 tracker and is made of segmented polycarbonate scintillator bars, interleaved with layers of lead.

Efficiency An event selection performance metric that measures the ratio of the number of selected signal events to the total number of true signal events.

Elecsim The ND280 data acquisition and electronics simulation used to simulate the detectors response to simulated particle tracks.

Fermi motion The motion of nucleons bound inside a nucleus.

FGD One of the two fine-grained subdetectors within ND280. FGD1 is a full-active polycarbonate scintillator tracker, whilst FGD2 also contains inactive layers of water.

FHC The forward horn current operation mode of the T2K beamline. In this mode positive hadrons are focussed into the beamline's decay tunnel to produce a neutrino beam that is predominantly made of neutrinos (as opposed to antineutrinos).

Fiducial cross section A cross section measured within a specific region of kinematic phase-space that the detector making the measurement in sensitive to.

Flux-integrated cross section A cross section that has not unfolded the effect of the incoming neutrino flux and is therefore experiment dependent.

FSI A within nuclear medium re-interaction of the hadronic final state following a primary interaction with a nuclear constituent. FSI stands for final state interactions. See 'SI' for a discussion of the distinction of FSI and SI.

FV The fiducial volume of a detector is the region in which it can reliably characterise an interaction (in a way that is free from edge effects within the detector).

GEANT A simulation software for the passage of particles through matter.

GENIE A widely-used neutrino-nucleus interaction simulation. GENIE is the primary simulation used for the NOνA and MINERνA experiments and a secondary simulation for T2K. More details can be found in [41].

GiBUU A simulation for the interaction of neutrinos, electrons, photons or hadrons, with nuclei in addition to heavy ion collisions. GiBUU contains an advanced implementation of FSI. More details can be found in [43, 44].

HMN/HMP The highest momentum negative/positive track in a selected event.

INGRID The on-axis near detector of the T2K experiment. INGRID's primary purpose is to measure the T2K neutrino beam profile but it is also capable of measuring neutrino-nucleus interaction cross sections.

IPS Stands for within phase-space and is the categorisation of neutrino interactions which are within the true phase-space defined for a fiducial cross-section measurement.

L-curve A method of identifying the optimal regularisation strength in a fit. The L-curve balances the impact of the regularisation on the fit (normalised penalty) against the compromised goodness of fit.

LFG A spectral function to describe Fermi motion. LFG stands for local Fermi gas and models nucleons as non-interacting fermions inside a potential defined by the local nuclear environment.

Lewellyn Smith (model) A widely-used model to describe CCQE neutrino-nucleus interactions. For more details see [46].

M_A The nucleon axial mass that appears in the Llewellyn Smith model of describing CCQE interactions (it also has an equivalent in the Rein-Seghal model of RES interactions - M_A^{RES}). M_A has been well-measured in neutrino-hydrogen/deuterium scattering and pion electroproduction experiments, but more recent results from neutrino-carbon/oxygen scattering experiments (such as MiniBooNE) measure a different value.

MiniBooNE A short baseline neutrino oscillation experiment at Fermilab which is also capable of measuring neutrino-nucleus interaction cross sections. MiniBooNE's interaction target is mineral oil (CH_2) and uses a neutrino beam with an average energy of 788 MeV [71]. MiniBooNE started taking data in 2002 and has since stopped.

MINERνA A dedicated neutrino-nucleus scattering experiment at Fermilab. MINERνA's active tracker region provides a hydrocarbon target but there are also a multitude of specialised heavier target regions. The 'NuMI' neutrino beam used by MINERνA is able to provide a range of different beam spectra, currently most MINERνA results use a neutrino beam with a peak energy of $\sim$ 3 GeV [152]. MINERνA was completed in 2010 and is currently still taking data.

NA61/SHINE A hadron production experiment at CERN. Targets are placed in a proton beamline and the ejected hadrons are measured. Since the beamline and targets are very similar to those used in the T2K primary beamline, NA61/SHINE provides powerful constraints on the principle uncertainties on the T2K flux prediction.

ND280 The off-axis near detector for the T2K experiment, designed to characterise the neutrino beam before any neutrino oscillations occur. ND280 is well suited to making neutrino-nucleus interaction cross-section measurements. The principle target material in ND280 is hydrocarbon and water but there are also specialised regions with heavier targets. The J-PARC beam provides ND280 with a neutrino beam with $\sim$ 0.6 GeV peak energy. ND280 took first data in 2009 and is still running.

NEUT The primary neutrino-nucleus interaction simulation used by T2K. In this thesis two NEUT simulations are used: NEUT 6B and NEUT 6D. NEUT 6B in generates events all over ND280 but has a bug in its FSI simulation, often causing nucleon FSI in neutrino-nucleus scattering events to be generated as if the interaction was with a heavier nuclear target than was simulated for the rest of the interaction. NEUT 6D contains no such bug, but is only generated in the central tracking region of ND280. For more details on NEUT see [40].

NIWG The T2K neutrino interactions working group. One of the NIWG's primary goals is to fit existing neutrino-nucleus scattering data using the NEUT used by T2K as the input simulation in order to obtain reasonable prior uncertainties on neutrino-nucleus scattering model parameters for the T2K oscillation analyses.

NOνA Alongside T2K, NOνA is the other currently running accelerator-based long baseline neutrino oscillation experiment. NOνA uses the Fermilab 'NuMI' neutrino beam to measure the oscillation of $\sim$ 2 GeV neutrinos over an 810 km baseline.

NUANCE A previously widely-used neutrino-nucleus interaction simulation. More details can be found in [45].

Nuisance parameter A parameter in a fit which alters the likelihood but is not directly of interest. Nuisance parameters are often marginalised over to find a fit result only as a function of the parameters of interest.

NUISANCE A newly-developed neutrino-nucleus interaction simulation comparison and tuning framework. See [149] for more details.

NuWro A neutrino-nucleus interaction simulation. See [42] for more details.

OOA A categorisation of neutrino interactions which fall outside the acceptance of the detector used to measure them. OOA stands for out of acceptance.

OOFV A categorisation of neutrino interactions which occur outside of the target detectors fiducial volume. OOFV stands for out of fiducial volume.

OOPS Stands for out of phase-space and is the categorisation of neutrino interactions which are outside of the true phase-space defined for a fiducial cross-section measurement.

PDF Stands for probability density function. A PDF in a function of a variable whose integral over some interval gives the probability of the variable value falling within that interval.

PDP Stands for parent decay point. In this thesis it is used to refer to a decay point of a T2K beam neutrino's parent (either a pion or a kaon) within the decay tunnel of the J-PARC neutrino beam.

P0D The most upstream central subdetector of ND280. The P0D is made of polycarbonate scintillator, drainable water bags and layers of heavier targets. P0D stands for π^0 detector.

PID Stands for particle identification and used in the context of discussing a detectors ability to determine a particle type from observables.

POT Stands for protons on target. This is the typical metric used to determine the amount of data taken by accelerator-based neutrino experiments. It refers to the number of protons that have impinged on the target of the neutrino beamline.

Purity An event selection performance metric that measures the ratio of the number of selected signal events to the total number of selected events.

RES A neutrino interaction mode in which a W- or Z-boson are exchanged with a nucleon to promote a delta baryon resonance state. The delta then quickly decays to produce a pion inside the nuclear medium of the target. RES stands for resonance scattering.

Regularisation A method of resolving the ill-posedness of a fitting problem. In this thesis an implementation of Tikhonov regularisation is used to impose a condition of smoothness on the extracted cross-section.

Rein-Seghal (model) A widely-used model to describe RES neutrino-nucleus interactions. For more details see [73].

RFG/GFG A spectral function to describe Fermi motion. RFG stands for relativistic Fermilab rmi gas and models nucleons as non-interacting fermions inside a potential defined by the whole nuclear environment. Since the local nuclear environment is not considered, this spectral function is sometimes known as a global Fermi gas (GFG).

RHC The reverse horn current operation mode of the T2K beamline. In this mode negative hadrons are focussed into the beamline's decay tunnel to produce a neutrino beam that is predominantly made of antineutrinos.

RPA Stands for random phase approximation and is a nuclear model to simulate the effects of the nuclear medium on the W- or Z-boson mediating a neutrino-nuclear interaction.

Spectral function A model to describe the momentum distribution of nucleons inside a nucleus.

SF Stands for Benhar spectral function, which is a spectral function to describe Fermi motion. SF allows for the existence or correlations between nucleons.

SI Secondary interactions of particles as they propagate through matter following an interaction. These are distinct from FSI: FSI happen inside the nuclear medium of the target and are thereby mediated by intra-nuclear forces (simulated by a neutrino interaction simulation); SI happen whilst the final state of an interaction is propagating through matter and are thereby mediated by inter-nuclear forces (simulated by GEANT).

Sideband A sample of an event selection designed to characterise the background with a minimal signal component. Sometimes also called a control region.

SK Super-Kamiokande, the far detector of the T2K experiment. SK is a large water-Cherenkov detector which is situated 295 km from the neutrino beam production point.

SMRD Stands for side muon range detector and is an outer detector of the T2K experiment made of coarsely separated scintillator bars.

SRC Stands for short range correlations, which are the initial-state pairs of correlated nucleons within the SF nuclear model.

STKI Stands for single transverse kinematic imbalance, referring to the kinematic imbalance between the outgoing lepton and the hadronic system in the plane transverse to the incoming neutrino in a neutrino-nucleus interaction.

STV Stands for the single transverse variables, which are observables that can characterise STKI.

Template weight The predominant fit parameters considered in this thesis. A template weight is a normalisation factor applied to events falling within both a particular STV bin and the IPS proton/muon kinematic binning (save for the one template weight to scale all OOPS events, integrated over all values of the STV).

TPC Stands for the ND280 time projection chambers. The three TPCs consist mostly of gaseous argon and are the principle tracking subdetectors at ND280.

Tracker (region of ND280) The central region of ND280 consisting of the FGDs and TPCs.

Unfolding The deconvolution of a reconstructed spectrum from detector acceptance and smearing.

Bibliography

[1] A. Sakharov, "Violation of CP invariance,C asymmetry, and baryon asymmetry of the universe", Soviet Physics Uspekhi,34(5):392 (1991), doi:10.1070/PU1991v034n05ABEH002497

[2] K. Abe et al. (T2K Collaboration), "Neutrino Oscillation Physics Potential of the T2K Experiment", Progress of Theoretical and Experimental Physics vol. 2015 (4) (2015), doi:10.1093/ptep/ptv031

[3] X.-G. Lu, L. Pickering, S. Dolan et al., "Measurement of nuclear effects in neutrino interactions with minimal dependence on neutrino energy", Phys. Rev. C 94, 015503 (2016), doi:10.1103/PhysRevC.94.015503

[4] S. Dolan for the T2K Collaboration, "Results and prospects of the latest $\nu_\mu CC0\pi$ analyses from the T2K near detectors", proceedings from NuFact 2016, arXiv:1703.04004 [hep-ex]

[5] S. Dolan for the T2K Collaboration, "Probing Nuclear Eects at the T2K Near Detector Using Single-Transverse Kinematic Imbalance", proceedings from Neutrino 2016, arXiv:1610.05077 [hep-ex]

[6] S. Dolan for the T2K Collaboration, "Probing Nuclear Eects at the T2K Near Detector Using Transverse Kinematic Imbalance", proceedings from NuPhys 2015, arXiv:1605.00179 [hep-ex]

[7] W. Pauli, "Letter of the 4th December 1930", Pauli Archive at CERN, `http://microboone-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=953;filename=pauli%20letter1930.pdf` (accessed April 2017)

[8] F. Reines and C. L. Cowan, "Detection of the Free Neutrino", Phys. Rev. vol. 92 pp. 830831 (1953), doi:10.1103/PhysRev.92.830

[9] E. Fermi, "Versuch Einer Theorie der -Strahlen: I", Zeitschrift fr Physik vol. 88 (3-4) pp. 161177 (1934), doi:10.1007/BF01351864

[10] G. Danby et al., "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos", Phys. Rev. Lett. vol. 9 pp. 3644 (1962), doi:10.1103/PhysRevLett.9.36

[11] M. L. Perl et al., "Evidence for Anomalous Lepton Production in e^+e Annihilation", Phys. Rev. Lett. vol. 35 pp. 14891492 (1975), doi:10.1103/PhysRevLett.35.1489

[12] K. Kodama et al., "Observation of Tau Neutrino Interactions", Physics Letters B vol. 504 (3) pp. 218 224 (2001), doi:10.1016/S0370-2693(01)00307-0

[13] ALEPH, DELPHI, L3, OPAL, and SLD Collaborations, and LEP Electroweak Working Group, and SLD Electroweak Group, and SLD Heavy Flavour Group, "Precision Electroweak Measurements on the Z Resonance", Physics Reports vol. 427 (56) pp. 257 454 (2006), doi:10.1016/j.physrep.2005.12.006

[14] R. Davis, D. S. Harmer, and K. C. Hoffman, "Search for Neutrinos from the Sun", Phys. Rev. Lett. vol. 20 pp. 12051209 (1968), doi:10.1103/PhysRevLett.20.1205

[15] J. N. Abdurashitov et al., "Solar Neutrino Flux Measurements by the Soviet-American Gallium Experiment (SAGE) for Half the 22-Year Solar Cycle", Journal of Experimental and Theoretical Physics vol. 95 (2), pp.181193 (2002), doi:10.1134/1.1506424

[16] M. Cribier et al., "Results of the Whole GALLEX Experiment", Nuclea Physics B - Proceedings Supplements vol. 70 (13) pp. 284 291 (1999), proceedings of the Fifth International Workshop on topics in Astroparticle and Underground Physics, doi:10.1016/S0920-5632(98)00438-1,

[17] Y. Fukua et al., "Measurements of the Solar Neutrino Flux from Super-Kamiokande's First 300 Days", Phys. Rev. Lett. vol. 81, pp. 1158 (1998), doi:10.1103/PhysRevLett.81.1158

[18] B. Pontecorvo, "Mesonium and anti-mesonium", Sov. Phys. JETP 6 (1957) 429, Zh. Eksp. Teor. Fiz. 33 (1957) 549-551 `http://www.jetp.ac.ru/files/pontecorvo1957_en.pdf` (accessed April 2017)

[19] B. Pontecorvo, "Neutrino Experiments and the Problem of Conservation of Leptonic Charge", Journal of Experimental and Theoretical Physics vol. 26 pp. 984988 (1968), `http://jetp.ac.ru/cgi-bin/e/index/e/26/5/p984?a=list` (accessed April 2017)

[20] Q. R. Ahmad et al. (SNO Collaboration), "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory", Phys. Rev. Lett. vol. 89 p. 011301 (2002), doi:10.1103/PhysRevLett.89.011301

[21] Y. Fukuda et al. (Super-Kamiokande Collaboration), "Evidence for Oscillation of Atmospheric Neutrinos", Phys. Rev. Lett. vol. 81 pp. 15621567 (1998), doi:10.1103/PhysRevLett.81.1562

[22] "The Nobel Prize in Physics 2015", Nobelprize.org, Nobel Media AB 2014 (2017), `http://www.nobelprize.org/nobel_prizes/physics/laureates/2015/` (accessed May 2017)

[23] Z. Maki, M. Nakagawa and S. Sakata, "Remarks on the Unified Model of Elementary Particles", Progress of Theoretical Physics. 28: 870 (1962), doi:10.1143/PTP.28.870

[24] V.N. Gribov and B. Pontecorvo, "Neutrino Astronomy and Lepton Charge", Phys. Lett. B28 493 (1969), doi:10.1016/0370-2693(69)90525-5

[25] E. Kh Akhmedov, "Neutrino Physics", Lectures given at Trieste Summer School in Particle Physics, June 7 July 9, 1999 (2000), arXiv:hep-ph/0001264

[26] K. Abe et al. (T2K Collaboration), "The T2K Experiment", Nucl. Instrum. Meth. A 659, 106 (2011), doi:10.1016/j.nima.2011.06.067

[27] D. Ayres et al. (NOvA Collaboration), "NOvA proposal to Build a 30 Kiloton Off-Axis Detector to Study Neutrino Oscillations in the Fermilab NuMI Beamline", Fermilab-Proposal-0929, arXiv:hep-ex/0503053 (2005)

[28] F. P. An et al. (Daya Bay Collaboration), "New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay", Phys. Rev. Lett. vol. 115 p. 111802 (2015), doi:10.1103/PhysRevLett.115.111802

[29] F. P. An et al. (Daya Bay Collaboration), "Independent Measurement of the Neutrino Mixing Angle 13 via Neutron Capture on Hydrogen at Daya Bay", Phys. Rev. D vol. 90 p. 071101 (2014), doi:10.1103/PhysRevD.90.071101

[30] Y. Abe et al. (Double Chooz Collaboration), "Ortho-Positronium Observation in the Double Chooz Experiment", Journal of High Energy Physics vol. 2014 (10) 32 (2014), doi:10.1007/JHEP10(2014)032

[31] Y. Abe et al. (Double Chooz Collaboration), "First Measurement of θ_{13} from Delayed Neutron Capture on Hydrogen in the Double Chooz Experiment" Physics Letters B vol. 723 (13) pp. 66 70 (2013), doi:10.1016/j.physletb.2013.04.050

[32] J. K. Ahn et al. (RENO Collaboration), "Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment", Phys. Rev. Lett. vol. 108 p. 191802 (2012), doi:10.1103/PhysRevLett.108.191802

[33] A. Gando et al. (KamLAND Collaboration), "Reactor On-Off antineutrino Measurement With KamLAND", Phys. Rev. D vol. 88 p. 033001 (2013), doi:10.1103/PhysRevD.88.033001

[34] The DUNE Collaboration, "Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 2: The Physics Program for DUNE at LBNF", arXiv:physics.ins-det/151206148 (2015)

[35] The Hyper-Kamiokande Working Group, "A Long Baseline Neutrino Oscillation Experiment Using J-PARC Neutrino Beam and Hyper-Kamiokande", Document submitted to 18th J-PARC PAC meeting in May 2014, arXiv:physics.ins-det/14124673 (2014)

[36] The LBNE Collaboration, "The Long-Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe", arXiv:hep-ex/13077335 (2013)

[37] The Hyper-Kamiokande Proto-Collaboration, "Physics potential of a long-baseline neutrino oscillation experiment using a J-PARC neutrino beam and Hyper-Kamiokande", PTEP 2015 053C02 (2015), doi:10.1093/ptep/ptv061

[38] Mark Scott for the NuPRISM, Hyper-K collaborations, "An Intermediate Water Cherenkov Detector at J-PARC", Proceedings of NuInt2015, Osaka, November 2015, arXiv:physics.ins-det/160301251 (2016)

[39] J.A. Formaggio and G.P. Zeller, "From eV to EeV: Neutrino Cross Sections Across Energy Scales", Rev. Mod. Phys. 84 1307-1341 (2012), doi:10.1103/RevModPhys.84.1307

[40] Y. Hayato, "A neutrino interaction simulation program library NEUT", Acta Physica Polonica B 40, 2477 (2009)

[41] C. Andreopoulos, et al., "The GENIE Neutrino Monte Carlo Generator", Nucl. Instrum. Meth. A614 87-104 (2010), doi:10.1016/j.nima.2009.12.009

[42] Tomasz Golan, Cezary Juszczak, and Jan T. Sobczyk, "Effects of final-state interactions in neutrino-nucleus interactions", Phys. Rev. C 86, 015505 (2012), doi:10.1103/PhysRevC.86.015505

[43] T. Leitner, O. Buss, L. Alvarez-Ruso, and U. Mosel, "Electron-nucleus and neutrino-nucleus scattering from the quasielastic to the resonance region", Phys. Rev. C 79 034601 (2009), doi:10.1103/PhysRevC.79.034601

[44] O. Buss, T. Gaitanos, K. Gallmeister, H. van Hees, M. Kaskulov, et al., "Transport-theoretical description of nuclear reactions", Phys. Rept. 512 1124 (2012), doi:10.1016/j.physrep.2011.12.001

[45] D. Casper et al., "The Nuance neutrino physics simulation, and the future", Nucl. Phys. Proc. Suppl. 112 161170 (2002), doi:10.1016/S0920-5632(02)01756-5

[46] C.H. Llewellyn Smith, "Neutrino Reactions at Accelerator Energies", Phys. Rept. 3 261-379 (1972), doi:10.1016/0370-1573(72)90010-5

[47] F. Halzen and A. D. Martin, "Quarks and Leptons", Wiley, New York (1984)

[48] O. Gayou et al. (Jefferson Lab Hall A Collaboration), "Measurement of GE_p/G_{M_p} in $ep \rightarrow ep$ to $Q^2 = 5.6GeV^2$, Phys. Rev. Lett. 88, 092301 (2002), doi:10.1103/PhysRevLett.88.092301

[49] R. Bradford, A. Bodek, H. Budd, and J. Arrington, "A new parameterization of the nucleon elastic form factors", Nucl. Phys. B Proc. Suppl. 159 127132 (2006), doi:10.1016/j.nuclphysbps.2006.08.028

[50] A. Bodek et al., "Vector and axial nucleon form factors: A duality constrained parameterization", Eur. Phys. J. C 53: 349 (2008), doi:10.1140/epjc/s10052-007-0491-4

[51] J. Garcia-Ramos et al., "Basic Concepts in Nuclear Physics: Theory, Experiments and Applications: 2015 La Rbida International Scientific Meeting on Nuclear Physics", Springer (2016)

[52] J. Nieves, I. R. Simo, and M. J. V. Vacas, "Inclusive charged-current neutrino-nucleus reactions", Phys. Rev. C 83, 045501 (2011), doi:10.1103/PhysRevC.83.045501

[53] M. Martini, M. Ericson, G. Chanfray, and J. Marteau, "Unified approach for nucleon knock-out and coherent and incoherent pion production in neutrino interactions with nuclei", Phys. Rev. C 80, 065501 (2009), doi:10.1103/PhysRevC.80.065501

[54] T. Golan, "Modeling nuclear effects in NuWro Monte Carlo neutrino event generator", Ph.D thesis, University of Wroclaw (2014), `http://neutrino.ift.uni.wroc.pl/files/phd_tomasz_golan.pdf` (accessed April 2017)

[55] S.K. Singh and E. Oset, "Inclusive quasielastic neutrino reactions in ^{12}C and ^{16}O at intermediate energies", Phys. Rev. C 48, 1246 (1993), doi:10.1103/PhysRevC.48.1246

[56] H. De Vries, C.W. De Jager and C. De Vries, "Nuclear charge-density-distribution parameters from elastic electron scattering", Atom. Data Nucl. Data Tabl. 36 pp. 495536 (1987), doi:10.1016/0092-640X(87)90013-1

[57] P.K.A. de Witt Huberts. "Proton spectral functions and momentum distributions in nuclei from high resolution (e, e-prime p) experiments". J.Phys. G16, pp. 507544 (1990), doi:10.1088/0954-3899/16/4/004

[58] D. Rohe et al., "Correlated strength in nuclear spectral function". Phys. Rev. Lett. 93, p. 182501. (2004), doi:10.1103/PhysRevLett.93.182501

[59] O. Benhar and A. Fabrocini, "Two nucleon spectral function in infinite nuclear matter", Phys. Rev. C62 034304 (2000), doi:10.1103/PhysRevC.62.034304

[60] C. Maieron, M. Martinez, J. Caballero, and J. Udias, "Nuclear model effects in charged current neutrino nucleus quasielastic scattering", Phys. Rev. C68 048501 (2003), doi:10.1103/PhysRevC.68.048501

[61] A. Meucci, C. Giusti, and F. D. Pacati, "Relativistic Greens function approach to charged current neutrino-nucleus quasielastic scattering", Nucl. Phys. A739 277290 (2004), doi:10.1016/j.nuclphysa.2004.04.108

[62] V. Pandey, N. Jachowicz, T. Van Cuyck, J. Ryckebusch, and M. Martini, "Low-energy excitations and quasielastic contribution to electron-nucleus and neutrino-nucleus scattering in the continuum random phase approximation", Phys. Rev. C 92, 024606 (2015), doi:10.1103/PhysRevC.92.024606

[63] L. Pickering, "Examining nuclear effects in neutrino interactions with transverse kinematic imbalance", JPS Conf. Proc. 12, 010032 (2016), doi:10.7566/JPSCP.12.010032

[64] S. Boyd et al, "Comparison of Models of Neutrino-Nucleus Interactions", AIP Conf. Proc. 1189, 60 (2009), doi:10.1063/1.3274191

[65] S. Dytman, "Neutrino event generators" AIP Conf.Proc. 896, 178184 (2007), doi:10.1063/1.2720468

[66] Patrick de Perio. "NEUT pion FSI", AIP Conf.Proc. 1405, 223228 (2011), doi:10.1063/1.3661590

[67] J. Nieves, I. R. Simo, and M. J. V. Vacas, "Inclusive charged-current neutrino-nucleus reactions", Phys. Rev. C 83 045501 (2011), doi:10.1103/PhysRevC.83.045501

[68] M. Martini, M. Ericson, G. Chanfray, and J. Marteau, "A unified approach for nucleon knock-out, coherent and incoherent pion production in neutrino interactions with nuclei", Phys. Rev. C80 065501 (2009), doi:10.1103/PhysRevC.80.065501

[69] A. Bodek, H. S. Budd, and E. Christy, "Neutrino quasielastic scattering on nuclear targets", Eur. Phys. J. C 71 (2011) 1726, doi:10.1140/epjc/s10052-011-1726-y

[70] G.T. Garvey, D.A. Harris, H.A. Tanaka, R. Tayloe, G.P. Zeller, " Recent advances and open questions in neutrino-induced quasi-elastic scattering and single photon production", Phys. Rep. 580 p1-45 (2015), doi:10.1016/j.physrep.2015.04.001

[71] A. A. Aguilar-Areval et al. (MiniBooNE Collaboration),"First Measurement of the Muon Neutrino Charged Current Quasielastic Double Differential Cross Section", Phys. Rev. D81, 092005 (2010), doi:10.1103/PhysRevD.81.092005

[72] J. Nieves, I. R. Simo, and M. V. Vacas, "The nucleon axial mass and the MiniBooNE quasielastic neutrinonucleus scattering problem", Phys. Lett. B 707, 72 (2012), doi:10.1016/j.physletb.2011.11.061

[73] D. Rein and L. M. Sehgal, "Neutrino-excitation of baryon resonances and single pion production", Annals of Physics 133 79153 (1981), doi:10.1016/0003-4916(81)90242-6

[74] D. Rein and L. M. Sehgal, "Coherent π^0 production in neutrino reactions", Nucl. Phys. B223 29 (1983), doi:10.1016/0550-3213(83)90090-1

[75] D. Rein and L. Sehgal, "PCAC and the deficit of forward muons in π^+ production by neutrinos", Phys. Lett. B657 (2007) 207209, doi:10.1016/j.physletb.2007.10.025

[76] D. Perevalov, "A measurement of the muon neutrino charged current quasielastic intereaction an a test of Lorentz violation with the MiniBooNE experiment", PhD thesis, Indiana University (2008), `https://www-boone.fnal.gov/publications/Papers/katori_thesis.pdf` (accessed July 2017)

[77] D. Perevalov, "Neutrino-nucleus neutral current elastic interactions measurement in MiniBooNE" PhD thesis, University of Alabama (2009), `https://www-boone.fnal.gov/publications/Papers/denis_thesis.pdf` (accessed April 2017)

[78] K. Abe et al. (T2K Collaboration), "Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target", Phys. Rev. D 91, 072010 (2015), doi:10.1103/PhysRevD.91.072010

[79] Costas Andreopoulos et al., "VALOR Run 1-7 Joint Fit Analysis with SK ν_e CC1pi^+ sample", Tech. Rep. 306, T2K (2016), `https://www.t2k.org/docs/technotes/306/` (accessed April 2017)

[80] M. Hartz, M. Scott and R. Shah, "Fake Data Studies with Cross Section Systematic Parameter Variations", Tech. Rep. 285, T2K (2016), `https://www.t2k.org/docs/technotes/285/` (accessed April 2017)

[81] P. Adamson et al. (NOvA Collaboration), "First measurement of electron neutrino appearance in NOvA", Phys. Rev. Lett. 116 no.15, 151806 (2016), doi:10.1103/PhysRevLett.116.151806

[82] P. Adamson et al. (NOvA Collaboration), "First measurement of muon-neutrino disappearance in NOvA", Phys. Rev. D93 no.5, 051104 (2016), doi:10.1103/PhysRevD.93.051104

[83] P. Adamson et al. (NOvA Collaboration), "Measurement of the Neutrino Mixing Angle θ_{23} in NOvA", Phys. Rev. Lett. 118, 151802 (2017), doi:10.1103/PhysRevLett.118.151802

[84] R. Gran et al. (K2K Collaboration), "Measurement of the quasielastic axial vector mass in neutrino interactions on oxygen", Phys. Rev. D 74, 052002 (2006), doi:10.1103/PhysRevD.74.052002

[85] V Lyubushkin et al. (NOMAD Collaboration), "A Study of quasi-elastic muon neutrino and antineutrino scattering in the NOMAD experiment", Eur.Phys.J. C63, 355-381 (2009), doi:10.1140/epjc/s10052-009-1113-0

[86] A.A. Aguilar-Arevalo et al. (MiniBooNE Collaboration), "First measurement of the muon antineutrino double-differential charged-current quasielastic cross section", Phys.Rev. D88 no.3, 032001 (2013), doi:10.1103/PhysRevD.88.032001

[87] K.A. Olive et al. (Particle Data Group), "Review of Particle Physics", Chinese Phys. C 38 090001 (2014), doi:10.1088/1674-1137/38/9/090001

[88] S. Barish, J. Campbell, G. Charlton, Y. Cho, M. Derrick, et al., "Study of neutrino interactions in hydrogen and deuterium. 1. Description of the experiment and study of the reaction neutrino $\nu_\mu + d \rightarrow \mu + p + p_s$", Phys. Rev. D16 3103 (1977), doi:10.1103/PhysRevD.16.3103

[89] D. Allasia, C. Angelini, G. van Apeldoorn, A. Baldini, S. Barlag, et al., "Investigation of exclusive channels in neutrino / antineutrino deuteron charged current interactions", Nucl. Phys. B343 285309 (1990), doi:10.1016/0550-3213(90)90472-P

[90] N. Baker, A. Cnops, P. Connolly, S. Kahn, H. Kirk, et al., "Quasielastic neutrino scattering: a measurement of the weak nucleon axial vector form-factor", Phys. Rev. D23 24992505 (1981), doi:10.1103/PhysRevD.23.2499

[91] G. Fanourakis, L. Resvanis, G. Grammatikakis, P. Tsilimigras, A. Vayaki, et al., "Study of Low-energy antineutrino Interactions on Protons", Phys. Rev. D21 562 (1980) doi:10.1103/PhysRevD.21.562

[92] T. Kitagaki, S. Tanaka, H. Yuta, K. Abe, K. Hasegawa, et al., "High-energy quasielastic $\nu_\mu + n \rightarrow \mu + p$ scattering in deuterium", Phys. Rev. D28 436442 (1983), doi:10.1103/PhysRevD.28.436

[93] L. Auerbach et al. (LSND Collaboration), "Measurements of charged current reactions of muon neutrinos on ^{12}C", Phys. Rev. C66 015501 (2002), doi:10.1103/PhysRevC.66.015501

[94] S. Bonetti, G. Carnesecchi, D. Cavalli, P. Negri, A. Pullia, et al., "Study of quasielastic reactions of neutrino and antineutrino in Gargamelle", Nuovo Cim. A38 260270 (1977), doi:10.1007/BF02730023

[95] N. Armenise, O. Erriquez, M. Fogli Muciaccia, S. Nuzzo, F. Ruggieri, et al., "Charged current elastic antineutrino interactions in propane", Nucl. Phys. B152 365375 (1979), doi:10.1016/0550-3213(79)90087-7

[96] S. Belikov, A. Bugorsky, V. S. Epshtein, N. Kalganov, L. Klimenko, et al., "Quasielastic neutrino and antineutrinos scattering: total cross sections, axial vector form-factor", Z. Phys. A320 625 (1985), doi:10.1007/BF01411863

[97] SKAT Collaboration, J. Brunner et al., "Quasielastic nucleon and hyperon production by neutrinos and antineutrinos with energies below 30 GeV", Z. Phys. C45 551 (1990), doi:10.1007/BF01556267

[98] G.A. Fiorentini, D.W. Schmitz, P.A. Rodrigues et al. (MINERμA Collaboration),"Measurement of Muon Neutrino Quasi-Elastic Scattering on a Hydrocarbon Target at $E_\nu \sim 3.5$ GeV", Phys. Rev. Lett. 111, 022502 (2013), doi:10.1103/PhysRevLett.111.022502

[99] K. Abe et al. (T2K Collaboration), "Measurement of the ν_μ CCQE cross section on carbon with the ND280 detector at T2K", Phys. Rev. D 92, 112003 (2015), doi:10.1103/PhysRevD.92.112003

[100] K. Abe et al. (T2K Collaboration), "Measurement of the ν_μ charged current quasielastic cross-section on carbon with the T2K on-axis neutrino beam", Phys. Rev. D 91, 112002 (2015), doi:10.1103/PhysRevD.91.112002

[101] C. Wilkinson et al., "Testing charged current quasi-elastic and multinucleon interaction models in the NEUT neutrino interaction generator with published datasets from the MiniBooNE and MINERνA experiments", Phys. Rev. D 93, 072010 (2016), doi:10.1103/PhysRevD.93.072010

[102] T. Walton, M. Betancourt et al. (MINERμA Collaboration), "Measurement of muon plus proton final states in ν_μ Interactions on Hydrocarbon at average E_ν of 4.2 GeV", Phys. Rev. D 91, 071301 (2015), doi:10.1103/PhysRevD.91.071301

[103] L. Aliaga et al. (MINERμA Collaboration), "Neutrino flux predictions for the NuMI beam", Phys. Rev. D 94, 092005 (2016), doi:10.1103/PhysRevD.94.092005

[104] G.D. Megias, et al., "Charged-current neutrino-nucleus reactions within the superscaling meson-exchange current approach", Phys. Rev. D 94, 093004 (2016), doi:10.1103/PhysRevD.94.093004

[105] T. Katori and M. Martini, "Neutrino-Nucleus Cross Sections for Oscillation Experiments", arXiv:1611.07770 [hep-ph] (2016), *section 3.2.3*

[106] K. Abe et al. (T2K Collaboration), "Measurement of double-differential muon neutrino charged-current interactions on C_8H_8 without pions in the final state using the T2K off-axis beam", Phys. Rev. D 93, 112012 (2016), doi:10.1103/PhysRevD.93.112012

[107] P.A. Rodrigues, J. Demgen, E. Miltenberger et al. (MINERμA Collaboration), "Identification of nuclear effects in neutrino-carbon interactions at low three-momentum transfer", Phys. Rev. Lett. 116, 071802 (2016), doi:10.1103/PhysRevLett.116.071802

[108] P.A. Rodrigues for the MINERμA Collaboration,"Identification of nuclear effects at low momentum transfer in MINERμA", NuInt15 contribution, `http://indico.ipmu.jp/indico/event/46/session/5/contribution/79/material/slides/0.pdf` (accessed April 17)

[109] X.-G. Lu, D. Coplowe, R. Shah, G. Barr, D. Wark, and A. Weber, "Reconstruction of energy spectra of neutrino beams independent of nuclear effects", Phys. Rev. D 92, 051302(R) (2015), doi:10.1103/PhysRevD.92.051302

[110] D. Allasia et al., "Transverse momentum of charged hadrons produced in neutrino and antineutrino deuterium charged current interactions", Z. Phys. C 27: 239 (1985), doi:10.1007/BF01556613

[111] K. Abe et al. (T2K Collaboration), "The T2K Neutrino Flux Prediction",Phys. Rev. D 87, 012001 (2013), doi:10.1103/PhysRevD.87.012001

[112] K. Abe et al. (T2K Collaboration), "Combined Analysis of Neutrino and Antineutrino Oscillations at T2K", Phys. Rev. Lett. 118, 151801 (2017), doi:10.1103/PhysRevLett.118.151801

[113] K. Abe et al. (T2K Collaboration), "Search for short baseline ν_e disappearance with the T2K near detector", Phys. Rev. D 91, 051102(R) (2015), doi:10.1103/PhysRevD.91.051102

[114] K. Abe et al. (T2K Collaboration), "Upper bound on neutrino mass based on T2K neutrino timing measurements", Phys. Rev. D 93, 012006 (2016), doi:10.1103/PhysRevD.93.012006

[115] B. Kirby, "Measurement of an Off-Axis Neutrino Beam Energy Spectrum", PhD thesis, University of British Columbia (2012), `http://www.t2k.org/docs/thesis/020/ubc2012fallkirbybrian` (accessed April 2017)

[116] J-PARC TDR, KEK-Report 2002-13 and JAERI-Tech 2003-044, `http://hadron.kek.jp/~accelerator/TDA/tdr2003/index2.html`, 2003 (accessed May 2017).

[117] G. Battistoni et al., "The FLUKA Code: Description and Benchmarking", AIP Conference Proceedings vol. 896 (1) pp. 3149 (2007), doi:10.1063/1.2720455

[118] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, "FLUKA: a Multi-ParticleTransport Code", Report No. CERN-2005-010; Report No. SLAC-R-773; Report No. INFN-TC-05-11 (2005), `http://slac.stanford.edu/pubs/slacreports/reports16/slac-r-773.pdf` (accessed April 2017)

[119] R. Brun et al., "GEANT: Detector Description and Simulation Tool", Technical Report No. CERN-W5013 (1994), `https://cds.cern.ch/record/1082634` (accessed April 2017)

[120] N. Abgrall et al. (NA61/SHINE Collaboration), "Measurements of Cross Sections and Charged Pion Spectra in Proton-Carbon Interactions at 31 GeV/c", Phys. Rev. C vol. 84 p. 034604 (2011), doi:10.1103/PhysRevC.84.034604

[121] N. Abgrall et al. (NA61/SHINE Collaboration), "Measurements of Production Properties of Positively Charged Kaons in Proton-Carbon Interactions at 31 GeV/c", Phys. Rev. C vol. 85 p. 035210 (2012), doi:10.1103/PhysRevC.85.035210

[122] M. Friend et al., "Flux Tuning and Uncertainty Updates for the 13a Flux", Tech. Rep. 217, T2K (2014), `https://www.t2k.org/docs/technotes/217` (accessed April 2017)

[123] P. A. Amaudruz et al., (T2K ND280 FGD Collaboration), "The T2K Fine-Grained Detectors", Nucl. Instrum. Meth. A 696, 1 (2012), doi:10.1016/j.nima.2012.08.020

[124] N. Abgrall et al., (T2K ND280 TPC collaboration), "The T2K Fine-Grained Detectors", Nucl. Instrum. Meth. A 637, 25 (2011), doi:10.1016/j.nima.2011.02.036

[125] G. Charpak, J. Derre, Y. Giomataris and P. Rebourgeard "MICROMEGAS, a multipurpose gaseous detector", Nucl. Instrum. Meth. A 560, 405 (2006), doi:10.1016/S0168-9002(01)01713-2

[126] C. Giganti, "Particle Identification in the T2K TPCs and study of the electron neutrino component in the T2K neutrino beam", PhD thesis, Institut de Recherche sur les lois Fondamentales de lUnivers (2010), `http://www.t2k.org/docs/thesis/005/thesis` (accessed August 2017)

[127] H. Bethe and J. Ashkin, "Passage of Radiations through Matter", Experimental Nuclear Physics Vol 1, ed. E. Segr, J. Wiley, New York, 1953, p. 253

[128] M. Nirkko, "Measurement of the K^+ production cross section from chargedcurrent ν_μ interactions in hydrocarbon at the T2K near detector", Tech. Rep. 278, T2K (2016), `https://www.t2k.org/docs/technotes/278` (accessed April 2017)

[129] S.Gomi et al, "Development and study of the multi pixel photon counter", Nucl. Instrum. Meth. A 581, 427 (2007), doi:doi.org/10.1016/j.nima.2007.08.020

[130] M. Yokoyama et al (KEK detector technology project photon sensor group), "Development of Multi-Pixel Photon Counters", arXiv:physics/0605241 [physics.ins-det]

[131] P. Sinclair, "Implementation of a multinucleon neutrino interaction simulation and comparison with T2K data", PhD thesis, Imperial College London (2014), `http://www.t2k.org/docs/thesis/054/thesis` (accessed April 2017)

[132] L. Pickering, PhD thesis, Imperial College London (2017), in preparation

[133] L. Monfregola, "Muon neutrino charged current quai-elastic interactions in the T2K off-axis near detector", PhD thesis, Universtitat de Valencia (2015), `http://www.t2k.org/docs/thesis/063/LauraMonfregolaThesis` (accessed April 2017)

[134] P. Martins, "ECal pi0 veto studies for P6B", Tech. Rep. 270, T2K (2016), `https://www.t2k.org/docs/technotes/270` (accessed April 2017)

[135] S. Bolognesi, "Best Practices for T2K Cross Section Measurements", Tech. Rep. 263, T2K (2016), `https://www.t2k.org/docs/technotes/263/` (accessed April 2017)

[136] P. Bartet-Friburg, "Mesure de la section efficace dinteraction des ν_μ parcourant charg sur du ^{12}C, sans production de pions dans l'tat final et en fonction de la multiplicit et de la cinmatique des particules produites, en utilisant le dtecteur proche de T2K.", PhD thesis, LPNHE Paris (2016), `http://www.theses.fr/2016PA066351` (accessed April 2017)

[137] S. Wilks, "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses", The Annals of Mathematical Statistics. 9: 6062 (1938), doi:10.1214/aoms/1177732360

[138] W. Feller, "An Introduction to Probability Theory and Its Applications", section 2.9, Vol. 1, 3rd ed. New York, Wiley (1968).

[139] G. Cowan, "Statistical Data Analysis", Oxford University Press, section 11.5.1 (1998)

[140] P. C. Hansen, "Analysis of Discrete Ill-Posed Problems by Means of the L-Curve", SIAM Rev., 34(4), 561580, doi:10.1137/1034115

[141] S. Schmitt, "TUnfold, an algorithm for correcting migration effects in high energy physics" Journal of Instrumentation, Volume 7 (2012), doi:10.1088/1748-0221/7/10/T10003

[142] K. Abe et al., "Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector", arXiv:1707.01048 [hep-ex] (2017)

[143] A. Bercellie et al., "Cross section parameters for 2014 oscillation analysis", Tech. Rep. 192, T2K (2014), `https://www.t2k.org/docs/technotes/192` (accessed April 2017)

[144] A. Higuera, A. Mislivec, et al. (MINERvA Collaboration), "Measurement of Coherent Production of $\pi^\pm$ in Neutrino and Anti-Neutrino Beams on Carbon from E_ν of 1.5 to 20 GeV", Phys. Rev. Lett. 113, 261802 (2014), doi:10.1103/PhysRevLett.113.261802

[145] A. Bercellie et al., "Implementation of additional NIWG cross section parameterizations for 2014 analyses", Tech. Rep. 193, T2K (2014), https://www.t2k.org/docs/technotes/193 (accessed June 2017)

[146] F. James and M. Roos, "Minuit: A System for function minimization and analysis of the parameter errors and correlations", Comput. Phys. Commun. 10(1975) 343367, doi:10.1016/0010-4655(75)90039-9

[147] S. Oser, "Elemental Composition of the FGD XY Modules", Tech. Rep. 091, T2K (2010), https://www.t2k.org/docs/technotes/091/ (accessed June 2017)

[148] S. Bolognesi, V. Galymov and C. Giganti, "Flux-integrated cross-section measurement of charged current events with no pions in the final state using the T2K near detector", Tech. Rep. 214, T2K (2014), https://www.t2k.org/docs/technotes/214/ (accessed April 2017)

[149] P. Stowell, C. Wret, C. Wilkinson, L. Pickering, et. al., "NUISANCE: a neutrino cross-section generator tuning and comparison framework.", 2017 JINST 12 P01016, doi:10.1088/1748-0221/12/01/P01016

[150] O. Hen, et. al. (CLAS Collaboration), "Measurement of transparency ratios for protons from short-range correlated pairs", Phys. Rev. B722, 63 (2013), doi:10.1016/j.physletb.2013.04.011

[151] C. Patrick, "Measurement of the Antineutrino Double-Differential Charged-Current Quasi-Elastic Scattering Cross Section at MINERνA", Ph.D thesis, Northwestern University (2016), http://lss.fnal.gov/archive/thesis/2000/fermilab-thesis-2016-04.pdf (accessed May 2017)

[152] L. Aliaga, et al. (MINERνA Collaboration),"Design, Calibration, and Performance of the MINERνA Detector", Nucl. Inst. and Meth. A743 130 (2014), doi:10.1016/j.nima.2013.12.053