

Measuring charged current neutrino interactions in the electromagnetic calorimeters of the ND280 detector

Mark Andrew Scott

High Energy Physics
Blackett Laboratory
Imperial College London

A dissertation submitted to Imperial College London
for the degree of Doctor of Philosophy

Abstract

This thesis presents a study of neutrino interactions within the electro-magnetic calorimeters (ECals) of the Tokai-to-Kamioka (T2K) off-axis near detector (ND280), using data collected from T2K Run II and Run III.

Neutrino oscillation physics is a rapidly advancing field, with the recent discovery of non-zero θ_{13} allowing the possibility of measuring CP-violation in the lepton sector. The current neutrino beam experiments must make precise measurements of the oscillation parameters and to do this require an improved understanding of neutrino interactions, which can only come through better measurements of neutrino cross-sections.

This thesis describes the development of a neutrino event selection using a boosted decision tree multi-variate analysis to separate interactions within the ND280 ECals from entering backgrounds. This is then used to provide a selection of neutrino event samples from each ECal module, which are inputs to a χ^2 fit that is used to extract the ν_μ -Pb charged current inclusive cross-section, which was found to be $\sigma_{\text{Pb}}^{\text{CC}} = 5.76_{-2.46}^{+2.65} \times 10^{-39} \text{ cm}^2 \text{ nucleon}^{-1}$.

Declaration

This analysis has been built upon the work of a large number of people over a long period of time and the purpose of this declaration is to make clear my personal contributions.

Chapter 1 and Chapter 2 are descriptions of the current state of neutrino physics and the T2K experiment respectively, and as such the information presented has been taken from a number of sources, which have been referenced appropriately.

Chapter 3 details the ND280 off-line software and event reconstruction. Everything presented is the work of the ND280 Software and ND280 Reconstruction working groups, except for the ECal Radon transform, which I developed during the course of this analysis.

The event selection described in Chapter 4 is entirely my own work, though it made use of the ROOT TMVA software suite, which is referenced accordingly.

Finally, the cross-section analysis presented in Chapter 5 is my own work, with contributions from the T2K Beam working group for the neutrino flux uncertainties, the T2K Neutrino Interactions working group, who provided the neutrino cross-section systematic uncertainties, and the T2K UK group who provided input for some of the detector systematic studies. The analysis also made use of the ν_μ charged current inclusive ND280 tracker event selection, which was provided by the ND280 ν_μ working group.

I hereby certify that all material presented here that is not my own work has been properly attributed.

Mark Scott

Acknowledgements

This thesis would not have been possible without the time and effort given by a number of people. I would like to thank my supervisor, Dr. Yoshi Uchida, for his help and support throughout my PhD. I would also like to thank Dr. Morgan Wascko for giving me the initial idea of making this measurement and our week on night shift together where we worked out how to make it happen. The T2K group at Imperial has been a good place to work over the last three and a half years and I have benefitted from the academic and social support from all the people who have been part of it.

The analysis performed here was directly dependent upon the hard work of the T2K collaboration, for which I am very grateful. I would like to thank the T2K Cross-Section working group for the input they provided for this work, and would particularly like to thank Prof. Alfons Weber, Dr. Kendall Mahn and Dr. Mark Hartz for their essential contributions.

Contents

List of Figures	7
List of Tables	16
1. Introduction	18
1.1. A brief history of neutrino oscillations	18
1.2. Neutrino oscillation theory	21
1.3. The current state of neutrino oscillation physics	23
1.4. Neutrino interaction cross-sections	24
1.4.1. Neutrino cross-section theory	25
1.4.2. Neutrino interactions on heavy nuclei	29
2. The T2K experiment	33
2.1. The T2K beamline	34
2.2. The ND280 detector suite	37
2.2.1. Multi-Photon Pixel Counter	37
2.2.2. The INGRID detector	39
2.2.3. The ND280 detector	40
2.3. The far detector	47
3. Event reconstruction and software	51
3.1. The ND280 software suite	51
3.1.1. The data processing chain	51
3.1.2. The Monte Carlo processing chain	52
3.2. ECal event reconstruction	56
3.2.1. ECal multi-pass Radon transform	61
4. Selecting neutrino interactions in the ND280 ECals	70
4.1. Data sample	71
4.1.1. Data quality requirements	71

4.1.2. Monte Carlo sample	73
4.2. A Monte Carlo study to select neutrino interactions in the ECals	74
4.2.1. Pre-selection of events	74
4.3. Boosted decision tree multivariate analyses	76
4.3.1. Creating a boosted decision tree	76
4.3.2. The boosted decision tree in this analysis	80
4.3.3. Data-MC comparisons	85
4.4. Final event selection	89
4.4.1. Final BDT performance	91
4.5. Event selection summary	97
5. Measuring the neutrino-lead charged current inclusive cross-section	104
5.1. Analysis method	104
5.2. Fitting to extract the cross-section	105
5.2.1. Samples used in this analysis	106
5.2.2. Systematics	111
5.2.3. Validating the fit performance	138
5.2.4. Fit applied to real data	143
5.2.5. Extracting the neutrino-lead charged current inclusive cross-section	149
6. Discussion and conclusions	152
6.1. Discussion of cross-section result	152
6.2. Conclusions	153
6.3. Outlook	154
Bibliography	156
A. Boosted decision tree input variables	160
B. ‘External’ population in barrel ECal samples	177

List of Figures

1.1.	Distribution of neutrino interactions at SK as a function of zenith angle	19
1.2.	The ratio of the observed $\bar{\nu}_e$ spectrum to the expectation for no oscillations as a function of $\frac{L}{E}$	20
1.3.	The allowed and excluded regions of parameter space for four types of neutrino oscillation	24
1.4.	Feynman diagrams showing QES examples of CC and NC interactions. The central ring indicates that, unlike the DIS case, the boson here interacts with the nucleon as a whole, rather than just a single quark	26
1.5.	Feynman diagrams showing (a) CC resonant pion production and (b) CC deep inelastic scattering.	26
1.6.	A summary of the current measurements of ν_μ CC interactions	27
1.7.	A comparison between the MiniBooNE and NOMAD ν_μ CCQE cross-section measurements	29
1.8.	The difference between the calculated and measured values of the electron-nucleus QES cross-section	30
1.9.	Measurements of the relative CC inclusive cross-sections for lead, iron, marble and polyethylene	31
1.10.	Measured ratio of the CC inclusive cross-section of ν_μ on lead to that on plastic	32
2.1.	The T2K experiment.	33
2.2.	Simulated neutrino energy spectra for CC interactions at SK	34

2.3.	The J-PARC accelerator complex	35
2.4.	An overview of the T2K neutrino beamline	36
2.5.	The T2K near detector hall	38
2.6.	The INGRID detector	39
2.7.	An example INGRID neutrino event	40
2.8.	An exploded view of the ND280 detector	41
2.9.	Simplified diagram of a TPC	43
2.10.	TPC measurements of the energy loss and momentum of particles	44
2.11.	A schematic of the PØD sub-detector	45
2.12.	SMRD construction image	48
2.13.	SK ID event displays	49
3.1.	A diagram showing the package structure of the ND280 software suite	52
3.2.	The predicted neutrino flux and the effect of kaon tuning	54
3.3.	Energy resolution of the DsECal energy fit as a function of energy	61
3.4.	Data-MC comparisons of control sample data to the appropriate particle gun MC	62
3.5.	The output from applying the described Radon transform to particle gun MC	63
3.6.	The output from applying the described Radon transform to neutrino beam MC	64
3.7.	The results from sequentially applying the Radon transform algorithm three times to a single ECal cluster, removing any tracks that are found at each stage.	65
3.8.	An event display showing the distribution of hits that created the Radon transform output	66

3.9. Difference between the true particle angle and that reconstructed by the Radon transform algorithm	67
4.1. The number of protons per pulse and the integrated number of POT over time	72
4.2. The ECal hit time relative to the beam trigger time	73
4.3. The number of hits in each cluster for (a) the DsECal and (b) the barrel ECal	75
4.4. The most upstream layer hit in the DsECal	75
4.5. The ROC curve for a number of different TMVA methods, showing the amount of background rejection achievable at a given signal efficiency.	77
4.6. A single decision tree taken from an example BDT	78
4.7. The number of different layers containing a hit for each ECal cluster	81
4.8. DsECal BDT classifier response	84
4.9. Barrel ECal BDT classifier response	85
4.10. DsECal ECal classifier cut tuning	85
4.11. Barrel ECal classifier cut tuning	86
4.12. Average Z position of all reconstructed clusters in the right barrel ECal	87
4.13. The truncated maximum ratio of clusters that pass the minimal pre-selection cuts	88
4.14. The PID Circularity of clusters that pass the minimal pre-selection cuts	88
4.15. The number of hits contained in clusters that pass the minimal pre-selection cuts	89
4.16. Histograms showing data-MC comparisons of the calculated (a) thrust and (b) χ^2_{MIP}	91
4.17. Histograms showing signal–background comparisons of the calculated (a) thrust and (b) χ^2_{MIP}	91

4.18. Histograms showing the DsECal and barrel ECal variables that display the greatest remaining difference between data and MC	92
4.19. The final DsECal and barrel ECal BDT responses	93
4.20. The selected events in the DsECal, plotted as a function of incident neutrino energy	95
4.21. The selected events in the barrel ECals, plotted as a function of incident neutrino energy	98
4.22. The purity and efficiency of the DsECal event selection versus neutrino energy.	100
4.23. The purity and efficiency of the DsECal event selection versus Q^2	101
4.24. The purity and efficiency of the barrel ECal event selection versus neutrino energy.	102
4.25. The purity and efficiency of the barrel ECal event selection versus Q^2 . . .	103
5.1. Selected CC candidate events from interactions in the most upstream FGD	107
5.2. The energy of all muon neutrinos that undergo a CC interaction in the top right and bottom left ECal modules	107
5.3. The BDT output from all clusters in the barrel ECals that pass the final pre-selection	108
5.4. Selected events from the ‘External’ population of barrel ECal modules . .	109
5.5. The BDT classifier output comparing data and MC	110
5.6. Fractional covariance matrix showing the uncertainties on the neutrino flux	112
5.7. The variation in the neutrino flux due to pion production uncertainties .	114
5.8. The nominal and shifted neutrino energy spectra in the bottom left barrel ECal module	115
5.9. The nominal and shifted neutrino energy spectra in the DsECal	116
5.10. The variation in the number of events selected from the MC sample coming from the uncertainties in the neutrino flux.	118

5.11. Matrices showing the change in the number of selected events due to uncertainties on the neutrino flux prediction	119
5.12. Covariance matrix showing the fractional uncertainties on the value of the NEUT cross-section parameters.	121
5.13. The variation in the total number of events selected due to uncertainties on the value of the main NEUT cross-section parameters.	121
5.14. Matrices showing the change in the number of selected events due to uncertainties on the NEUT cross-section parameters	122
5.15. The change in the neutrino energy spectrum due to 1σ shifts in the value of M_A^{QE}	123
5.16. The charge deposited in each hit from a sample of cosmic rays passing through the DsECal	125
5.17. The spread of MPVs from cosmic data collected over T2K Run III, fitted with a Gaussian curve to show the 1σ variation.	126
5.18. Covariance matrix created by considering the effects of variation in the ECal charge response on the sub-detector event samples.	126
5.19. Average number of pixel-equivalents as a function of the number of incident photons from the nominal MC	127
5.20. Average number of pe as a function of the number of incident photons from the corrected MC	128
5.21. Covariance matrix created from the MPPC saturation systematic uncertainty	129
5.22. Covariance matrix created from the TiO_2 thickness uncertainty	130
5.23. Covariance matrix created from the systematic uncertainty on the scintillator bar size	130
5.24. Covariance matrix created from the systematic uncertainty on the lead sheet size	131
5.25. Estimation and tuning of MC noise rate	132
5.26. Covariance matrix created from the uncertainty on the rate of noise hits in the ECals	133

5.27. The combined detector systematic uncertainty covariance matrix.	133
5.28. A covariance matrix generated from two independent samples of nominal MC, showing the intrinsic variation within the samples.	134
5.29. Fractional covariance matrix showing the effect of detector uncertainties on the FGD event selection.	135
5.30. The total fractional covariance matrices from evaluating the different systematic error sources	136
5.31. The distribution of the fitted values and pulls of the lead population normalisation parameter	140
5.32. Fitted values and pulls of the four normalisation parameters when applying parameter set 1	140
5.33. Fitted values and pulls of the four normalisation parameters when applying parameter set 2	141
5.34. Fitted values and pulls of the four normalisation parameters when applying parameter set 3	141
5.35. The fitted parameter values when fitting the signal-enriched or background-enriched samples alone	142
5.36. The minimum χ^2 value from 10,000 fits to fake data generated using parameter set 1	143
5.37. The number of events selected in data compared to the number selected in the nominal MC across the sub-detector samples	145
5.38. The number of events selected in data compared to the number selected in the fitted MC across the sub-detector samples	145
5.39. The value of the χ^2 function as a function of R^{Pb} , R^{C} and R^{External} . . .	146
5.40. The number of events selected in data compared to the number selected in the fitted MC when fitting the left ECal module	148
5.41. The minimum χ^2 value across a range of values of the R^{Pb} and R^{C} parameters	148
5.42. The simulated neutrino flux for T2K Run II and Run III	150

A.1. The mean charge of all hits in the cluster.	160
A.2. The charge skew of hits in the cluster.	161
A.3. The RMS of the charge of hits in the cluster.	161
A.4. The summed charge of all hits in the cluster.	161
A.5. The number of ECal layers containing a hit from the cluster.	162
A.6. The layer closest to the centre of the ND280 that contains a hit from the cluster.	162
A.7. The layer furthest from the centre of the ND280 that contains a hit from the cluster.	162
A.8. Pointing X axis.	163
A.9. Pointing Y axis.	163
A.10. Pointing Z axis.	163
A.11. Maximum perpendicular bar hit.	164
A.12. Maximum parallel bar hit.	164
A.13. Minimum bar hit.	164
A.14. PID Angle.	165
A.15. PID ShowerAngle.	165
A.16. Thrust value.	165
A.17. Thrust X axis.	166
A.18. Thrust Y axis.	166
A.19. Thrust Z axis.	166
A.20. Maximum hits per layer.	167
A.21. PID Circularity.	167
A.22. Truncated maximum ratio.	167
A.23. Front back ratio.	168

A.24.Distance parameter of first Radon line in first view of cluster.	168
A.25.Distance parameter of second Radon line in first view of cluster.	168
A.26.Distance parameter of third Radon line in first view of cluster.	169
A.27.Distance parameter of first Radon line in second view of cluster.	169
A.28.Distance parameter of second Radon line in second view of cluster.	169
A.29.Distance parameter of third Radon line in second view of cluster.	170
A.30.Angle of first Radon line in first view of cluster.	170
A.31.Angle of second Radon line in first view of cluster.	170
A.32.Angle of third Radon line in first view of cluster.	171
A.33.Angle of first Radon line in second view of cluster.	171
A.34.Angle of second Radon line in second view of cluster.	171
A.35.Angle of third Radon line in second view of cluster.	172
A.36.Percentage of available hits used in first Radon line in first view of cluster.	172
A.37.Percentage of available hits used in second Radon line in first view of cluster.	172
A.38.Percentage of available hits used in third Radon line in first view of cluster.	173
A.39.Percentage of available hits used in first Radon line in second view of cluster.	173
A.40.Percentage of available hits used in second Radon line in second view of cluster.	173
A.41.Percentage of available hits used in third Radon line in second view of cluster.	174
A.42.Number of hits in the first view of the cluster.	174
A.43.Number of hits in the second view of the cluster.	174
A.44.The total number of hits in the cluster.	175
A.45. χ^2_{MIP}	175
A.46.The average Z position and length of clusters in the barrel ECal modules.	175

A.47. The X-Y position of clusters in the DsECal.	176
B.1. Top right barrel ECal module ‘External’ population.	177
B.2. Top left barrel ECal module ‘External’ population.	178
B.3. Left side barrel ECal module ‘External’ population.	178
B.4. Bottom right barrel ECal module ‘External’ population.	178
B.5. Bottom left barrel ECal module ‘External’ population.	179

List of Tables

1.1. The best fit values and 1σ uncertainties on the three flavour neutrino oscillation parameters	25
3.1. The performance of the multi-pass Radon algorithm for finding tracks in 2D clusters in the barrel ECals.	66
3.2. The 3D matching performance of the multi-pass Radon transform.	68
3.3. The 3D matching performance of the multi-pass Radon transform for the first track	68
4.1. The input variables for the DsECal and barrel ECal BDTs.	82
4.2. The results of applying the DsECal event selection to MC, broken down at each cut point.	86
4.3. The results of applying the barrel ECal event selection to MC, broken down at each cut point.	86
4.4. The stricter pre-selection cuts for the DsECal and barrel ECal event selections	90
4.5. The final set of input variables to the barrel and downstream ECal BDTs.	92
4.6. The results of applying the final DsECal event selection to MC, broken down at each cut point. The pre-selection cut values can be found in Table 4.4.	94
4.7. The results of applying the final barrel ECal event selection to MC, broken down at each cut point. The pre-selection cut values can be found in Table 4.4.	94

5.1. The fractional systematic uncertainty on the selected ECal samples, broken down into the contribution from each source of systematic error.	137
5.2. The three population normalisation parameter sets used to test the χ^2 fit.	138
5.3. The total number of events selected in the different sub-detector samples for both data and MC.	144
5.4. The fitted values of the four normalisation parameters with their errors from the fitting.	144
5.5. The fitted values of the three normalisation parameters found when fitting each of the barrel ECal modules separately with the DsECal and FGD samples.	147

Chapter 1.

Introduction

The Standard Model (SM) of Particle Physics is one of the most successful theories in science, explaining a number of observed experimental results at its inception and making precise predictions that were later confirmed. Since the discovery of a Higgs-like boson at the LHC [1, 2] we have observed all of the particles that the Standard Model contains, and all the measured properties of these particles agree with the predictions it makes. However, the SM does not currently contain a mass term for the neutrino. The quantum mechanical phenomenon of neutrino oscillations therefore stands as the only experimentally verified result that goes beyond the Standard Model.

1.1. A brief history of neutrino oscillations

The first observation of the effect of neutrino oscillations was made by R. Davis [3] at the Homestake Solar Neutrino Experiment in 1968. Davis was searching for solar neutrinos produced in the sun, using the reaction $\text{Cl}^{37}(\nu, e^-)\text{Ar}^{37}$. The total rate of neutrino interactions was expected to be $2.0 \times 10^{-35} \text{ s}^{-1}$, primarily from B^8 decay. The experiment did not manage to detect any neutrino interactions above their measured background so put an upper limit on the neutrino interaction rate of $3.0 \times 10^{-36} \text{ s}^{-1}$, almost an order of magnitude below the value predicted using the standard solar model (SSM). This deficit was then confirmed in 1990 by the Kamiokande-II experiment [4], which observed the ratio of the measured B^8 neutrino flux to the predicted as $0.46 \pm 0.05(\text{stat}) \pm 0.06(\text{syst})$. During this period, the SAGE [5] and GALLEX [6] experiments were making measurements of solar neutrinos from both B^8 decay and the proton-proton fusion reaction chain. Again, both experiments saw a deficit compared to the SSM prediction, with the SAGE

experiment showing that this deficit had to exist in both the proton-proton reaction and B^8 decay.

There were a number of possible explanations for the observed solar neutrino deficits, and though neutrino oscillations were perhaps the preferred explanation they had still not been experimentally confirmed. In 1988 the Kamiokande-II [7] detector published their first analysis of atmospheric neutrinos, where they had collected 277 fully contained events. Their measured electron neutrino flux agreed with that predicted, whereas the measured muon neutrino flux was only $59 \pm 7\%$ of that expected. This atmospheric neutrino anomaly was confirmed by the IMB [8], Soudan2 [9] and the Super-Kamiokande (SK) [10] experiments, with SK publishing a landmark paper in 1998 describing an observed deficit of muon neutrinos that was dependent upon the zenith angle of the events (Fig. 1.1). This paper demonstrated that the flux of muon neutrinos from the atmosphere varied with the distance the neutrinos had travelled, suggesting that this disappearance was caused by neutrino oscillations.

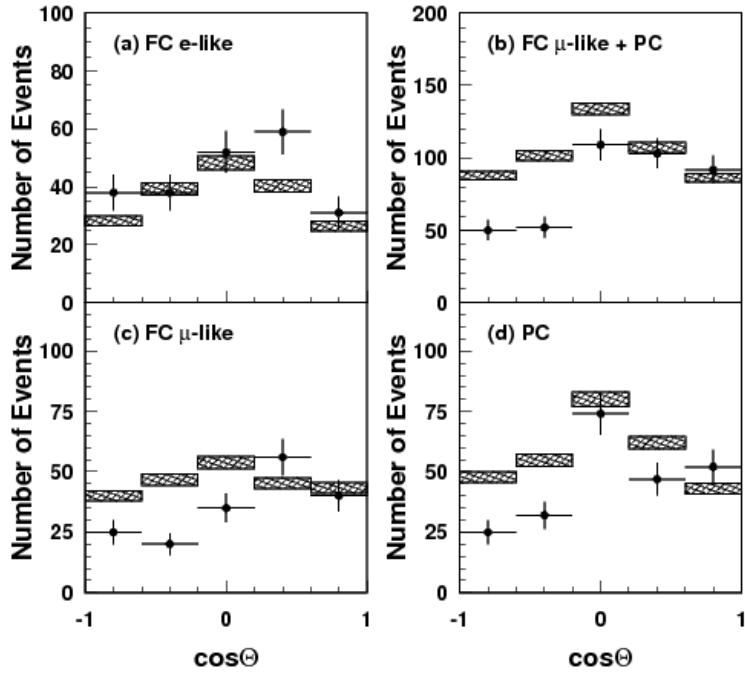


Figure 1.1.: The number of reconstructed neutrino interactions at SK as a function of zenith angle, data points, compared to the number predicted, the hashed rectangles. A deficit in μ -like events can be seen for negative values of $\cos \Theta$ [10].

Further evidence for oscillations came from the Sudbury Neutrino Observatory (SNO) experiment in 2002 [11]. At this time SNO was a water Cherenkov detector using

deuterium as a target and was able to measure charged current (CC), neutral current (NC) and elastic scattering (ES) interactions. Both charged and neutral current neutrino interactions will be defined in detail in the following sections of this thesis, but the key to the SNO experiment was that the CC interactions were only sensitive to electron neutrinos, whilst the NC and ES processes were sensitive to all three neutrino flavours. The measured rate of ES interactions at SNO agreed with the high precision measurements made by the SK experiment, and by comparing the neutrino flux derived from CC events at SNO to the flux from ES events at SK they demonstrated that the solar neutrino flux must be made up in part by at least one other active neutrino flavour.

The final confirmation of neutrino oscillations came from the KamLAND [12] experiment, which used liquid scintillator to detect electron anti-neutrinos from nuclear reactors. KamLAND calculated the ratio of the observed $\bar{\nu}_e$ spectrum to that expected without neutrino oscillations, which is shown in Fig 1.2. The variation in this ratio clearly shows both the disappearance and subsequent reappearance of $\bar{\nu}_e$ s as a function of their energy, proving beyond doubt that neutrino oscillations occur.

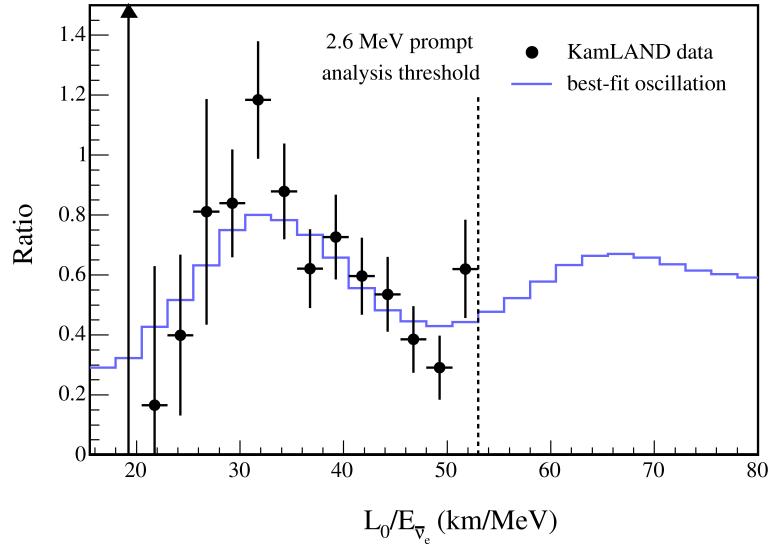


Figure 1.2.: The ratio of the observed $\bar{\nu}_e$ spectrum to that expectation with no oscillations as a function of $\frac{L}{E}$ [12]. The data points show the ratio measured by the KamLAND experiment, using $\bar{\nu}_e$ s from a number of nuclear reactors surrounding the experiment, and the blue line is the best-fit to that data assuming there are neutrino oscillations.

1.2. Neutrino oscillation theory

The theory of neutrino oscillations can be constructed using one underlying assumption: that the observed neutrino flavours are in fact a linear combination of freely propagating mass states. These mass eigenstates are then related to the flavour eigenstates by

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle, \quad (1.1)$$

where ν_α are the flavour eigenstates, ν_i the mass states and $U_{\alpha i}$ an element of the mixing matrix describing their relation to one another.

One can then state that the wavefunction of a given flavour state at a given time and distance, $|\phi(x, t)\rangle$ can be expressed as

$$|\phi(x, t)\rangle = \sum_i U_{\alpha i} |\nu_i\rangle e^{-i\mathbf{p}_\nu \cdot \mathbf{x}_\nu} \quad (1.2)$$

with \mathbf{p}_ν the four momentum of the neutrino and \mathbf{x}_ν its four vector position.

If the neutrino mass is small compared to its momentum then $E_i \simeq p_\nu + \frac{M_i^2}{2p_\nu}$. Having $M_i \ll p_\nu$ requires that the neutrino be travelling at approximately c , so, in natural units, its position and time become equivalent. Eqn. 1.2 can then be re-written as

$$|\phi(x, t)\rangle = \sum_i U_{\alpha i} |\nu_i\rangle e^{-i(\frac{M_i^2}{2p_\nu})x}. \quad (1.3)$$

The neutrino mass states can then be re-expressed as flavour states using the relation

$$|\nu_i\rangle = \sum_\beta U_{\beta i}^* |\nu_\beta\rangle, \quad (1.4)$$

which then gives

$$|\phi(x, t)\rangle = \sum_\beta \sum_i U_{\alpha i} e^{-i(\frac{M_i^2}{2p_\nu})x} U_{\beta i}^* |\nu_\beta\rangle. \quad (1.5)$$

From Eqn. 1.5 the probability of observing a specific neutrino flavour at a distance x from the neutrino's production point, $P(\nu_\alpha \rightarrow \nu_\beta, x)$, is given by

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta, x) &= \left[\sum_j U_{\alpha j}^* e^{i(\frac{M_j^2}{2p_\nu})x} U_{\beta j} \right] \left[\sum_i U_{\alpha i} e^{-i(\frac{M_i^2}{2p_\nu})x} U_{\beta i}^* \right] \\ &= \sum_{i, j} U_{\alpha i} U_{\alpha j}^* U_{\beta i}^* U_{\beta j} e^{i\frac{M_i^2 - M_j^2}{2p_\nu}x}. \end{aligned} \quad (1.6)$$

Eqn. 1.6 shows that the probability of observing a specific neutrino flavour changes as a function of the mass squared splitting between the neutrino mass eigenstates, the neutrino momentum and the distance it has travelled. It also demonstrates that observing oscillations requires that there is at least one massive neutrino and that there must be non-trivial mixing between the neutrino mass and flavour eigenstates. As neutrinos are ultra relativistic, $\frac{M_i^2 - M_j^2}{2p_\nu}x$ can be expressed as $\frac{\Delta m_{ij}^2}{2E_\nu}L$, where $\Delta m_{ij}^2 = M_i^2 - M_j^2$, E_ν is the neutrino energy and L the distance the neutrino has travelled.

The matrix that describes this non-trivial mixing is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. This matrix can be parameterised by three mixing angles, θ_{12} , θ_{23} and θ_{13} , and either a single complex phase, δ , if neutrinos are Dirac particles, or the phase δ alongside two additional complex phases if neutrinos are Majorana particles [13]. It has been shown [13] that these additional Majorana phases do not affect the oscillation probability derived above, so these phases can be ignored when discussing neutrino oscillations. Using the above parameterisation, the PMNS matrix can be written as

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (1.7)$$

where $c_{ij} = \cos(\theta_{ij})$, $s_{ij} = \sin(\theta_{ij})$ and θ_{ij} is the mixing angle between the i and j mass states. This parameterisation neatly splits the PMNS matrix into three rotations, which can be roughly associated with the neutrino sources that can be used to measure them. Explicitly, atmospheric and beam neutrinos provide sensitivity to θ_{23} , beam and reactor neutrinos are used to measure θ_{13} , and solar neutrinos probe θ_{12} . These associations are due to the dominant term in the neutrino oscillation probability in each case but this is not the only factor that determines the sensitivity of individual experiments. The ratio of an experiment's baseline to the energy of the observed neutrinos and the mass

squared splitting between the two neutrino states govern the oscillation probability, which will ultimately determine the sensitivity of the experiment to both the dominant and sub-dominant mixing angles.

1.3. The current state of neutrino oscillation physics

Since the discovery of atmospheric neutrino disappearance by SK there has been a huge increase in the number of experiments studying neutrino oscillations, leading to more and better measurements of the oscillation parameters. To illustrate this, Fig. 1.3 shows the experimental constraints on the mass squared splittings and mixing angles as of September 2011. This shows that KamLAND, SK and SNO together give a strong constraint on both θ_{12} and Δm_{12}^2 whilst the combination of SK and MINOS constrain θ_{23} and $|\Delta m_{23}^2|$. Fig. 1.3 also shows whether an individual experiment is sensitive to the sign of the mass squared splittings, which appears as an asymmetry about $\tan \theta = 1$. The solar neutrino experiments are able to determine the sign of Δm_{12}^2 due to matter effects in the sun, a process called the Mikheyev, Smirnov, Wolfenstein (MSW) effect that is explained in greater detail in Ref. [13].

In 2011 the Tokai to Kamioka (T2K) collaboration published the first indication that θ_{13} was non-zero [14], observing 6 electron neutrino interactions in an off-axis muon neutrino beam when only 1.5 ± 0.3 events were expected assuming $\theta_{13} = 0$. This early result was quickly confirmed by the Daya Bay [15] and RENO [16] collaborations, who both saw the disappearance of electron anti-neutrinos from the flux produced by nuclear reactors with a 5.2σ and 4.9σ significance respectively.

Including these results, our current knowledge of the three flavour neutrino oscillation parameters is summarised in Table 1.1, along with their 1σ uncertainties.

From Eqn. 1.7 it can be seen that the discovery of non-zero θ_{13} opens up the possibility of measuring the CP violating phase δ , perhaps leading to an observation of CP violation in the lepton sector. This, along with the sign of Δm_{23}^2 , is one of the remaining unknowns in neutrino physics, and making measurements of these parameters is the goal of the current and next generation of neutrino oscillation experiments.

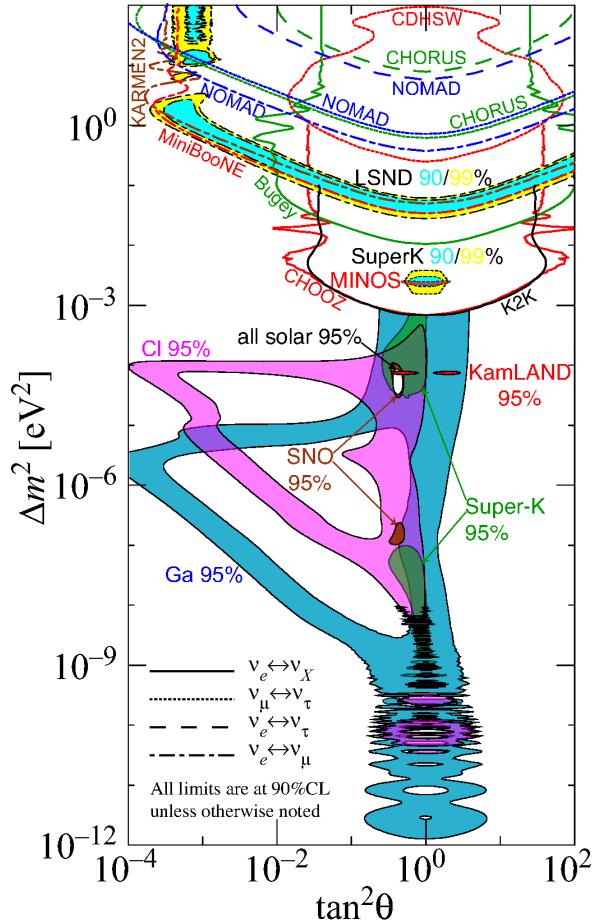


Figure 1.3.: The allowed and excluded regions of parameter space for four types of neutrino oscillation as a function of the mass splitting, Δm^2 and oscillation angle, $\tan^2 \theta$. The shaded regions show the allowed contours for the oscillation parameters as measured by the indicated experiments. The dashed lines show the regions of parameter space that have been excluded in the search for different neutrino oscillation modes.

1.4. Neutrino interaction cross-sections

Neutrino oscillation studies have now entered the precision era, with high statistics experiments measuring the oscillation parameters to the percent level. The discovery of large θ_{13} also means that experiments will have to consider the full three neutrino oscillation probability in future analyses and as a consequence will need a better knowledge of the other oscillation parameters, in particular θ_{23} .

All experiments measure the rate of neutrino interactions, which is given by the neutrino flux multiplied by the neutrino interaction cross-section. In almost all recent

Oscillation Parameters	
Parameter	Best Fit ($\pm 1\sigma$)
Δm_{21}^2 [10 $^{-5}$ eV 2]	7.58 $^{+0.22}_{-0.26}$
$ \Delta m_{31}^2 $ [10 $^{-3}$ eV 2]	2.35 $^{+0.12}_{-0.09}$
$\sin^2 \theta_{12}$	0.306 $^{+0.018}_{-0.015}$
$\sin^2 \theta_{23}$	0.42 $^{+0.08}_{-0.03}$
$\sin^2 \theta_{13}$	0.0251 ± 0.0034

Table 1.1.: The best fit values and 1σ uncertainties on the three flavour neutrino oscillation parameters, combining results from atmospheric, solar and reactor neutrino data [13]. The θ_{13} value quoted is the PDG average of the Daya Bay, RENO and Double CHOOZ results.

results, the largest source of error has come from uncertainties on the neutrino cross-sections. To achieve the necessary reduction in these uncertainties requires precise, repeated measurements of the relevant cross-sections using a range of neutrino energies, flavours and target materials.

1.4.1. Neutrino cross-section theory

Neutrinos interact with matter through the weak nuclear force, exchanging either a W or Z boson with the target nucleus. The exchange of a W boson is known as a ‘charged current’ (CC) interaction, shown in Fig. 1.4a, whilst Z boson exchange gives ‘neutral current’ (NC) interactions, in Fig. 1.4b. NC interactions are characterised by the incoming neutrino appearing in the final state, whilst CC interactions result in the production of a charged lepton whose flavour matches that of the neutrino in the initial state. For both charged and neutral current interactions there are a variety of different interaction modes, whose relative cross-sections depend upon the neutrino energy. Elastic (ES) and quasi-elastic scattering (QES) dominate at neutrino energies below 1 GeV, and describe the neutrino interacting with an entire nucleon, with this nucleon (and possibly others) being ejected from the target nucleus. Above this energy the target nucleon can be excited to a resonant state, which then decays to produce a variety of mesons. These resonant interactions (RES) give a final state composed of nucleons and mesons and are the dominant process for neutrinos with an energy of a few GeV. At about 5 GeV and

above deep inelastic scattering (DIS) has the largest cross-section, a process characterised by the neutrino scattering from an individual quark within the nucleon. DIS produces a hadronic shower as the final state due to the energy transferred to the nucleon. These three process are illustrated in Fig. 1.4a, Fig. 1.5a and Fig. 1.5b using CC interactions.

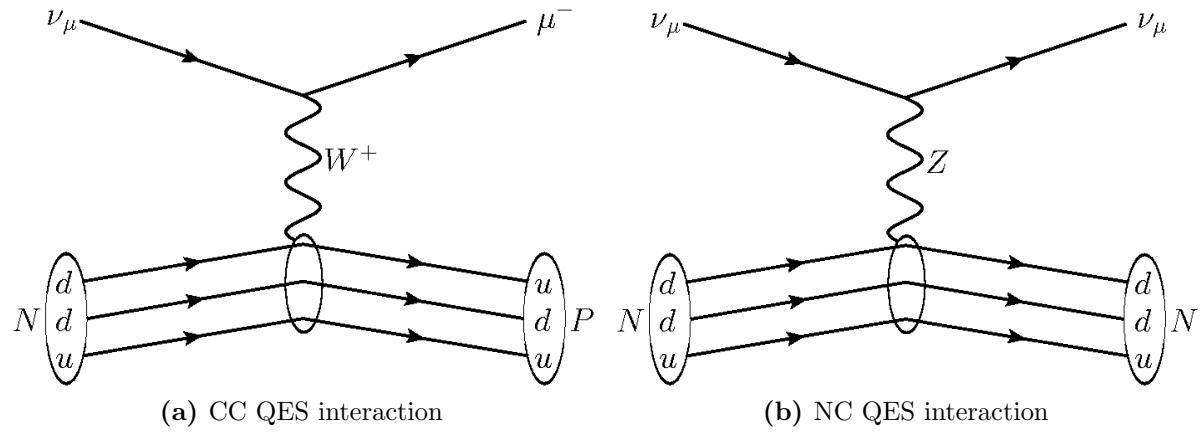


Figure 1.4.: Feynman diagrams showing QES examples of CC and NC interactions. The central ring indicates that, unlike the DIS case, the boson here interacts with the nucleon as a whole, rather than just a single quark

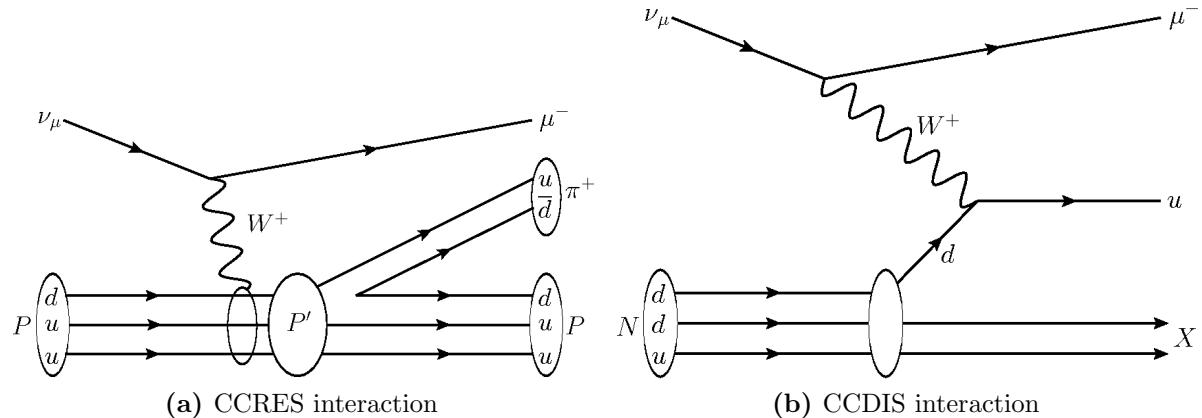


Figure 1.5.: Feynman diagrams showing (a) CC resonant pion production and (b) CC deep inelastic scattering.

CC neutrino interactions are the most common signal channel in neutrino oscillation experiments, since the lepton they produce is easy to detect and provides information about the neutrino flavour and energy. A subset of these, charged current quasi-elastic (CCQE) interactions, allow the reconstruction of the initial neutrino energy from the final state lepton kinematics alone and as such are of particular interest.

There exist adequate [17] theoretical models of these interactions, but no universal description of the transitions between the QES, RES and DIS regions. Also, the importance of nuclear effects is only recently being addressed, which have an effect on both the kinematics of the interaction itself and the possible final state particles that are observed. Unfortunately, this lack in the theory is matched by a lack of data at these energy scales. Fig. 1.6 shows measurements of the neutrino CC cross-section, split into the three different interaction modes. The large uncertainty on the experimental measurements around 1 GeV can be clearly seen.

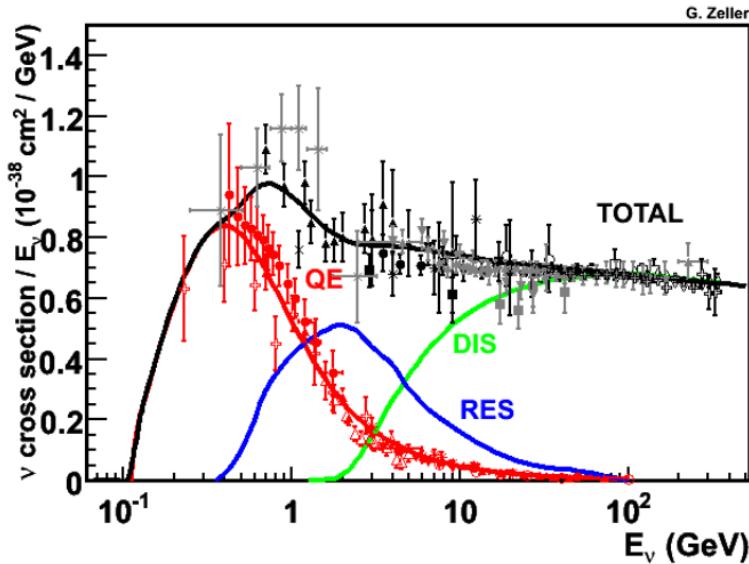


Figure 1.6.: A summary of the current measurements of ν_μ CC cross-sections, coming from a range of experiments looking at a range of neutrino energies and interaction modes [17]. The cross-section is well measured at neutrino energies above 10 GeV but shows large uncertainties in the 1 GeV region.

The paucity of data in this energy region is now being addressed by a number of experiments, both those dedicated to cross-section measurements and the near detectors of the long baseline experiments. Recently, the MiniBooNE experiment produced a measurement of the muon neutrino CCQE cross-section, using the world's largest sample of ν_μ CCQE events at 1 GeV [18]. At the same time the collaboration used this event sample to extract the value of M_A^{QE} , the quasi-elastic axial mass. M_A^{QE} is a parameter in the CCQE interaction models that has a direct effect on the interaction cross-section, and comes from the parameterisation of the nucleon structure. This parameterisation gives three dominant nucleon form factors: two vector form factors that have been well measured by electron scattering data and a single axial-vector form factor $F_A(Q^2)$. A

dipole form is used for the axial-vector form factor in the majority of neutrino interaction models, which gives $F_A(Q^2)$ as

$$F_A(Q^2) = \frac{F_A(0)}{(1 + Q^2/M_A^2)^2}. \quad (1.8)$$

In Eqn. 1.8 Q^2 is the four-momentum transferred from the neutrino to the target nucleon and $F_A(0)$ is the axial-vector form factor at $Q^2 = 0$, which has been well determined by neutron beta-decay. Eqn. 1.8 shows that M_A^{QE} affects the Q^2 distribution of neutrino interactions, a distribution that can be calculated from the final state lepton kinematics of CCQE events. Using their large, pure sample of CCQE events MiniBooNE was able to fit the Q^2 distribution of these interactions and extract a value of M_A^{QE} . Before MiniBooNE, the best measurement of M_A^{QE} came from the NOMAD experiment [19], whose CCQE sample covered a neutrino energy range of 3–100 GeV. NOMAD found that $M_A^{\text{QE}} = 1.05 \pm 0.02(\text{stat}) \pm 0.06(\text{syst})$ GeV, in good agreement with the then current world average of $M_A^{\text{QE}} = 1.026 \pm 0.021$ GeV. In contrast, MiniBooNE measured $M_A^{\text{QE}} = 1.35 \pm 0.17$ GeV at an average neutrino energy of 800 MeV. The MiniBooNE result did however agree well with results published by the K2K [20] and MINOS [21] experiments. The MiniBooNE and NOMAD data is shown in Fig. 1.7, which plots the measured CCQE cross-section as a function of reconstructed neutrino energy and shows the results from the cross-section model using the two values of M_A^{QE} .

In the MiniBooNE paper they say that they are making a measurement of $M_A^{\text{effective}}$ rather than M_A^{QE} , which is an important distinction. All current oscillation experiments use carbon, oxygen or argon nuclei as the neutrino target rather than free nucleons. This means that any cross-section measurement is measuring an effective M_A^{QE} in a particular nuclear environment, rather than the fundamental parameter. In many cases, the nuclear environment is modelled in the cross-section simulations as a relativistic Fermi gas (RFG) [23], which provides the nucleon binding energy and Fermi momentum. The combination of the RFG model and dipole form factor is simple and able to describe the data well at higher neutrino energies, but the recent cross-section measurements from the above experiments show that there is a strong tension between the theory and the data around 1 GeV. These simple theoretical models are unlikely to accurately describe neutrino interactions and Fig. 1.6 shows that more experimental evidence is needed to truly understand what is occurring.

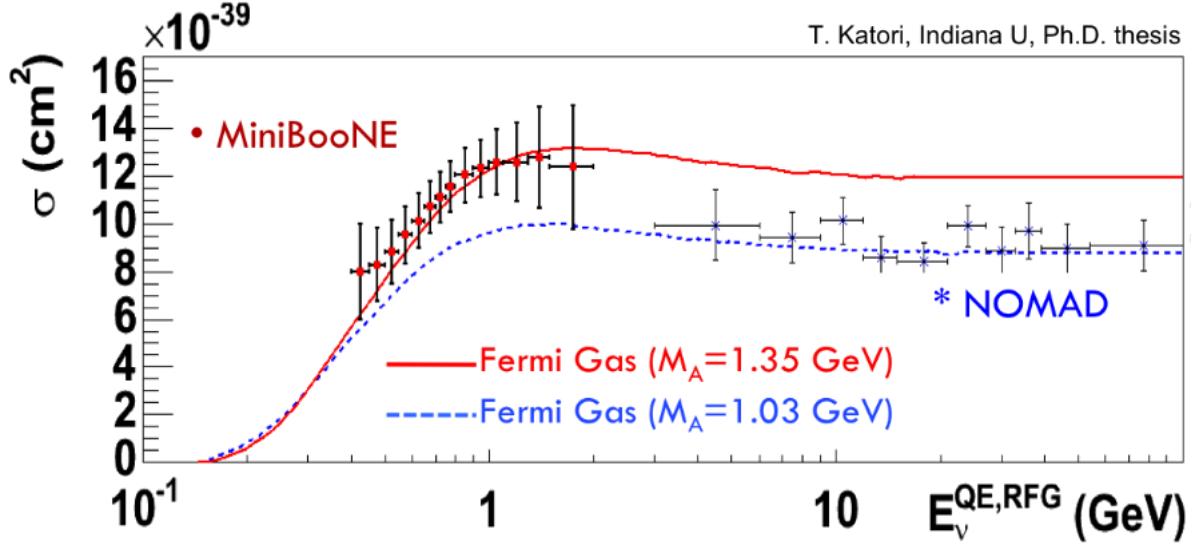


Figure 1.7.: A comparison between the MiniBooNE and NOMAD ν_μ CCQE cross-section measurements as a function of neutrino energy. The red line shows the predicted cross-section using $M_A = 1.35$ GeV whilst the blue shows the cross-section using the value of M_A measured by NOMAD [22].

1.4.2. Neutrino interactions on heavy nuclei

As mentioned above, the current neutrino interaction models do not accurately reflect what is seen in the data. The RFG model has been tested against high precision electron scattering data and seems to perform well at calculating the nucleon Fermi momentum, binding energy and QES cross-section [24]. On the other hand, a more recent paper [25] has shown that the RFG model does not model the electron QES cross-section particularly well in the region where the electron does not transfer much momentum to the nucleon. The percentage difference observed at the maximum cross-section is shown in Fig. 1.8, demonstrating the poor agreement and also highlighting the worsening of that agreement as the atomic mass of the nucleus increases.

The T2K collaboration uses two main neutrino interaction generators, NEUT [26] and GENIE [27], and in both it is only the RFG model that is affected by the atomic mass of the target nuclei. The RFG models used by the generators both provide the nucleon Fermi momentum (p_f) and binding energy (ϵ_B), which depend upon both the atomic mass (A) and atomic number (Z) of the target nucleus, and it is these parameters that affect the calculated cross-sections. As an added difficulty, p_f and ϵ_B also modify the secondary interactions of any hadrons liberated from the nucleus, further changing what is seen by experiments.

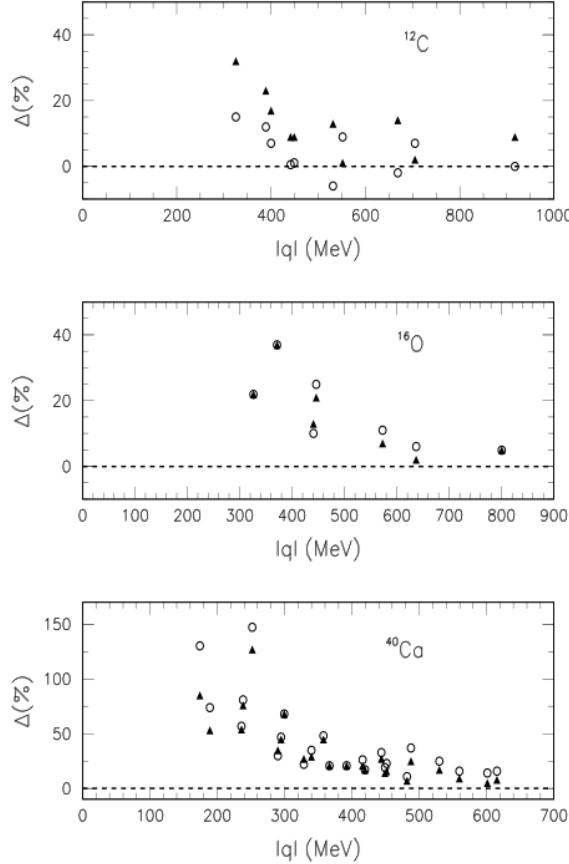


Figure 1.8.: The difference between the calculated and measured values of the electron-nucleus QES cross-section for carbon, oxygen and calcium as a function of three-momentum transfer $|q|$. The filled triangles come from using the relativistic Fermi gas model whilst the open circles correspond to using the plane wave impulse approximation [25].

There are currently only two experimental measurements of the interaction cross-section for neutrinos on lead nuclei, both of which are at neutrino energies above 2 GeV and are measurements of the ratio of the cross-section on lead to that on carbon. The first was performed by the CHORUS collaboration [28] using the CERN SPS neutrino beam, who were investigating the Z/A dependence of the total CC cross-section. Their results are shown in Fig. 1.9, where the measurement on lead has $Z/A = 0.4$ and the other points, from left to right, come from measurements made on iron, marble and polyethylene respectively. The CERN SPS beamline produced a wide-band neutrino beam whose average energy was 27 GeV, so the event sample was dominated by DIS interactions.

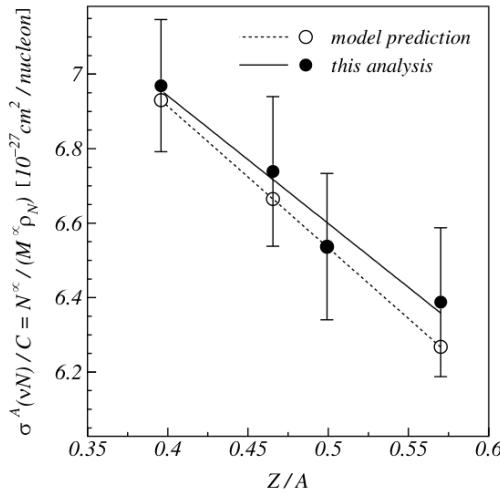


Figure 1.9.: Measurements of the relative CC inclusive cross-sections for lead, iron, marble and polyethylene plotted as a function of Z/A . The black points are the measured result whilst the open circles represent the values expected from their cross-section model. The solid and dashed lines are the best linear fit to the data and model points [28].

Fig. 1.9 also shows the expected linear scaling of the cross-section in Z/A according to the cross-section model they describe, and it can be seen that the data points are in agreement with this. This is not unexpected, since Fig. 1.6 shows that current cross-section theory is a good description of the observed data in the DIS region.

The second, much more recent measurement came from the MINERνA collaboration [29] in 2012. The MINERνA detector sits in the NuMI beamline and aims to measure neutrino cross-sections and final state effects across a range of nuclear targets. The detector itself abuts the MINOS near detector, which provides a measurement of the momentum of muons produced in CC interactions. This is then combined with a calorimetric measurement of the remaining energy deposited in the MINERνA detector to give an estimate of the initial neutrino's energy. MINERνA has presented the ratio of the ν_μ CC inclusive lead cross-section to that on plastic scintillator alongside a comparison of the number of events seen in data compared to the number in their Monte Carlo (MC) simulation. This result is shown in Fig. 1.10, with data covering the neutrino energy range from 2–22 GeV. Across most of this range there is no significant disagreement between both data and MC and between the predicted cross-section ratio and observed rate, though there is some indication of a difference at the lowest neutrino energies.

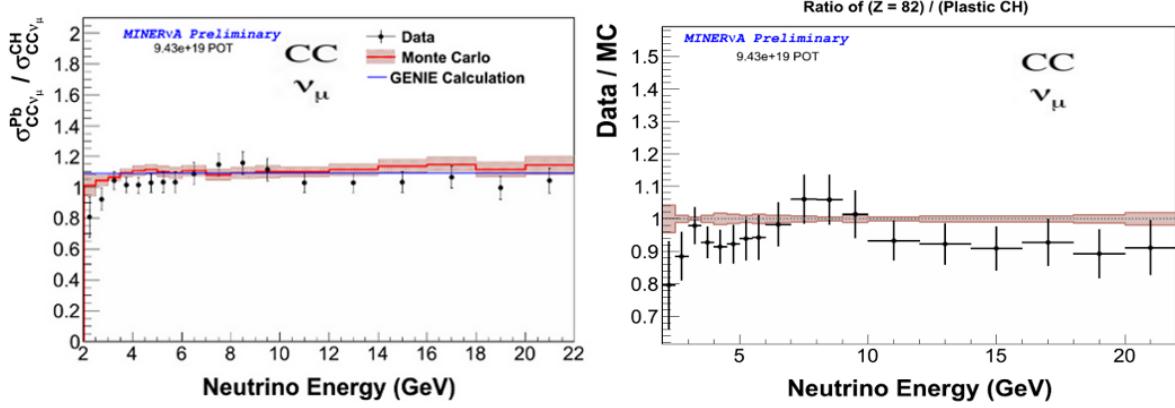


Figure 1.10.: Measured ratio of the ν_μ CC inclusive cross-section on lead to that on plastic (left) between 2 GeV and 22 GeV, alongside the measured ratio in data divided by that present in the Monte Carlo (right) across the same energy region [29].

It is widely understood that the next generation of neutrino oscillation experiments need an improved understanding of neutrino cross-sections, and that to achieve this more cross-section measurements are needed across a range of nuclear targets. This thesis describes the measurement of the ν_μ CC cross-section on lead nuclei using the electromagnetic calorimeters from the near detector of the T2K experiment. This result will be the first absolute measurement of this cross-section and the first measurement at all in the 1 GeV neutrino energy region.

Chapter 2.

The T2K experiment

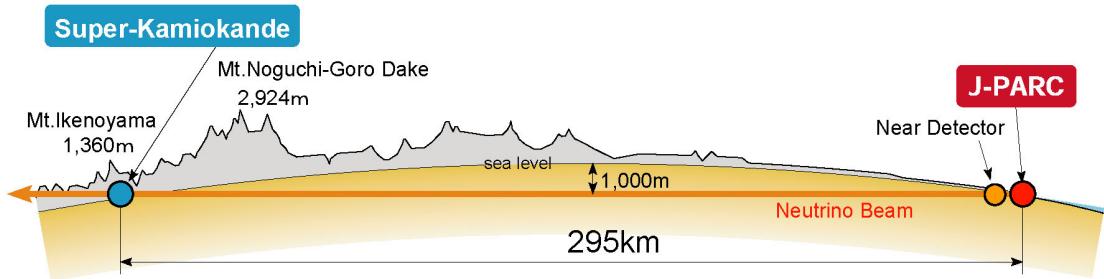


Figure 2.1.: The T2K experiment.

The Tokai-to-Kamioka (T2K) experiment is a long baseline neutrino oscillation experiment that is using a ν_μ beam to precisely measure the parameters governing ν_μ oscillation. The experiment's primary goals are to observe the appearance of ν_e in the ν_μ beam, providing information on the parameter θ_{13} , and to measure the parameters governing ν_μ disappearance, $\sin^2 2\theta_{23}$ and Δm_{23}^2 , to a precision of $\delta \sin^2 2\theta_{23} \sim 0.01$ and $\delta \Delta m_{23}^2 \sim 10^{-4} \text{ eV}^2$ respectively.

The experiment has three main components; a neutrino production beamline situated at the Japanese Proton Accelerator Research Complex (J-PARC) in Tokai, a suite of near detectors, also at J-PARC but sited 280 m away from the neutrino beamline target station and the Super-Kamiokande water Cherenkov detector 295 km distant in Kamioka. The T2K experiment is the first to employ the off-axis technique, whereby the SK detector and one of the near detectors are both 2.5° off the neutrino beam axis. This modifies the neutrino energy spectrum they experience, as shown in Fig. 2.2, increasing the neutrino flux at the oscillation maximum for ν_μ disappearance and reducing the flux

at higher energies, resulting in a narrow band beam. This increases the expected number of signal events at SK and also reduces the expected background events due to high energy, unoscillated neutrinos.

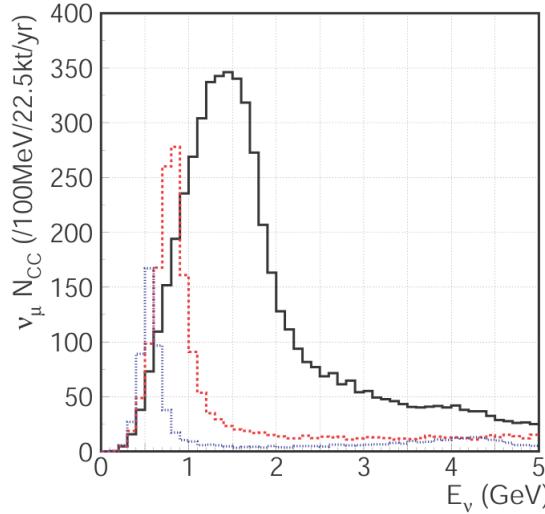


Figure 2.2.: Simulated neutrino energy spectra for CC interactions at SK. The solid histogram is the expected spectra when SK lies one degree off the neutrino beam axis while the dashed and dash-dotted histograms correspond to two and three degrees off-axis respectively. [30]

In this chapter a short description of each section of the T2K experiment is presented, with a more in-depth look at the ND280 detector that is used in this analysis. There is also an overview of the latest neutrino oscillation results from the T2K experiment and a discussion of how this study can lead to improvements in these measurements.

2.1. The T2K beamline

J-PARC is a newly constructed accelerator facility consisting of three accelerators: a linear accelerator (LINAC), a rapid-cycling proton synchrotron (RCS) and the main ring synchrotron (MR). The accelerator complex is shown in Fig. 2.3, with the three accelerators marked alongside the different experimental halls. At the LINAC H^- ions are accelerated up to 180 MeV before being converted to H^+ ions by charge-stripping foils at the point of injection to the RCS. The RCS accelerates these ions up to 3 GeV with a 25 Hz cycle rate. The RCS holds two bunches per cycle, which are either injected into the MR or supplied to the muon and neutron beamline for use by the Material and Life Sciences facility also at J-PARC. The MR holds up to eight proton bunches which

can be extracted at two points. The slow extraction point provides high energy protons to experiments in the hadron hall, whilst the fast extraction deflects the proton beam to the neutrino beamline for use by T2K. At the fast extraction point the entire eight bunches are extracted within a single turn by five kicker magnets, forming a single beam spill. Each bunch is confined within a 100 ns window, and each bunch within a spill is separated by 500 ns. This timing information is linked to SK through a GPS system, allowing beam neutrinos to be selected based on timing alone.

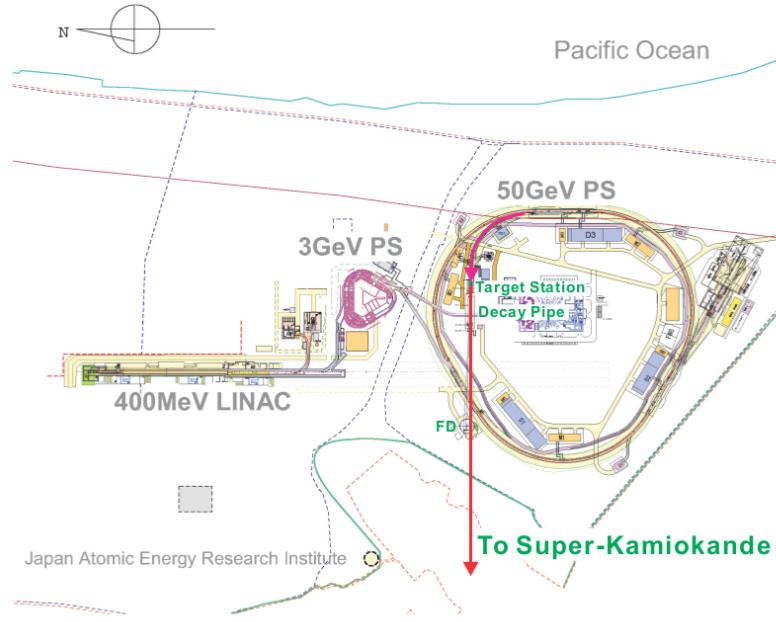


Figure 2.3.: The J-PARC accelerator complex, with the three main accelerators highlighted. The accelerator labelled as the 3 GeV PS is the RCS and the accelerator labelled as 50 GeV PS is the MR [30].

A schematic of the neutrino beamline is shown in Fig. 2.4. The beamline is composed of two sections, the primary beamline, which transports the proton beam from the MR and points it toward Kamioka, and the secondary beamline. As in all neutrino beam experiments, stable and well tuned beam operation is vitally important to produce both a high intensity and consistent neutrino beam. To do this, the primary beamline contains a large number of beam monitors to measure the profile, intensity, position and beam loss of the proton beam. These are described in more detail in [31]. The secondary beamline contains the proton target station followed by a 96 m decay volume and ends in a beam dump and muon monitor. The proton target station contains the proton target itself, an optical transition radiation monitor (OTR) and three magnetic focussing horns. The OTR monitors the proton beam direction and profile directly in front of the target, using

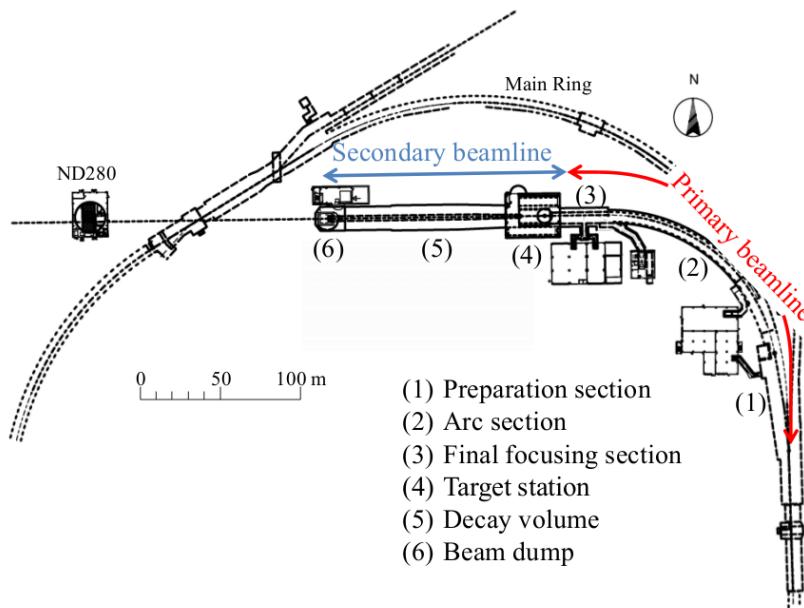


Figure 2.4.: An overview of the T2K neutrino beamline, indication the components of the primary and secondary beamlines [31].

the transition radiation produced by the protons as they pass through a titanium alloy film. The target is a 1.9 interaction length (91.4 cm long), 2.6 cm diameter graphite rod, surrounded by a 2 mm thick graphite sleeve and a 0.3 mm thick titanium casing and is cooled by helium gas flowing between each layer. Protons interact within the graphite, creating large numbers of secondary mesons, the majority of which are pions. The target sits within the inner conductor of the first magnetic horn, which collects these mesons. The second and third horns then focus these mesons to produce a meson beam. When operating at their design current of 320 kA (they are currently running at 250 kA) a field of 2.1 T is created inside the horns and the subsequent beam focussing increases the neutrino flux at SK by a factor of 16 compared to operation at 0 kA.

As the secondary mesons travel through the decay volume they produce neutrinos, primarily through the following processes

$$\begin{aligned}\pi^+ &\rightarrow \mu^+ + \nu_\mu, \\ K^+ &\rightarrow \mu^+ + \nu_\mu, \\ K^+ &\rightarrow \pi^0 + e^+ + \nu_e,\end{aligned}$$

and

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu.$$

This results in a neutrino beam largely composed of muon neutrinos from the decay of positive pions, with a small contamination from $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$, and a peak energy of approximately 600 MeV.

At the end of the decay volume is the beam dump, formed from 3.17 m of graphite and 2.4 m of iron. The beam dump absorbs almost all the surviving particles at the end of the decay length, with only muons above 5 GeV managing to reach the muon monitors located behind it. Since the majority of the neutrinos in the beam were produced in two-body decays alongside a muon, recording the muon direction gives a reliable measure of the neutrino beam direction. The muon monitor is designed to measure the neutrino beam direction to an accuracy of 0.25 mrad and monitor the neutrino beam intensity to a precision of 3%.

It is not possible to record the total number of neutrinos produced in the neutrino beamline. Instead, the number of protons that collide with the graphite target, called the number of protons on target (POT), is recorded. The integrated POT count over a period of time is then used as a proxy for integrated neutrino flux over that same period.

2.2. The ND280 detector suite

The ND280 detector suite sits 280 m downstream of the graphite target and contains two detectors, the INGRID on-axis detector and the ND280 detector, which sits 2.5° off-axis, on the line between the target and the far detector.

2.2.1. Multi-Photon Pixel Counter

Multi-Pixel Photon Counters (MPPCs) are used throughout both the INGRID and ND280 detectors to detect scintillation light produced by particles depositing energy in plastic scintillator. MPPCs were chosen due to both spatial constraints and their ability to operate in a magnetic field, since the ND280 magnetic field is essential for realising the detector's physics goals. The MPPCs are multi-pixel avalanche photodiodes, with 667 pixels arranged in a 26×26 square (9 pixels are lost to fit an electrode on the

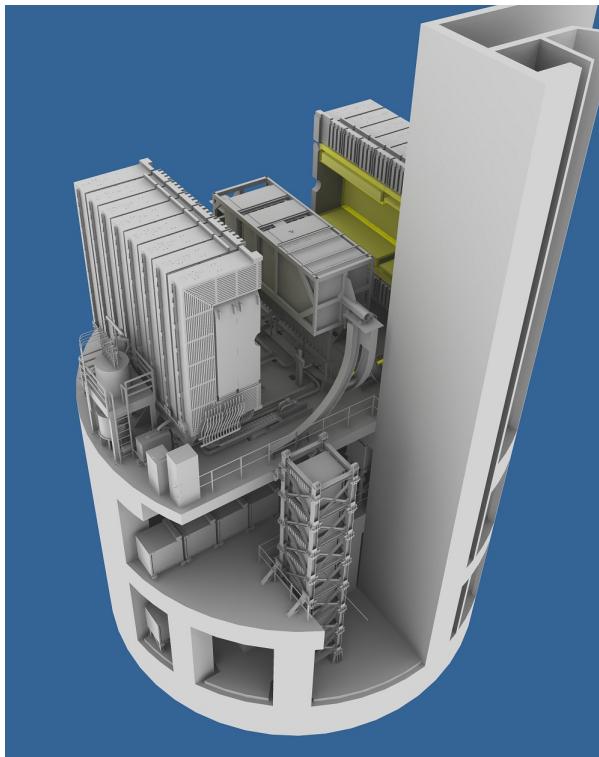


Figure 2.5.: The T2K near detector hall, showing the ND280 off-axis detector on the top level, the INGRID horizontal module one level down and the INGRID vertical module spanning the lower two levels [31].

diode) each held at between 0.8–1.5 V above their breakdown voltage, giving a gain of approximately 1×10^6 . When a photon strikes an individual pixel a Geiger avalanche is created. The single pixel then gives a binary response, either having been struck or not, that is decoupled from the number of photo-electrons created by the incident photon. The total MPPC output is then the sum of the number of pixels that fired within a given time period, which is proportional to the number of photons observed.

The ND280 scintillating detectors are formed from bars of plastic scintillator, with a wavelength shifting (WLS) optical fibre threaded through the bar centre. Any scintillation photon created in the plastic is collected by the WLS fibre and transported to one end of the bar. The WLS fibre shifts the wavelength of these captured photons, changing their spectrum to match the spectrum the MPPCs are best able to detect. Further information and detailed performance analysis for the MPPCs can be found in [32–34].

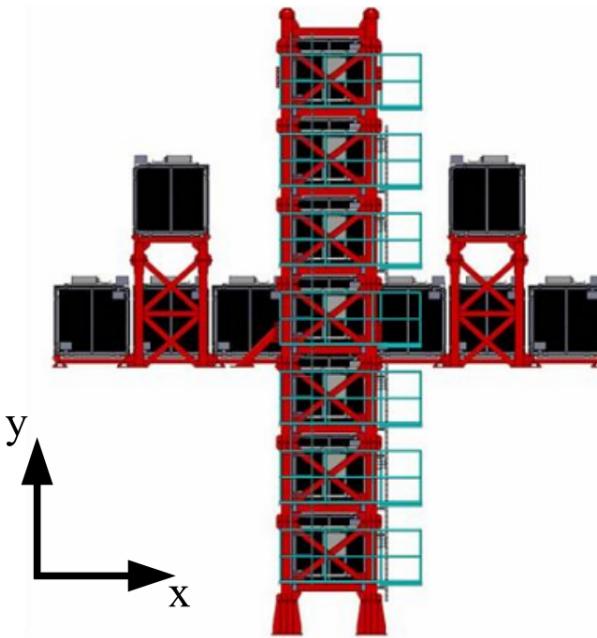


Figure 2.6.: The INGRID detector, showing the vertical, horizontal and off-axis modules [31].

2.2.2. The INGRID detector

The Interactive Neutrino GRID (INGRID) detector is composed of 16 identical modules, arranged as shown in Fig. 2.6. These modules are spaced so that each one samples a different part of the neutrino beam, with the beam centre passing through the centre of the cross. With this constraint, INGRID can monitor an off-axis angle from 2° to 2.5° , where deviation outside this would lead to a less precise measurement of the beam direction. Each module is made from 9 sheets of iron and 11 tracking scintillator planes, with each scintillator plane containing 24 horizontal and 24 vertical bars of plastic scintillator, read out by MPPCs. Each module has a total iron mass of 7.2 t so that, at the nominal beam intensity, there are enough neutrino interactions in each module to monitor the beam direction on a day-by-day basis. Using the number of neutrino interactions in each module the beam centre is known to a precision of better than 10 cm, which equates to 0.4 mrad. An example neutrino event is shown in Fig. 2.7. The two off-axis modules are used to check the axial symmetry of the neutrino beam. Each INGRID module is surrounded on the top, bottom and both sides by veto planes, composed of 22 scintillating bars. These planes are used to veto charged particles entering the module from the outside.

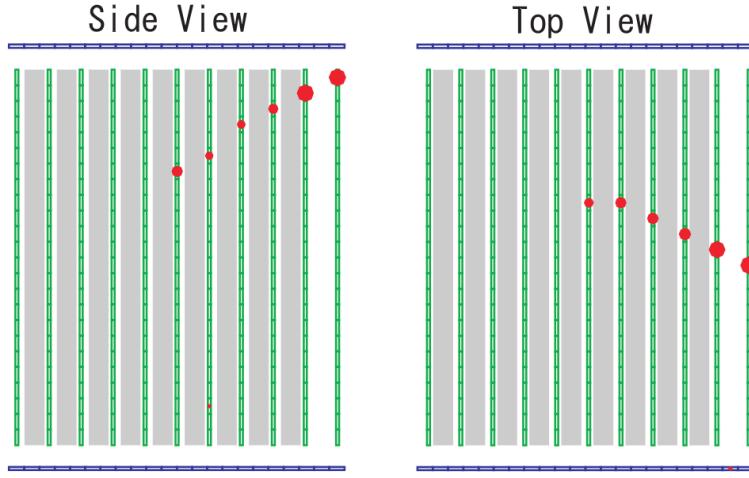


Figure 2.7.: An example INGRID neutrino event. The green cells indicate the position of scintillator bars, the grey areas correspond to the iron sheets and the blue cells show the veto scintillator planes. Here a neutrino entered from the left of the module, interacted in the central iron plane producing a charged particle that deposited energy in the remaining scintillator planes, indicated by the red circles [31].

The proton module is an extra module that has been added to INGRID, placed in the centre between the vertical and horizontal arms of the cross in Fig. 2.6. This module does not contain any iron planes, and uses smaller scintillating bars to give much finer granularity detector. This module is designed to measure both the proton and muon produced in the charged current quasi-elastic interaction channel for comparisons to the beamline MC and the current neutrino generators.

2.2.3. The ND280 detector

The second of the two near detectors is the ND280, situated above and to one side of the INGRID detector, 2° off the neutrino beam axis. The task of the ND280 is to measure the flux, flavour content and energy spectrum of the neutrino beam seen by the far detector. Alongside this, the detector is ideally placed to precisely measure a number of neutrino interaction cross sections in the 1 GeV energy range. As shown in Fig. 2.8, the ND280 detector comprises of a number of sub-detectors enclosed within the UA1 magnet. This magnet produces a 0.2 T field in the centre of the detector, allowing the reconstruction of charge and momentum from the curvature of particles produced in neutrino interactions. Interleaved within the magnet yoke is the side muon range

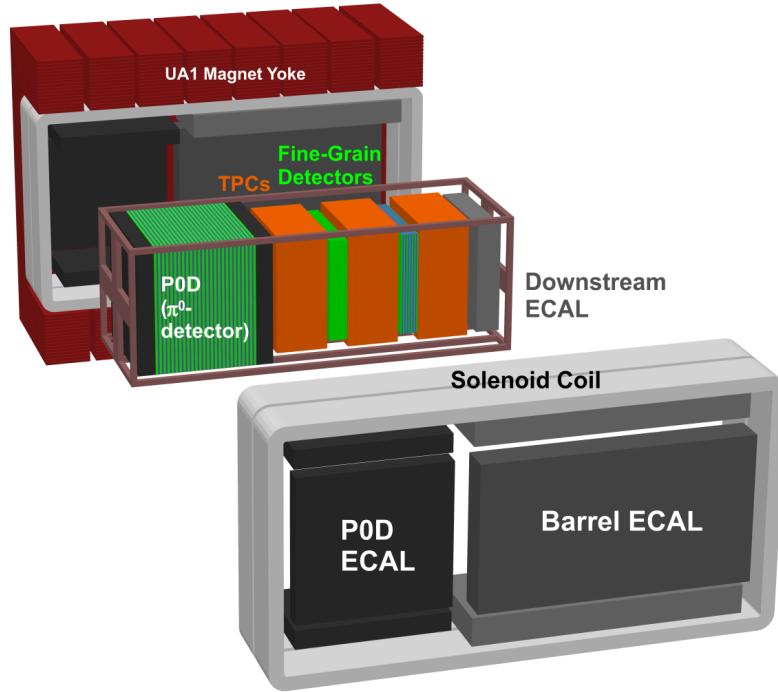


Figure 2.8.: An exploded view of the ND280 detector, showing the positions of the separate sub-detectors [31].

detector, which surround the electromagnetic calorimeters that are affixed to the inner face of the magnet. These in turn contain the tracker, which consists of the π^0 detector, three time projection chambers, two fine grained detectors and the downstream ECal.

The fine grain detectors

The fine grain detectors (FGDs) are the primary neutrino target of the ND280 and made from layers of scintillator bars. These bars, 9.61 mm wide by 9.61 mm deep and 1864.3 mm long, are made from polystyrene scintillator. Each bar has a thin reflective coating of TiO_2 and a WLS fibre through the centre. One end of the fibre is mirrored by the vacuum deposition of aluminium, the other connects to an MPPC and the associated readout electronics. The first FGD has 30 layers of 192 bars, whose orientation alternates between horizontal and vertical with each new layer. The scintillator provides a 1.1 t target mass and the alternating bars allow the reconstruction of particles produced in neutrino interactions. The FGD also has a light injection system (LIS) mounted alongside the readout electronics, which is used to directly calibrate the detector with the magnet

closed. This LIS flashes the exposed end of the WLS fibre, testing the MPPC saturation, charge response and non-linearity.

The second FGD uses the same materials and electronics as the first, yet has only 14 layers of scintillator grouped together in pairs of alternate orientation. These 7 pairs are separated by 6 layers of hollow, thin walled, corrugated polycarbonate 2.5 cm thick sheets, one end of which has been sealed with polyurethane sealant. These layers are filled with water and kept below atmospheric pressure, so that if a leak should develop air is sucked into the layer rather than water leaving and entering the FGD. By comparing the neutrino interaction rates in the two FGDs a measurement of the neutrino interaction cross-section on water can be obtained and compared directly to that measured at the far detector.

Unlike the other scintillating detectors, the FGD MPPCs are readout by customised ‘AFTER’ ASICs boards. These store the digitised photosensor waveforms at a rate of 50 MHz and have been designed to allow continuous readout from the detector during a beam spill. The continuous readout ensures that the detector registers Michel electrons from muon decay and combined with the MPPC waveforms gives a timing resolution of 2 ns [35].

As mentioned above, the FGDs are the principle neutrino target of the ND280 detector. They are completely active, allowing accurate reconstruction of neutrino interaction vertices and measurements of vertex activity, which are necessary when looking for exclusive final states. The FGDs also have the best timing resolution, and are used as the base component when particles tracks are reconstructed through multiple sub-detectors.

The time projection chambers

Either side of both FGD detectors are the three time projection chambers (TPCs) which provide high resolution tracking of charged particles. Combined with the ND280 magnetic field, the TPCs provide charge identification and momentum measurements. Finally, the ionisation left by these particles can be combined with their measured momentum to separate the different types of particle very effectively.

An individual TPC consists of an inner and outer box, with the outer box measuring 2.3 m by 2.4 m by 1.0 m. The inner box is filled with an argon based drift gas, Ar:CF₄:iC₄H₁₀ (95:3:2), while the outer box is filled with CO₂ as an insulator. The inner box walls have a strip pattern machined onto their copper coating, which produces a

uniform electric field when combined with the central cathode. A simplified diagram of a TPC is shown in Fig. 2.9.

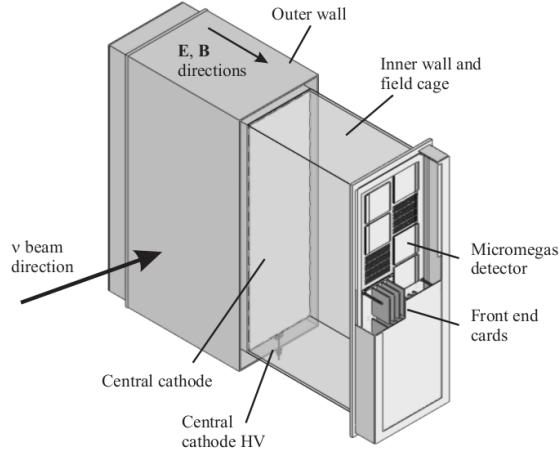


Figure 2.9.: Simplified diagram of a TPC, showing the micromegas detectors and electronics, the central cathode and neutrino beam direction [31].

When a charged particles passes through the detector, it creates ionisation electrons in the TPC gas. In the electric field, which is precisely aligned with the ND280 magnetic field, these electrons drift away from the cathode and towards one of the TPC readout planes. These readout planes are formed from micromegas detectors, which first amplify then measure the ionisation produced. The micromegas pads are segmented in the vertical and horizontal direction, which is combined with the measurement time and the known ionisation drift velocity to produce a complete 3D image of the charged particle tracks within the TPC. From this image the momentum of the particle is calculated from the track curvature and the known magnetic field. To perform particle identification, the track is split into segments, each of which corresponds to a single column of micromegas pads. For each segment the energy loss per unit length is calculated, and a truncated mean of all the segments used as the final result. A plot of this energy loss against measured particle momentum is shown in Fig. 2.10, which clearly shows the differing response to different particle types.

The π^0 detector

The π^0 detector (PØD) is the most upstream detector in the central tracker region of the ND280. It is formed from layers of scintillator bars, brass and lead sheets and HDPE water target bags. These bags can be emptied, allowing the determination of

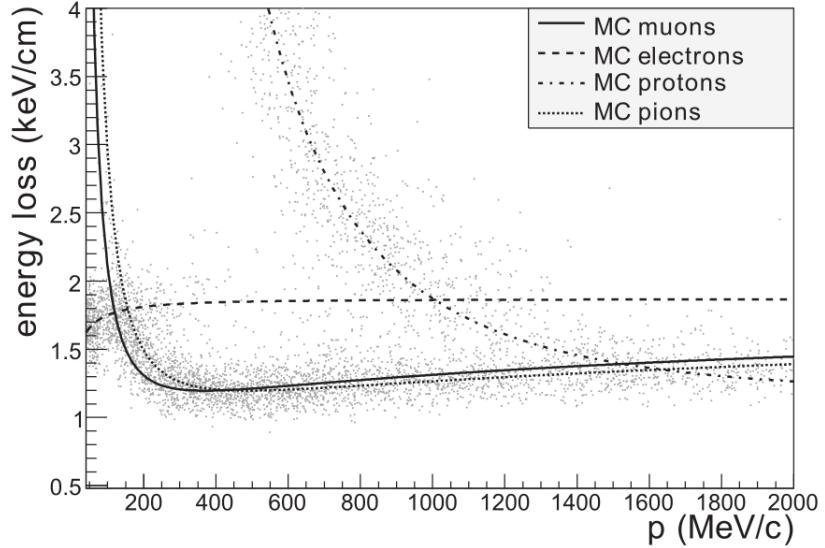


Figure 2.10.: Data points show one TPC’s measurements of the energy loss and momentum of particles produced in neutrino interactions in the ND280 detector. The overlaid curves show the expected values for different particle species, calculated using Monte Carlo simulations [31].

neutrino interaction cross sections on water using a subtraction method. As in the FGD, the scintillator bar layers are grouped in pairs of opposite orientation, providing 3D reconstruction of charged particle tracks. The primary purpose of the PØD is to measure the production cross section of π^0 particles through the neutral current interaction $\nu_\mu + N \rightarrow \nu_\mu + N + \pi^0 + X$ on a water target. This interaction is one of the major backgrounds to the ν_e appearance measurement at the far detector, so constraining it with the ND280 will lead to a much improved result. To do this, the PØD has a central water target region, composed of scintillator bars, brass sheets and water bags, capped at either end by calorimetry regions that contain only lead sheets and scintillator layers. This is shown in more detail in Fig. 2.11. The calorimetry sections improve the containment of electromagnetic showers and also provide veto regions to remove entering particles. The PØD uses the same scintillator, WLS fibre and MPPCs as the other detectors of the ND280, with the exception that the scintillator bars used in the PØD have a triangular cross section. Each pair of scintillator layers has 134 vertical bars and 126 horizontal bars, each of which is readout at a single end by an MPPC, in the same manner as the FGD.

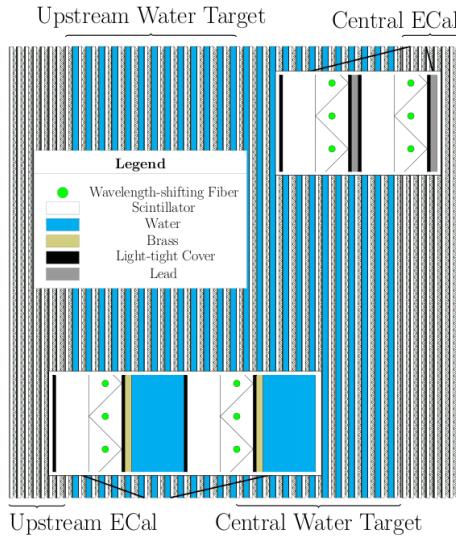


Figure 2.11.: A schematic of the PØD sub-detector showing the triangular cross-section scintillator bars, the ECals and water target, and the different metallic absorber sheets [31].

The electromagnetic calorimeters

The electromagnetic calorimeters (ECals) have three main sections —the downstream ECal (DsECal), barrel ECals and PØD ECals as shown in Fig. 2.8. The ECals are tracking, sampling calorimeters that give near-hermetic coverage for particles entering or leaving the inner tracking region of the ND280. The ECals were built to both help identify and measure the energy of photons created in the inner detectors and in particular to aid the reconstruction of π^0 particles from interactions in the FGD detectors. Alongside this the ECals provide complementary particle identification information to the TPCs and help with the full reconstruction of events.

Each ECal module has a near identical design, composed of alternating layers of scintillator bars and lead sheets. All the scintillator bars have a $4.0\text{ cm} \times 1.0\text{ cm}$ cross section, with a central $2.0\text{ mm} \times 3.0\text{ mm}$ elliptical hole running through the bar centre, and were extruded at the same Fermilab facility. They are formed from polystyrene doped with 1% PPO (2,5-DiPhenylloxazole) and 0.03% POPOP (1,4-Bis(5-phenyl-2-oxazolyl)benzene). The PPO is the primary scintillator, the photons from which cause the secondary scintillation of the POPOP, which acts as a wavelength shifter to produce photons with an emission peak of 420 nm. This emission is matched to the absorption spectrum of the WLS fibre threaded down the centre of the bar, which is centred at a wavelength of 430 nm. The WLS fibre transports the captured photons to an MPPC at

one end of the bar, shifting their wavelength in the process as described in 2.2.1. Each ECal module is sandwiched between two carbon fibre plates and held in an aluminium frame, which gives structural support and secures each module in place. On the side wall of each ECal are thinner aluminium plates that hold the readout electronics, water cooling, dry-air circulation systems and the power distribution bars.

The DsECal has 34 layers of scintillator bars, interleaved with 33 layers of 1.75 mm thick lead absorber to give $10.6 X_0$ thickness and a total mass of 4.80 t. Each scintillator layer contains 50 2.04 m long bars, each of which is readout at both ends by an MPPC. The bars alternate orientation by 90° between each layer, with bars in the most upstream layer running horizontally, then vertically in the next layer and so on. This results in each ECal module having two 'views', each corresponding to a signal in either the horizontal or vertical layers. Information from both views can then be combined to create 3D reconstructed objects, and is also used individually to provide more information about the event.

There are six barrel ECal modules which are symmetric about the vertical plane through the centre of the ND280. Each side has two identical top and bottom modules and a larger side module, each of which contains 31 scintillator layers and 30 layers of 1.75 mm thick lead. This equates to a thickness of $9.7 X_0$ and was enforced by the size of the UA1 magnet. In the barrel modules the scintillator bars are organised into layers parallel or perpendicular to the neutrino beam direction, with the layer closest to the centre of the ND280 perpendicular to the beam. This gives the same two views and reconstruction ability as the DsECal, the only difference being that the perpendicular bars have only one instrumented end, the other having been mirrored using the vacuum deposition of aluminium. The two side barrel modules have 96 bars in each perpendicular layer and 57 in each parallel layer, giving a total mass of 9.21 t. The four top and bottom modules also have 96 bars in their perpendicular layers but only 38 bars in the layers parallel to the beam. Each of these modules has a mass of 6.62 t.

Finally, there are six PØD ECal modules, whose layout echoes that of those in the barrel ECal. These modules surround the PØD detector and are required to detect photons that escape the PØD without showering and to help capture showers that are only partly contained within the PØD. They also act as vetos against entering particles and provide some particle identification, particularly for minimally ionising particles. Unlike the other ECals, the PØD ECal modules have only six layers of scintillator bars, each lying parallel to the neutrino beam. These are interspersed with five layers of 4 mm thick lead, giving a total thickness of $3.6 X_0$. In the four top and bottom modules each

scintillator layer contains 38 bars, whilst the two side modules contain 69 bars per layer. Each bar is 2.34 m long and readout at a single end by an MPPC.

The DsECal was constructed in 2008 and shipped to CERN in 2009 for commissioning and cosmic ray and test beam data taking. This was used to characterise the detector response to electrons, pions and muons covering a momentum range of 0.3 GeV - 5.0 GeV, produced by the CERN T9 beamline. The DsECal was then installed at the ND280 site for the first T2K physics run, with the barrel and PØD ECals installed in October 2010 for the second physics run.

The side muon range detector

The side muon range detector (SMRD) is created from 440 scintillator modules, placed in the air gaps of the return yoke of the UA1 magnet. The SMRD tracks muons produced at a high angle to the neutrino beam and provides measurements of their momentum. It also functions as part of the cosmic trigger for the ND280 and helps identify neutrino interactions that occur in the magnet iron.

The UA1 magnet is formed from 16 C-shaped flux return yokes, grouped in pairs to form 8 rings enclosing the ND280, labelled 1 - 8 from upstream to downstream. Each yoke is made from 16 iron plates, giving 15 air gaps in each yoke. The SMRD has three layers of scintillator modules on the top and bottom of each yoke and at least three layers along the sides. In yoke pair 6 there are four scintillator modules whilst in pairs 7 and 8 there are six scintillator modules, giving better detection in the forward direction. These modules are made up of either four or five scintillator counters, depending upon the size of the air gap in different parts of the flux return yoke. All the scintillator counters are readout from a single end using the same combination of WLS fibre and MPPC as the other scintillator detectors of the ND280. The SMRD modules are much larger than the scintillator bars in the rest of the ND280, so the WLS fibre follows an S-shaped groove through the module, to maximise the capture of scintillation light. This is shown in Fig. 2.12.

2.3. The far detector

T2K uses the Super-Kamiokande detector as its far detector, which is situated 295 km from J-PARC under Mt. Ikenoyama, as shown in Fig. 2.1. Built in a cavern 1 km underground,

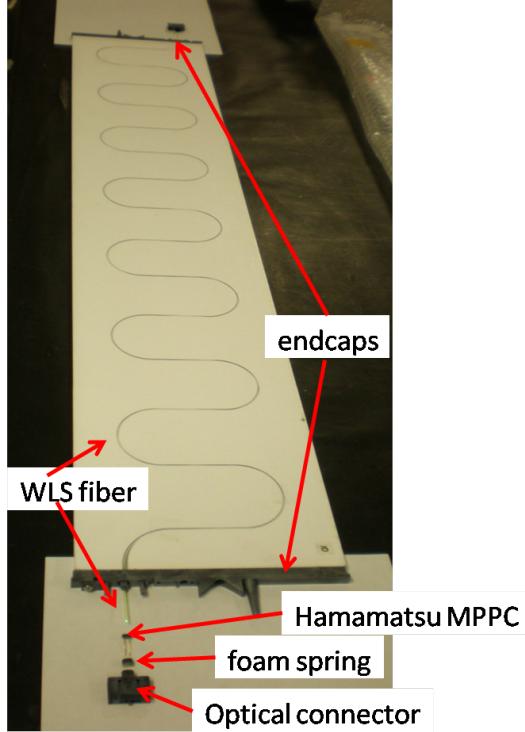


Figure 2.12.: SMRD construction image showing the WLS fibre and MPPC [31].

the detector is a cylinder containing 50 kTons of ultra pure water surrounded by ~ 13000 photomultiplier tubes (PMTs). These tubes detect the Cherenkov light produced by charged particles from neutrino interactions. This light is used to find the neutrino interaction vertex, from which the neutrino flavour and energy can be reconstructed. It is this ability to accurately measure the neutrino flavour that allows T2K to search for the appearance of ν_e in the ν_μ beam and to precisely measure the parameters governing ν_μ disappearance.

As mentioned earlier, the SK experiment is famous for providing the first evidence for atmospheric neutrino oscillations back in 1996. Since then there have been many upgrades to the detector to improve performance, but even so the detector's long running operation means that it is well understood. The energy scale has been calibrated to the percent level and the MC simulation matches data also to the percent level.

The detector itself is composed of an inner and an outer volume, separated by a stainless steel cylinder. The inner detector (ID) measures 33.8 m in diameter and 36.2 m in height and contains 11,129 inward facing 50 cm PMTs. The outer detector surrounds the ID with a depth of 2 m of water and has 1,885 outward facing 20 cm diameter PMTs on its inner wall. The OD is sparsely populated with phototubes due to its original

purpose of providing an active veto for cosmic ray muons. Even so, it provides an almost 100% efficient veto for these entering backgrounds [31]. By using the beam timing window, it is also possible to identify neutrino interactions in the OD. From volume considerations alone it is expected that there will be a similar number of interactions in the OD compared to the ID and these can be used as further checks of the MC simulation and beam flux. The ID itself has 40% PMT coverage, with each PMT having a combined quantum and collection efficiency of 20%. Cherenkov light is emitted in a cone along the direction of flight of charged particles with energies above a certain threshold. These photons reach the ID wall in the shape of a ring, which is used to reconstruct the event.

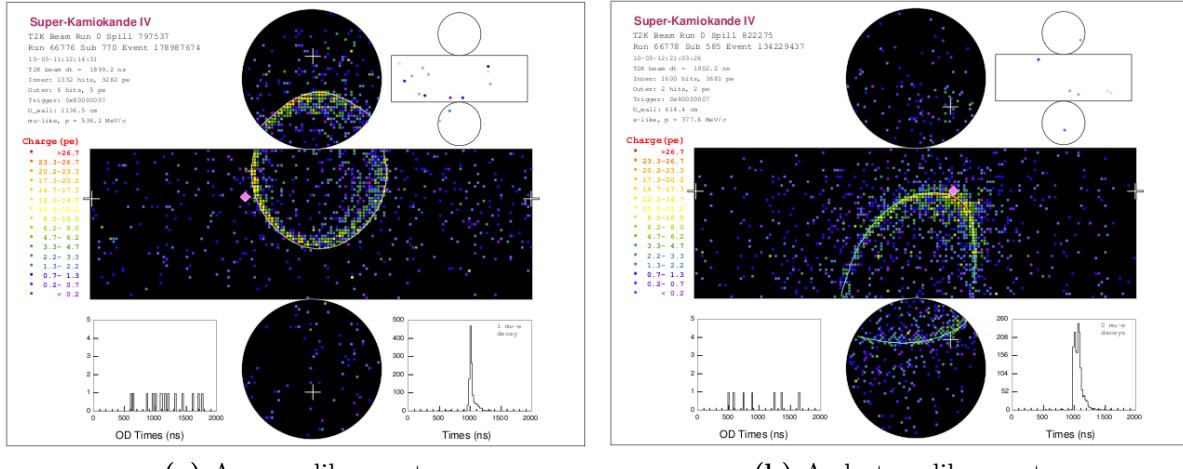


Figure 2.13.: SK ID event displays, showing the PMT hits and fitted rings for (a) a muon-like and (b) an electron-like event [31].

Charge current quasi-elastic interactions (CCQE) are the primary channel used for both ν_μ disappearance and ν_e appearance studies. Ignoring final state interactions, these events produce a proton and their respective charged lepton in the detector. Due to its mass, the muon will suffer little change in its motion when travelling through the detector. The Cherenkov light therefore will form a sharp, distinct ring when it reaches the PMTs. Electrons, in contrast, will frequently scatter and so follow a more convoluted path through the detector. In addition, at the energies expected from the T2K beam, the electron will also create electromagnetic showers. More information on the event reconstruction and analysis at SK can be found in [36]. The combination of these two phenomenon results in the Cherenkov light from the electron forming a "fuzzy" ring when it reaches the sides of the ID. Examples of the two types of Cherenkov ring are shown in Fig. 2.13. By reconstructing the number of Cherenkov rings produced and identifying the species of particle that created the ring, the SK detector can record both

the energy spectrum of ν_μ and ν_e and the number of each neutrino species in the T2K beam. From the number of interactions alone it is possible to extract the oscillation parameters that T2K has been designed to measure [14], and in future analyses with higher statistics the neutrino energy spectrum will be used to give improved results.

Chapter 3.

Event reconstruction and software

The analysis described here has been performed on the final output of the ND280 software suite, and therefore it is essential to fully understand this output. In this section the overall framework of the ND280 software is described, with the relevant sections of the event reconstruction explained in greater detail.

3.1. The ND280 software suite

The same set of software tools is used to process both the data and MC, and is formed from a number of separate packages that each perform a specific task. The main components of the software are shown in Fig. 3.1.

3.1.1. The data processing chain

Data is recorded at the ND280 detector in the MIDAS [37] format and both saved to tape and distributed across the GRID. These files are then converted into the oaEvent format, which is based upon the standard ROOT classes and written specifically for the ND280 software, by the oaUnpack package.

The events then pass through the calibration chain contained within the oaCalib package. This holds separate calibration routines for each of the ND280 sub-detectors, alongside global calibrations for common components, such as the MPPCs. The calibration constants used for this are stored in a MySQL database, with different constants stored for different experimental running periods.

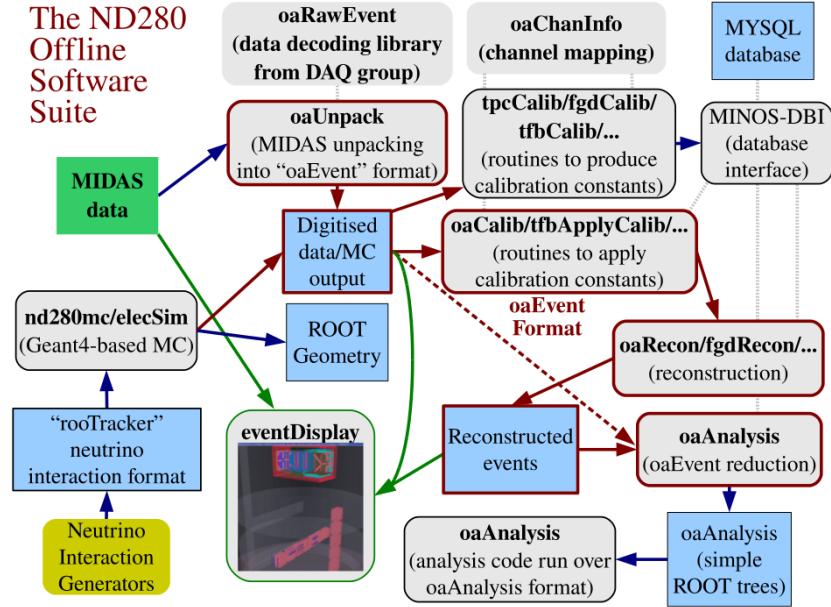


Figure 3.1.: A diagram showing the package structure of the ND280 software suite [31].

Once the data has been calibrated it is passed to the ND280 reconstruction algorithms, which attempt to reconstruct tracks, showers and neutrino interaction vertices. Each sub-detector has its own, internal, reconstruction algorithms that produce objects within that sub-detector. These are then combined by the global reconstruction, using the RECPACK [38] toolkit, to form objects spanning multiple sub-detectors. The ECal reconstruction algorithms are described in more detail later in this chapter, and more information on the other reconstruction algorithms can be found in Ref. [39].

The final part of the software chain is the data reduction and analysis package oaAnalysis. This reduces all of the data from the reconstruction to extract physics objects and their properties, removing any surplus information. This produces a much smaller file in a ROOT-based format that can be used by analysers without the need for the rest of the ND280 software suite.

3.1.2. The Monte Carlo processing chain

From the calibration steps onward the Monte Carlo (MC) processing chain is identical to that followed by the data, but using specific MC calibration constants rather than those used by the data processing. Before this, though, the neutrino beam needs to be modelled and the neutrino trajectories tracked. The neutrino interactions must then be simulated

and the paths of any subsequent charged particles calculated. Finally, the response of the detector and electronics to these particles has to be reproduced accurately.

Neutrino flux simulation

The JNUBEAM [40] simulation has been created by the T2K experiment in order to recreate the J-PARC neutrino beam. JNUBEAM simulates the proton beam impinging upon the T2K graphite target then follows the particles produced through the remainder of the T2K beamline. The decay of the particles to neutrinos is calculated and information recorded for those neutrinos whose path intersects with either the ND280 or SK detectors, producing a file of neutrino vectors.

JNUBEAM performs this using a combination of the FLUKA [41] hadron production model, Geant3/GCALOR [42], fits to data from the NA61/SHINE [43, 44] experiment and previous experimental measurements. FLUKA is used to model the primary and secondary interactions of the incident protons with the T2K target, with Geant3 taking over to propagate the subsequent particles to the decay volume, where the decays are calculated using the latest Particle Data Group [13] information on branching fractions and decay rates.

Once this initial simulation has been created, the results are tuned to match the NA61 data. The NA61 experiment has measured the charged pion [43] and charged kaon [44] spectra produced from colliding 31 GeV protons with a graphite target. The experiment has collected data using both a thin graphite target, with a thickness of 4% of a nuclear interaction length, and with a replica of the T2K target. In the analysis presented here only the thin target data has been used to tune the MC, but the results from the replica target will be used in future versions of the T2K flux simulation. The flux prediction at the ND280 is shown in Fig. 3.2a, and a plot showing the effect of the kaon tuning on the ν_μ and ν_e flux at the ND280 is shown in Fig. 3.2b. The tuning is applied both as an input to the MC and as a weighting of events applied during analysis. This allows the latest data to be included in the flux tuning without requiring a complete MC processing whenever more data becomes available. The MC used in this analysis was created using the 11a neutrino flux simulation, with the 11b version 3.2 flux tuning applied.

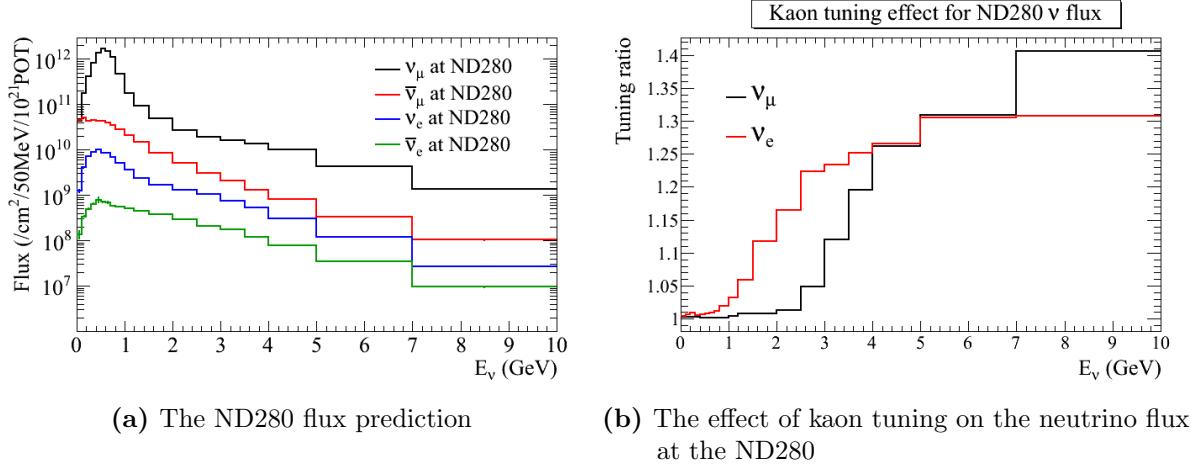


Figure 3.2.: The predicted neutrino flux (a), and the effect of kaon tuning on that flux (b), at the ND280 detector. 3.2a shows the expected flux for the four neutrino species present in the neutrino beam as a function of neutrino energy. This has been normalised to 1×10^{21} POT, showing that the ν_μ component is dominant. 3.2b shows the ratio of the predicted flux with kaon tuning applied to that predicted without tuning, for both the ν_μ and ν_e components of the neutrino beam.

Neutrino interaction generators

Neutrino interaction generators are used to simulate neutrinos interacting with nuclei and include various theoretical and empirical models to describe the possible physical processes that can occur. The T2K experiment uses both the NEUT and GENIE neutrino event generators, but the published oscillation results use the NEUT generator, which is also used in this analysis. The steps described for the NEUT generator below are also present in GENIE, but different theoretical models are implemented and the results tuned to different data sets.

The NEUT generator has two external inputs, the neutrino vector file produced by the beam MC and a ROOT-based simulation of the ND280 detector geometry. Combined with its own neutrino interaction models, the generator uses this information to create neutrino interactions throughout the ND280. Initially, for a given nucleus, NEUT uses the Moniz-Smith relativistic Fermi gas model [23] to describe the nucleons it contains. Together with the neutrino interaction processes within the generator this is used to calculate the probability of each interaction mode occurring on that nucleus. Using the initial neutrino vector and the ND280 geometry, NEUT calculates the interaction probability across all the nuclei along the neutrino's path. Random number generators

are then used to decide whether an interaction occurs, the position it occurs at, and the interaction mode itself.

Once this has been done, if required, the formation of primary hadronic showers is simulated. This uses internal NEUT code for hadronic systems with energies below 2 GeV and the PYTHIA/JetSet algorithms for those with energies above 2 GeV. The internal algorithms generate showers directly from experimental data and were developed because PYTHIA/JetSet does not describe the data well at these low energies.

After these steps, NEUT has generated all the primary particles created by the neutrino interaction, which must now be propagated from the interaction point to the edge of the nucleus. This is done using a cascade model, where any primary hadrons are moved in steps through the nucleus. At each step the probability of the hadron interacting with the nuclear environment is calculated, the result of which determines whether an interaction is simulated. This is repeated for all primary hadrons until they either interact or exit the nucleus.

The neutrino generator outputs a set of final state particle vectors alongside information about the neutrino interaction that created them. This then passed on to the ND280 detector MC software package, which simulates the detector's response to these particles.

ND280 detector simulation

The ND280 detector simulation handles the creation of the ND280 detector geometry and particle propagation through that geometry, using ROOT [45] and Geant4 [46]. The ND280 geometry is created from the design specifications of the different ND280 sub-detectors and records all of the necessary properties of the materials used in their construction. This information is used by Geant4 to correctly simulate the energy deposition, scattering and decay of particles as they pass through the detector. The simulation provides a list of energy deposits in the active regions of the detector, which are then processed by the electronics simulation software. At the same time the neutrino beam timing structure is applied to the neutrino interactions, so that the correct number of interactions are included in each beam bunch and each bunch is correctly separated from the others within the simulated beam spill.

ND280 electronics simulation

The final part of the ND280 MC simulates the response of the detector's active regions to energy deposited within them and reproduces the response of the detector electronics. In particular, the drift of charged particles in the TPC and the response of the scintillating bars and WLS fibres is calculated to provide what is actually readout by the electronics. The electronics simulation then mimics the effect of the various readout boards on the ND280 sub-detectors, including noise and saturation effects, before digitising the signals, resulting in an output that has the same format as the data recorded by the ND280.

At this point, the MC can now be processed by the same software chain as the data, which was described at the beginning of this chapter.

3.2. ECal event reconstruction

The analysis described in this thesis is based solely on reconstructed ECal objects, so the ECal reconstruction is described here in some detail.

As mentioned in Sec. 2.2.3 the ECal is a tracking sampling calorimeter formed from layers of scintillator bars and lead sheets. Each scintillator bar has a 10 mm by 40 mm rectangular cross-section, which defines the granularity of the detector and in turn informs many choices in the reconstruction algorithms. The ECal reconstruction takes place in a number of steps:

1. Hit preparation
2. Basic clustering
3. Combine clustering
4. Expand clustering
5. 3D matching
6. 3D hit reconstruction
7. Low energy photon reconstruction
8. Energy fitting and particle identification

Hit preparation

The initial stage of the ECal reconstruction takes partially calibrated hits and prepares them for the later algorithms.

The hits arrive separated by module and are then split into separate time bunches. This is done by ordering them in time before searching for any gaps of greater than 50 ns between two adjacent hits. The hits are then filtered according to the sensor they occurred on. If two hits came from sensors on either end of the same ECal bar and are close in time then they are combined to form a reconstructed hit, otherwise they are left as a single hit. For reconstructed hits, the relative time difference between the hits, assuming a constant speed of light within the WLS fibre, is used to roughly calculate the position of the reconstructed hit along the length of the bar.

Once this is done the hit is calibrated to remove the effects on the hit charge of light attenuation along the WLS fibre and to convert that charge into minimally ionising particle equivalent units (MEU). A minimally ionising particle, such as a high energy muon, should deposit a similar, small, amount of energy in each bar that it passes through. The scale factor to convert a hit's charge into MEU is calculated using cosmic data, providing a standard candle for charge measured in the detector. By applying this calibration to each bar in turn, and by measuring the response to cosmic rays relatively often, the response of each bar to a given amount of deposited energy can be kept constant even if other factors have changed.

Basic clustering

The basic clustering is a nearest neighbour aggregation algorithm that produces 2D clusters separately for each view of the ECal detector. To form the clusters, the highest charge hit within a 30 ns window is selected. Using this as a seed, hits are added to the nascent cluster provided they are within the time window and they are within two layers and one bar of a hit currently in the cluster. The final clusters must contain at least three hits and are then passed on to the next clustering algorithm alongside any remaining hits that have not been clustered.

Combine clustering

The next stage of the clustering chain is the combine clustering algorithm, which attempts to re-connect single tracks or showers that produced two basic clusters.

The basic cluster containing most hits is selected as the primary cluster, since it contains the most information, and the principle direction calculated using charge weighted principle component analysis (PCA). The algorithm then calculates the smallest distance between the principle axis of the primary cluster and any other basic cluster present. If the smallest distance between two clusters is less than 80 mm, and the difference between the charge weighted average hit time of the clusters is less than 40 ns, they are combined into a single cluster.

The above process is then repeated, with this new cluster used as the primary cluster, until all possible clusters have been combined. These clusters are then passed onto the Expand clustering algorithm with any basic clusters that were not used.

Expand clustering

This clustering algorithm takes the unmatched hits from the basic clustering and attempts to add them to the clusters produced by the combine clustering algorithm.

The algorithm looks at each unmatched hit individually, trying to match it to any of the existing clusters. For each cluster the primary and secondary axes are calculated using PCA, along with the 1σ width of the cluster along the two axes. The distance along these axes of the unmatched hit from the cluster centre is then calculated, in units of σ , before being added in quadrature to give the matching weight.

This matching weight is calculated for the same unmatched hit using each of the available clusters. The hit is then matched to the cluster providing the lowest matching weight, provided that that weight is below 80 and that the hit is within 40 ns either side of the earliest and latest hits within the cluster. This process is repeated for each unmatched hit to form the final 2D clusters.

3D matching

The 3D matching combines the final 2D clusters into 3D objects, using two different likelihoods to do so. The first likelihood uses a track from the tracker (FGDs and TPCs)

as a seed, propagating this track to the inner face of the nearest ECal module with the RECPACK tools. It then examines every possible pair of 2D clusters in that module that both contain hits in the inner four layers of the ECal. The average position of the two clusters on the inner face of the module is calculated along with the uncertainty on that position. Then, for both position and time, the absolute difference between the clusters and the seed is found, Δx and Δt . The uncertainty on the position difference is calculated as the quadrature sum of the cluster and seed position uncertainties, whilst the uncertainty on the time difference is taken to be 100 ns. These uncertainties are used to create two Gaussian distributions, with a mean of zero, in position and time (G_x and G_t respectively).

The likelihoods are then calculated as follows

$$\mathcal{L}_x = \frac{G_x(\Delta x | 0, \sigma_x)}{G_x(0 | 0, \sigma_x)} \quad (3.1)$$

$$\mathcal{L}_t = \frac{G_t(\Delta t | 0, \sigma_t)}{G_t(0 | 0, \sigma_t)}, \quad (3.2)$$

and the seeded likelihood is simply

$$\mathcal{L}_{\text{seeded}} = \mathcal{L}_x \times \mathcal{L}_t. \quad (3.3)$$

The second matching likelihood uses information from each possible pair of 2D clusters, regardless of their position relative to the inner face of the detector. For each pair the ratio of the total charge in one compared to the other, Q_{ratio} , is calculated then compared to a log-normal distribution generated from MC truth, giving $\mathcal{L}_{Q_{\text{ratio}}}$. The difference in layer number, Δ_{layer} , between the innermost layer of each cluster is also found and compared to a probability density function, also created using MC, which gives $\mathcal{L}_{\Delta_{\text{layer}}}$. The matching likelihood is then

$$\mathcal{L}_{\text{matching}} = \mathcal{L}_{Q_{\text{ratio}}} \times \mathcal{L}_{\Delta_{\text{layer}}}. \quad (3.4)$$

If a pair of clusters has both $\mathcal{L}_{\text{matching}}$ and $\mathcal{L}_{\text{seeded}}$, the final likelihood is the product of the two, otherwise it is just $\mathcal{L}_{\text{matching}}$. The two clusters giving the maximum value of the likelihood are then matched, provided this likelihood is greater than $1e^{-5}$. This

continues until all 2D clusters have been considered for matching, and produces a set of fully reconstructed 3D objects.

3D hit reconstruction

Once 3D objects have been created the information from each ECal view can be combined to give a 3D position to each hit within the object. To do this, the object hits are first split according to which view they came from. Then, the charge weighted average position of the hits contained within each layer is calculated. To reconstruct the 3D position of a hit, in the opposing view a linear fit is performed to the average hit positions from the four layers surrounding the hit in question. This fit provides the position of the hit in its unknown coordinate.

Low energy photon reconstruction

The low energy photon reconstruction is the final stage in the reconstruction chain that can create new ECal clusters. It does this by examining all the unused hits and unmatched 2D clusters left after the previous steps. Considering each unmatched cluster in turn, if any unused hit is in time with the cluster and in an adjacent layer then it is matched. The timing cut requires that the unused hit time be less than 10 ns before the unmatched cluster start time and less than 10 ns after the cluster end time. The layer cut requires that the hit be within one layer of the first layer in the cluster and one layer of the last layer in the cluster. This algorithm attempts to reconstruct low energy photons from π^0 decay, which rarely have enough energy to create three hits in each ECal view.

Energy fitting and particle identification

The ECal was designed to identify and measure the energy of electromagnetic showers, particularly those created by photons from π^0 decay. The electromagnetic energy fit is performed first and provides an estimate of the cluster energy from 25 MeV–20 GeV. It does this by minimising a likelihood function based on the total cluster charge, the charge RMS and the charge skew. Particle gun MC was used to generate a spline relating each of these reconstructed quantities to the true particle energy. The minimiser produces an estimate of the cluster energy and from the splines receives the expected values of the

reconstructed charge variables. By minimising the distance between the measured charge variables and those calculated from the splines, an accurate estimate of the particle's energy is returned. Fig. 3.3 was created from data collected by the DsECal during the CERN testbeam, where particles with a known energy were impinged upon the DsECal, and shows the width of the ECal energy fit distribution divided by the square root of the incident particle's energy as a function of that energy.

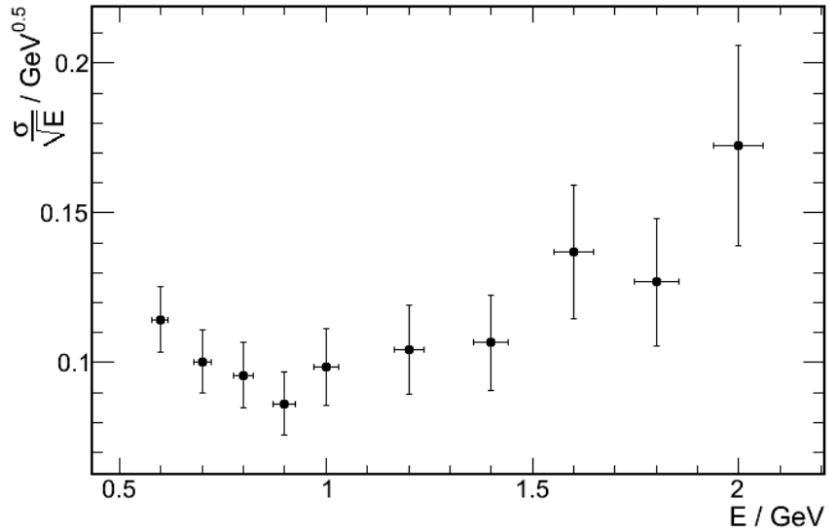


Figure 3.3.: The energy resolution of the DsECal energy fit as a function of true particle energy

The results of this energy fit are then used by the particle identification (PID) routines, which currently separates particles into those that are track-like and those that are shower-like. It does this using a multi-layer perceptron neural network from the ROOT TMVA package [47], which was again trained using electron and muon particle gun MC. The output of this algorithm was then compared to data collected by the DsECal at a CERN beam test [48] and to data collected in-situ at the ND280 hall. The results of this are shown in Fig. 3.4b for testbeam electrons, and in Fig. 3.4a for in-situ muons.

3.2.1. ECal multi-pass Radon transform

The current ND280 reconstruction identifies neutrino interaction vertices in the FGDs and PØD and then looks for tracks or showers in the other sub-detectors. This ignores any interactions that occur in the ECal, the most massive of the fully instrumented sub-detectors. The large mass of the ECal means that there are many more neutrino

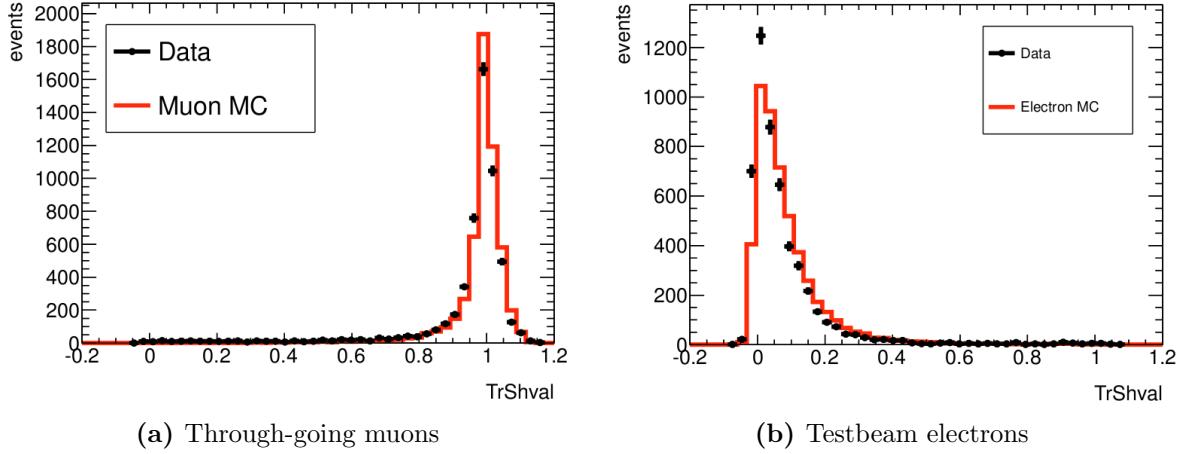


Figure 3.4.: Data-MC comparisons of control sample data to the appropriate particle gun MC for (a) through-going muon data collected at the ND280 hall and (b) electron data collected at the CERN testbeam.

interactions within it than occur in the FGD, and these interactions create backgrounds for analysers trying to select FGD vertices.

To help identify neutrino vertices in the ECals a multi-pass Radon transform was developed. A Radon transform is a common pattern recognition technique, used in many fields, where analysers try to find geometric shapes in distributions of data points. The interaction with the highest cross-section in the T2K neutrino beam is the CCQE interaction, mentioned in Chap. 1, which produces a muon and typically a proton in the final state. The ECal reconstruction algorithms will combine the hits from these particles into a single cluster, which will then be identified as a shower. Applying a Radon transform to the hits in this cluster can show that it was formed from a number of individual tracks, rather than a single shower, which can then be used to help identify the cluster as a neutrino vertex.

In high energy physics the literature describes the use of a Fuzzy Radon transform [49]. This paper defines the measured hit density $\rho(x)$ as

$$\rho(x) = \int \rho_p(x) D(p) dp, \quad (3.5)$$

where x is a set of measured parameters describing the hits, p is the set of parameters that describe possible particle tracks, $D(p)$ is a function describing the tracks that created the measured hits and $\rho_p(x)$ is the detector response to a particle with parameters p .

The Fuzzy Radon transform, $\tilde{D}(p)$, is the inversion of this integral, where one takes a known hit density and response function to find $D(p)$

$$\tilde{D}(p) = \int \rho(x)\rho_p(x)dx. \quad (3.6)$$

Local maxima of $\tilde{D}(p)$ then give the parameters p that describe the tracks that created the observed hit density.

It is not possible to carry out the integral in Equ. 3.6 analytically, so the above algorithm is simplified to become solvable.

To begin with, the transform is applied to each view of the ECal separately, so the set of hit parameters x contains only the 2D hit positions. Secondly, the detector response is modelled as a delta function for each bar that a particle with parameters p would travel through. Finally, the potential tracks are assumed to be straight lines, so can be described by a distance, d and an angle, θ . These changes mean that, for each point in the $d - \theta$ space of possible tracks, the value of $\tilde{D}(d, \theta)$ is given by counting the number of hits that that line intersects. When applied to a reconstructed ECal cluster, the distance d is defined as the minimum separation of the potential Radon line from the charge-weighted centre of the cluster and θ is defined as the angle made by the Radon line to the local X axis.

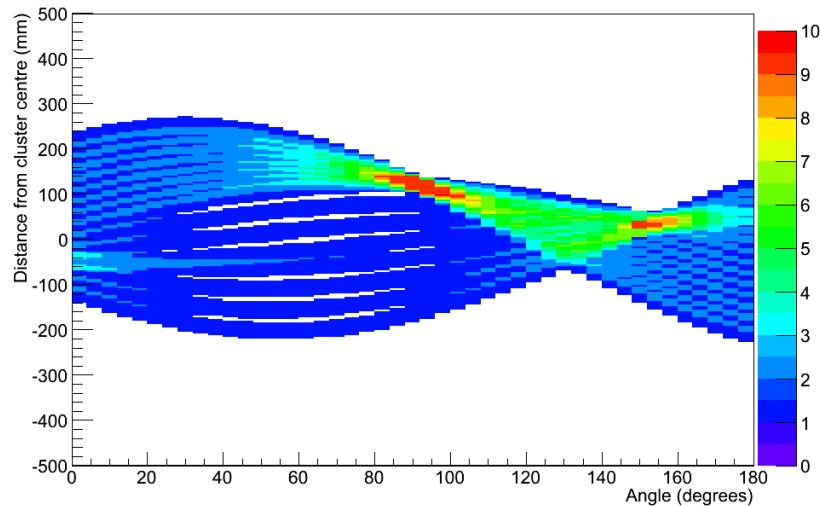


Figure 3.5.: The output from applying the described Radon transform to particle gun MC. Each cell corresponds to a single Radon line, with the cell height indicating the number of hits that line crossed

Fig. 3.5 shows a histogram of $\tilde{D}(d, \theta)$ across the $d - \theta$ space, with steps of 12 mm in d and 4° in θ . This output came from applying the transform to particle gun MC where a muon and proton had been fired from the same point within a barrel ECal module but with an angular separation of 60° . The two particle tracks can be clearly distinguished from one another, one travelling at 90° to the local X axis and the other at 150° . Fig. 3.6 shows the same plot but with the transform applied to a simulated neutrino interaction rather than the idealised particle gun. In this situation it was required that at least two charged particles were created in the interaction and that both deposited hits within the same ECal cluster. The particle tracks are now much less clearly defined, so whilst it is simple to find the global maximum it is much more difficult to find the second or third maxima.

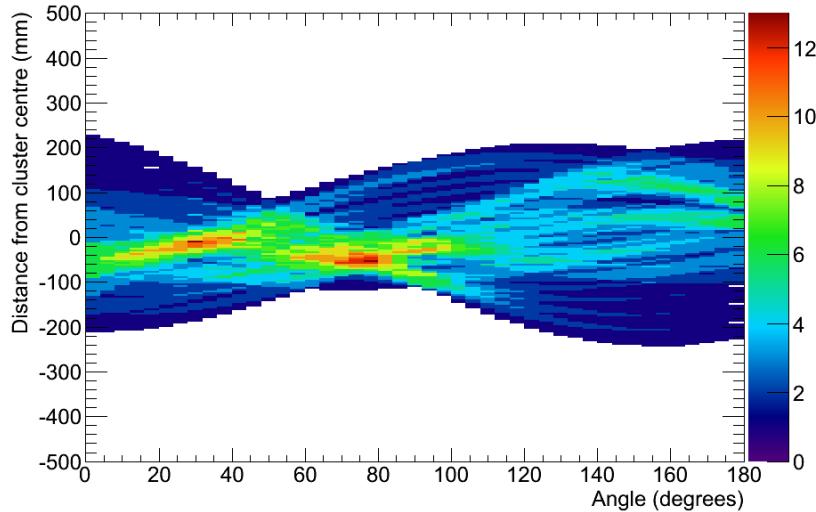
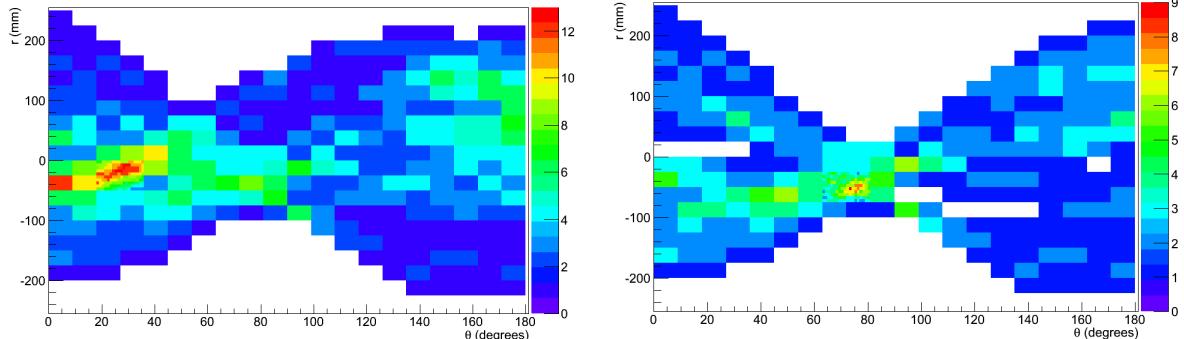


Figure 3.6.: The output from applying the described Radon transform to neutrino beam MC. The result shown here came from an ECal cluster formed by a DIS event creating two pions and a proton that deposited energy within the ECal.

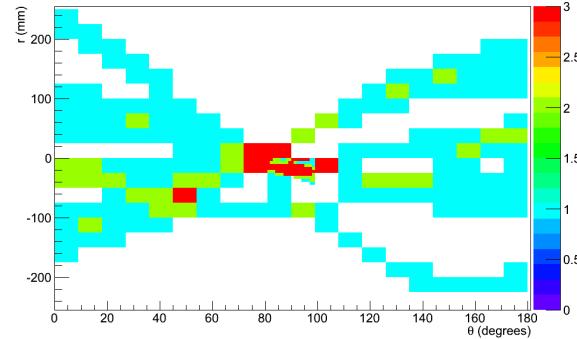
To allow multiple tracks to be identified, the Radon transform is applied more than once to the same 2D cluster. For each iteration, when the best Radon line has been found, the hits that are intersected by that line are removed from the algorithm. The transform is then applied to the remaining hits, and this time the global maximum corresponds to the secondary maximum seen in the output of the first iteration. This process is repeated up to three times, and a track identified if the Radon line passes through at least three hits. This is a similar method to that used at Super-Kamiokande to count the number of Cherenkov rings from a neutrino interaction [36].

Radon transforms can be computationally expensive, since the entire $d-\theta$ space must be covered each time the transform is applied. To mitigate the effect of performing the transform three times, a fast scan of the parameter space is performed first, using steps of 25 mm in d and 9° in θ . A second, fine grained, scan is then applied to the $d-\theta$ cell that contained the maximum from the fast scan, using steps of 5 mm and 1°. A comparison was made between this method and performing a fine grained scan over the entire space, and the Radon lines found agreed between the two methods in all cases.

The histograms in Fig. 3.7 show the output of the Radon algorithm applied three times to a single 2D cluster, which is shown in the event display in Fig. 3.8. The three particle tracks that form this cluster can be clearly seen, and have been correctly found by the multi-pass Radon transform.



(a) The first application of the Radon algorithm (b) The second application of the Radon algorithm



(c) The third application of the Radon algorithm

Figure 3.7.: The results from sequentially applying the Radon transform algorithm three times to a single ECal cluster, removing any tracks that are found at each stage.

To categorise the performance of the algorithm it was necessary to perform a hand scan of each event, since in many cases the truth information was insufficient to determine

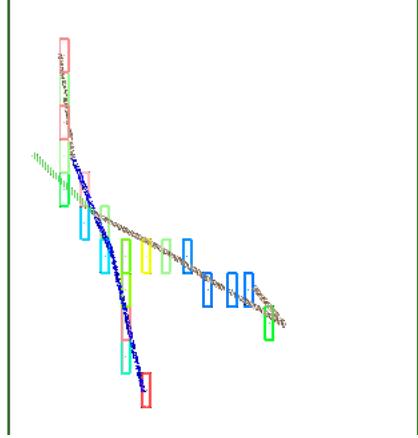


Figure 3.8.: An event display showing the distribution of hits that created the Radon transform output in Fig. 3.7. The hits are represented as rectangular cells whose colour indicates the measured charge. The MC trajectories of the three charged particles that created these hits are also shown, with two the red lines corresponding to two charged pions and the blue line to a proton. The final dashed green line shows the trajectory of the initial neutrino that created this cluster.

	Number of tracks found		
	1 Track	2 Tracks	3 Tracks
Number found	34	46	30
Number correctly found	32	31	26
Number missed	1	4	3
Purity	94%	67%	87%
Efficiency	97%	89%	90%

Table 3.1.: The performance of the multi-pass Radon algorithm for finding tracks in 2D clusters in the barrel ECals.

whether a true particle track had been found. Table 3.1 shows the results from hand-scanning 127 2D clusters.

At this stage there is no category for clusters formed from more than three tracks or from showers, so they have been added to the ‘3 Tracks’ sample. Also, the ECals are dominated by single, through-going tracks, so after sufficient single track clusters had been found only clusters with two or more tracks were hand scanned. The biggest source of mis-classification arises from the degree of curvature of single tracks. If a minimally

ionising particle travels a significant distance within the ECal then the approximation used in the algorithm, that the particles are travelling in straight lines, is no longer valid. This causes the algorithm to incorrectly fit the curved track with two straight lines.

As mentioned above, there were numerous cases where it was not possible to associate a clear track in the detector to the truth particle that created it. Bearing this in mind, to get an idea of the resolution of the reconstructed direction of the Radon tracks, the difference in angle between the primary Radon track and the truth trajectory that contributed most hits to the ECal cluster is shown in Fig. 3.9

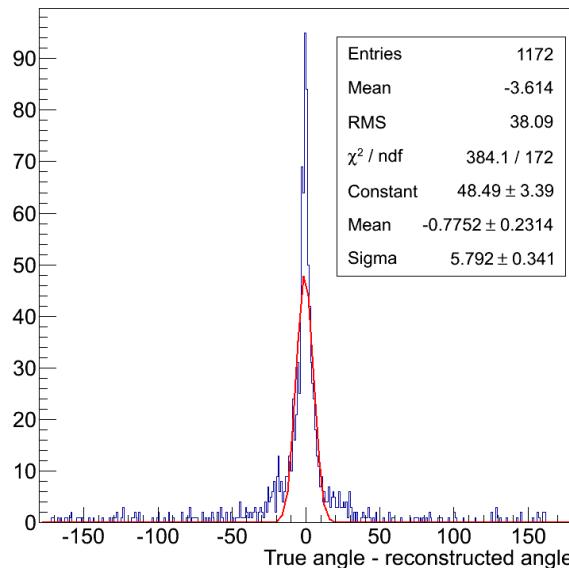


Figure 3.9.: The difference between the angle that the primary truth particle made to the local X axis and the angle found by the Radon transform. The red line shows a Gaussian fit to the central peak, to give an indication of the direction resolution of the algorithm.

A Gaussian fit was performed to the central peak of the distribution in Fig. 3.9, the width of which is used as a measure of the algorithm's direction resolution, which is found to be 5.8° . Each ECal scintillator bar is 40 mm wide and each module is approximately 400 mm deep. Using these numbers, for particles travelling perpendicularly to the ECal front face the maximum achievable direction resolution is 5.7° , showing that the achieved resolution is quite acceptable.

Once the transform has been applied to both views, the 2D Radon tracks it creates are then combined to form 3D tracks, with the primary Radon track in one view matched to the primary track in another view, the second to the second and so on. The performance of this matching is shown in Table 3.2, when the algorithm was applied to 1306 3D

		Number of 3D tracks found		
		1 Track	2 Tracks	3 Tracks
Number matched		1306	114	40
Number correctly matched		1041	19	1
Percent correct		80%	17%	3%

Table 3.2.: The 3D matching performance of the multi-pass Radon transform.

		Total number of tracks found from both views				
		2 Tracks	3 Tracks	4 Tracks	5 Tracks	6 Tracks
Number matched		1013	156	57	34	46
Number correctly matched		861	119	30	13	18
Percent correct		85%	76%	53%	38%	39%

Table 3.3.: The 3D matching performance of the multi-pass Radon transform for the first track. The 3D clusters have been split according to the total number of 2D tracks they contain.

ECal clusters. The table shows the total number of clusters that were found with a given number reconstructed 3D tracks, and also shows the number of clusters where each 2D track was correctly matched to its partner in the opposite view. The matching

performance looking at the first track alone is shown in Table 3.3, split according to the total number of 2D Radon tracks in the 3D cluster. The number matched category here shows the total number of clusters containing the specified number of 2D tracks, whilst the correctly matched category is the number of clusters where the primary 2D Radon track was correctly matched to its partner in the opposite view. These results show that

the primary Radon track is correctly reconstructed in the majority of cases provided there are less than 5 2D Radon tracks within the cluster. Table 3.2 and Table 3.3 show that the second and third Radon tracks are only very rarely correctly matched to their partner, indicating that matching Radon tracks based upon their number of hits alone is not sufficient. This matching could be significantly improved in future versions of this algorithm by looking at the charge of the 2D Radon tracks, since the spatial distribution of the charge in 2D tracks created by the same particle should be very similar in each

ECal view. This could be used in the construction of a matching likelihood in a similar manner to that described earlier in Sec. 3.2, along with information on the relative position and number of hits in each track.

Having considered the 2D and 3D performance of the algorithm, it was decided that the 2D Radon transform algorithm results would be used in this analysis. There is a lot of scope for further developments of the algorithm, by fitting curved rather than straight tracks, developing a robust classification for showers or by improving the 3D matching of the 2D Radon results. The algorithm described above provides information that can be used to classify ECal clusters as ‘vertex-like’, rather than either ‘shower-like’ or ‘track-like’, and further development will improve this.

Chapter 4.

Selecting neutrino interactions in the ND280 ECals

This analysis presents a measurement of the CC inclusive cross-section for ν_μ interactions with lead nuclei in the ND280 ECals. Making this measurement requires a sample of neutrino events whose vertex is within an ECal. The ECals are constructed from alternating layers of uninstrumented lead sheet and layers of scintillator bars, as described in Chap. 2. Neutrino interactions that occur in the lead sheets are indistinguishable from those that occur in the scintillator bars, so the event selection presented here does not attempt to differentiate between them. The signal is therefore defined as a ν_μ induced charged current interaction that occurs within a lead or scintillator layer of one of the ND280 ECals. There are two main types of background that must be identified by the selection. The most obvious background is made up of minimally ionising particles that pass straight through an ECal, leaving a clear track, labelled as ‘through-going’ backgrounds. The second background is created by charged particles entering and then stopping within an ECal module. These stopping backgrounds mimic the charge deposition expected from signal interactions and are labelled as ‘signal-like’ backgrounds. Some of these signal-like events form an irreducible background in this analysis, and are included in the cross-section extraction fit described in Chap. 5.

The ECals can be divided into two categories by their orientation to the neutrino beam. The DsECal lies perpendicularly to the beam, so particles from neutrino interactions will primarily be travelling at right angles to the DsECal face. The barrel ECals all lie parallel to the neutrino beam, so neutrino induced particles generally will be passing along the plane of the barrel ECal layers. These geometric differences motivate the development of two separate selections. Both selections use a number of pre-selection cuts to select

reconstructed events of good quality, remove through-going backgrounds and to remove regions where the data is not well modelled by the MC. Separate multi-variate boosted decision trees are then used to select the signal interactions.

A detailed description of these two selections is presented and their performance evaluated in the following sections. It should be noted that both selections were initially optimised on Monte Carlo and then applied unchanged to the data sample in order to remain unbiased.

4.1. Data sample

The T2K data taking is divided into a number of different periods, known as physics runs, each corresponding to roughly one year of operation. Between each run improvements are made to the J-PARC proton beam, usually resulting in an increased beam power for each subsequent run. The data used in this analysis were collected between November 2010 and June 2012 from two physics runs—Run II, from November 2010 until March 2011, and Run III from March 2012 until June 2012. The number of protons per pulse and the integrated number of POT are shown in Fig. 4.1 for both Run II and III. Before November 2010, in Run I, only the DsECal was installed and functioning at the ND280. This analysis is not statistics limited, unlike the oscillation analyses, so including the data from Run I would not reduce any statistical uncertainty on the final result and would require extra systematics to be evaluated. For this reason it is not included in the analysis.

4.1.1. Data quality requirements

There are a number of requirements that each beam spill must pass in order to be classified as good quality. Firstly, the neutrino beamline monitor data are checked to ensure that each component of the beamline is operating within its nominal parameters and that the proton beam properties are as expected. These checks are described in detail in [50] but in short require that

1. Each hardware component is working correctly
2. All horns currents are within 5 kA of the mean current

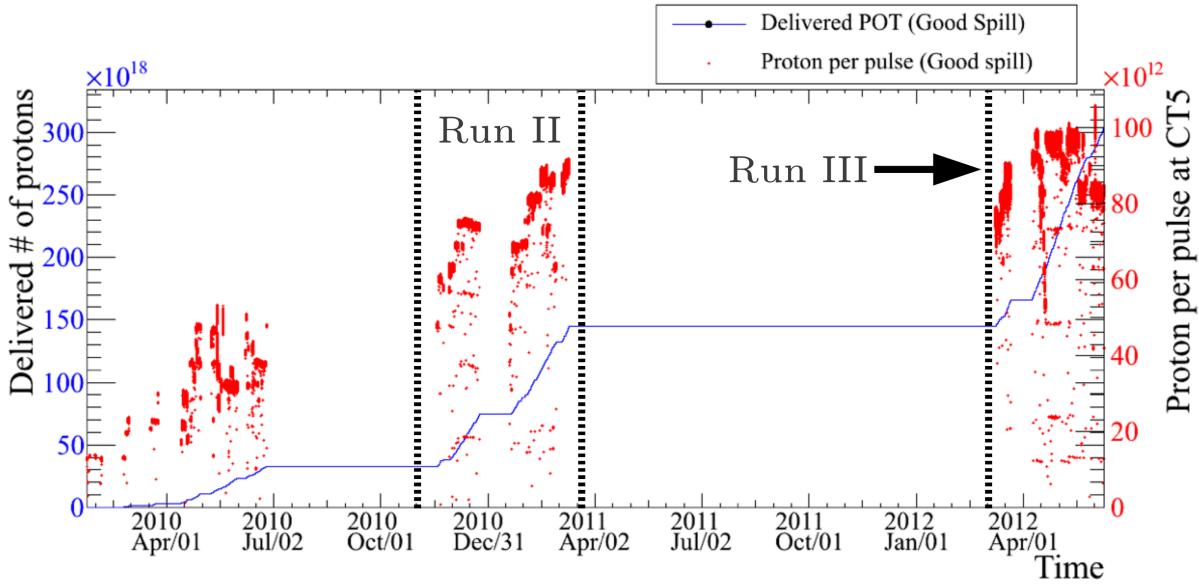


Figure 4.1.: The number of protons per pulse and the integrated number of POT over time, showing physics runs II and III, as measured by the fifth current transformer (CT5) beam monitor.

3. The beam angle is within 1 mrad of the mean value
4. The total muon yield measured by the muon monitor is within 5% of the mean value

This information is summarised by a boolean flag that highlights every spill where these requirements are not met.

At the ND280 detector a dedicated data quality group examines the sub-detector outputs on a spill by spill basis, providing another boolean flag to indicate when each sub-detector is collecting good quality data. This selection only considers reconstructed objects from the ECal sub-detectors, so only requires that the ECal data is of good quality, rather than the entire ND280 detector.

The ECal data quality requirements are described in Ref. [51], with the main requirement being that the ECal electronics are working as expected. Checks are performed every 3 hours, due to the time needed to process the raw data. As explained in Chap. 2, the MPPCs produce a discrete amount of charge for each pixel that fires, so by plotting the recorded charge from a single channel, multiple peaks are seen that correspond to the number of pixels that have fired. The pedestal value of a channel is defined as the amount of charge recorded by the MPPC when no pixel has fired, and the gain defined as the difference in charge between the central values of neighbouring peaks. The measured

gain and pedestal values are used to calibrate the data taken over each three hour period, but are also checked by eye to identify significant deviations in the gain of each MPPC or in the pedestal charge. If these requirements are not satisfied that beam spill is rejected.

Looking at beam triggers alone, the data quality group check the hit time relative to the beam trigger time. Taking into account offsets from the readout electronics, the beam trigger timing can be mapped onto the electronics readout cycle and compared to the data. By doing this it can be seen in Fig. 4.2 that all ECal hits occur near the expected beam timing window and do not coincide with the electronics reset periods, so the complete charge is recorded for each hit.

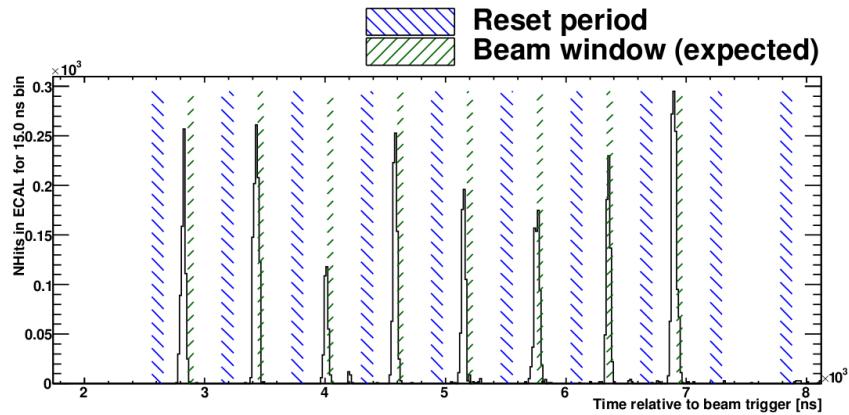


Figure 4.2.: The ECal hit time relative to the beam trigger time, showing the expected beam window and the electronics reset periods. The vast majority of hits occur during the expected beam windows. The reset periods are 100 ns long and the electronics integration period between two reset periods is 480 nm long [51].

The total data used in this analysis corresponds to 2.665×10^{20} protons on target (POT), of which 2.588×10^{20} POT pass these basic quality cuts.

4.1.2. Monte Carlo sample

The NEUT neutrino interaction generator was used to produce 5.330×10^{20} POT of neutrino beam MC that was used in this analysis. Of this, 2.170×10^{20} POT was generated using a simulated proton beam power of 120 kW to match the average power achieved in physics Run II, whilst the remainder was simulated with a proton beam power of 178 kW corresponding to the Run III average proton beam power. At these powers we expect an average of 7.99 neutrino interactions in the ND280 detector per spill during Run II and 9.46 per spill during Run III. This corresponds to roughly one

interaction per bunch across the entire ND280 detector in both runs, so the chance of event pile-up affecting the selected events is negligible.

Combined with this is 1.795×10^{20} POT of neutrino beam MC with neutrino interactions generated in the ND280 detector pit and in the surrounding substrate rather than in the ND280 detector itself. This sample is referred to as *sand* MC in the remainder of this thesis. In an ideal situation an equal amount of sand and beam MC would be used, but unfortunately this quantity of sand MC was not available at the time of writing. Even so, the sand MC used produced 228,000 clusters in the ND280 ECals and fully covers the expected kinematic space of particles from these sand interactions. It is therefore appropriate to use this smaller sample and scale the results to match that expected from the total integrated data POT being analysed.

4.2. A Monte Carlo study to select neutrino interactions in the ECals

4.2.1. Pre-selection of events

The barrel and downstream ECal pre-selections were designed to remove any ECal cluster that was easily distinguishable from the signal described above while at the same time retaining the vast majority of that signal.

As mentioned in Chap. 3, whilst at least six hits are needed to form a fully reconstructed object, many clusters with fewer hits are created by the low energy photon reconstruction algorithm. To remove these, a cut is made on both the barrel ECal and DsECal samples requiring at least six hits in the final reconstructed object. The number of hits in each cluster is shown in Fig. 4.3 for both the DsECal and barrel ECals, with events separated into signal, through-going background and signal-like background categories. The histograms in Fig. 4.3 are area-normalised, with the two background categories having a combined area equal to that of the signal category, to allow the shape of each category to be clearly seen.

In addition to this, in the DsECal selection a cut is made requiring that there are no hits in the first layer of the DsECal, layer 0. This acts to veto any particle entering from a sub-detector further upstream and Fig. 4.4 shows the distribution of signal and

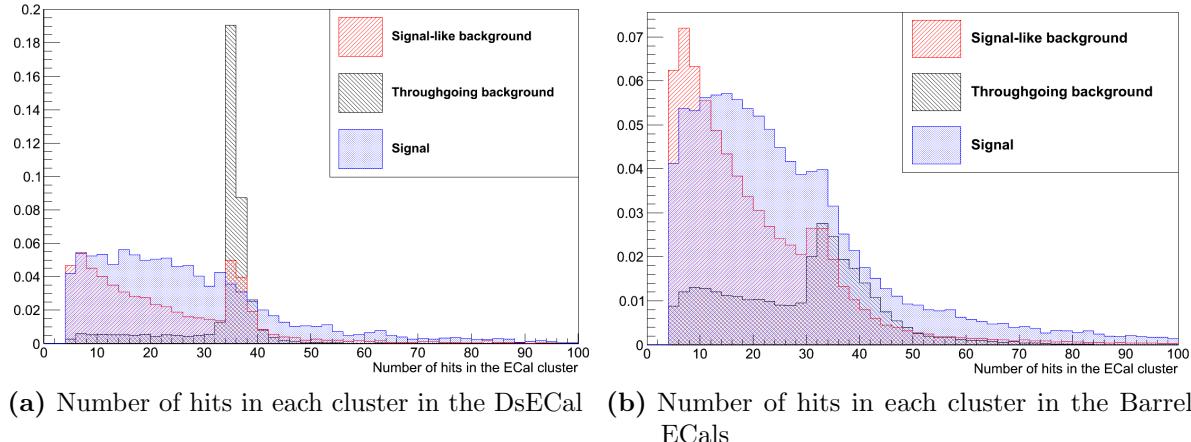


Figure 4.3.: The number of hits in each cluster for (a) the DsECal and (b) the barrel ECal, split into three categories: signal (in blue), through-going background (in black) and signal-like background (in red).

background events relative to this cut. A similar cut is not applied to the barrel ECal samples due to the differences in orientation described at the start of this chapter.

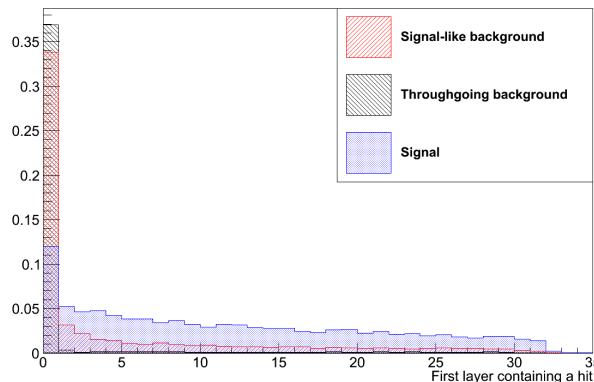


Figure 4.4.: The most upstream layer hit in the DsECal split into three categories: signal (in blue), through-going background (in black) and signal-like background (in red).

These pre-selections produce samples of fully reconstructed events, with the majority of the through-going background cut from the DsECal sample. A multivariate classifier is then used to distinguish signal from background in these samples.

4.3. Boosted decision tree multivariate analyses

When trying to divide events into signal and background categories, better separation is generally achieved by using a discriminant formed from some combination of other discriminating variables. Methods to create these discriminants are known as multivariate analysis (MVA) techniques, and they achieve better separation by taking into account the correlations between different variables, in contrast to cut based selections. MVAs are often used to select rare events, such as the decay of a Higgs boson, but are also useful when the individual discriminating variables show little separation between the signal and background events, as is the case here. Performing a cuts based analysis would both limit the achievable efficiency and purity of the selection and potentially produce samples biased toward certain neutrino energies or lepton kinematic regions.

In the past artificial neural networks (ANNs) have been used for this purpose but recently boosted decision trees (BDTs) have become more prevalent. An ANN is a pattern recognition algorithm that uses a network of nodes, arranged in layers, to separate signal from background. An input layer is associated to the input variables and an output layer provides the final discriminant, with a number of hidden layers in between. ANNs provide good separation but must be carefully tuned and trained to ensure they do not suffer from over-training. BDTs are capable of providing as good, if not better, separation than ANNs and are both simpler to use and very robust [52]. In this analysis the TMVA package[47] was used to both train and apply MVA methods, and the performance of the BDT method was compared to that of the ANN. This comparison is shown in Fig. 4.5, where it can be seen that the ANN and BDT give almost identical separation across the range of signal efficiencies. Since the BDT is easier to train and more robust, it was chosen for this selection.

4.3.1. Creating a boosted decision tree

A decision tree is a simple MVA that uses a series of simple cuts to classify events as signal or background, and by itself is a weak discriminator. A boosted decision tree is created by applying a boosting algorithm to a large number of decision trees. This combines and weights the response of the decision trees, transforming many weak classifiers into a single powerful discriminator.

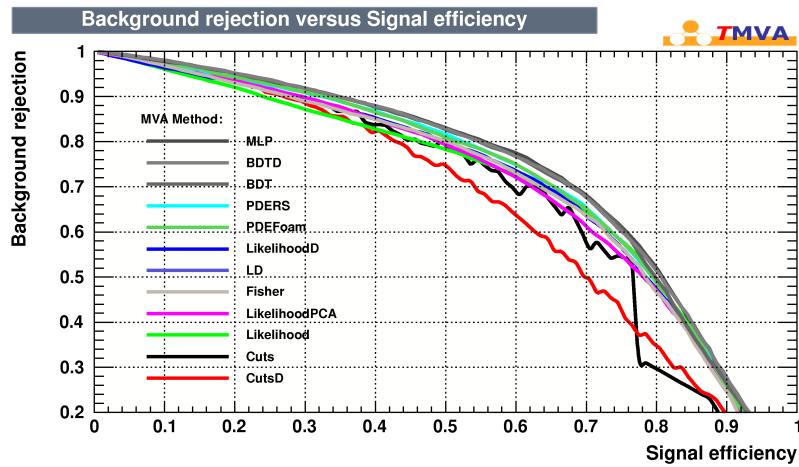


Figure 4.5.: The ROC curve for a number of different TMVA methods, showing the amount of background rejection achievable at a given signal efficiency.

Growing a decision tree

To create a single decision tree one must first have a set of signal and background events with a number of variables that can be used to discriminate one from the other. This is the training sample. Then, for each variable, scan over all the possible values it can take, calculating the signal and background purity for all events to the left or right of that value. This splitting creates two ‘branches’, one signal-like the other background-like. The splitting value is chosen to give the maximum separation between signal and background, determined using the Gini index, G . This is defined as

$$G = \left(\sum_{i=1}^n W_i \right) P(1 - P), \quad (4.1)$$

where W_i is the weight of event i and P is the purity of signal events in the current branch. Initially all the events have a weight of 1. To achieve the best split, the sum of the Gini index of each branch must be minimised. In practice this is equivalent to maximising the quantity C , given by

$$C = G_{\text{Parent}} - (G_{\text{Left Branch}} + G_{\text{Right Branch}}), \quad (4.2)$$

where G_{Parent} is the Gini coefficient of the parent events, $G_{\text{Left Branch}}$ is the Gini coefficient of the events that fall to the left of the splitting value and $G_{\text{Right Branch}}$ the Gini coefficient of those to the right.

The variable and splitting value that give the best separation are chosen and form the first ‘node’ of the decision tree. Applying this cut then produces two branches, one enriched with signal events the other with background. For each branch the process described above is repeated, choosing the best variable to separate the events in that branch into new signal and background branches. This process is repeated until either the final branch, known as a ‘leaf’, contains too few events, the leaf contains purely signal or background events, or the requisite number of branches have been created. Each leaf is then categorised as signal-like or background-like according to the signal purity of the events it contains. Any background event on a signal leaf, or vice versa, is misclassified. An example decision tree is shown in Fig. 4.6

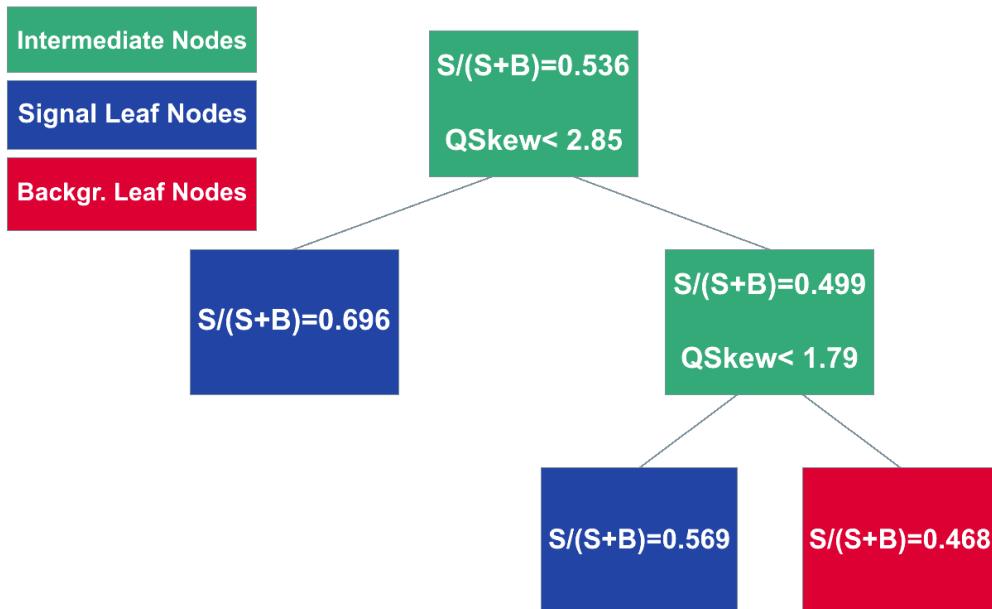


Figure 4.6.: A single decision tree taken from an example BDT. The signal-like and background-like leaves are highlighted, with the signal significance of each leaf shown, and the cuts used to create each branch appearing in the parent node.

Creating a decision tree is a simple process that is very similar to using conventional cuts, except that no event is completely removed in the above algorithm, so no information is lost. Unfortunately decision trees by themselves are not stable, and small fluctuations in the training sample can create big differences in the decision tree. This problem is removed by boosting.

Boosting a decision tree

There are a number of different boosting algorithms currently available, but the one used in this analysis is the AdaBoost algorithm [53]. This operates by increasing the weight of all misclassified events. The error, E_m , of a tree, m , is defined as

$$E_m = \frac{\sum_{i=1}^{n_{\text{Mis-classified}}} W_i}{\sum_{i=1}^{n_{\text{Total}}} W_i}, \quad (4.3)$$

the sum of the event weights, W_i , for all the misclassified events divided by the sum of all the event weights. The score of the tree, α_m , can then be calculated as:

$$\alpha_m = \beta \ln \frac{1 - E_m}{E_m}, \quad (4.4)$$

where β is a constant defined by the user. The weight for each misidentified event is then increased so that

$$W_i^{\text{new}} = W_i e^{\alpha_m}, \quad (4.5)$$

and the weights for all events renormalised to give the same total. These new weights are then used to grow a second tree, which should now correctly classify some of the events misclassified by the old tree assuming one of the input variables provides some separating power for these events. This process is repeated a number of times to produce a forest of boosted decision trees. Each event is passed through every tree in the forest, and the final BDT classifier output, $T(x_i)$, is given by

$$T(x_i) = \frac{1}{N_{\text{tree}}} \sum_{m=1}^{N_{\text{tree}}} \alpha_m T_m(x_i), \quad (4.6)$$

where x_i is the i th event and $T_m(x_i)$ is equal to 1 if the event ends up in a signal leaf or -1 if it ends up on a background leaf.

Boosting the decision trees and then using the weighted sum of the tree responses to classify events not only increases the separating power of the classifier but also increases its robustness. Another useful feature, compared to ANNs, is the way extra input variables are handled. In a BDT, variables are selected by their ability to separate the weighted events and then a simple cut is applied. An additional variable is only selected

when it provides some separation between the events, and is otherwise ignored. This is not the case for the ANN, where the additional variable will be added to the input layer and can affect the output even if it provides no separating power.

4.3.2. The boosted decision tree in this analysis

The BDT used in this analysis was grown with 800 trees, a maximum of 8 final leaves per tree and with $\beta = 0.3$. The BDT was trained using a sample of 1.06×10^{20} POT of beam MC, with neutrino interactions simulated throughout the magnet and tracking regions of the ND280, and with an equal amount of sand MC. These MC files were generated using the NEUT generator described in Chap. 3, and split into signal and background samples. These samples were randomly split in two when growing the BDT, with half of each being used to train the BDT and the other half used to test the final tree.

Input parameters

Input variables were identified by comparing the distribution of signal and background events within that variable. This highlighted variables that provided some discriminating power between signal and background, an example of which is shown in Fig. 4.7. The list of variables used in the DsECal and barrel ECal BDTs is in Table 4.3.2 whilst plots showing the signal separation of these variables can be found in Appendix A. More details on the individual variables can be found in Ref. [54, 55].

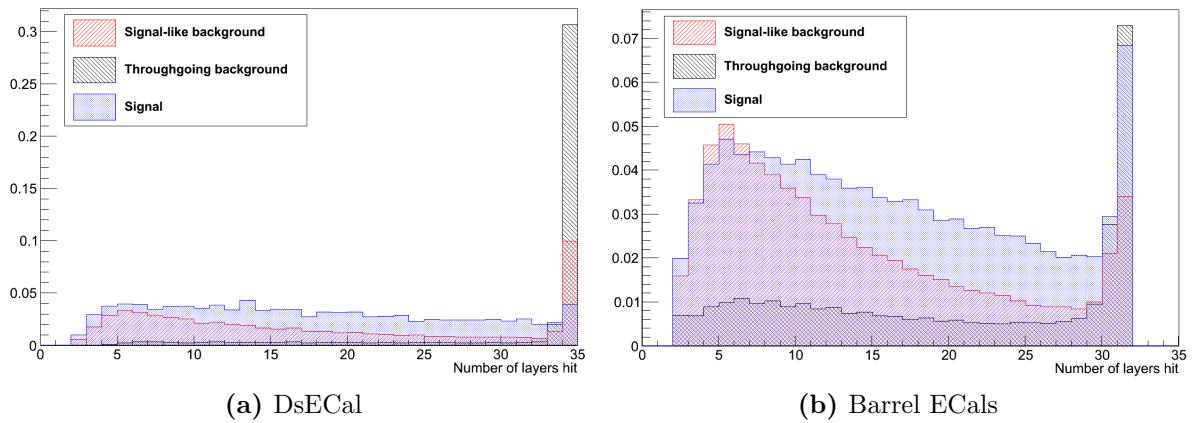


Figure 4.7.: The number of different layers containing a hit for each cluster in (a) the DsECal and (b) the barrel ECal, split into three categories: signal (in blue), through-going background (in black) and signal-like background (in red). The histograms have been area normalised so that the through-going background and signal-like background combined have the same area as the signal histogram.

Table 4.1.: The input variables for the DsECal and barrel ECal BDTs.

Variable	Description
Charge Mean	The mean charge of all hits in the cluster
Charge Skew	The charge skew of hits in the cluster
Charge RMS	The RMS of the charge of hits in the cluster
Charge Sum	The summed charge of all hits in the cluster
Number of layers hit	The number of ECal layers containing a hit from the cluster
Most upstream layer hit	The layer closest to the centre of the ND280 that contains a hit from the cluster
Most downstream layer hit	The layer furthest from the centre of the ND280 that contains a hit from the cluster
Pointing Direction	The charge weighted alignment of hits along the cluster's principle axis
Maximum perpendicular bar hit	Bar closest to the outermost edge of the ECal module that contains a hit from the cluster, for bars perpendicular to the neutrino beam
Maximum parallel bar hit	Bar closest to the outermost edge of the ECal module that contains a hit from the cluster, for bars parallel to the neutrino beam
Minimum bar hit	Bar closest to the innermost edge of the ECal module that contains a hit from the cluster
PID Angle	The angle between the cluster's direction and the axis normal to the ECal module's face, constructed to be less than 90°

Continued on Next Page...

Variable	Description
PID ShowerAngle	The opening angle of the shower, defined as the arctan of the ratio of the shower's width to its length, as calculated by principle component analysis (PCA)
Thrust axis	The thrust axis of the cluster, defined in Ref. [56]
Thrust	The charge weighted alignment of the cluster hits with the cluster's thrust axis
Maximum hits per layer	The maximum number of hits in a single layer, minus one
PID Circularity	The circularity of the cluster, with a perfectly spherical cluster having a circularity of zero and a straight line having a circularity of one
Truncated Max Ratio	The highest and lowest charge hits are removed from the cluster, then the total charge per layer calculated. The ratio of the minimum layer charge to the maximum gives the truncated max ratio
Front Back Ratio	The ratio of the total charge in the outermost quarter of the cluster to the innermost quarter, where the cluster has been divided into four equal lengths along the principle PCA axis
2D Radon track angle	The 2D Radon track angle defined in Chap.3, with a separate variable for each of the six possible Radon tracks
2D Radon track distance	The 2D Radon track distance defined in Chap.3, with a separate variable for each of the six possible Radon tracks
Radon hits used	The percentage of the available hits that are added to each 2D Radon track
Number of hits in each view	The number of hits in each 2D view of the cluster

BDT performance

The BDT performance is monitored using the tools provided by the TMVA package. In Fig. 4.8 the DsECal BDT classifier is shown for both signal events, in blue, and background events, in red. The points show the BDT response from the training sample whilst the filled histograms show the response from the testing event sample. The training sample output agrees well with the testing sample output, and a good separation is achieved between signal and background events. This plot is shown for the barrel ECal BDT in Fig. 4.9, which also shows good agreement and separation.

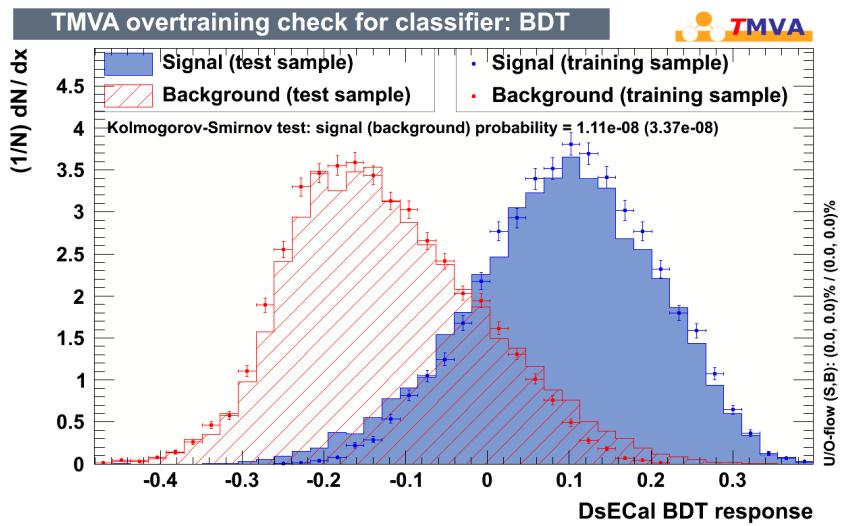


Figure 4.8.: DsECal BDT classifier response for both the training sample (points) and testing sample (histogram). Signal events are shown in blue and background events are in red.

The ROOT TMVA package provides tools to optimise the cut value applied to the classifier. By scaling the background and signal to match the numbers expected in the MC after the pre-selection cuts, a cut value of -0.05 maximises the significance of the DsECal selection, shown in Fig. 4.10. The significance is defined here as the number of selected signal events divided by the square root of the total number of events selected. An equivalent plot for the barrel ECal is shown in Fig. 4.11, where the optimum cut value is -0.08 .

These two selections were applied to an independent sample of 4.3×10^{19} POT of beam and sand MC, the results of which are shown in Table 4.2 for the DsECal and Table 4.3 for the barrel ECal.

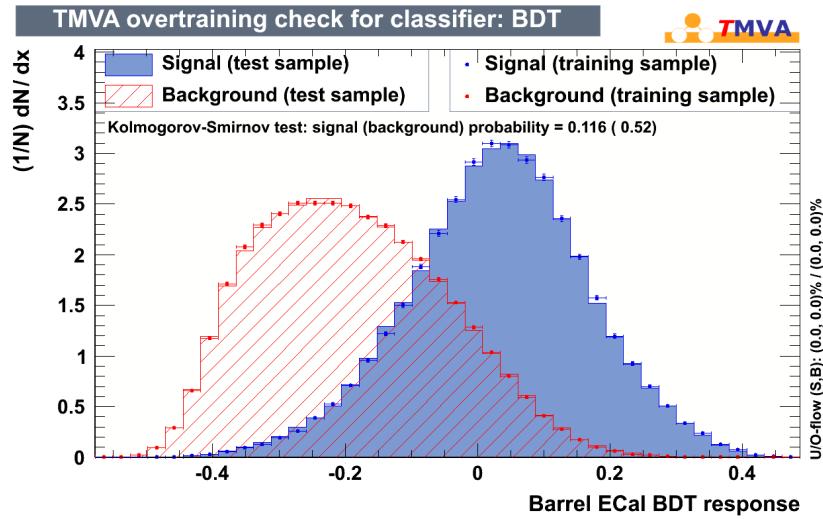


Figure 4.9.: Barrel ECal BDT classifier response for both the training sample (points) and testing sample (histogram). Signal events are shown in blue and background events are in red.

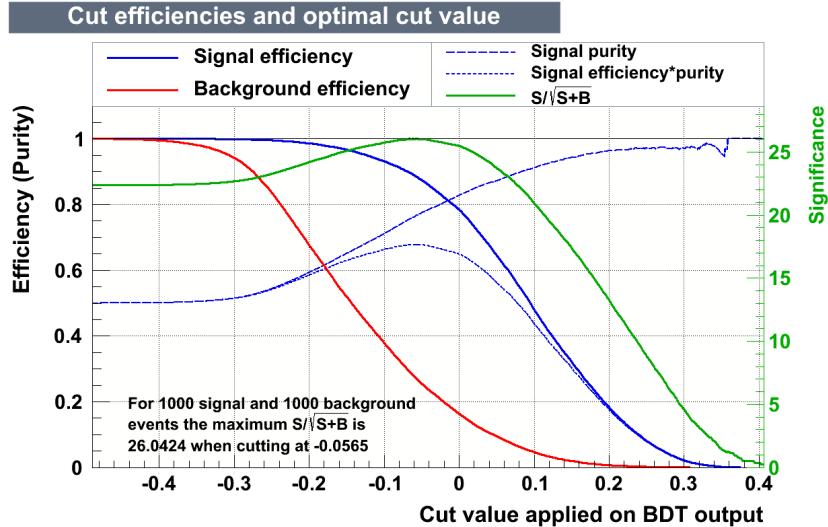


Figure 4.10.: DsECal ECal classifier cut tuning, where the ratio of signal to background events has been altered to match that expected after the pre-selection cuts

4.3.3. Data-MC comparisons

The pre-selection cuts and optimised BDT cuts described above were decided upon based on MC alone. The pre-selection was minimal, selecting well reconstructed ECal clusters and cutting out obvious entering backgrounds from the DsECal, whilst the BDTs used every variable that provided some separation between signal and background events.

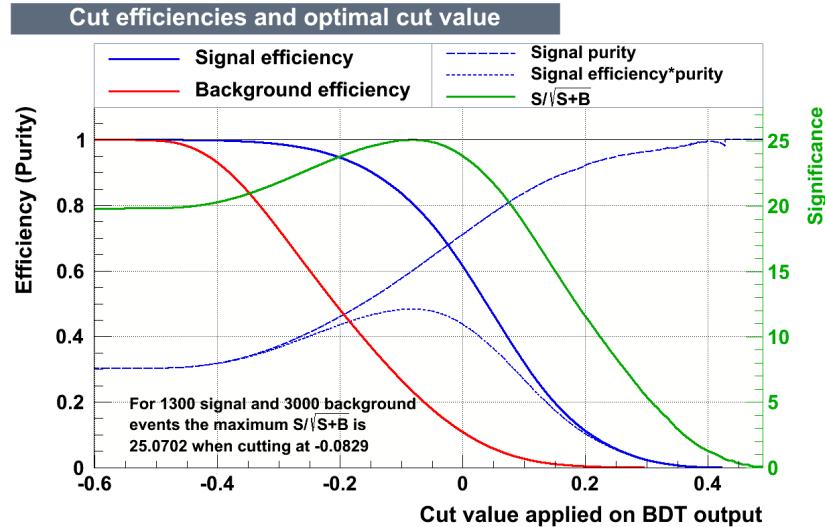


Figure 4.11.: Barrel ECal classifier cut tuning, where the ratio of signal to background events has been altered to match that expected after the pre-selection cuts

Cut	No. Signal	No. Background	Purity (%)	Efficiency (%)
Initial MC	13716	57817	19.2	100.0
No. of hits ≥ 6	13130	54886	19.3	95.7
Most upstream layer hit > 0	11568	14674	44.1	84.3
BDT response > -0.05	10239	3669	73.6	74.7

Table 4.2.: The results of applying the DsECal event selection to MC, broken down at each cut point.

This produced two ideal selections that show what can be achieved with the current ECal reconstruction.

Cut	No. Signal	No. Background	Purity (%)	Efficiency (%)
Initial MC	135041	318332	29.8	100.0
No. of hits ≥ 6	129469	295600	30.5	95.9
BDT response > -0.08	91597	49374	65.0	67.8

Table 4.3.: The results of applying the barrel ECal event selection to MC, broken down at each cut point.

A comparison was made between the BDT input variable distributions in beam MC and beam data, to show that the MC used to train the BDTs was describing the observed data well. The comparison used 1.3×10^{20} POT of sand and beam MC and 1.40×10^{20} POT of Run III beam data, with the resulting variable distributions normalised to have unit area before being compared.

Fig. 4.12 shows the average Z position (the mean position of the cluster along the beam axis) of all reconstructed clusters in the right barrel ECal, for both MC, Run II (3.0×10^{19} POT) and Run III data. There is a clear difference in the shape of the distributions, with a sharp spike in Run III data at 1400 mm that is not in either of the other two samples. The collaboration is still investigating this discrepancy, but it appears to be due to two dead front end electronics boards. The boards were damaged between Run II and Run III, but the detector simulation was not updated to reflect this. These two boards control up to 128 readout channels and losing these channels will alter the properties of reconstructed clusters in the right ECal for Run III data. As a consequence, the right ECal was removed from this analysis.

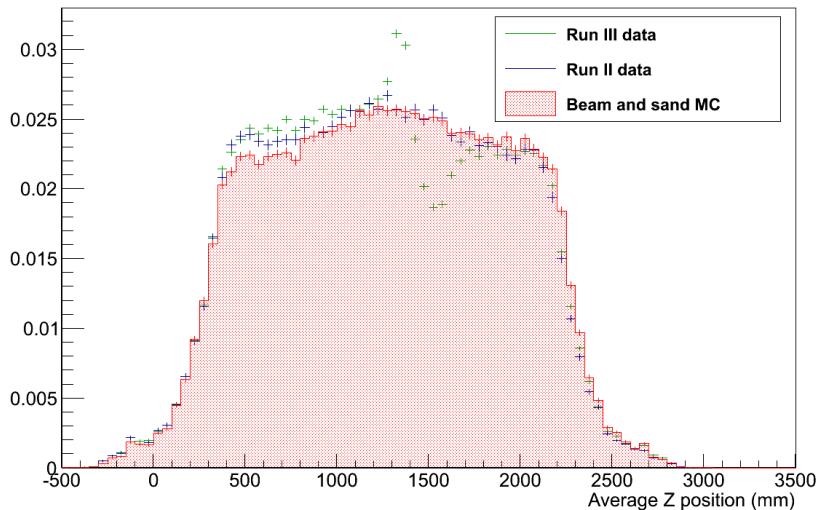


Figure 4.12.: Average Z position of all reconstructed clusters in the right barrel ECal. Histograms with unit area are shown for MC (red), Run II data (blue points) and Run III data (green points)

When applying the minimal pre-selection cuts a number of differences were observed between the shape of the BDT input variable distributions in data and MC. An example of these differences is shown in Fig. 4.13, where the truncated maximum ratio of each cluster that passes the pre-selection cuts is shown for both data and MC. These differences can be seen in both the barrel and downstream ECal samples and appear to be caused

by an excess of track-like clusters in the ECal modules. This is shown by Fig. 4.14, which plots the circularity of each ECal cluster for both data and MC separately for the barrel and downstream ECals. Both Fig. 4.14a and Fig. 4.14b show an excess of data events with a high circularity, corresponding to track-like clusters.

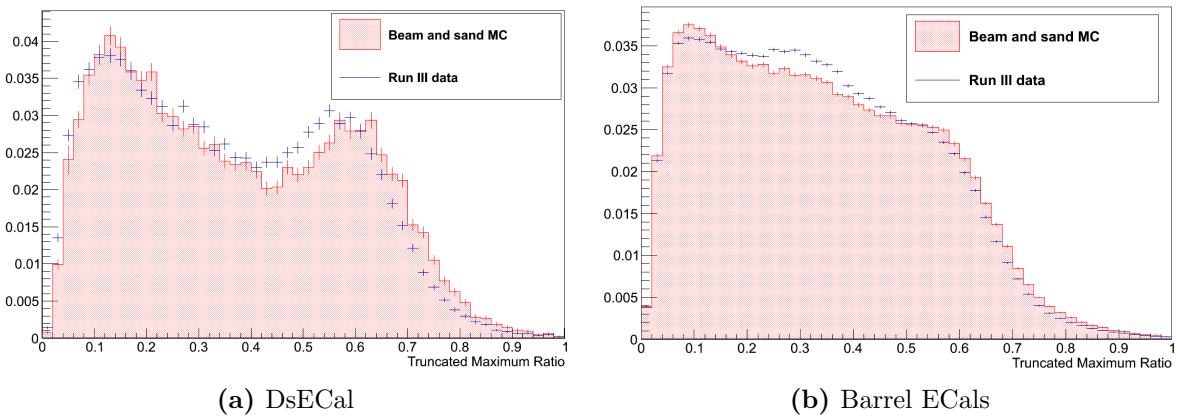


Figure 4.13.: The truncated maximum ratio of clusters that pass the minimal pre-selection cuts in (a) the DsECal and (b) the barrel ECal. The filled red histogram shows MC whilst the blue points show Run III data, both normalised to have unit area.

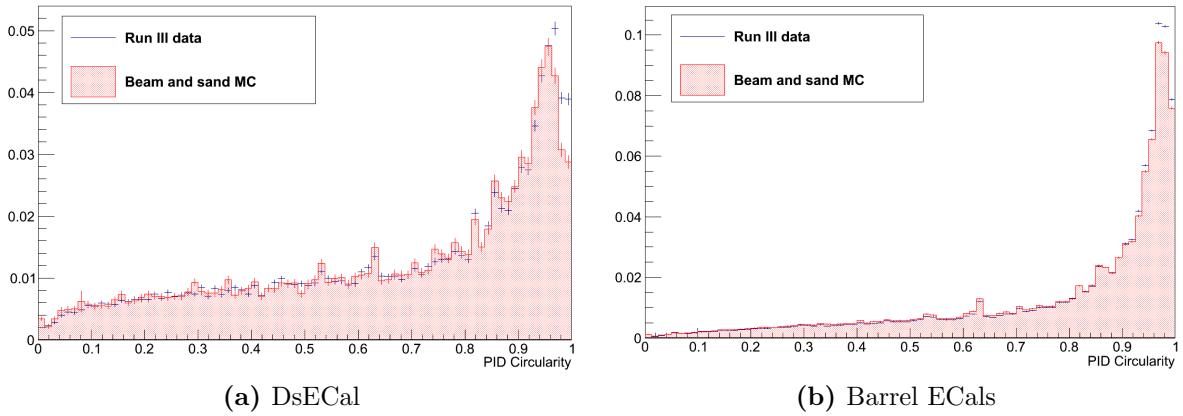


Figure 4.14.: The PID Circularit of clusters that pass the minimal pre-selection cuts in (a) the DsECal and (b) the barrel ECal. The filled red histogram shows MC whilst the blue points show Run III data, both normalised to have unit area.

Along with shape-only comparisons of the data and MC a POT normalised comparison was performed to check for gross differences in event rates. Two illustrative plots are shown in Fig. 4.15, showing the data-MC comparisons of the number of hits in each reconstructed cluster after applying the minimal pre-selection cuts. Fig. 4.15a shows

the distributions in the bottom left ECal module, the closest to the beam centre, whilst Fig. 4.15b shows the same distributions in the top right ECal module, the module furthest from the beam centre. There is a clear difference between the two, with an excess of data in the bottom left ECal compared to an excess of MC in the top right ECal.

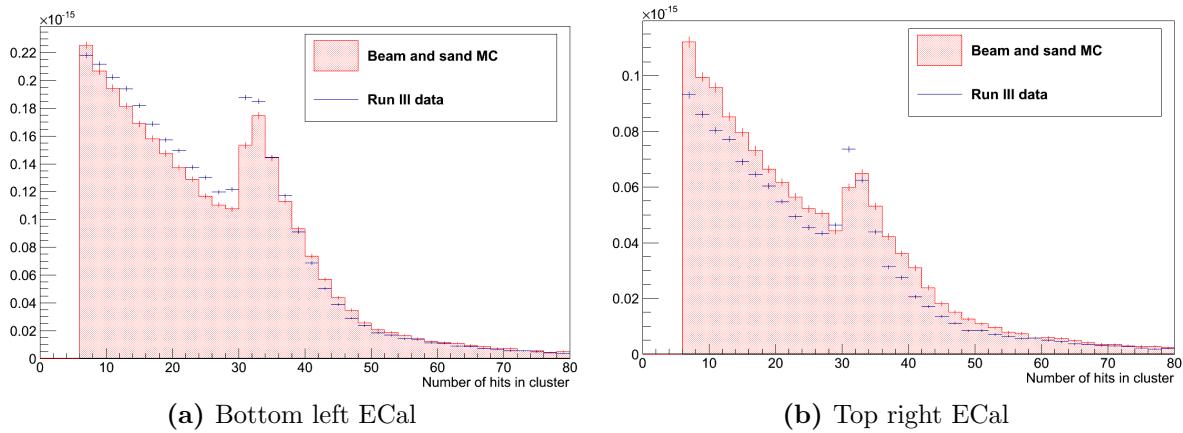


Figure 4.15.: The number of hits contained in clusters that pass the minimal pre-selection cuts in (a) the bottom left barrel ECal and (b) the top right barrel ECal. The filled red histogram shows MC whilst the blue points show Run III data, both normalised to show the number of expected events per POT.

From these data-MC comparisons it can be seen that there are beam-induced, track-like clusters present in the data that are not being modelled by the MC. These are most likely being created by muons from neutrino interactions outside the central tracking region of the ND280 detector. This region may not be correctly modelled in the ND280 detector simulation, the sand MC may not be correct or there could be a mis-modelling of the neutrino beam shape by the neutrino beamline simulation. Investigating this goes beyond the scope of this thesis, but it must be understood for future analyses.

4.4. Final event selection

The observed shape differences between the data and MC must be minimised in order to have confidence in the output of any multi-variate classifier. To do this, more stringent pre-selection cuts, which are listed in Table 4.4, were used to remove major areas of difference. These cuts both reduce the target volume of the ECal modules, by requiring no hits in any of the bars at the edges of the module, and cut out very straight tracks. The new cuts also require at least three hits in each ECal view, to remove low energy

	DsECal	Barrel ECals
Number of hits in each view	> 2	> 2
χ^2_{MIP}	> 100	> 100
Maximum perpendicular bar hit	< 49	< 94
Maximum parallel bar hit		< 57 or < 37
Minimum bar hit	> 0	> 0
Thrust	< 0.99825	< 0.99825
Most upstream layer hit		> 0

Table 4.4.: The stricter pre-selection cuts for the DsECal and barrel ECal event selections. The exact value used for the maximum parallel bar cut depends upon the barrel ECal module containing the cluster, but it is the outermost bar in each case.

photon clusters more effectively. To cut out the straightest tracks, cuts are made against the thrust and χ^2_{MIP} variable, shown in Fig. 4.16 for the DsECal. The χ^2_{MIP} is calculated as

$$\chi^2_{\text{MIP}} = \sum_{i=1}^{\text{NLayers}} \frac{(Q_i - Q^{\text{MIP}})^2}{\sigma_{Q^{\text{MIP}}}^2}, \quad (4.7)$$

where NLayers is the number of layers in the ECal module, Q_i is the summed charge in layer i from the cluster, Q^{MIP} is the mean charge deposited in a single layer by a minimally ionising particle and $\sigma_{Q^{\text{MIP}}}$ is the error on that quantity, both calculated from particle gun MC.

Fig. 4.16 shows the data-MC comparisons for the thrust and χ^2_{MIP} variables while Fig. 4.17 shows the same variables broken down according to whether the events are signal or background. Removing any event with a thrust above 0.99825 corresponds to removing the highest three bins in Fig. 4.16a and Fig. 4.17a, which show that this will remove the worst of the data-MC discrepancy whilst preserving the majority of the signal. Similarly for the χ^2_{MIP} variable in Fig. 4.16b and Fig. 4.17b, events with a χ^2 below 100 are removed.

After applying these stricter pre-selection cuts only those variables whose shape agreed between data and MC were chosen as inputs for newly trained DsECal and barrel ECal

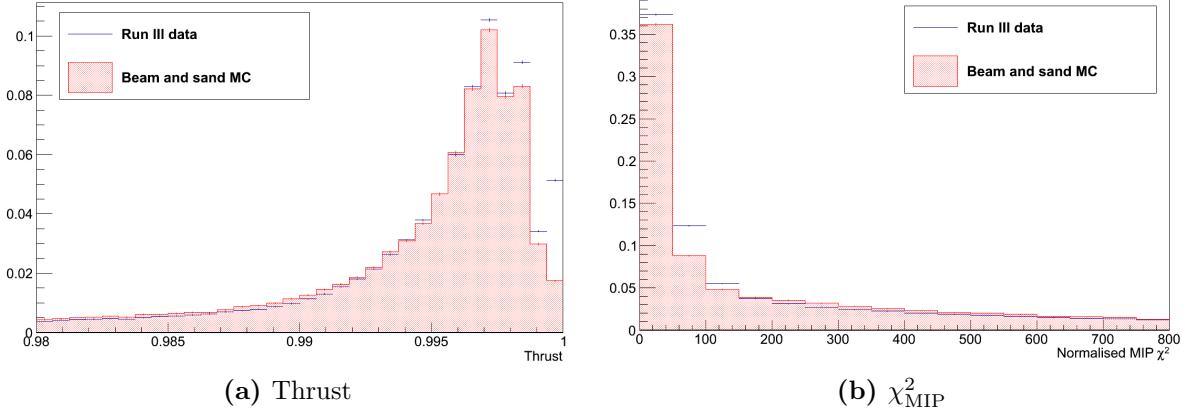


Figure 4.16.: Histograms showing data-MC comparisons of the calculated (a) thrust and (b) χ^2_{MIP} for clusters in the DsECal. The filled red histogram shows MC whilst the blue points show Run III data, both normalised to have unit area.

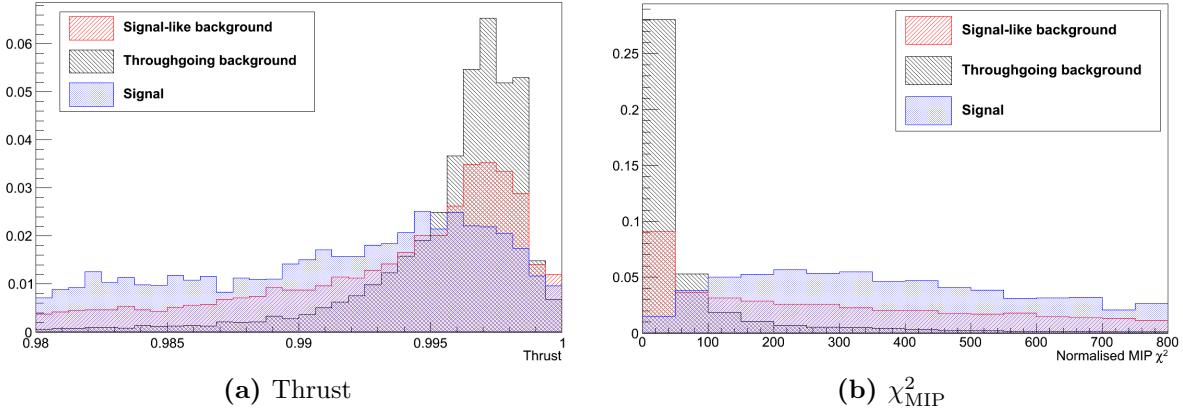


Figure 4.17.: Histograms showing signal–background comparisons of the calculated (a) thrust and (b) χ^2_{MIP} for clusters in the DsECal. The dotted blue histogram shows signal interactions, the black histogram shows through-going backgrounds and the red histogram shows signal-like backgrounds. The two background histograms have a combined area equal to the signal histogram.

BDTs. These variables are listed in Table 4.5, with those showing the greatest remaining data-MC disagreement in the DsECal and barrel ECal selections shown in Fig. 4.18.

4.4.1. Final BDT performance

A new pair of BDTs for the downstream and barrel ECals were trained on the MC sample described in Sec. 4.3.2, using the stricter pre-selection and input variables described

DsECal	Barrel ECals
Charge Skew	Charge Skew
Maximum perpendicular bar hit	Maximum perpendicular bar hit
Minimum bar hit	Minimum bar hit
Thrust	Thrust
PID Circularity	PID Circularity
2D Radon track angle	2D Radon track angle
2D Radon track direction	2D Radon track direction
Number of hits in the cluster	Number of hits in the cluster
Number of layers hit	PID Angle
X position of cluster	Object length
Y position of cluster	Average Z position of cluster

Table 4.5.: The final set of input variables to the barrel and downstream ECal BDTs.

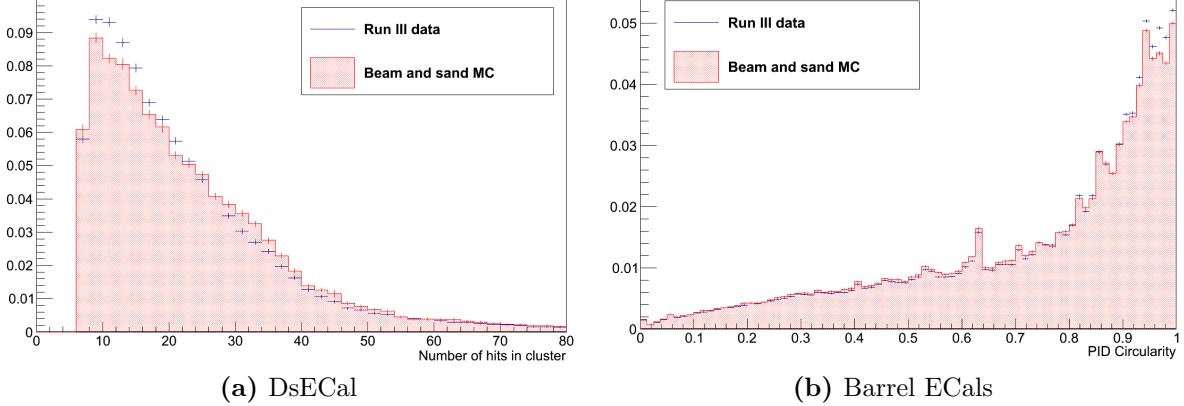


Figure 4.18.: Histograms showing the DsECal and barrel ECal variables that display the greatest remaining difference between data and MC. Fig. 4.18a shows the number of hits in DsECal clusters while Fig. 4.18b shows the PID circularity of barrel ECal clusters. The filled red histogram are MC whilst the blue points are Run III data, both normalised to have unit area.

in Sec. 4.4. The final classifier output is shown in Fig. 4.19, showing signal events in blue and background events in red, with the points indicating the BDT response to the training sample and the filled histograms the response to the testing sample. Again the

training sample agrees with the testing sample for both the barrel and downstream BDTs, but the separation between the signal and background events is much less pronounced than that seen previously. This is due to both the restricted number of input variables, giving the BDTs less information to separate the events, and the stricter pre-selection cuts removing a greater amount of the easily identifiable through-going background.

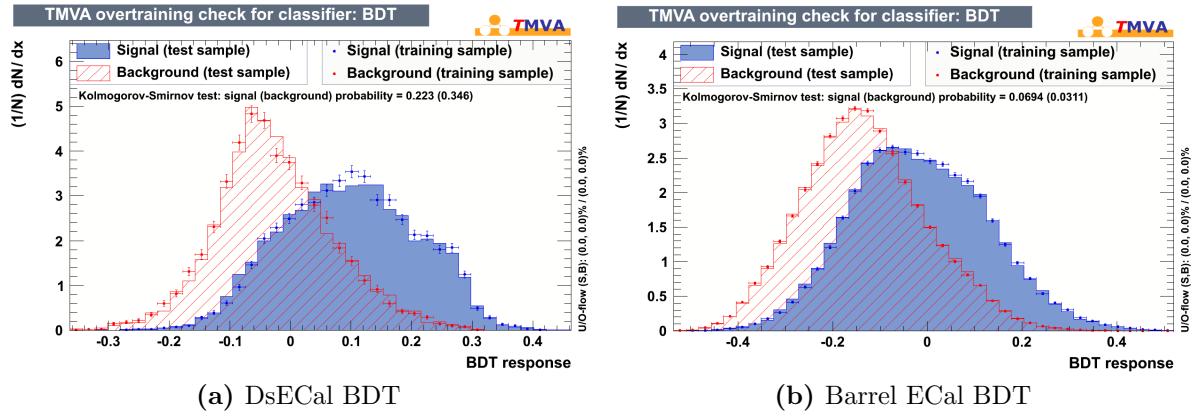


Figure 4.19.: The BDT response from the training (points) and testing (solid histogram) MC samples for (a) the DsECal and (b) the barrel ECal BDTs. The samples have been split into signal, in blue, and background, in red, which have been scaled to have equal area.

Using the same cut optimisation tool as described previously, the cut that maximises the significance of the DsECal signal selection was found to be at -0.13 , and at -0.17 for the barrel ECal signal selection. These two selections were applied to the MC described in Sec. 4.1.2, the results of which are shown in Table 4.6 for the DsECal and Table 4.7 for the barrel ECal.

The selected events are shown in Fig. 4.20 and Fig. 4.21 for the DsECal and barrel ECals respectively, plotted as a function of the incident neutrino energy. Figs. 4.20a and 4.21a show the location of the neutrino vertex that created the selected ECal cluster, demonstrating that the majority of selected events occurred in the lead or scintillator layers of the ECals. These plots also show that one of the main backgrounds in both selections is caused by particles coming from interactions in the magnet. Figs. 4.20b and 4.21b are broken down according to the identity of the nucleus that was involved in the neutrino interaction, and both plots confirm the previous findings that the majority of selected events came from interactions on lead and carbon nuclei.

Cut	No. Signal	No. Background	Purity (%)	Efficiency (%)
Initial sample	197949	1009643	16.4	100.0
No. of hits per view	179006	937811	16.0	92.2
Most upstream layer hit	156123	198428	44.0	80.4
χ^2_{MIP}	150068	178524	45.7	77.3
Minimum bar hit	128686	131191	49.5	66.3
Maximum perpendicular bar hit	111349	85044	56.7	57.3
Thrust	109100	82237	57.0	56.2
BDT response > -0.13	107134	72256	59.7	55.1

Table 4.6.: The results of applying the final DsECal event selection to MC, broken down at each cut point. The pre-selection cut values can be found in Table 4.4.

Cut	No. Signal	No. Background	Purity (%)	Efficiency (%)
Initial sample	1558520	4229420	26.9	100.0
No. of hits per view	1425380	3697540	27.8	91.5
χ^2_{MIP}	1324690	2721810	32.7	85.0
Minimum bar hit	1145490	2124880	35.0	73.5
Maximum parallel bar hit	999764	1750180	36.4	64.2
Maximum perpendicular bar hit	903426	1481530	37.9	58.0
Thrust	845561	1333270	38.8	54.3
BDT response > -0.17	735206	832443	46.9	47.2

Table 4.7.: The results of applying the final barrel ECal event selection to MC, broken down at each cut point. The pre-selection cut values can be found in Table 4.4.

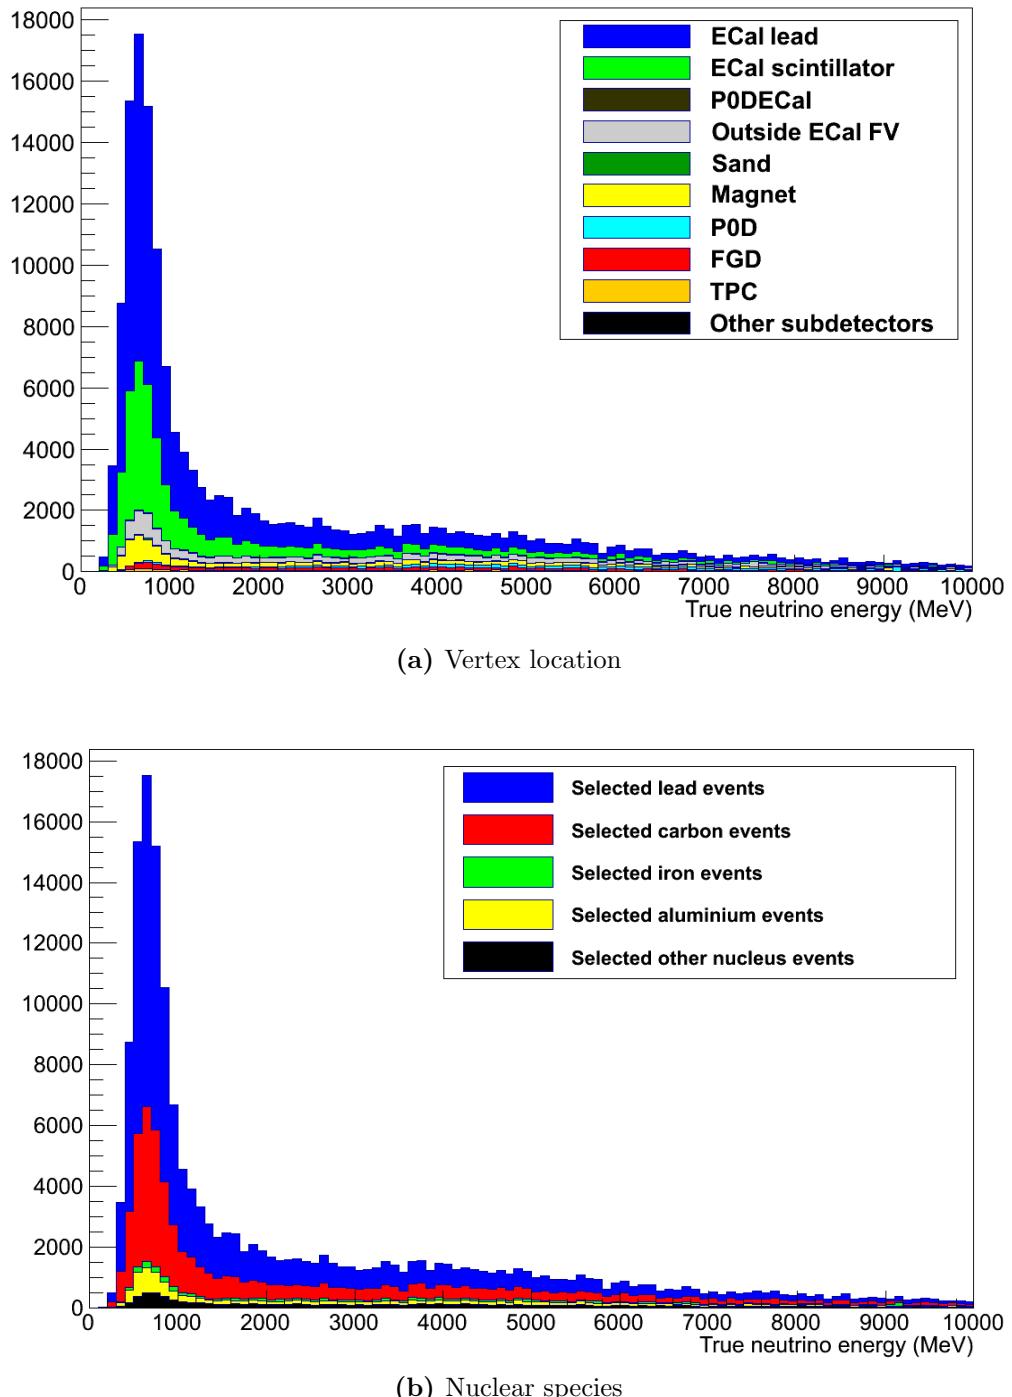


Figure 4.20.: The selected events in the DsECal, plotted as a function of incident neutrino energy. The distributions have been broken down by (a) the location of the neutrino vertex and (b) the species of nucleus that the neutrino interacted with.

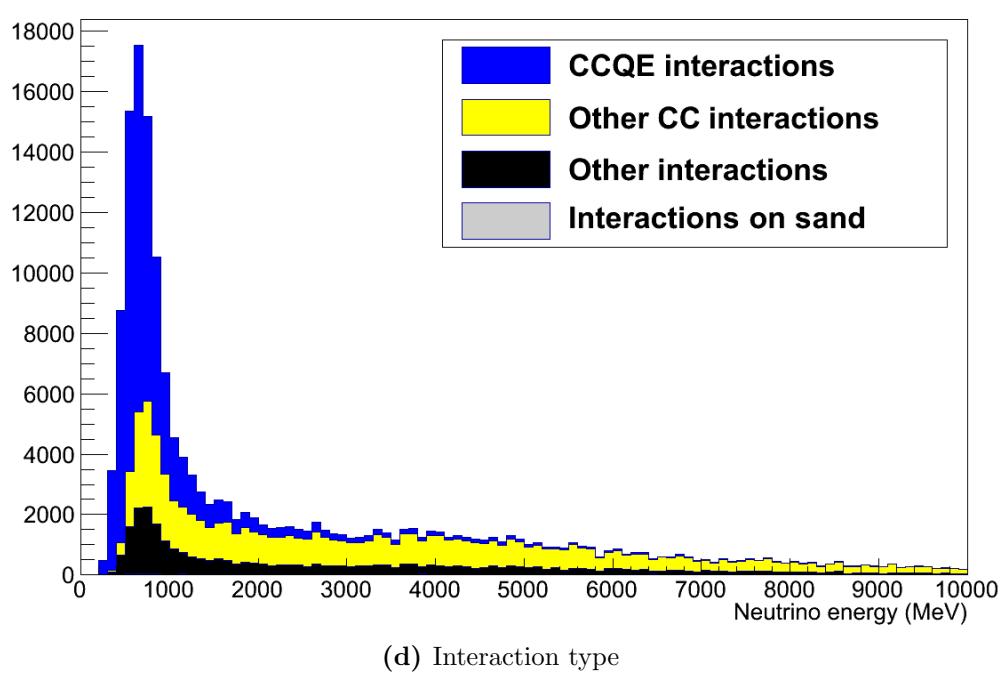
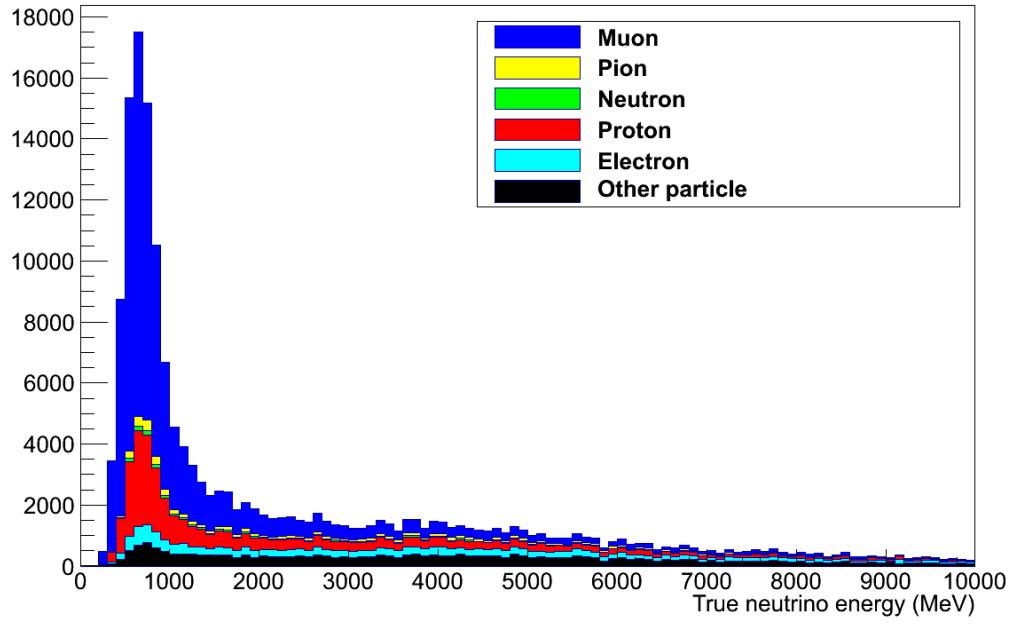


Figure 4.20.: The selected events in the DsECal, plotted as a function of incident neutrino energy. The distributions have been broken down by (c) the particle that deposited most energy in the selected ECal cluster, and (d) the type of neutrino interaction that was selected.

In Fig. 4.20c and 4.21c the selected events are broken down by the species of particle that deposited most energy in the ECal cluster. For both the barrel and downstream

ECal samples the selected clusters are created predominantly by muons and protons, which is expected since the T2K neutrino beam is largely composed of muon neutrinos and the greatest contribution to the interaction cross section at these energies comes from the CCQE channel. The final two plots, Figs. 4.20d and 4.21d, show the type of interaction that occurred at the neutrino vertex, indicating that the selected clusters were mainly created by CC interactions, as desired.

The purity and efficiency of selecting signal events versus the selected neutrino energy is shown in Figs. 4.22 and 4.24 for the downstream and barrel ECal modules respectively. The same quantities are shown in Figs. 4.23 and 4.25 as a function of Q^2 , which is defined as

$$Q^2 = -q^2, \quad q = p^\ell - p^\nu, \quad (4.8)$$

where p^ℓ is the four-momentum of the outgoing lepton and p^ν the four-momentum of the incoming neutrino. These figures show that no single part of the neutrino energy spectrum is being selected with a greater efficiency than any other part, nor any part of the Q^2 spectrum. This is important, since this thesis describes an inclusive analysis over the entire energy range available and should not be biased towards any particular part of it.

4.5. Event selection summary

This chapter has described an effective method of selecting neutrino interactions within the lead and scintillator layers of the ND280 ECals. Use of a boosted decision tree multivariate technique provided the power to separate signal interactions from background, with an ideal selection achieving 73.6% purity and 74.7% efficiency in the DsECal. Differences between the observed data and the ND280 Monte Carlo simulation restricted the number of variables that could be used by the boosted decision trees, but even so the final DsECal selection had a purity of 59.7% and an efficiency of 55.1% for selecting charged current muon neutrino interactions within the ECal lead and scintillator layers.

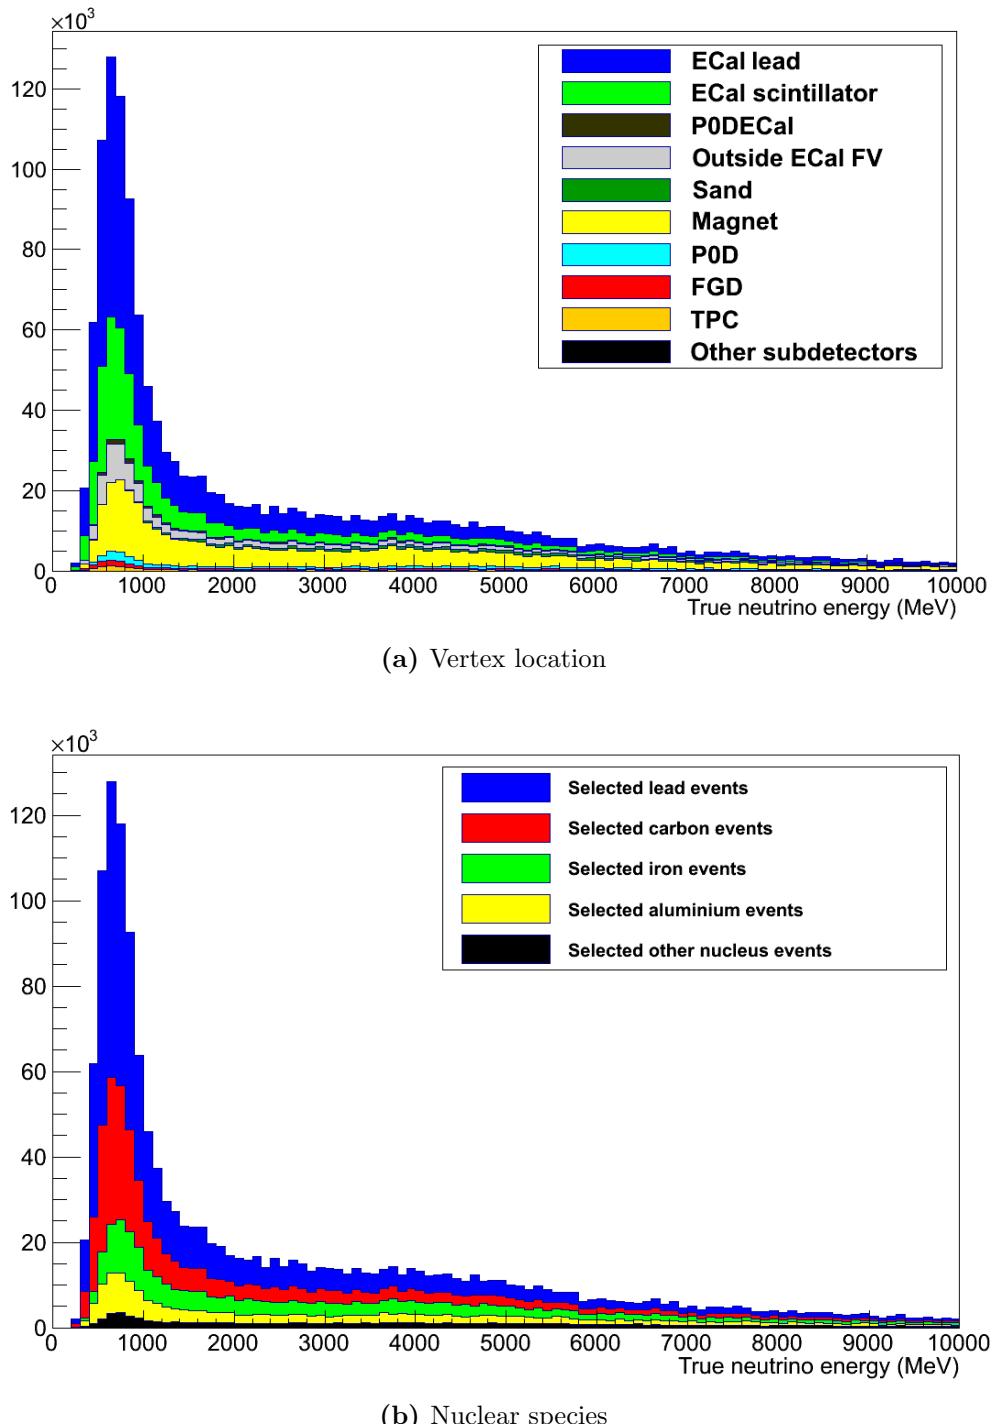


Figure 4.21.: The selected events in the barrel ECals, plotted as a function of incident neutrino energy. The distributions have been broken down by (a) the location of the neutrino vertex and (b) the species of nucleus that the neutrino interacted with.

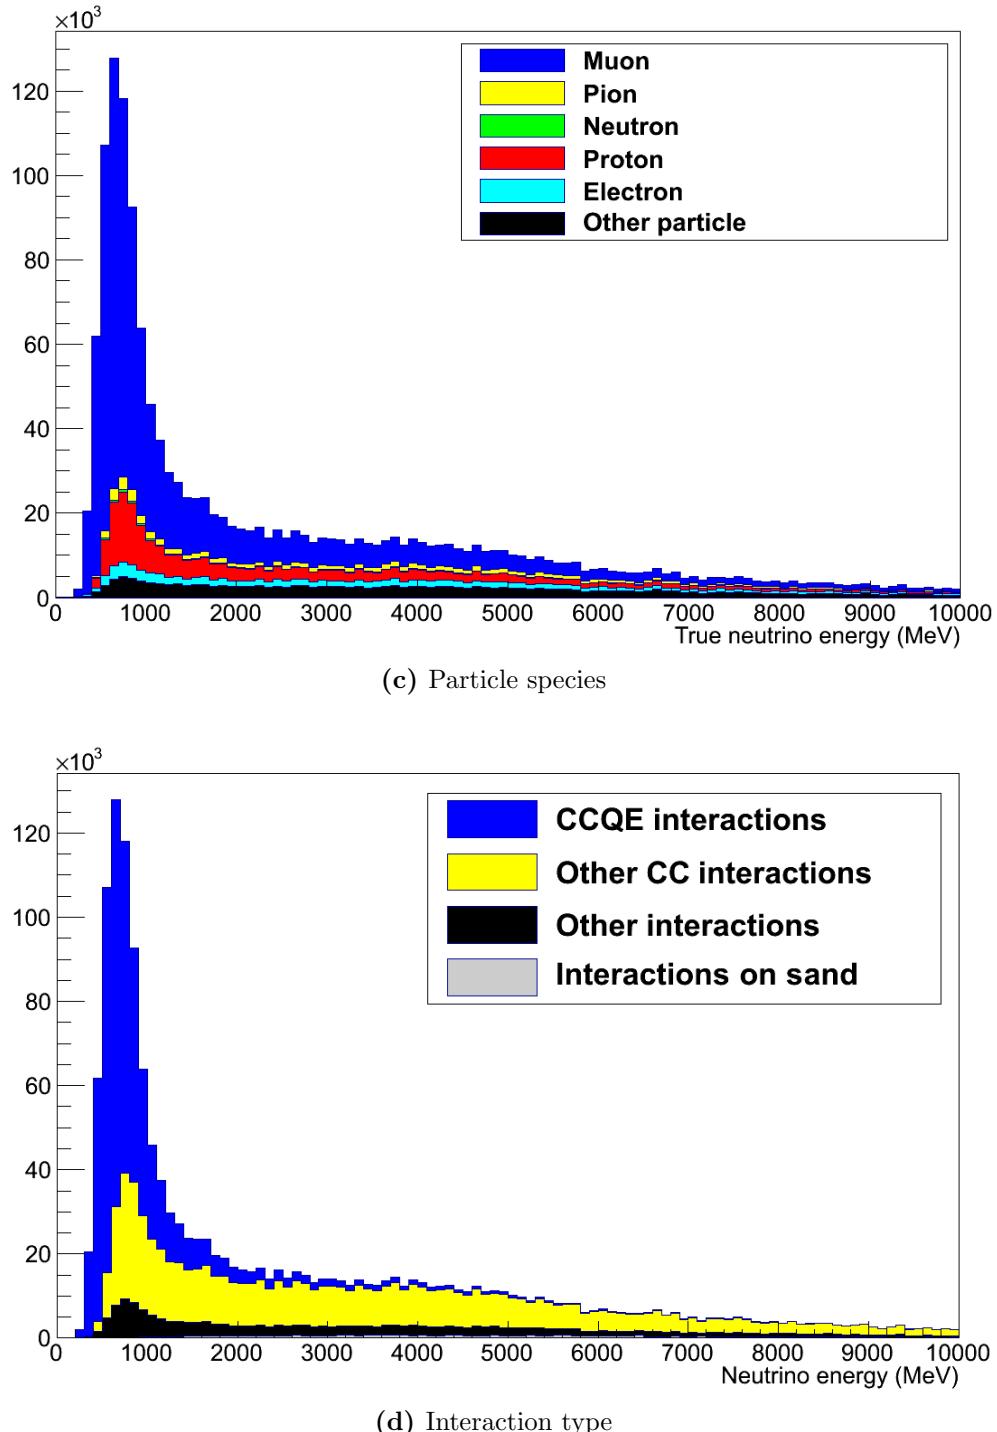


Figure 4.21.: The selected events in the barrel ECals, plotted as a function of incident neutrino energy. The distributions have been broken down by (c) the particle that deposited most energy in the selected ECal cluster, and (d) the type of neutrino interaction that was selected.

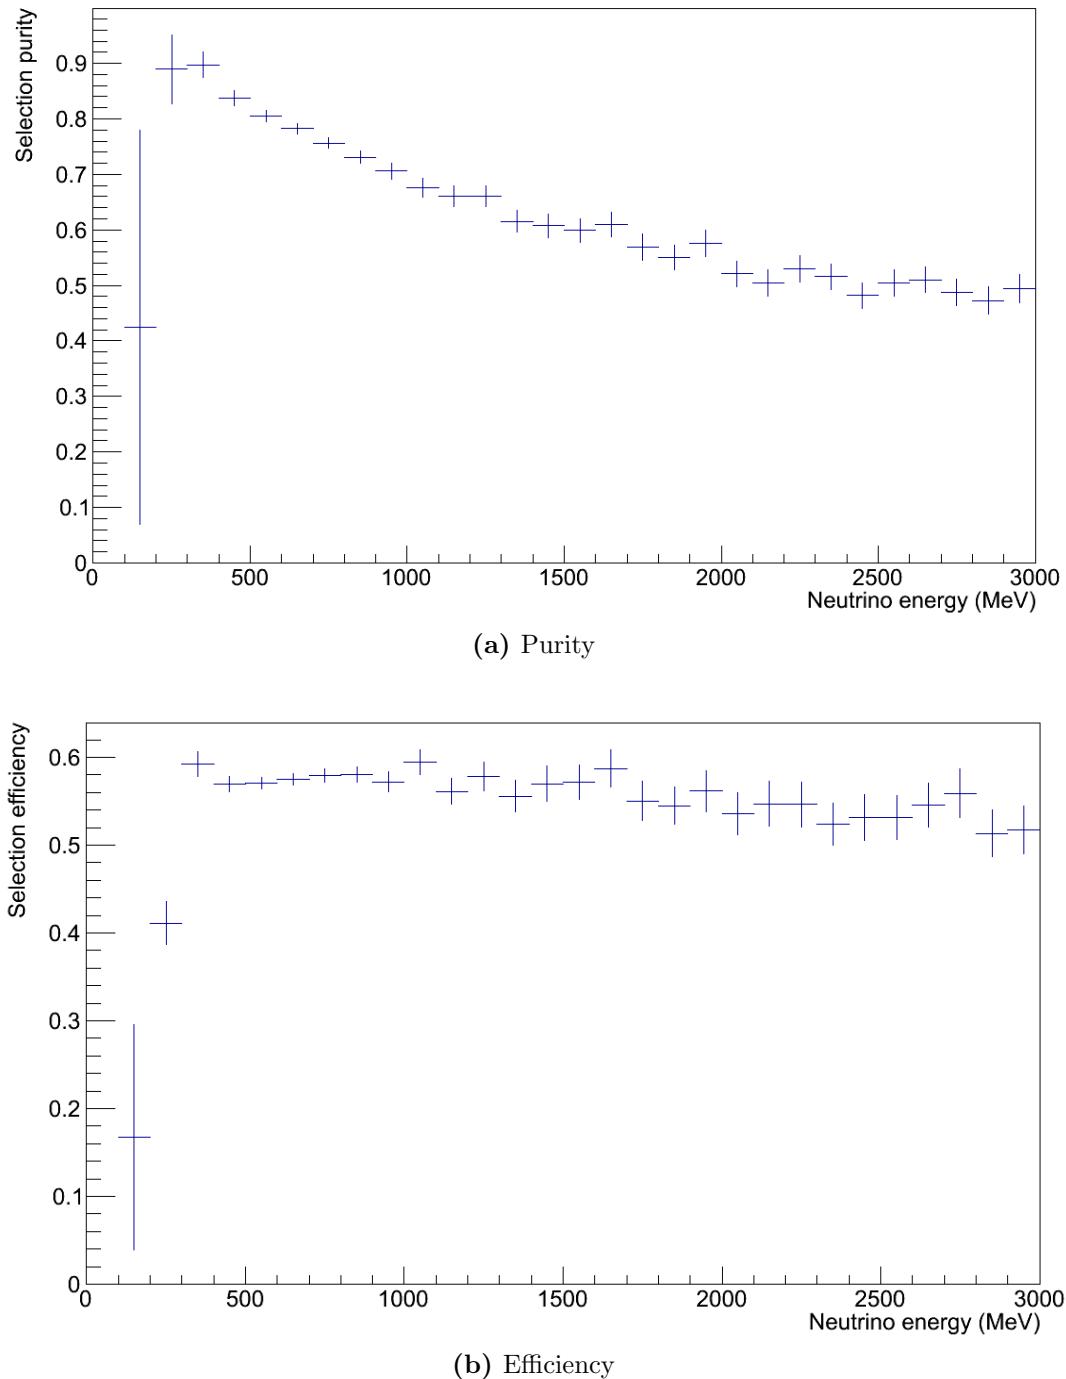


Figure 4.22.: The purity and efficiency of the DsECal event selection versus neutrino energy.

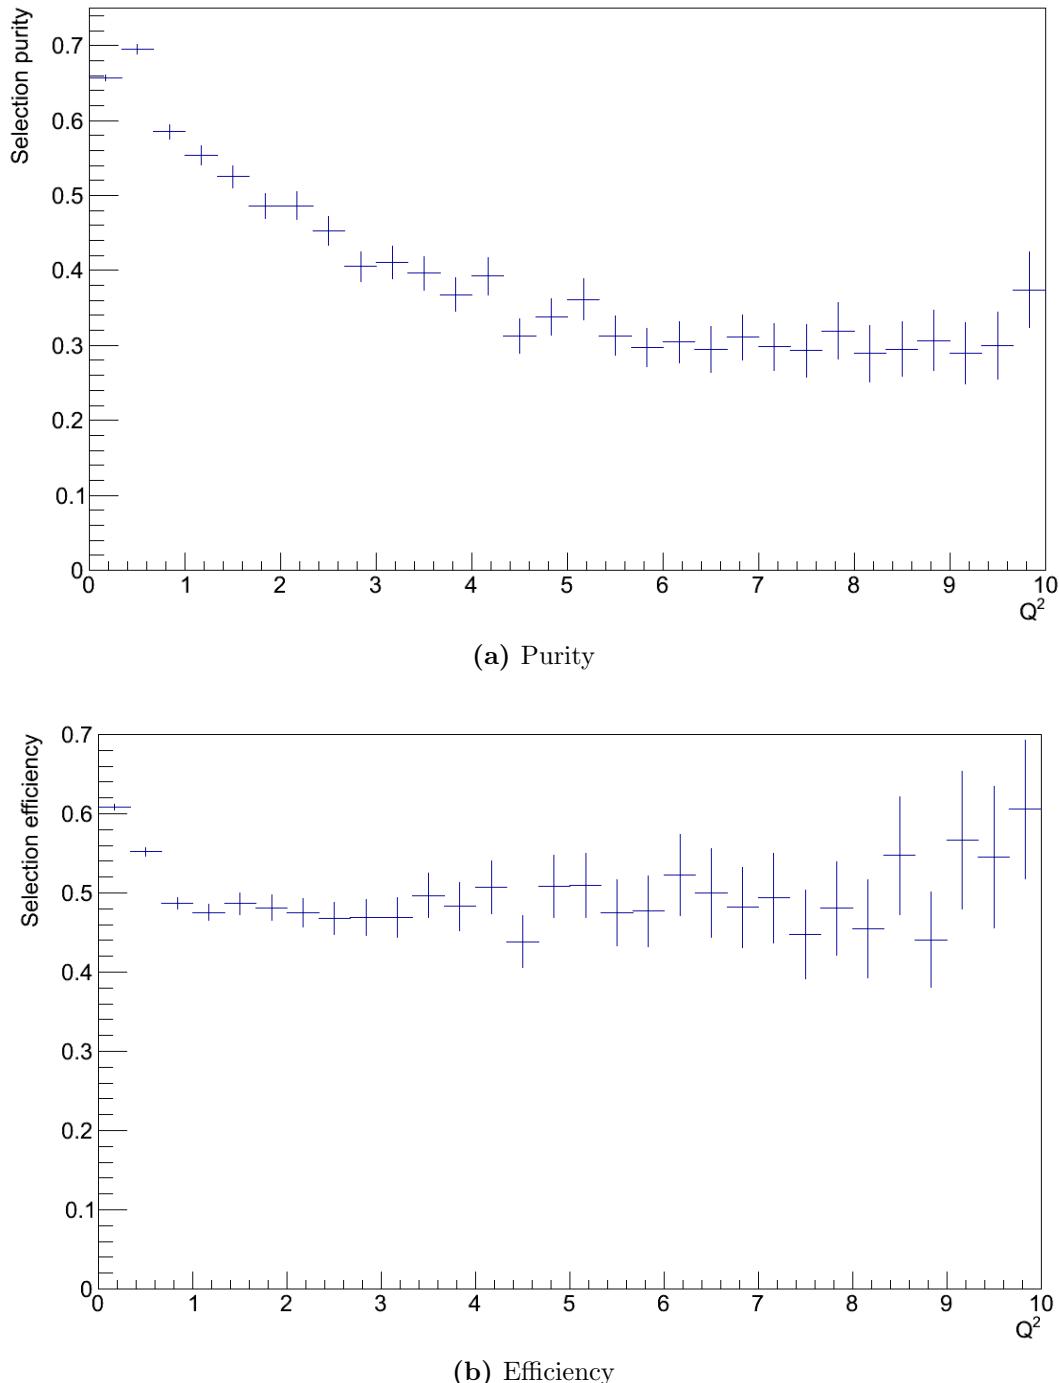


Figure 4.23.: The purity and efficiency of the DsECal event selection versus Q^2 .

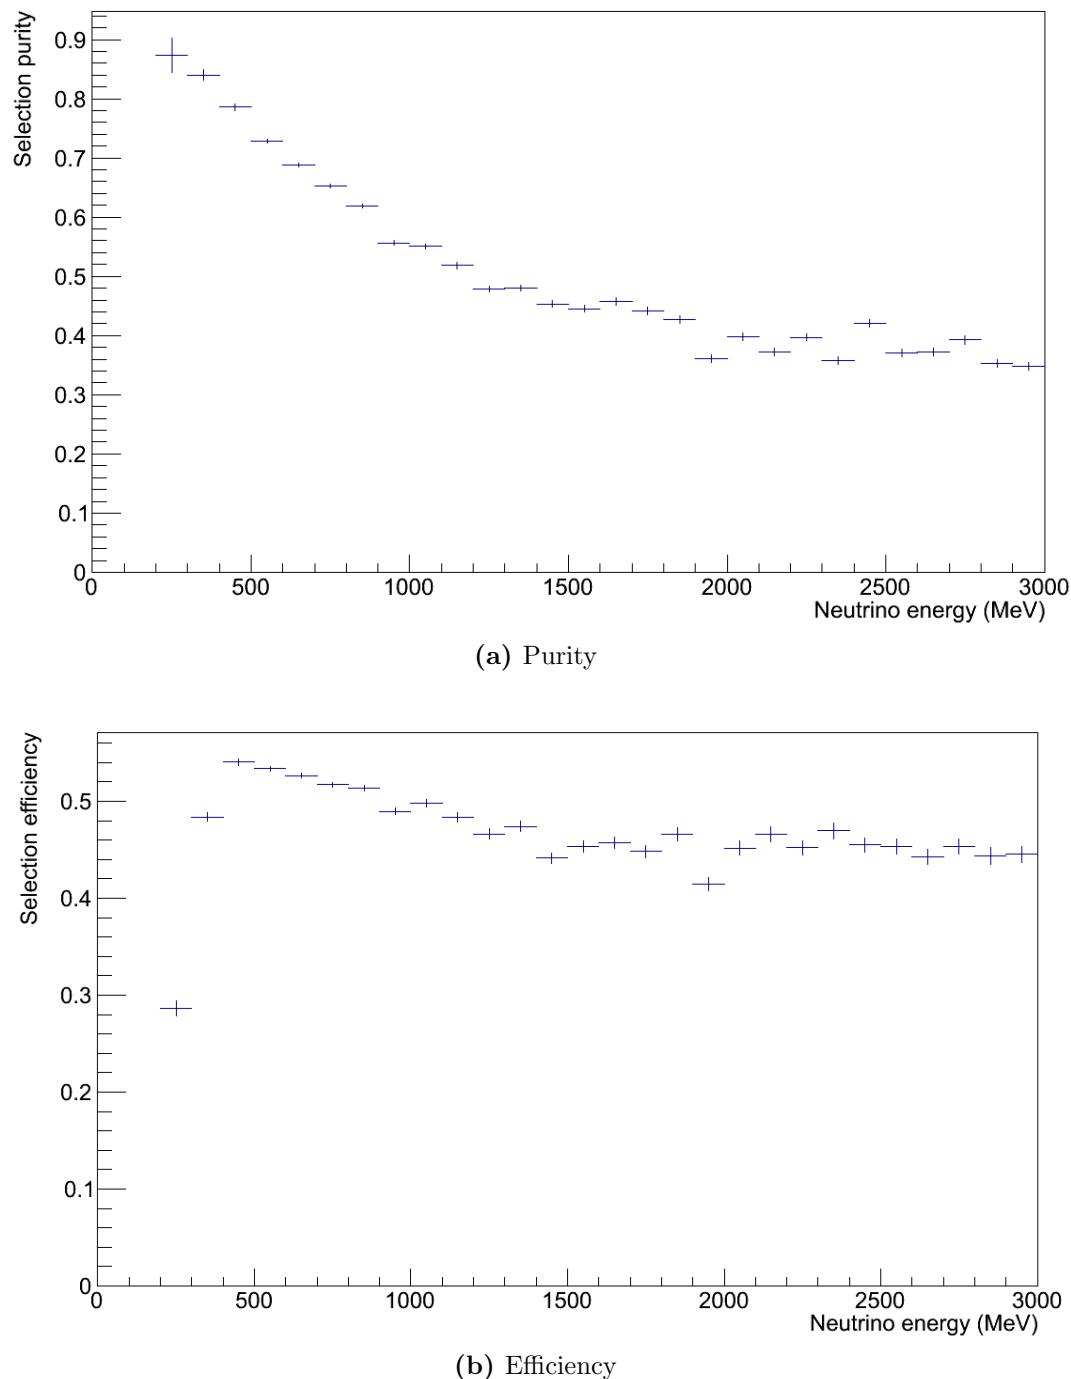


Figure 4.24.: The purity and efficiency of the barrel ECal event selection versus neutrino energy.

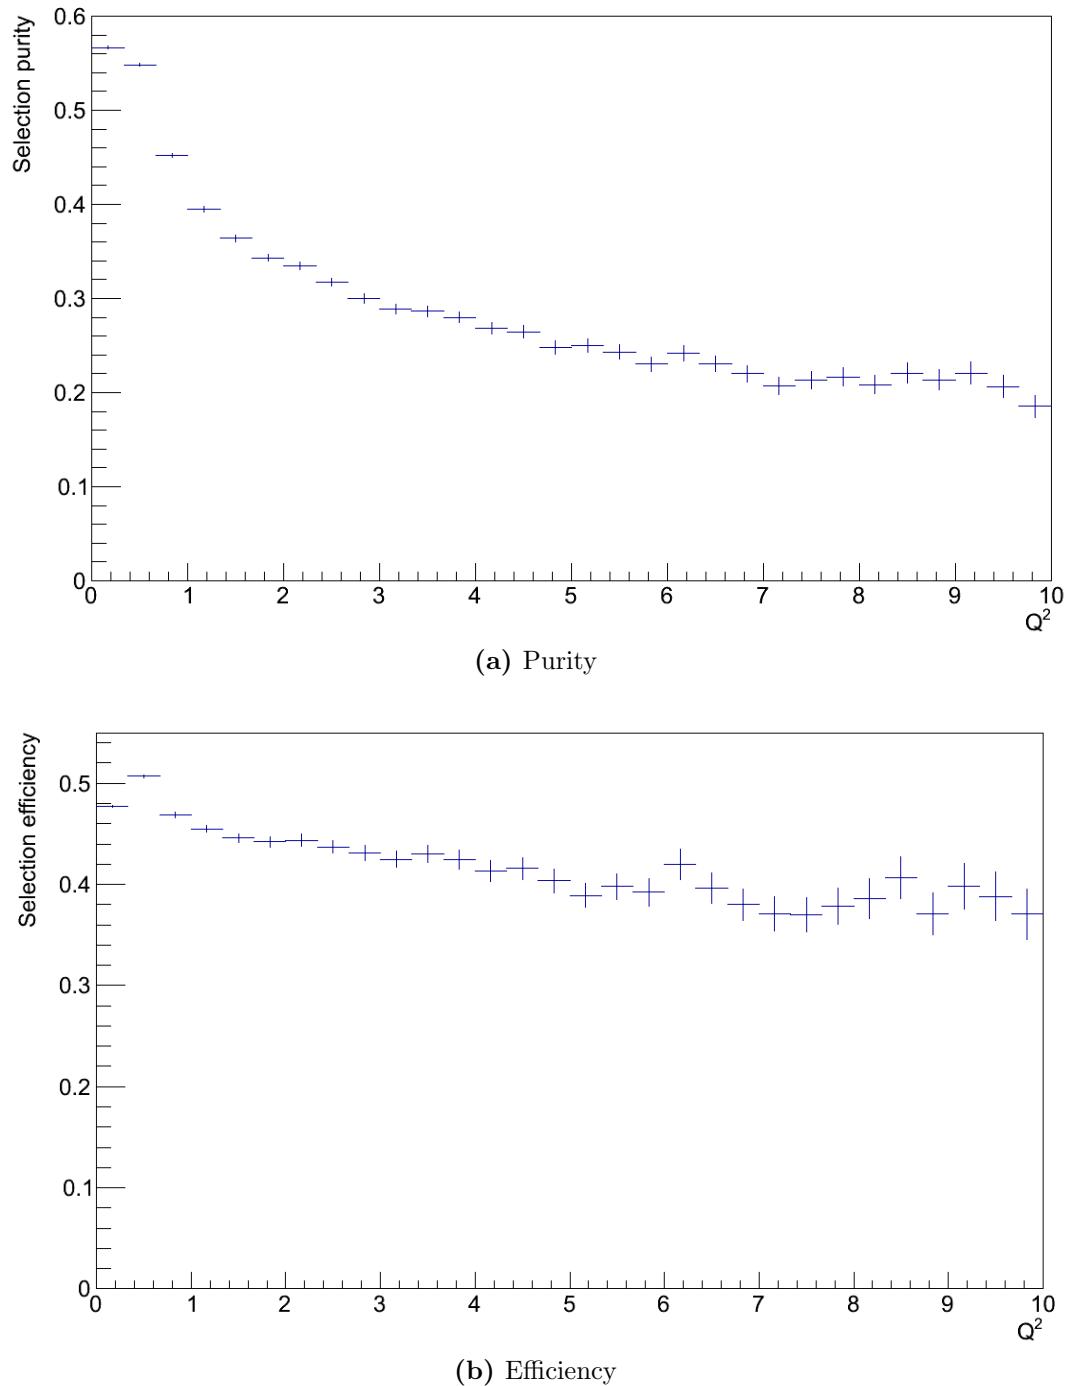


Figure 4.25.: The purity and efficiency of the barrel ECal event selection versus Q^2 .

Chapter 5.

Measuring the ν_μ -Pb charged current inclusive cross-section

The aim of this analysis is to extract the ν_μ -Pb charged current inclusive cross-section from a sample of neutrino interactions in the ND280 ECals. This chapter describes the method used to do this and presents the results of the analysis.

5.1. Analysis method

The analysis method used here was based upon similar analyses performed by the K2K [57] and SciBooNe [58] experiments. In both these analyses multiple data samples are fit simultaneously using MC templates. The content of these templates is varied using a number of normalisation parameters, and it is these normalisation parameters that contain the physics the analysers are trying to extract.

The ND280 detector is particularly well suited to performing analyses of this type, since it is formed from multiple sub-detectors, each of which contains a different target material, or is more sensitive to a different neutrino interaction mode. These offer a number of samples that can be used to constrain backgrounds contained in the main event selection, where the constraint can come either from the sample itself or from correlations between different samples.

5.2. Fitting to extract the cross-section

This analysis minimises a χ^2 function in order to fit the MC samples to the data. This χ^2 is defined as

$$\chi^2 = \sum_{i, j}^{\text{No. of samples}} \Delta N_i^T \cdot V_{ij}^{-1} \cdot \Delta N_j, \quad (5.1)$$

with V being the covariance matrix for both systematic and statistical uncertainties in each sample and

$$\Delta N_i = N_i^{\text{data}} - N_i^{\text{exp}}, \quad (5.2)$$

$$\begin{aligned} N_i^{\text{exp}} = & (R^{\text{Pb}} n_i^{\text{Pb}} + R^{\text{C}} n_i^{\text{C}} + R^{\text{Pb}} n_i^{\text{Pb,other}} \\ & + R^{\text{C}} n_i^{\text{C,other}} + R^{\text{Other}} n_i^{\text{Other}} + R^{\text{External}} n_i^{\text{External}}), \end{aligned} \quad (5.3)$$

where the subscript i runs over each sub-detector sample in the fit. N_i^{data} here is the number of events observed in sample i in the data and N_i^{exp} is the number of events expected in sample i from the MC. In Eqn. 5.3, R^X is the normalisation parameter for population X whilst n_i^X is the number of events from population X in sample i , taken directly from the MC.

The fitter is allowed to vary the R^X parameters with no prior constraints, minimising the χ^2 by altering the number of expected events in each sample to match the data. The four R^X parameters were chosen to cover the different types of events that are present in the samples. The R^{Pb} parameter alters the normalisation for CC neutrino interactions on lead nuclei and R^{C} controls the normalisation for CC neutrino interactions on carbon nuclei. The remaining parameters are used to decouple the background events from the CC interactions in the ECals. The R^{External} parameter scales any selected cluster whose interaction vertex occurred in either the magnet yoke or the material surrounding the ND280 detector. Finally, R^{Other} is used to scale any event that does not fall into the previous categories.

The fit is performed using the ROOT TMinimiser framework and the Minuit2 algorithms within that, due to both its widespread use and ease of implementation.

5.2.1. Samples used in this analysis

In this analysis the event selections described in Chap. 4 are used to select samples of CC neutrino interactions on both lead and carbon nuclei in the DsECal and five of the barrel ECal modules, since the right barrel ECal module is not being used in this analysis. Unlike in Chap. 4, where CC neutrino interactions on both lead and carbon were classed as signal, the signal events referred to from now on are CC neutrino interactions on lead alone, since the aim of this analysis is to measure the ν_μ -Pb cross-section.

It was shown in Chap. 4 that these samples contain the signal events alongside backgrounds from interactions on carbon, from interactions in the UA1 magnet, PØD and Tracker, and from interactions in the material surrounding the ND280 detector (hereafter referred to as *sand muons* due to the major constituent of the substrate surrounding the ND280 pit and the particle type that reaches the detector).

To constrain the interactions on carbon nuclei a selection of CC interactions occurring in the FGD is included, taken from the T2K flux-averaged inclusive CC cross-section measurement [59]. Fig. 5.1 shows the FGD event selection when applied to MC and data versus the momentum of the reconstructed muon. The filled histograms show the signal and background interactions taken from the NEUT generator whilst the black points show the selected data events. The selection has a purity of 73.9% for selecting CC ν_μ interactions on carbon nuclei, and is described in detail in Ref. [59].

The downstream ECal will be the main signal sample in the fit, which will be used to extract the cross-section result after the fit. As shown in Chap. 4 the DsECal provides the highest purity and efficiency for selecting neutrino CC interactions on lead nuclei whilst still retaining a large number of events, making it the best sample to use for extracting the cross-section.

As the ND280 is very close to the neutrino production point, the neutrino beam width is the same order of magnitude as the XY plane of the detector. This means that, on average, each of the barrel ECal modules is at a different off-axis angle to each of the others, so will experience different beam intensities and different neutrino energy spectra, as shown in Fig. 5.2 for the top right and bottom left modules. The different locations of the ECal modules also mean that the separate event samples will contain different amounts of beam related backgrounds. By including the barrel ECal modules separately in the fit these differences allow the neutrino flux and beam related backgrounds in the DsECal sample to be constrained.

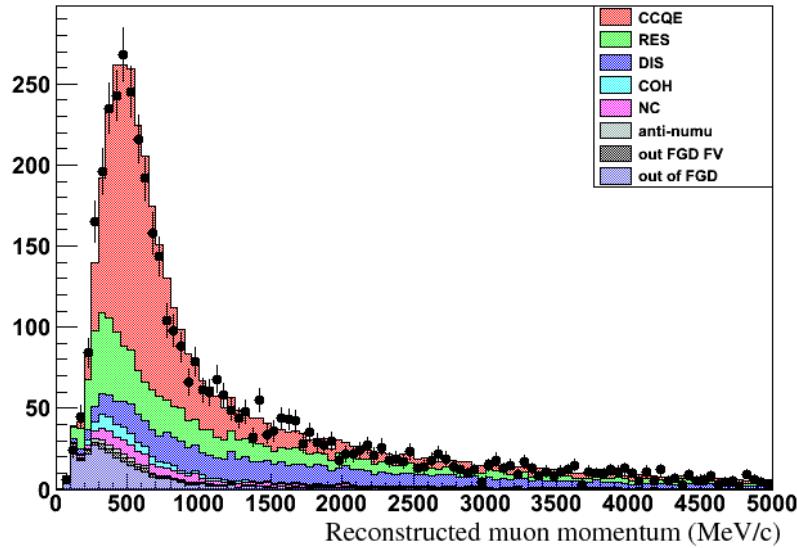


Figure 5.1.: Selected CC candidate events from interactions in the most upstream FGD. The NEUT MC prediction has been broken down into signal (labelled NEUT) and background categories, whilst the total GENIE prediction is shown by the blue line.

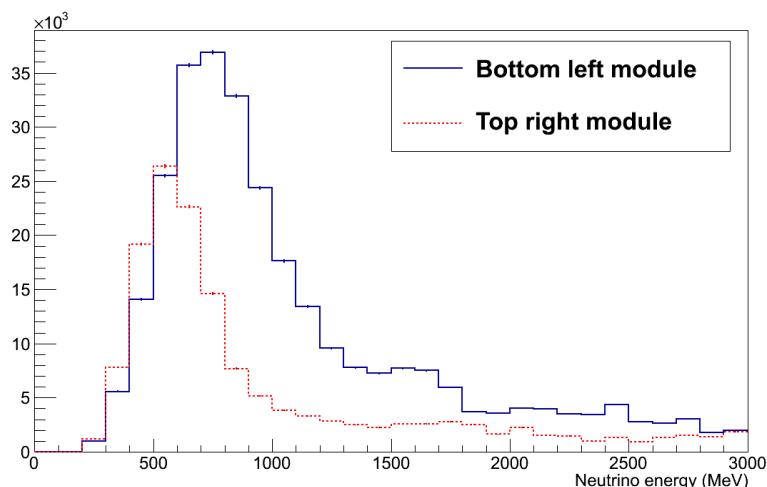


Figure 5.2.: The energy of all muon neutrinos that undergo a CC interaction in the top right (solid, red line) and bottom left (dashed, blue line) ECal modules, demonstrating the the differing normalisation and the different position and width of the neutrino beam peak.

Signal- and background-enriched samples

As explained above, the barrel ECals were included in this analysis to put a constraint on backgrounds coming from interactions in the magnet and from sand muons. To further improve their ability to do this both signal and background-enriched samples were created for each barrel ECal module. The signal-enriched samples add further constraint on the lead cross-section normalisation whilst the background-enriched samples provide a better constraint on the external backgrounds.

The background-enriched samples come from the same data set as the signal-enriched samples described in Chap. 4, and are made up of all the events that fail the BDT cut applied to the barrel ECal clusters, shown in Fig. 5.3 for the MC sample. The BDT distribution has been broken down to show the contribution from each event population in the χ^2 fit, with the ‘External’ category formed from the ‘Sand’ and ‘Magnet’ histograms from Fig. 5.3 and the ‘Other’ category formed from the ‘NC’ and ‘Other’ histograms.

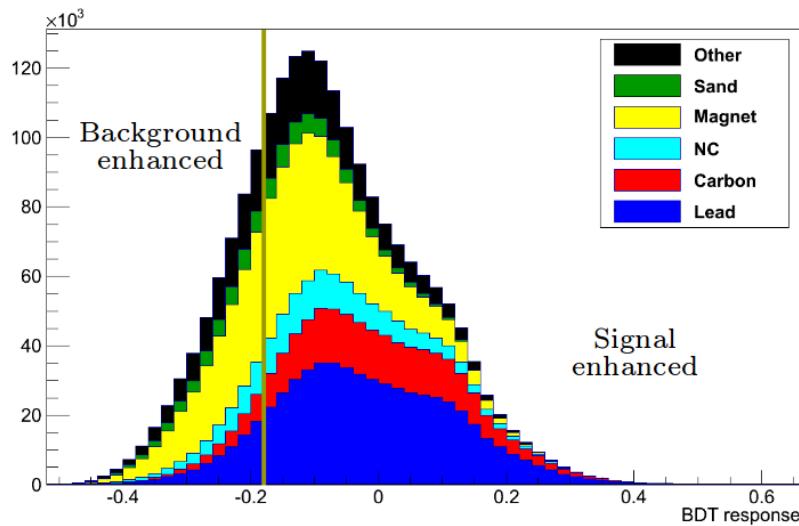


Figure 5.3.: The BDT output from all clusters in the barrel ECals that pass the final pre-selection detailed in Sec. 4.4 when applied to the MC described in Chap. 4. The cut value is shown by the vertical line, and the background- and signal- enhanced samples indicated.

The ECal samples can also be broken down to show the contribution in each from the backgrounds they are designed to constrain. Fig. 5.4 shows the ‘External’ population of the top right and left side barrel ECal modules in both the signal-enriched samples. These distributions have been split according to interaction type and plotted against the true energy of the neutrino involved. Similar plots for the remaining ECal event samples

can be found in Appendix B. Fig. 5.4 demonstrates that the ‘External’ component of the two barrel ECal samples both cover a similar range of neutrino energies and types of neutrino interaction, so are both sensitive to, and therefore able to constrain, the ‘External’ normalisation parameter. It is important to check this, since it is the small differences between the sub-detector samples that will provide the constraint on the ‘External’ normalisation. If, for example, one module was only selecting low energy CCQE interactions from the magnet, then it would be providing an incorrect constraint on the ‘External’ background in the other samples, which also contain sand muons and higher energy DIS interactions.

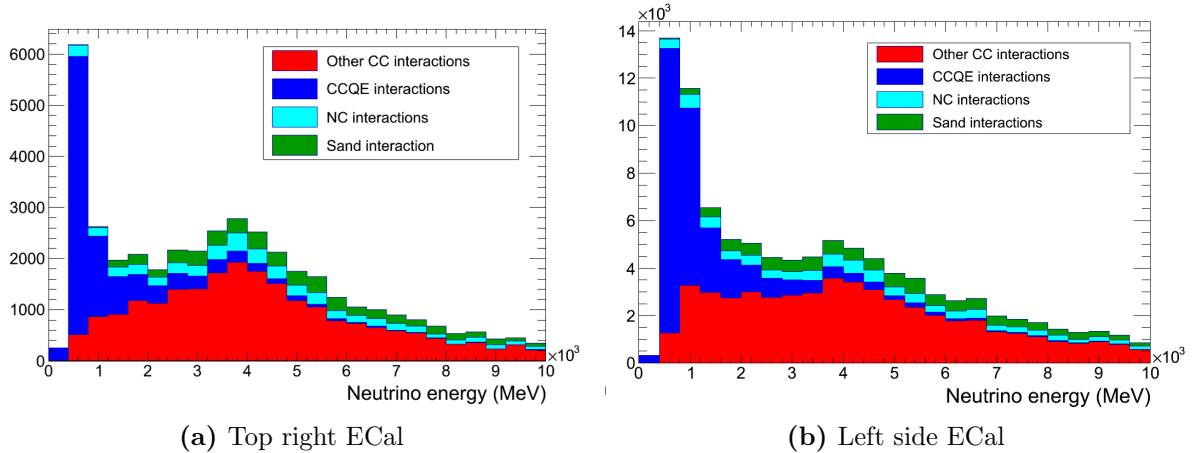


Figure 5.4.: Selected events from the ‘External’ population of (a) the top right ECal and (b) the left side ECal signal-enriched samples, broken down by interaction type.

Including the background and signal enhanced event selections from each of the barrel ECal modules with the signal selection from the DsECal and the FGD events provides a total of 12 event samples which can be used to extract the ν_μ -Pb cross-section.

BDT output data-MC comparisons

Before fitting the data, a comparison of the data and MC BDT output was performed for both the barrel and downstream ECals. Fig. 5.5 shows the results of this comparison, with the MC normalised to the data POT and split into the fit populations described above. For both distributions there is a clear excess of data events in the background-like region and an excess of MC events in the signal-like region, a fact that should be borne out in the fit results.

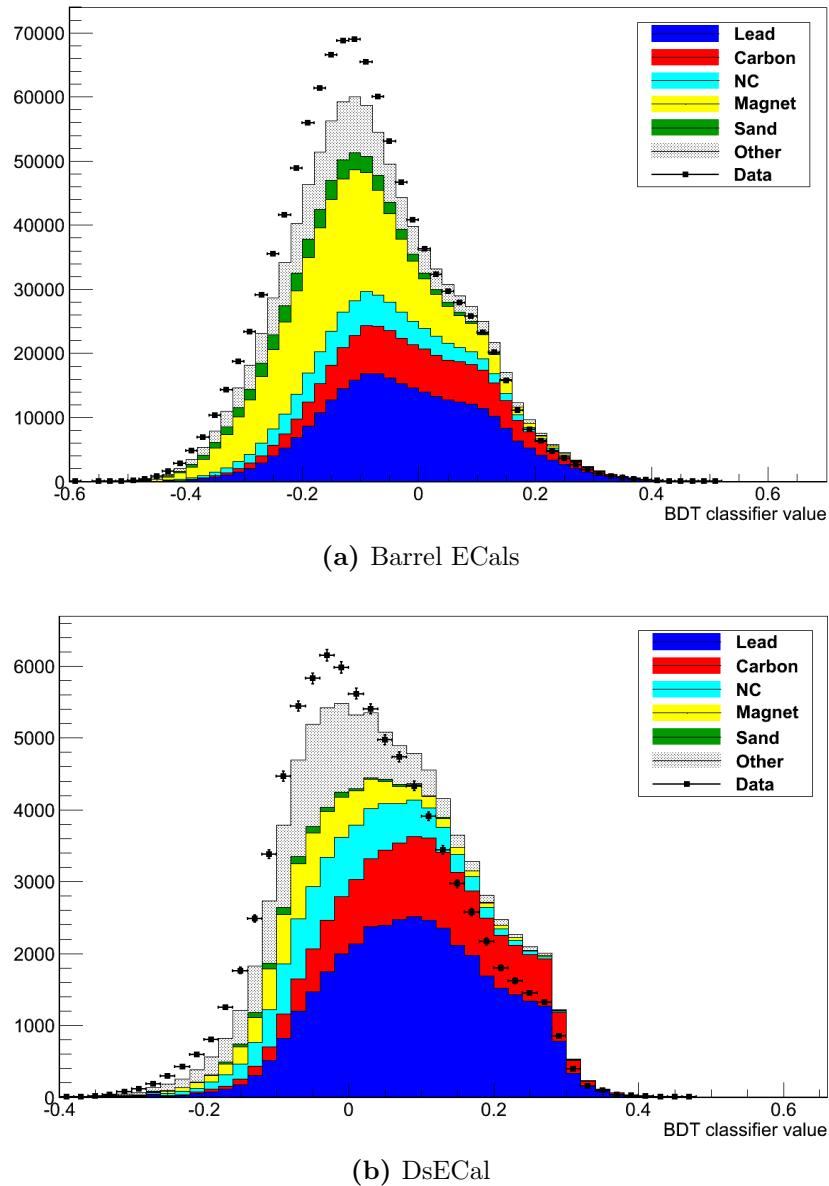


Figure 5.5.: The BDT classifier output comparing data to MC, normalised to the data POT, for both the combined barrel and downstream ECals. The MC has been split according to the populations in the χ^2 fit, with the ‘Sand’ and ‘Magnet’ categories making up the ‘External’ population.

5.2.2. Systematics

As in any analysis, the result described here depends upon a good understanding of the different systematic error sources that can affect it. In this experiment there are three categories under which the systematic errors can be grouped:

1. Flux systematics—any systematic coming from uncertainties on the neutrino flux
2. Cross-section systematics—any systematic from the current uncertainty on the neutrino interaction cross-sections
3. Detector systematics—any systematic due to the detector hardware or software

Eqn. 5.3 shows that every selected event is scaled by one normalisation term. This means that during the χ^2 fit the overall MC normalisation is free to float and it is only the shape difference between the data and MC that will constrain the fitted parameters. Therefore the covariance matrix in Eqn. 5.1 must only take into account the shape differences introduced to the samples by the different systematic source.

In each case the systematic errors in question are presented as fractional covariance matrices between each of the different sub-detector samples, with the background-enriched samples labelled as ‘Reverse’. They are then simply added together to form the total systematic error covariance matrix. These covariance matrices have all been calculated using a 4.9×10^{19} POT sample of beam and sand MC, which provides approximately 150,000 selected events.

Flux systematics

For any neutrino beam experiment, a good understanding of the neutrino flux is vital. On the T2K experiment a dedicated group of people are working on this, using a range of data: from the NA61/SHINE experiment at CERN, from previous experiments, from the proton beam monitors around the neutrino beam production station and from the INGRID detector.

The NA61/SHINE collaboration provided T2K with their measured meson angular and momentum distributions and also with their final error on those distributions. The beam working group then evaluate the effects of these uncertainties on the neutrino flux seen by the different T2K detectors. This systematic effect is combined with those calculated from the measured uncertainties of the proton beam, the magnetic focussing

horns and from differences in the production of secondary nucleons between the beam MC and previous experiments. The final output is a fractional covariance matrix, shown in Fig. 5.6, for neutrinos at the ND280 detector, binned by neutrino flavour and energy. More detail on the systematic uncertainties can be found in Ref. [50].

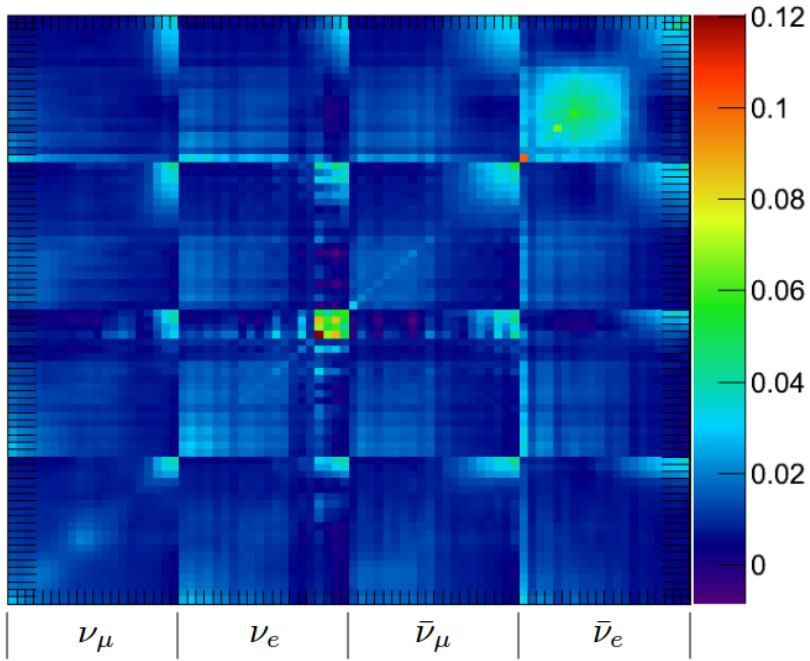


Figure 5.6.: Fractional covariance matrix showing the uncertainties on the neutrino flux. The matrix is divided into 4 groups of 20 bins, where each set of 20 bins spans the neutrino energy range and there is one set of 20 bins per neutrino species.

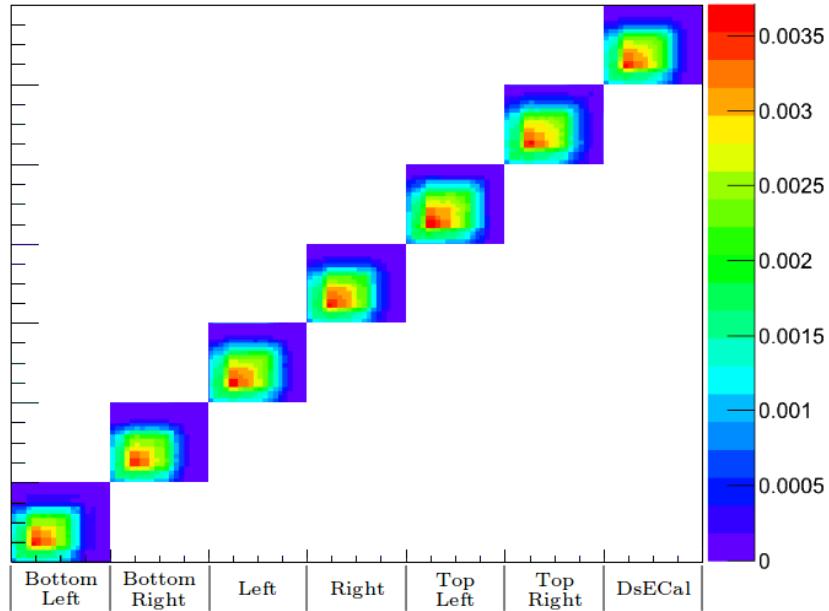
This matrix was generated for use by the official T2K oscillation analyses, which use a sample of neutrino interactions in the FGD detectors as a constraint on the neutrino flux and cross-sections. As a consequence of this, the flux systematic matrix had only been evaluated over the face of the FGD detectors rather than the whole ND280. The measurement made in this analysis, although based on events in the DsECal, also has input from neutrino interactions in the barrel ECals, so it was necessary to validate the use of this matrix for those samples. It was not possible to re-make the matrix in its entirety, so validation focussed on the pion production and beam direction uncertainties. It should be noted that, since the DsECal sits behind the FGDs, the errors evaluated for the DsECal are the same as those evaluated for the FGD and so can be used for comparison.

The effect of the pion production uncertainties was evaluated separately for each barrel ECal module, producing an individual fractional covariance matrix for each module,

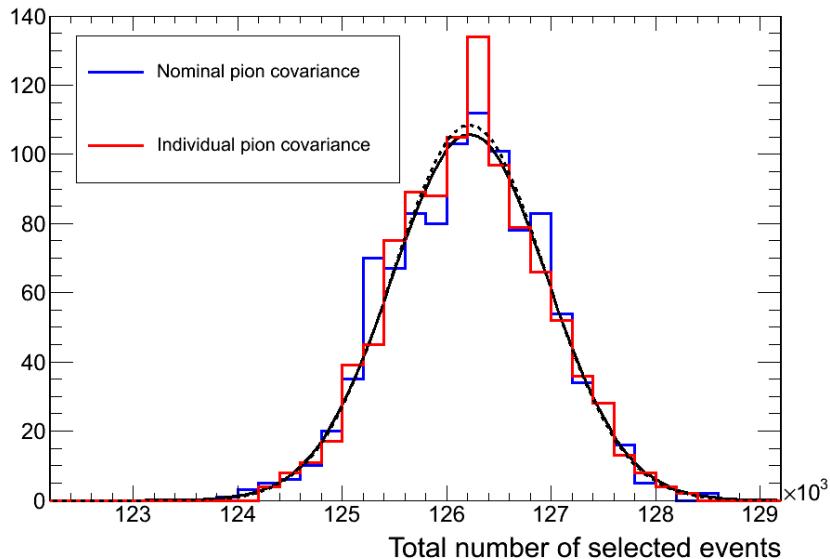
shown in Fig. 5.7a. One thousand random throws were taken from this matrix for a sample of selected ECal neutrino events, and the variation in the total number of events selected is shown in Fig. 5.7b. This was then compared to the same distribution, but with the random throws taken only from the fractional covariance matrix generated across the front face of the FGD, the nominal matrix, also shown in Fig. 5.7b. The two histograms in Fig. 5.7b have both been fitted with Gaussian distributions, which turn out to be almost identical. This demonstrates that using the measured pion production uncertainties from the tracker region of the ND280 gives the same uncertainty on events in the barrel regions as is given by evaluating the pion production uncertainties separately for each of the barrel ECal modules.

The systematic effect of the neutrino beam direction uncertainty was also evaluated separately for each of the ECal modules and compared to the uncertainty from the FGD plane. This was done by shifting the beam centre by ± 12 cm, the variation measured by the INGRID detector. The beam was shifted simultaneously in both the vertical and horizontal directions, moving the ND280 detector either further on- or off-axis by ± 17 cm. A ratio of the shifted neutrino spectrum to the nominal was calculated for each ECal module and compared to that seen for the DsECal. The ratio for the bottom left module is shown in Fig. 5.8b as an example of the maximum deviation seen, whilst the original spectra are shown in Fig. 5.8a. The same plots for the DsECal are shown in Fig. 5.9.

When creating the plots in Figures 5.8 and 5.9, neutrinos were generated across the front face of the ND280 detector by the JNUBEAM software. The neutrinos that passed through the XY face of a given ECal module were used to evaluate the error caused by the beam direction uncertainties, so the DsECal sample was much larger than the barrel ECal samples due to its larger front face. This does not cause any problems at the neutrino beam peak, but towards the tails of the distribution statistical effects disrupt the comparison, creating larger differences. By focussing on the events around the peak neutrino energy, it can be seen that the beam direction uncertainty changes the neutrino spectrum of the DsECal by about 2%, whereas for the bottom left ECal module the difference is around 5%. This difference comes primarily from a shift in the neutrino spectrum, rather than a change in the total number of neutrinos seen. It was found that shifting the beam position results in a 1.7% change in the number of neutrinos passing through the DsECal and a 2.5% change in those passing through the barrel ECals. Fig. 4.24 shows that the selection efficiency in the barrel ECals is changing slowly at the neutrino energy peak (around 650 MeV), so shifts in the neutrino energy spectrum due

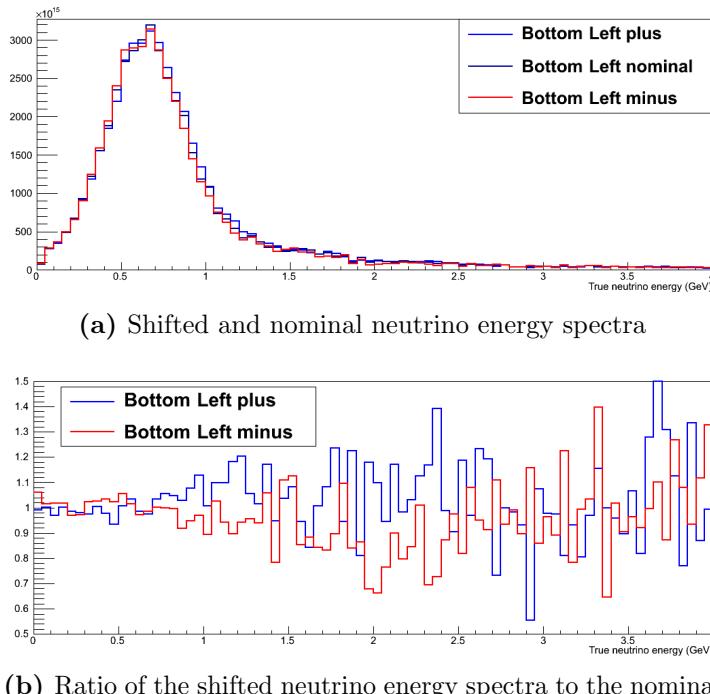


(a) Pion production uncertainty covariance matrix



(b) Total number of selected MC events after reweighting for the pion production errors

Figure 5.7.: The variation in the neutrino flux due to pion production uncertainties. Fig. 5.7a is a fractional covariance matrix with one set of 20 neutrino energy bins for each ECal module. Fig. 5.7b shows the variation in the number of events selected when random throws are taken from such a covariance matrix. Throws using the nominal pion production covariance matrix alone produce the blue distribution while throws from matrix in Fig. 5.7a gave the red distribution. The nominal and individual distributions have been fit with two Gaussian functions (the solid and dashed black lines respectively).



(b) Ratio of the shifted neutrino energy spectra to the nominal

Figure 5.8.: The nominal and shifted neutrino energy spectra in the bottom left barrel ECal module, where the shifted histograms show the effect of moving the neutrino beam centre 12 cm towards the ND280 in both the vertical and horizontal directions (blue) and 12 cm away from the ND280 (red).

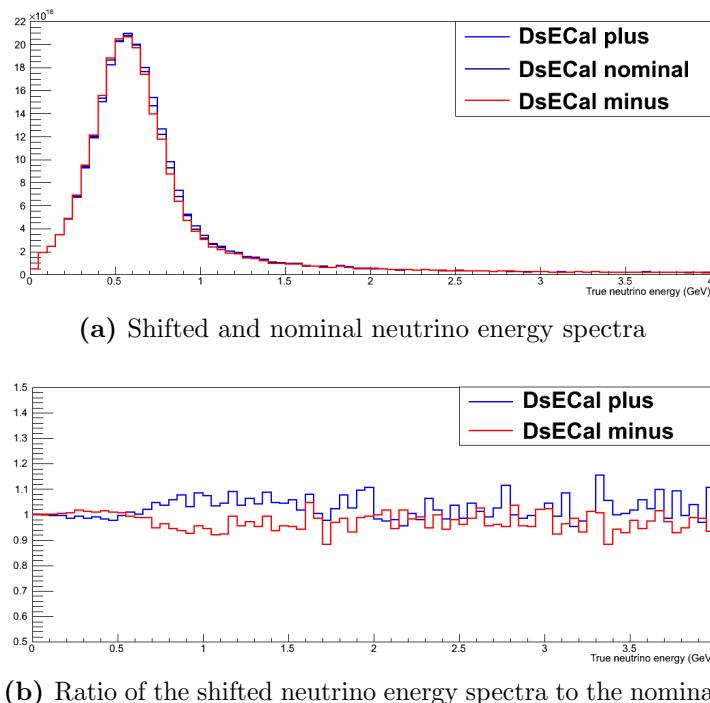


Figure 5.9.: The nominal and shifted neutrino energy spectra in the DsECal, where the shifted histograms show the effect of moving the neutrino beam centre 12 cm towards the ND280 in both the vertical and horizontal directions (blue) and 12 cm away from the ND280 (red).

to the beam position uncertainty should have a small effect on the final sample. This leaves the 0.8% difference between the DsECal and bottom left ECal as the remaining error, which is negligible when compared to the overall beam uncertainty.

This brief investigation into the flux uncertainties in the barrel regions of the ND280 has shown no significant difference in the calculated flux uncertainties between the barrel and the downstream ECals. It is therefore appropriate to use the existing flux covariance matrix, shown in Fig. 5.6, to describe the uncertainties on the number of events selected in the ND280 ECals.

To see the effect of the flux systematics on the events used in this analysis, a set of parameters b_i is created to vary the tuned input flux

$$\Phi_i^{\text{varied}} = b_i \Phi_i^{\text{tuned}}, \quad (5.4)$$

where i identifies the neutrino energy-flavour bin from the flux covariance matrix and the b_i parameters can be generated from the fractional covariance matrix via Cholesky decomposition

$$b_i = \sum_{j=0}^{80} L_{ij}^T u_j. \quad (5.5)$$

Here the u_j is a random draw from a Gaussian, with mean zero and a width of one, for the neutrino energy-flavour bin j of the flux covariance matrix. L_{ij}^T is the transpose of the lower triangular matrix from the Cholesky decomposition of the flux covariance matrix. The flux covariance matrix shows the fractional change in each neutrino energy-flavour bin, so the b_i parameters are used to weight the selected neutrino events to provide the varied MC.

The process above describes a single throw from this covariance matrix, generating a single set of new weights for the selected neutrino events. To assess the full impact of these uncertainties, 1000 random throws were performed and the resultant number of selected events recorded for each of the ECal and FGD samples, producing the distribution in Fig. 5.10 where the total number of selected events has been plotted for each throw.

The thrown event samples, N_a^{throw} , are then renormalised so that the total number of events in both the thrown ($N_{\text{thrown}}^{\text{total}}$) and nominal ($N_{\text{nominal}}^{\text{total}}$) samples is the same, giving

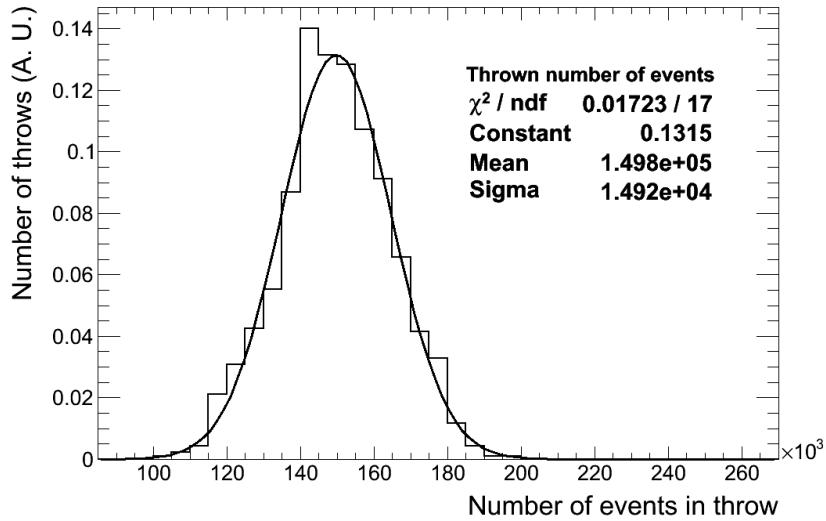


Figure 5.10.: The variation in the number of events selected from the MC sample coming from the uncertainties in the neutrino flux.

the renormalised thrown sample, $N_a'^{\text{throw}}$, as

$$N_a'^{\text{throw}} = N_a^{\text{throw}} \frac{N_{\text{nominal}}^{\text{total}}}{N_{\text{thrown}}^{\text{total}}}. \quad (5.6)$$

The differences between the thrown sample and the nominal MC, $\Delta N_a^{\text{shift}}$ is then found

$$\Delta N_a^{\text{shift}} = N_a'^{\text{throw}} - N_a^{\text{nominal}}, \quad (5.7)$$

where N_a^{nominal} is the nominal number of events in sub-detector sample a . The fractional covariance matrix element v_{ab} can then be written as

$$v_{ab} = \frac{\Delta N_a^{\text{shift}} \cdot \Delta N_b^{\text{shift}}}{N_a^{\text{nominal}} \cdot N_b^{\text{nominal}}}, \quad (5.8)$$

with a and b running over the 11 ECal detector samples and the FGD sample. This matrix is created for each throw to preserve the correlations between samples. The final fractional covariance matrix for the flux systematic uncertainties is the mean average of the matrices created by each throw, and is shown in Fig. 5.11, alongside the correlation matrix.

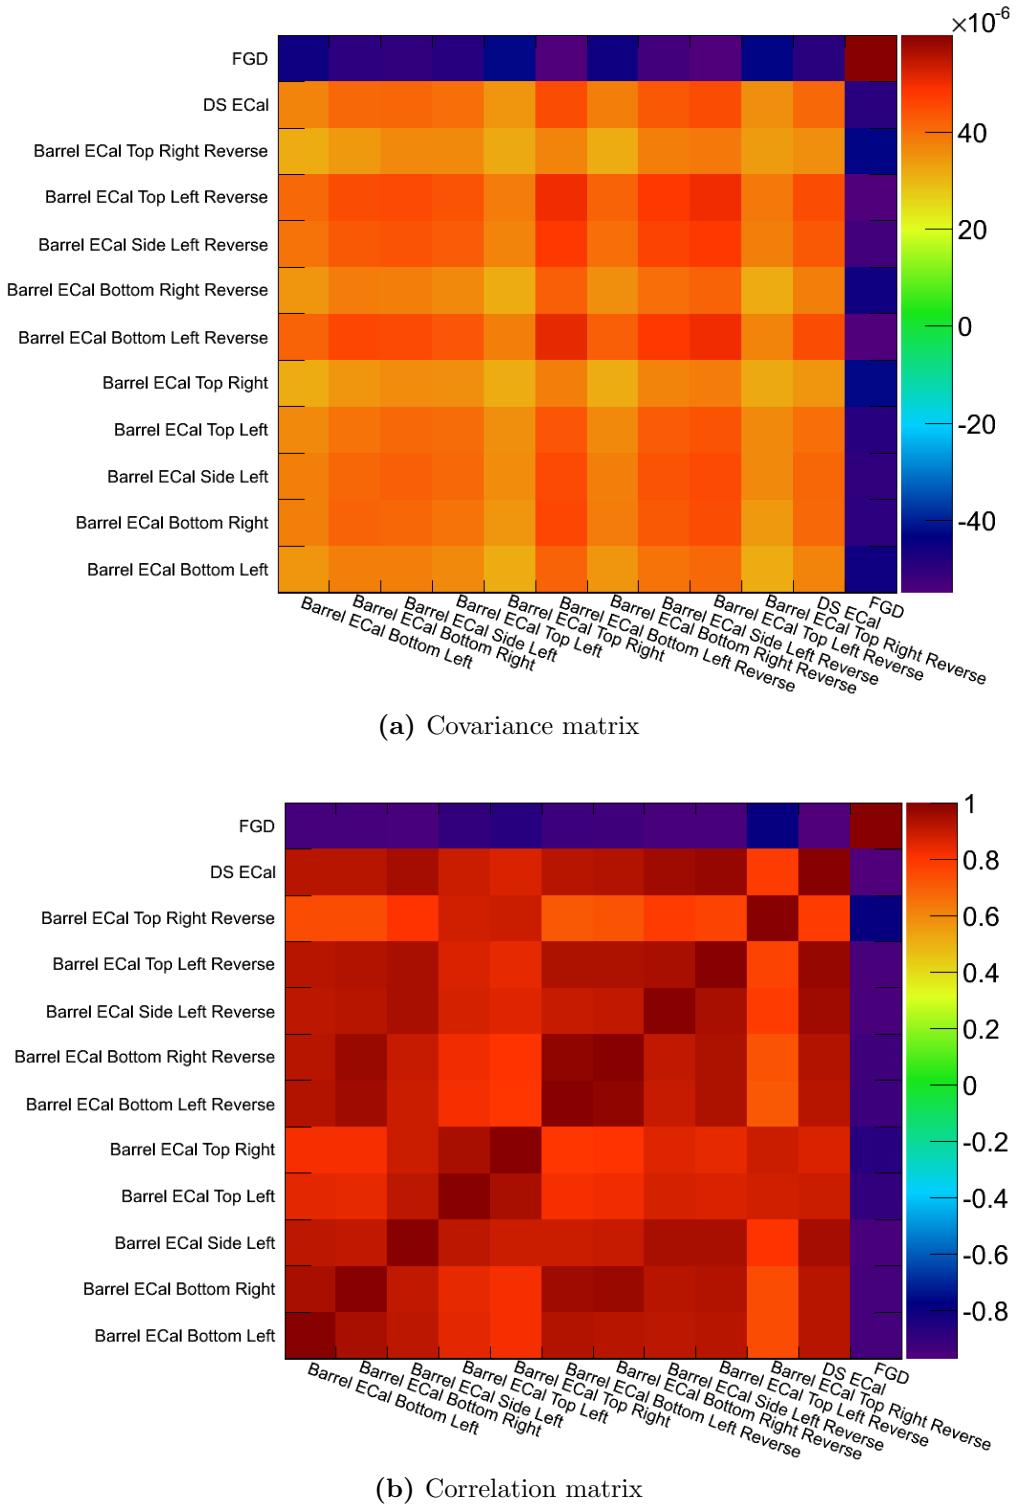


Figure 5.11.: Matrices showing (a) the uncertainty and (b) the correlations between the uncertainties on the number of events selected in each sub-detector sample due to uncertainties on the neutrino flux prediction.

Cross-section systematics

As discussed in Chap. 1, the current picture of neutrino cross-sections is incomplete, particularly at the neutrino energies seen by the ND280 detector. The neutrino interaction models used by current neutrino interaction generators contain a number of tunable parameters that have come from experimental results. The uncertainties on these measured parameters affect the neutrino interaction cross-sections in the T2K MC, therefore directly altering the number of events that are expected in the ND280. Alongside this, there are some competing models that describe what is actually happening in the neutrino interaction, giving rise to systematic uncertainties from the choice of model. At T2K there is a dedicated neutrino interactions working group (NIWG) tackling these issues.

The NIWG provides uncertainties on the different parameters used in the neutrino interaction generators by performing fits to both external and ND280 data samples. It also provides response functions that relate the changes in these parameters to changes in the events selected at both the near and far T2K detectors [60]. Uncertainties on the cross-section parameters are given as a covariance matrix, presented visually in Fig. 5.12, containing the absolute error on each parameter along with the correlations between the parameters, where these are known.

The effect of these uncertainties on the event selection from Chap. 4 was analysed in a similar fashion to the flux systematics. Cholesky decomposition of the cross-section matrix combined with a vector of draws from a random Gaussian distribution was used to generate a single random throw of the parameter errors. In contrast to the flux systematic, this is a throw of the errors rather than a throw of event weights. The thrown errors were added to the nominal parameter values to provide a set of thrown cross-section parameters. The T2KReWeight [61] software takes these new parameters and produces a set of event weights from them, using the response functions provided by the NIWG. Using this machinery 1000 random throws from the cross-section covariance matrix were made and applied to the selected neutrino events. The variation this produces in the total number of selected events is shown in Fig. 5.13.

A fractional covariance matrix for each throw was produced in the same way as described for the flux systematic and is shown in Fig. 5.14 alongside the corresponding correlation matrix.

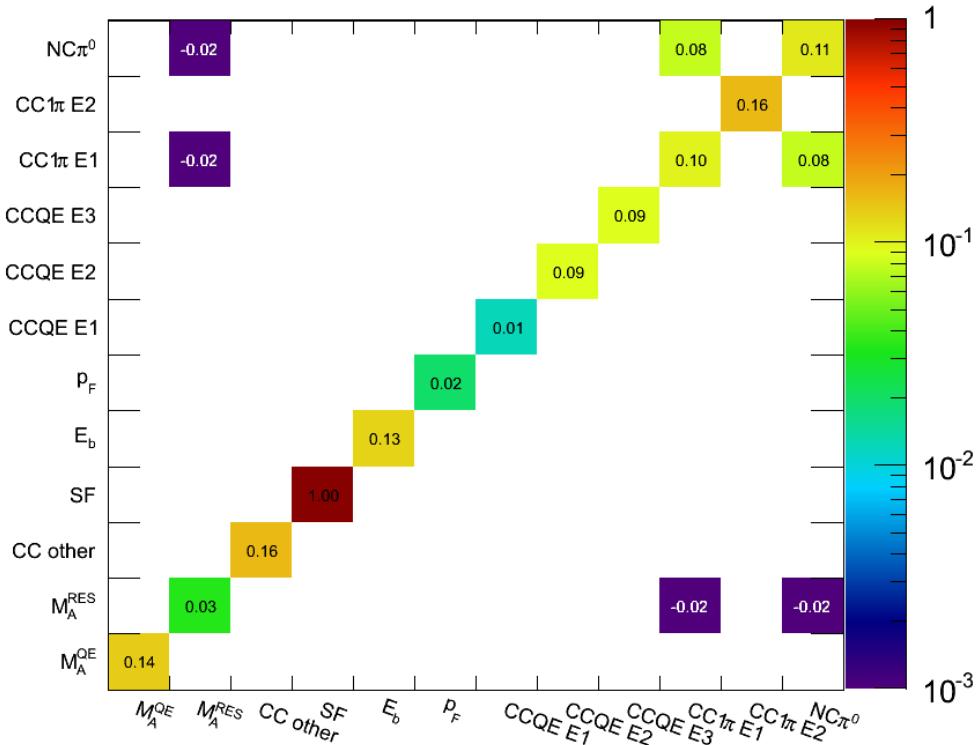


Figure 5.12.: Covariance matrix showing the fractional uncertainties on the value of the NEUT cross-section parameters.

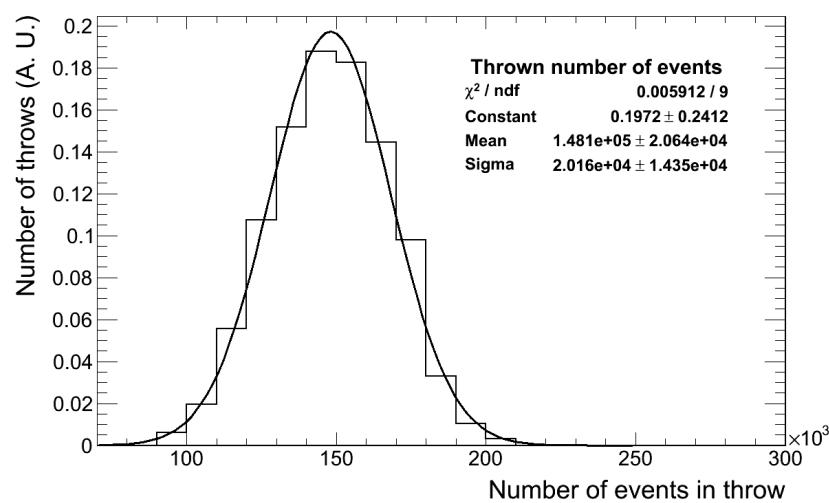


Figure 5.13.: The variation in the total number of events selected due to uncertainties on the value of the main NEUT cross-section parameters.

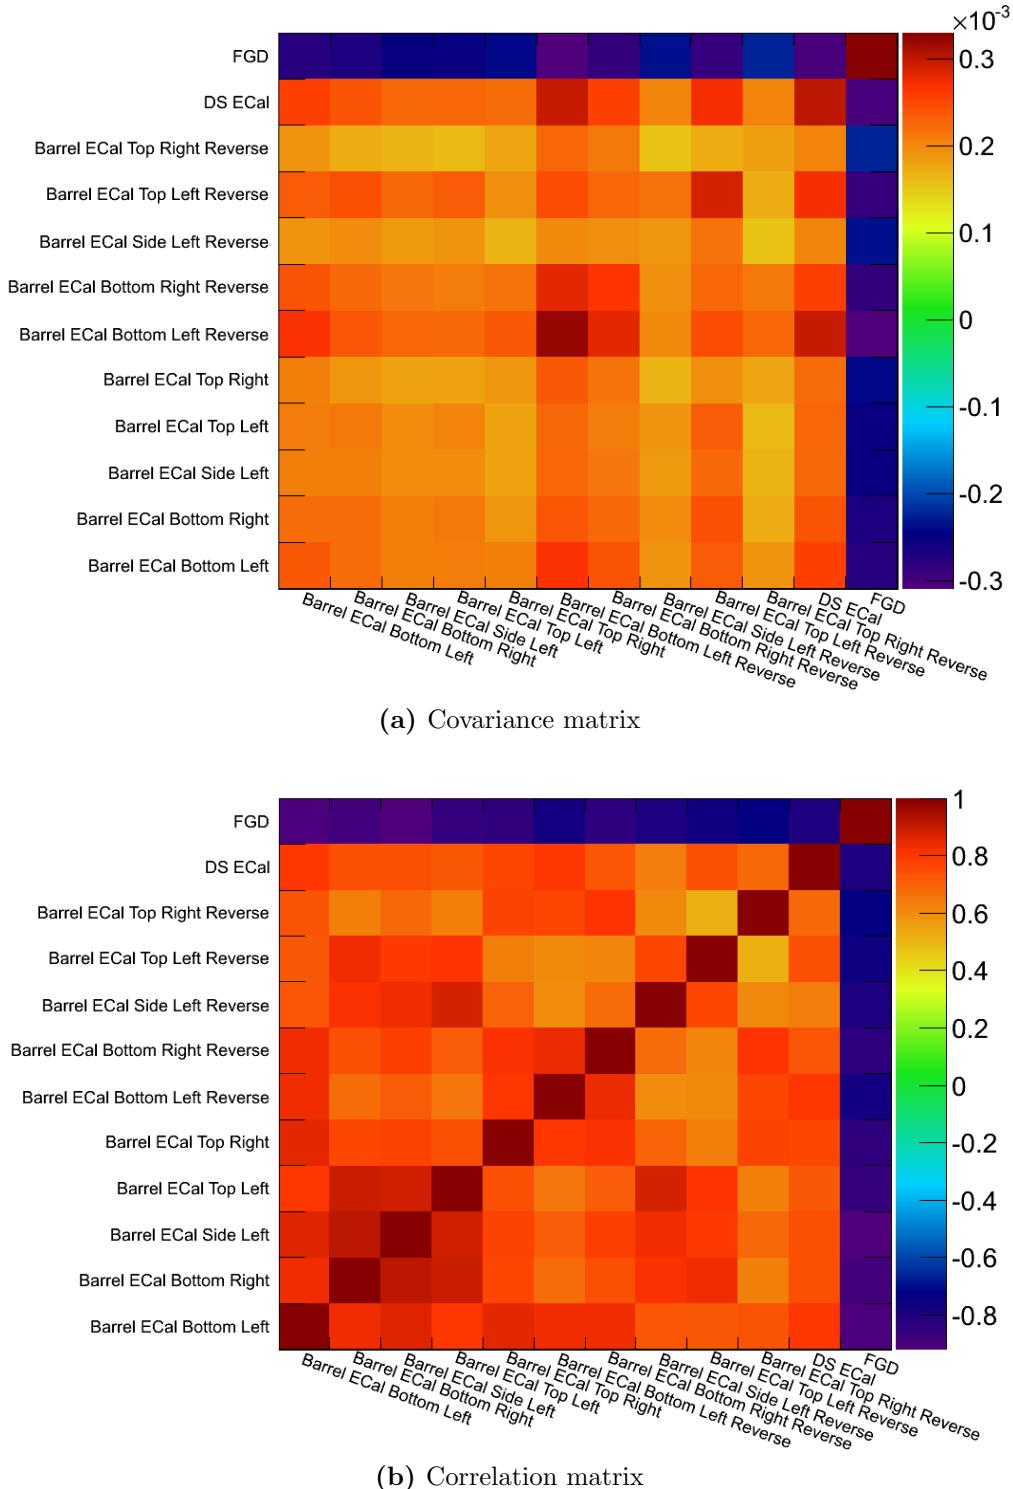


Figure 5.14.: Matrices showing (a) the uncertainty and (b) the correlations between the uncertainties on the number of events selected in each sub-detector sample due to uncertainties on the NEUT cross-section parameters.

As a cross check, each of the individual cross-section parameters was varied by itself and the resultant matrices summed. Random throws were generated as above, correctly taking into account the parameter correlation, but only the variation in the parameter of interest was applied using T2KReWeight. The final summed matrix was found to be identical to the one in Fig. 5.14a. As a second cross-check, the neutrino energy distribution was plotted for the nominal value of M_A^{QE} and $M_A^{\text{QE}} \pm 1\sigma$, shown in Fig. 5.15. M_A^{QE} primarily governs the rate of CCQE interactions, which is the main interaction mode below about 2 GeV. As mentioned in Chap. 1, M_A^{QE} appears in the dipole term of the axial-vector form factor, so from Eqn. 1.8 an increase in M_A^{QE} should increase the neutrino interaction cross-section and a decrease in M_A^{QE} should reduce it. Fig. 5.15 shows this expected behaviour, giving further confidence that the event re-weighting is being applied correctly.

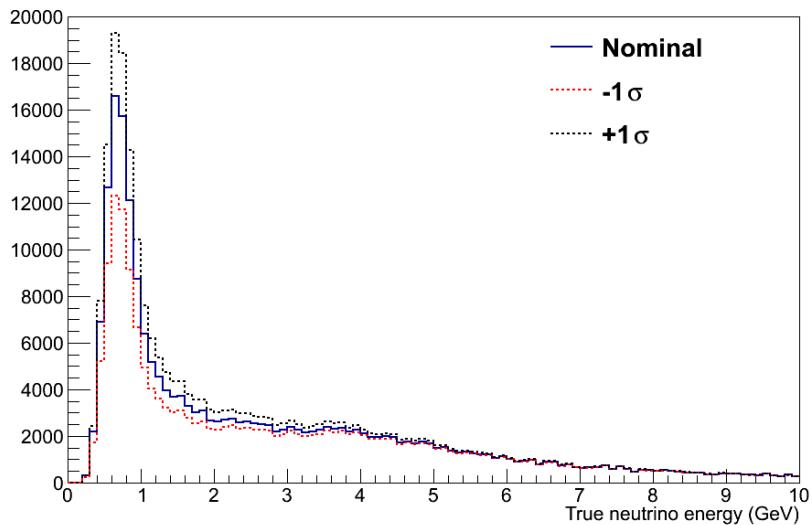


Figure 5.15.: The change in the neutrino energy spectrum due to 1σ shifts in the value of M_A^{QE} .

The NIWG also provide uncertainties on the final state interactions (FSI) of the neutrino induced pions, kaons, etas, protons and neutrons. These FSI effects will produce mixing between the different CC interaction final states, causing a CCQE event to produce a π^+ for example. This analysis uses an inclusive CC selection to give event samples, so FSI uncertainties will have a negligible effect compared to the underlying cross-section parameter and flux systematics and are not included in the analysis.

Detector systematics

The systematics presented here are only relevant to the selected ECal events, not those in the FGD sample, which will be discussed later. For all the detector systematics, a MC sample with the systematic variation applied was generated and the event selection described in Chap. 4 applied. For each ECal module, the difference between the number of events selected in the systematically varied MC and the nominal MC is taken as the effect of the systematic, which is then divided by the number of events selected in the nominal MC to give the fractional effect. Again, the varied number of events in each sample is renormalised so that the total number of events across the varied samples is equal to the total number of events in the nominal MC. The measured effect on each ECal module is then combined with that from the other modules to create a fractional covariance matrix.

In this analysis the event selection only uses simple, low level variables whose values are determined solely by the position and charge of the hits in the ECal. As described in Chap. 3, the unknown coordinate of a hit is calculated using the charge and known coordinates of the surrounding hits. The known hit coordinates come from the physical position of the bar containing that hit, so are unaffected by anything except the probability that a hit is created in the first place. In the ECal MPPC calibration a cut is applied requiring every hit to have at least 3.5 photon equivalent units of charge before it is considered for reconstruction. This means that only variations in the hit charge before reconstruction is performed, or uncertainties in the ECal construction that affect the amount of energy deposited by a particle in a scintillator bar, will affect the probability that a hit is created. An understanding of the majority of detector systematics therefore only requires an understanding of the variation in hit charge.

The argument above neatly splits the detector systematics into two categories, the first being the variation of the detector charge response and the second being discrepancies between the real detector and the MC simulation of the detector that could give differences in the amount of charge deposited by ionising particles.

The effect of the detector charge response on the number of selected events is estimated by applying a smearing to the calibrated charge of each ECal hit in the beam MC. This smearing is applied by taking a random draw from a Gaussian with a mean of zero and a width equal to the charge resolution of the ECal, then adding this number to the calibrated hit charge.

Cosmic rays were used to calculate the charge resolution of the ECals, since at ground level these are composed almost exclusively of MIPs. Plotting the charge of each hit in a sample of cosmic ray tracks produces a characteristic curve, as shown in Fig. 5.16, which can be fit with a Landau-Gaussian distribution. This provides the most probable value (MPV) for the charge in a single hit created by a MIP.

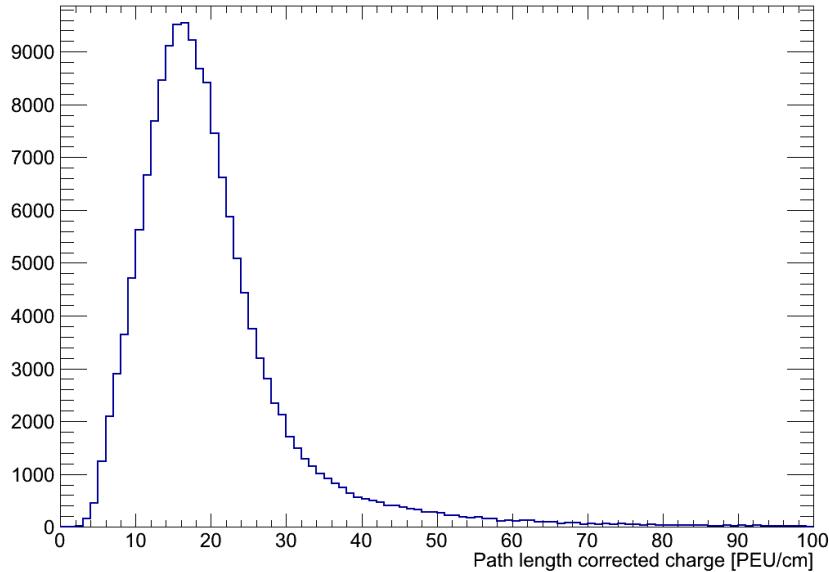


Figure 5.16.: The charge deposited in each hit from a sample of cosmic rays passing through the DsECal, where the charge has been corrected for the path length of the cosmic ray within each scintillator bar.

This curve is identical for all MIPs, regardless of their energy, and is used as a standard candle in charge calibration. By looking at the change in the MPV of this curve over time a measurement of the variation of the ECal charge response can be made. Cosmic muon data was collected in between neutrino beam spills throughout the T2K physics runs, so will give the correct charge resolution for neutrino beam data. Fig. 5.17 shows the MPV measured from cosmic data samples taken from T2K Run III, fitted with a Gaussian. The most probable amount of charge deposited by a cosmic muon does not change significantly over this time period, and the width of the Gaussian fit has been taken as the ECal charge variation. The fractional covariance matrix generated for this systematic is shown in Fig. 5.18.

The second category of systematics looks to quantify the differences between the data and MC, which come from a number of places. Following on from the above discussion of the charge resolution, the MC simulation of the MPPC response to incident photons

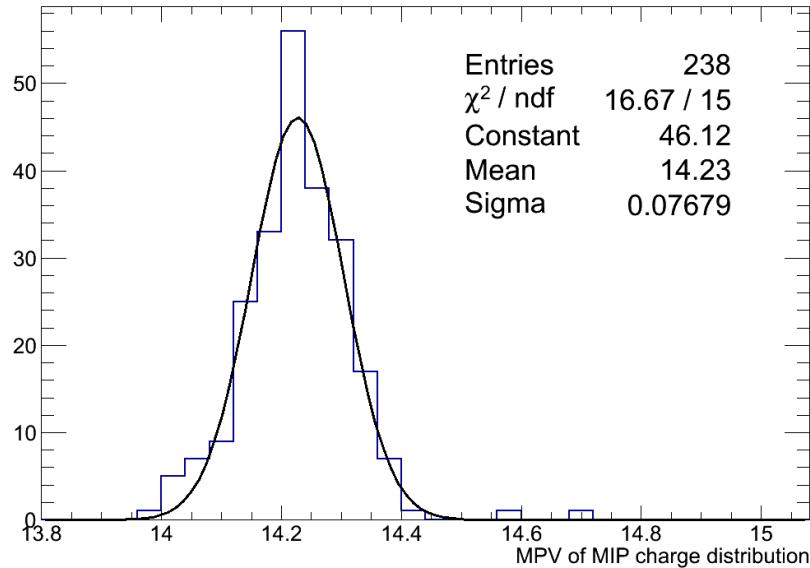


Figure 5.17.: The spread of MPVs from cosmic data collected over T2K Run III, fitted with a Gaussian curve to show the 1σ variation.

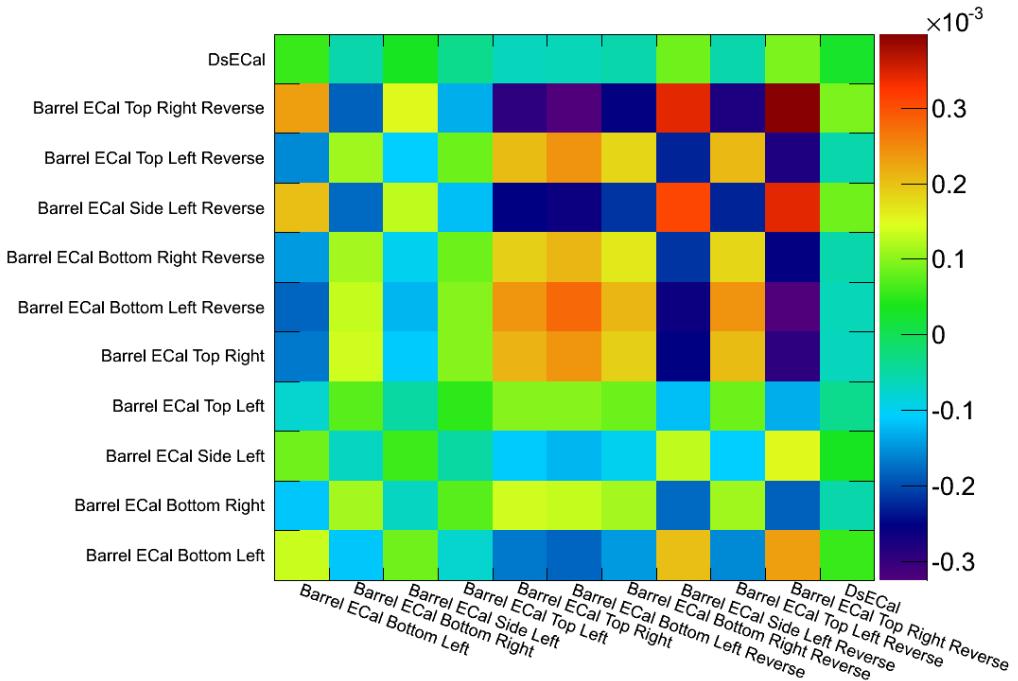


Figure 5.18.: Covariance matrix created by considering the effects of variation in the ECal charge response on the sub-detector event samples.

does not match the data presented in Ref. [62]. The data in Ref. [62] was taken from a testbench setup, using LEDs to illuminate a WLS fibre at a set distance from an MPPC connected to it. The MPPC response was then characterised in a controlled environment similar to that of the ECals in the ND280 pit. Fig. 5.19 shows the saturation curve recorded in Ref. [62] alongside the saturation curve that is present in the ND280 electronics simulation. The plot relates the number of photons hitting the MPPC active surface to the number of photo-electrons recorded by the sensor, showing the saturation of the sensor when large numbers of photons strike it. The discrepancy between the data and MC as the MPPCs saturate can be clearly seen.

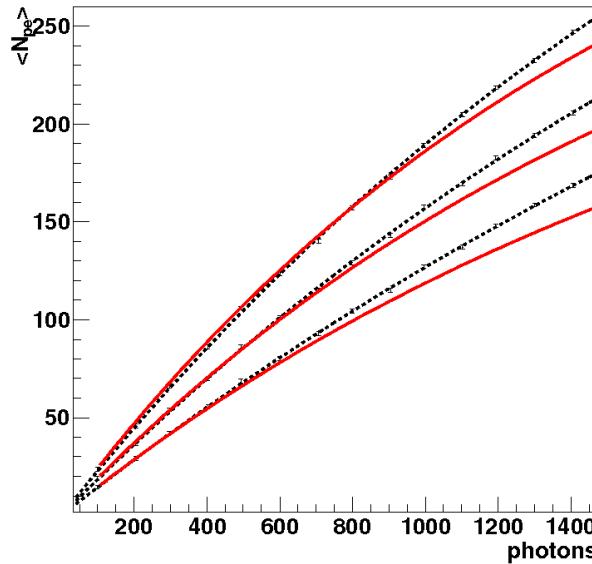


Figure 5.19.: Average number of pixel-equivalents (pe)—the number of pixels that fire—as a function of the number of photons hitting the active surface of the MPPC, corrected for the MPPC photo-detection efficiency. The solid curves show the response encoded in the MC whilst the dashed line is a fit to the data points, which were taken from Ref. [62]. The three separate curves represent MPPCs at three different overvoltages.

Fig. 5.20 shows the same plots as Fig. 5.19, but in this case the MC response has been corrected to better match the data. The correction linearly reduces the simulated MPPC output according to the number of incident photons, reducing it by 0% up to 500 photons and then by up to 10% at 1400 photons. Applying this correction brings the simulated MPPC response in line with the measured response and it will become part of the future ND280 software. At this time the correction is not present by default, so its affect will be covered by a systematic.

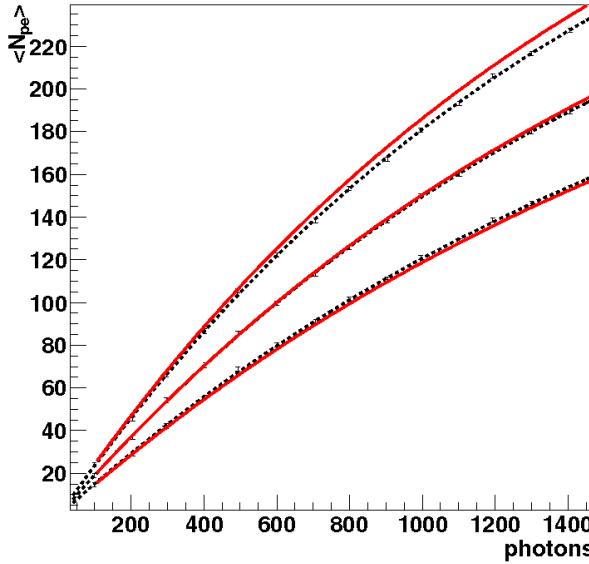


Figure 5.20.: Average number of pe as a function of the number of photons hitting the active surface of the MPPC, corrected for the MPPC photo-detection efficiency. The solid curve is the corrected MC response whilst the dashed line is a fit to the data points, which were taken from Ref. [62]. The three separate curves represent MPPCs at three different overvoltages.

Systematically varied MC was generated by applying this correction at the elecSim stage of processing. After applying the event selection to both the varied and nominal MC, the fractional covariance matrix in Fig. 5.21 was generated.

Potentially, the largest sources of systematic error could come from differences between the ECal component dimensions in the MC and those components in real life. The ECals are simulated in the MC according to the design specification used during their construction, so any difference will come from uncertainties on the dimensions of the ECal components. To affect this analysis, those differences have to change the volume of the active or dead materials within the ECal volume. There are only three components where any difference could produce a noticeable effect: the thickness of the TiO_2 coating on the scintillator bars, the volume of the bars themselves and the volume of the lead sheets. The other materials used in the ECal construction were considered, but the uncertainties on the masses involved are all much less than 1% of the total detector mass, so will have a negligible effect. Unless otherwise stated, each component's volume uncertainty was calculated using Ref. [31].

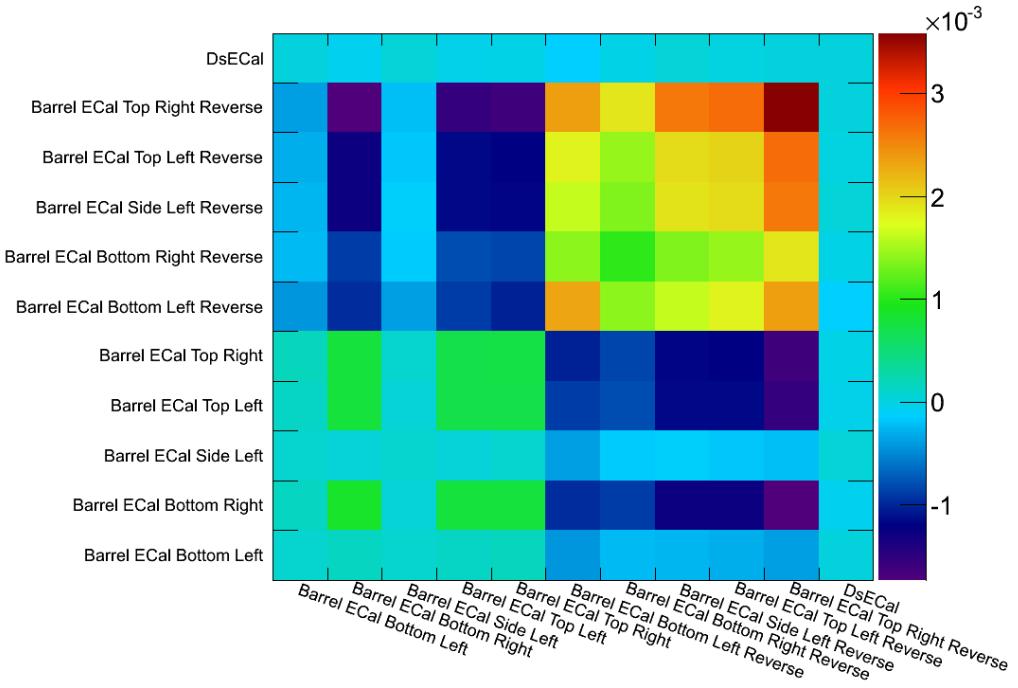


Figure 5.21.: Covariance matrix created by comparing the sub-detector event samples selected from nominal MC to those selected from MC with the MPPC saturation correction applied.

All of these systematics were evaluated by changing the ECal simulation in the nd280mc software package to produce varied MC. Each change was applied one by one and the change in the ECal’s simulated mass recorded and compared to calculations by hand, to confirm both that the change had been applied and that it had been applied correctly.

The ECal scintillator bars have a 0.25 mm thick TiO_2 coating in order to contain scintillation light within the bar. From the bar extrusion process, the uncertainty in this thickness is 0.13 mm. This coating does not scintillate itself, so increasing its thickness by over 50% will increase the amount of dead material a charged particle passes through as it penetrates an ECal module. The fractional covariance matrix created from this varied MC is shown in Fig. 5.22.

The ECal bars themselves have a 40 mm by 10 mm cross-section, but could be up to 0.4 mm smaller in both width and height. The bars were cut to length during construction, giving a 0.1 mm uncertainty on the bar length. Since the bar cross-section can only be smaller than the nominal, the bar size was reduced in the varied MC to give the smallest possible volume. This systematic produces the fractional covariance matrix shown in Fig. 5.23.

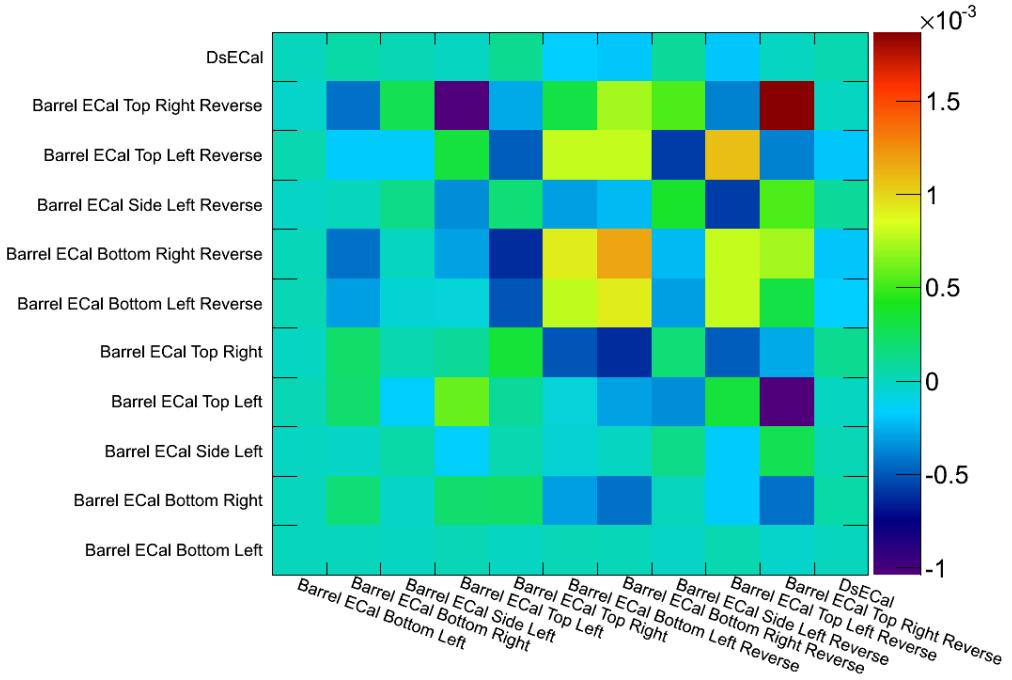


Figure 5.22.: Covariance matrix created by comparing the sub-detector event samples selected from nominal MC to those selected from MC where the thickness of the TiO_2 coating on the ECal scintillator bars had been increased by 50%.

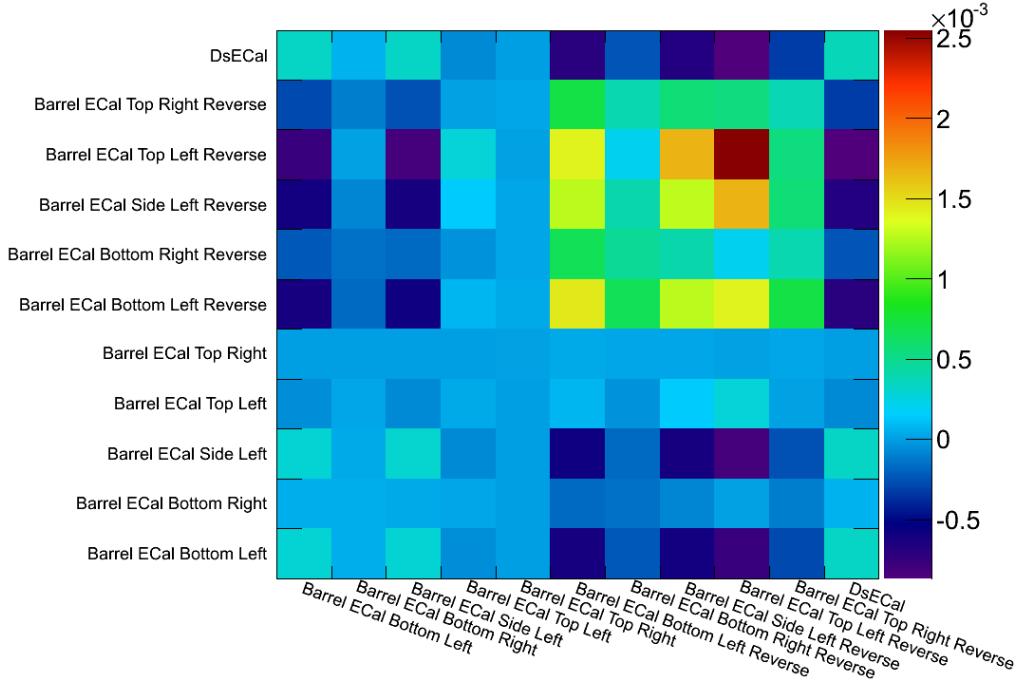


Figure 5.23.: Covariance matrix created by comparing the sub-detector event samples selected from nominal MC to those selected from MC where the size of the ECal scintillator bars had been reduced by up to 0.4 mm in each dimension.

The final ECal component to be varied is the lead sheeting. The total length and width of the sheeting that covers each layer is known to within 3 mm, whilst the thickness of the sheets has a 0.1 mm uncertainty from the production process. During construction, checks were made to ensure the lead layers completely covered the scintillator bars in each layer, so the 3 mm uncertainty enters as a 1.5 mm uncertainty on the overhang of the lead from the scintillator at each of the four sides of the layer. Applying this change produces the fractional covariance matrix in Fig. 5.24.

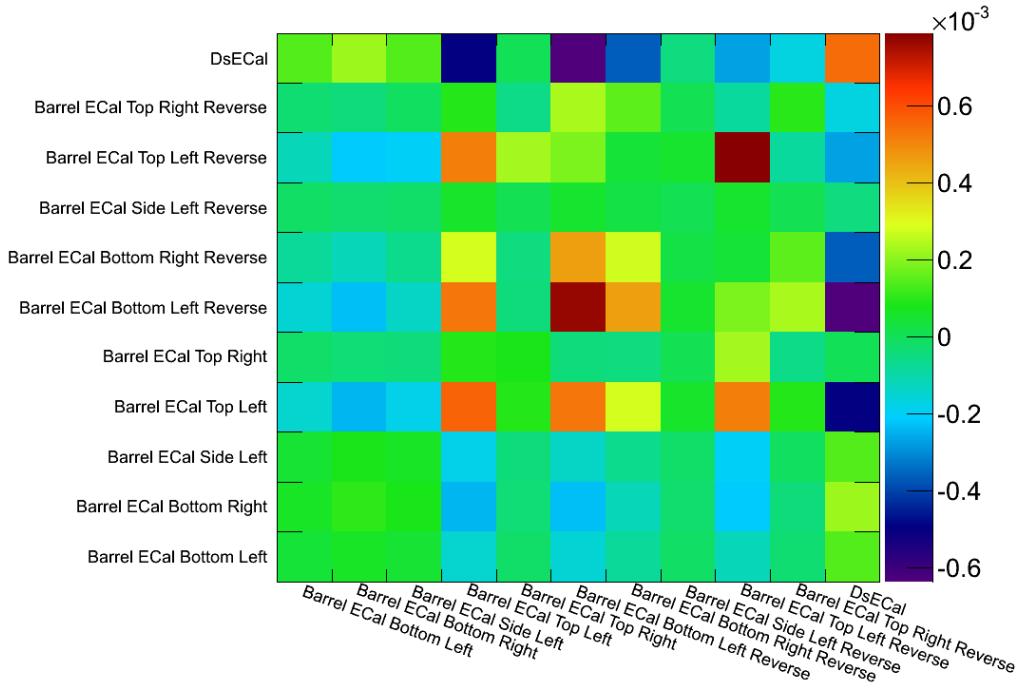


Figure 5.24.: Covariance matrix created by comparing the sub-detector event samples selected from nominal MC to those selected from MC where the volume of the lead sheets in the ECals had been increased according to their production uncertainties.

The final detector systematic comes from the difference between the number of noise hits present per event in the neutrino beam data and that simulated in the MC. In each MC event, the combined number of noise hits in all ECal modules was counted immediately before any clustering took place. To estimate the number of noise hits in a cosmic data event, the total number of hits present immediately before clustering began was counted and then the number of hits that went on to form fully reconstructed 3D tracks was subtracted from this total. The result of this estimation when applied to cosmic MC is shown in Fig. 5.25a, alongside the number of true noise hits in the same events, showing that the estimate is accurate. This method was used to estimate the number of noise hits per event in cosmic data, shown in Fig. 5.25b for a single inter-spill

cosmic dataset taken during T2K Run III. The MC simulation was then tuned to produce the same number of noise hits as seen in the data, the results of which are also shown in Fig. 5.25b. The fractional difference in the number of selected events in the tuned MC compared to the nominal is used to quantify the effect of this systematic, creating the covariance matrix shown in Fig. 5.26.

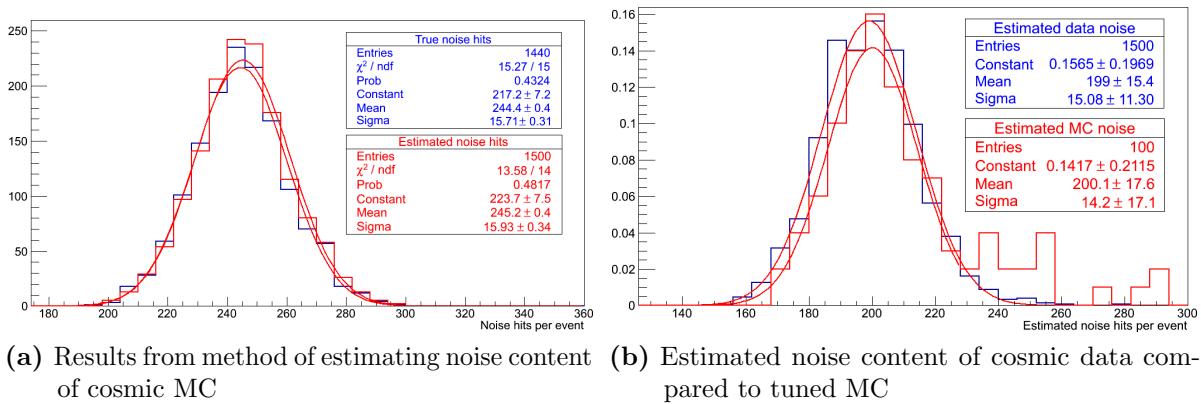


Figure 5.25.: Fig. 5.25a shows the number of noise hits per event from a sample of cosmic MC, as calculated using MC truth information (blue) and the estimate from counting the number of fully reconstructed hits (red), both fit with Gaussian curves. Fig. 5.25b shows the number of estimated noise hits per event in cosmic data (blue) and in the tuned MC (red), showing that the tuned MC noise rate agrees with that observed in the data.

The ECal detector systematics are summarised in Table 5.1 and the combined fractional covariance matrix they produce together is shown in Fig. 5.27.

Since the MC generation itself is intrinsically random, there are differences in the number of selected events per module even between two independent sets of nominal MC. Fig. 5.28 shows a covariance matrix generated from these two sets of nominal MC, including the normalisation differences between them. The top left signal-enriched selection shows a 3.6% difference from one dataset to the other. This variation will be present in each varied MC sample, and will hide the effect of smaller systematics. To mitigate the effects of this, two covariance matrices were generated for each detector error source, using two independent samples of systematically varied MC and nominal MC. An average of the two fractional covariance matrices was then taken as the final effect of the systematic in question. Ideally, a response function would be generated for each reconstructed ECal variable by generating multiple samples of varied MC for each systematic error source. The processing time needed to create these response functions made this approach prohibitive for the analysis described here, but doing this would

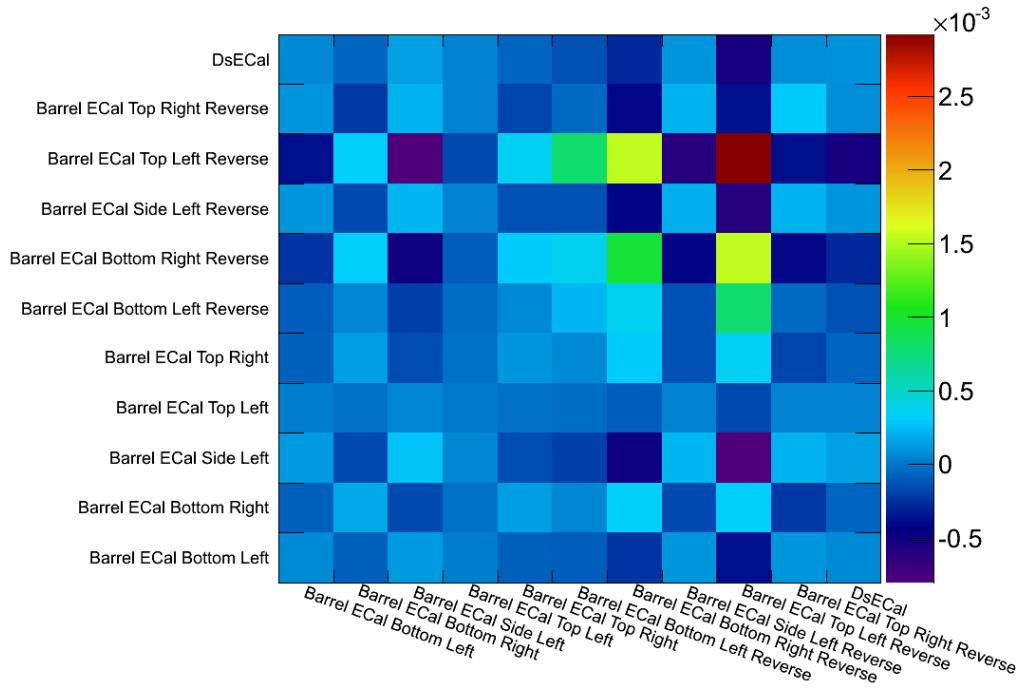


Figure 5.26.: Covariance matrix created by comparing the sub-detector event samples selected from nominal MC to those selected from MC in which the noise rate in the ECals has been tuned to data.

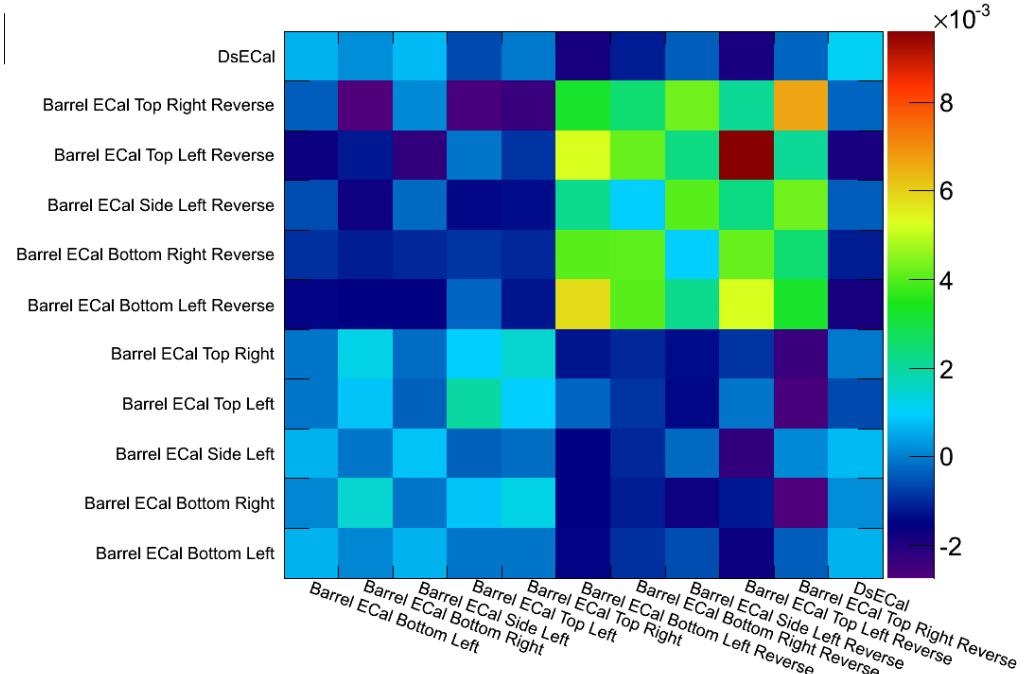


Figure 5.27.: The combined detector systematic uncertainty covariance matrix.

remove the effect of the intrinsic MC variation and allow the detector systematics to be evaluated in the same way as the flux and cross-section systematics, which would be valuable in the future.

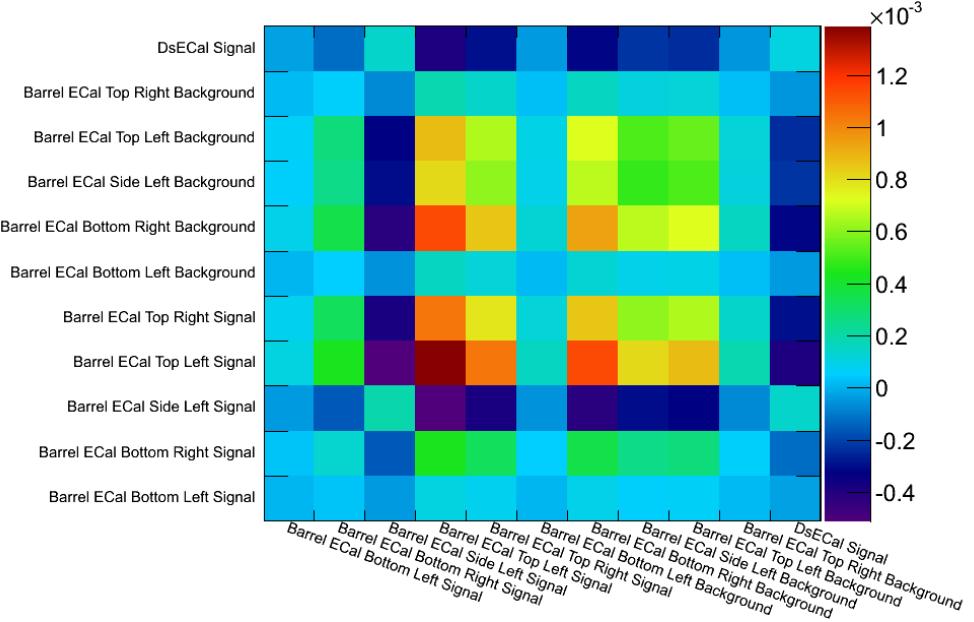


Figure 5.28.: A covariance matrix generated from two independent samples of nominal MC, showing the intrinsic variation within the samples.

The intrinsic MC variation described above will still be present in each individual detector systematic covariance matrix, and will therefore introduce some double-counting of errors into the final uncertainty. The covariance presented in Fig. 5.27 is therefore a conservative estimate of the true detector error but, as shown in Table 5.1, the detector systematic uncertainties are still the smallest contributor to the overall error.

FGD detector systematics

The FGD detector systematics for the CC-inclusive selection are described extensively in [63]. The systematics have been provided in the form of a fractional covariance matrix, shown in Fig. 5.29, that is split into 20 bins of reconstructed lepton momentum and angle and then separated into events passing the CCQE-like selection and those that do not, giving 40 bins in total.

This analysis is using the total number of FGD charged current events and ignoring the $p\theta$ information. To calculate the FGD systematic uncertainty, the diagonal from the fractional covariance matrix is used to provide the error on each $p\theta$ bin and then the

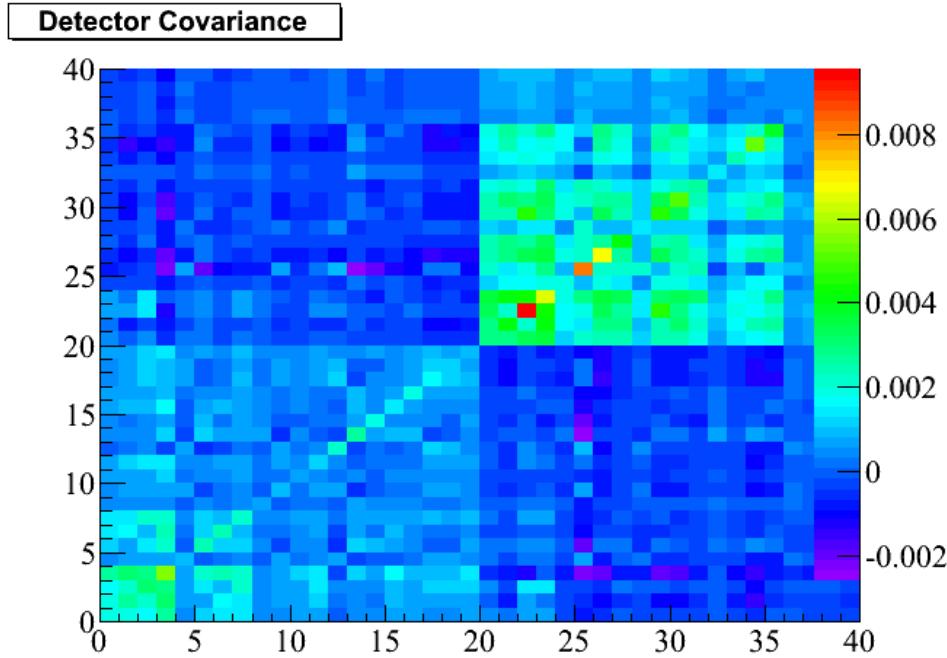


Figure 5.29.: Fractional covariance matrix showing the effect of detector uncertainties on the FGD event selection.

mean average error for the selected MC events calculated. This average error comes to 3.22%, which is independent from the ECal detector systematics and as such is included in the detector systematics covariance matrix with no correlations.

Summary of systematic errors

In this analysis the major sources of systematic error have been evaluated and presented above, with covariance matrices showing their effect on the shape of the sub-detector event samples. These systematic source were also evaluated without renormalising the varied MC in order to see their effect on the sample normalisation. The combined normalisation and shape fractional errors from each systematic source are summarised in Table 5.1.

To produce the total systematic error covariance matrix for this analysis, the detector, flux, and cross-section shape-only matrices are summed, resulting in the matrix in Fig. 5.30a, which will be used in the χ^2 fit discussed at the beginning of this chapter. Alongside this, Fig. 5.30b shows the covariance matrix generated including the normalisation uncertainty in each systematic error source.

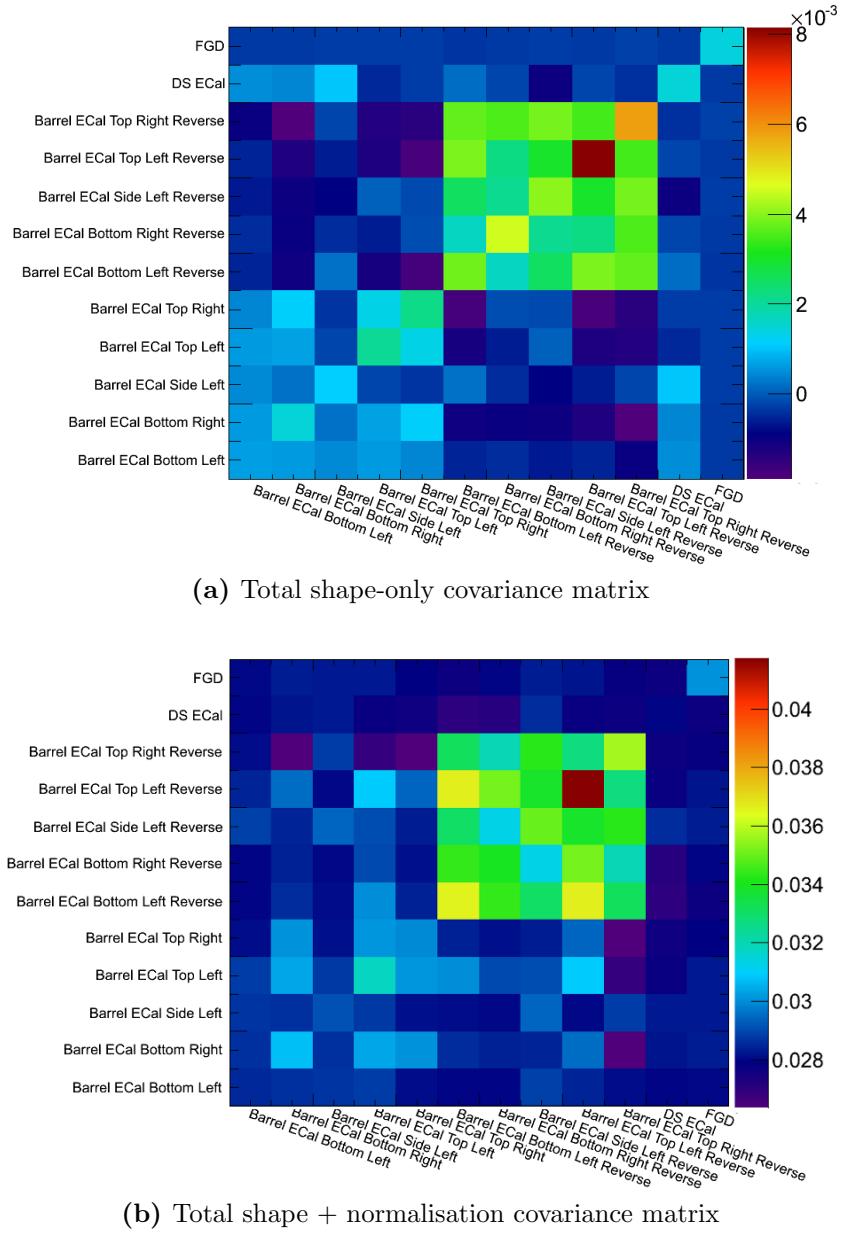


Figure 5.30.: The total covariance matrices calculated looking at (a) the shape-only effects and (b) both the shape and normalisation effects of the different systematic error sources.

Systematic source	Largest error (%)	Average error (%)
Charge response	2.7	1.6
MPPC saturation	5.2	3.5
TiO ₂ coating	3.6	2.2
Scintillator bar dimensions	7.2	4.2
Lead sheet dimensions	2.1	1.2
Noise simulation	3.6	1.7
Detector systematics total	10.1	6.4
Flux uncertainties	10.3	10.3
Cross-section uncertainties	13.5	13.2
Total	19.7	17.9

Table 5.1.: The fractional systematic uncertainty on the selected ECal samples, broken down into the contribution from each source of systematic error.

This is the first T2K analysis to examine neutrino interactions within the ND280 ECals and the first to include the sand muon MC when looking at interactions outside the ND280 tracker. As a consequence, the systematic errors presented are not definitive but are instead robust, conservative estimates of the true errors. In future analyses there are a number of improvements that could be made that would result in improved error estimates.

For the detector systematics, the majority have been calculated using a small number of shifted MC samples or data from a small period of running. Work is now underway to evaluate these systematics over a larger dataset and produce response functions in the same manner as the cross-section uncertainties. This would both improve the estimate of the error and allow future analysers to easily re-weight their selected events and evaluate the effect of the error. The largest detector errors come from the construction tolerances of the ECal modules. A detailed investigation into the ECal mass, in the same manner as that done for the FGDs [63], would be able to greatly constrain these uncertainties. Finally, both the cross-section and flux systematic uncertainties are being constantly updated by the working groups concerned, with upcoming data from the ND280 detector being used to further reduce these systematic errors.

Parameter set	R^{Pb}	R^{C}	R^{External}	R^{Other}
1	1.0	1.0	1.0	1.0
2	2.0	1.0	1.0	1.0
3	0.9	1.0	1.2	1.1

Table 5.2.: The three population normalisation parameter sets used to test the χ^2 fit.

5.2.3. Validating the fit performance

The χ^2 fit described at the start of this chapter has been validated by applying it to toy MC datasets, using the fractional covariance matrix derived above and including the effect of statistical errors. The sub-detector MC templates were created using the selected MC events in each sub-detector sample, broken down according to the position and species of the target nucleus. These templates were generated from the same MC samples described in Chap. 4 and scaled to represent the 2.588×10^{20} POT of data that were analysed.

To test the fitter, three sets of population normalisation parameters, R^{Input} , were used to generate fake datasets, summarised in Table 5.2. One set of parameters represented the nominal MC, with each normalisation set to unity, the second introduces a large change to the lead cross-section normalisation and the final set has a small shift applied to each of the normalisation parameters. These inputs scale the appropriate population in each sub-detector sample MC template

$$n_i^{\text{X Input}} = R^{\text{X Input}} n_i^{\text{X MC}}, \quad (5.9)$$

where i runs over each sub-detector sample, X represents the populations within that sample, $R^{\text{X Input}}$ is the input normalisation of population X , $n_i^{\text{X MC}}$ is the number of events in population X of sample i in the nominal MC and $n_i^{\text{X Input}}$ is the scaled number of events in population X of sample i . The total input number of events for each sample, N_i^{Input} , is given by summing the input populations within that sample

$$N_i^{\text{Input}} = \sum^{\text{X}} n_i^{\text{X Input}}. \quad (5.10)$$

To generate the toy datasets, a random throw was taken from the total fractional covariance matrix, shown in Fig. 5.30a, using the Cholesky decomposition method described earlier in this chapter

$$b_i = \sum_{j=0}^{\text{N Samples}} L_{ij}^T u_j. \quad (5.11)$$

Here, L_{ij} is the ij th component of the lower diagonal matrix from the decomposition of the fractional covariance matrix, u_j is a draw from a Gaussian of width 1 and mean 0 while i and j run over each sample in the analysis. The vector produced, \mathbf{b} , holds the fractional variation that must be applied to each sample to create a single random throw. This throw is then passed to the χ^2 fitter, which returns the best fit normalisation parameters it finds.

Ten thousand throws for a given set of input parameters were performed to construct the fitted parameter and parameter pull distributions. An example of these distributions is shown in Fig. 5.31a and Fig. 5.31b for R^{Pb} from parameter set 1. The fit parameter distribution is fitted with a Gaussian curve, the mean of which gives the final fitted parameter value. The pull, δ_i^X , on parameter R_i^X for each throw i is given by

$$\delta_i^X = \frac{R_i^X \text{ Fit} - R_i^X \text{ Input}}{\sigma_i^X \text{ Fit}}, \quad (5.12)$$

where $R_i^X \text{ Fit}$ is the fitted value of parameter X and $\sigma_i^X \text{ Fit}$ is the error on the fitted parameter value. This error is calculated by taking either the positive or negative MINOS error produced by the fit, according to whether the fitted normalisation is below or above the input value. Fig. 5.31a shows that the fitted value agrees with the input parameter whilst Fig. 5.31b demonstrates that the MINOS errors truly reflect the uncertainty on the fit and that the fit is unbiased.

Fig. 5.32a shows the fitted parameter value and error for each of the four parameters of input parameter set 1, from which it can be seen that the fitter is correctly returning the input value for each parameter. Fig. 5.32b shows the pull distribution mean and σ values for each parameter, showing that in each case the fit is unbiased and is reporting the correct uncertainty. Figures 5.33 and 5.34 show these same plots for the second and third set of input parameters, and in each case the fitter returns the correct input parameter values with the expected errors and only a very slight bias.

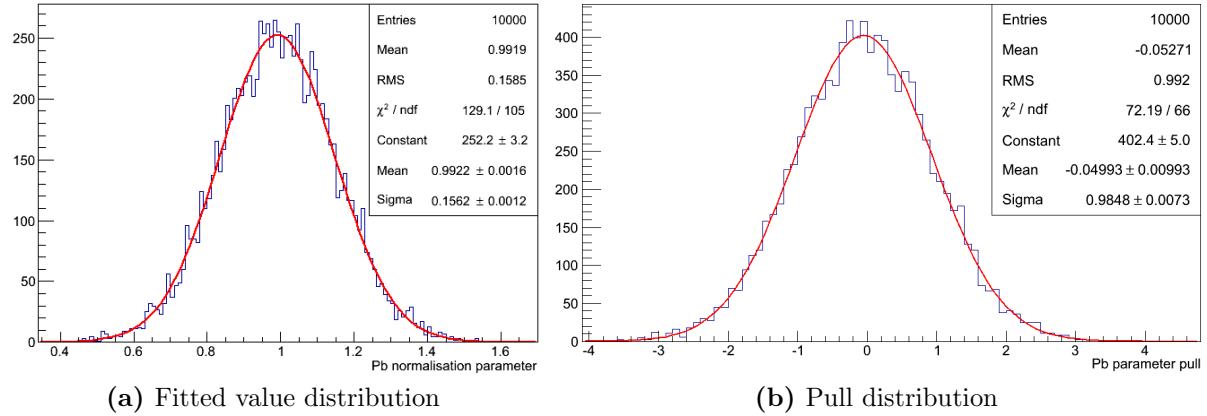


Figure 5.31.: The distribution of (a) the fitted parameter value and (b) the parameter pull of R^{Pb} , collected from 10,000 fake data throws using parameter set 1

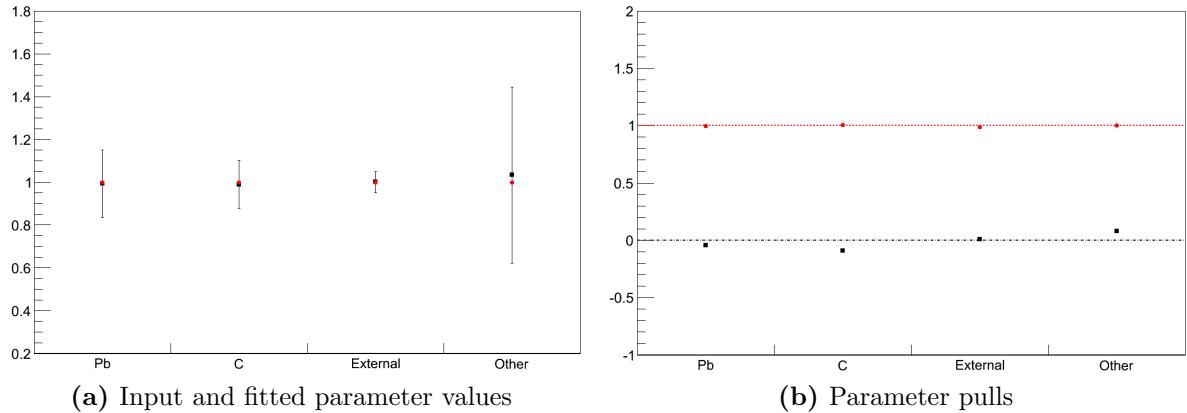


Figure 5.32.: Fig. 5.32a shows mean the fitted value for the four normalisation parameters (black points) compared to the input parameter values (red points). Fig. 5.32b shows the mean values (black points) and widths (red points) of the four parameter pull distributions after fitting each with a Gaussian. The fake data used in these fits comes from applying parameter set 1.

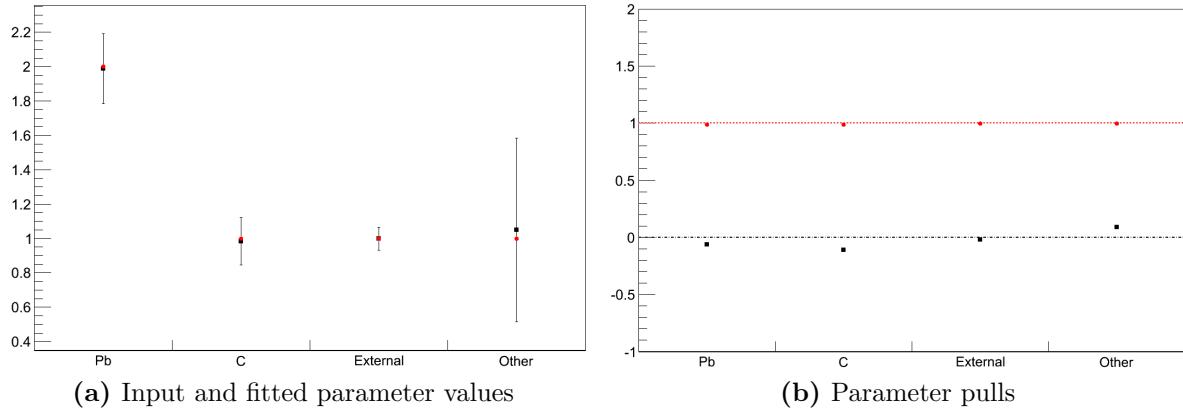


Figure 5.33.: Fig. 5.33a shows mean the fitted value for the four normalisation parameters (black points) compared to the input parameter values (red points). Fig. 5.33b shows the mean values (black points) and widths (red points) of the four parameter pull distributions after fitting each with a Gaussian. The fake data used in these fits comes from applying parameter set 2.

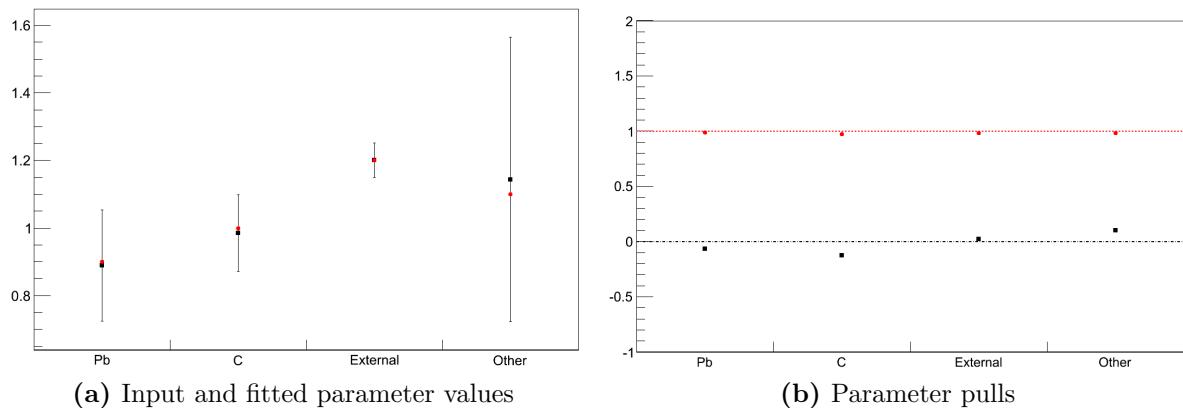


Figure 5.34.: Fig. 5.34a shows mean the fitted value for the four normalisation parameters (black points) compared to the input parameter values (red points). Fig. 5.34b shows the mean values (black points) and widths (red points) of the four parameter pull distributions after fitting each with a Gaussian. The fake data used in these fits comes from applying parameter set 3.

The fit uncertainty on R^{Pb} is approximately 0.16 for parameter sets 1 and 3, which represent realistic values of the lead normalisation, rising to 0.2 for parameter set 2. To quantify the effect of the barrel ECal signal- and background-enriched samples the fit was performed on throws from the nominal input parameters with either the signal- or background-enriched samples removed. It is necessary to include one set of barrel ECal samples in the fit, since without them there would be four free parameters and only the DsECal and FGD samples acting as constraints, resulting in too many degrees of freedom in the fit. The parameter fit results when the signal-enriched barrel ECal samples have been included are shown in Fig. 5.35a, and the results achieved when fitting the background-enriched sample are in Fig. 5.35b.

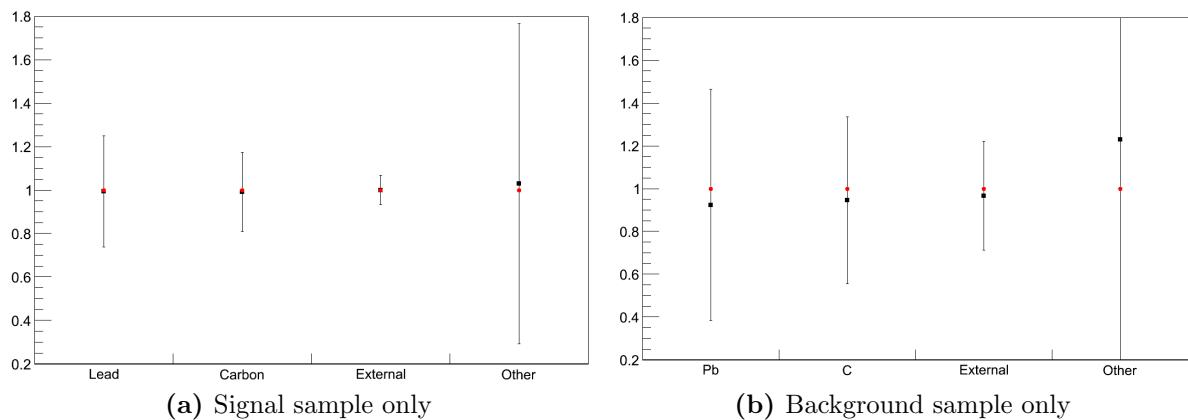


Figure 5.35.: The fitted values (black points) and input parameter values (red points) for the four normalisation parameters from parameter set 1. Fig. 5.35a shows the results achieved when only the signal-enriched sub-detector samples are fit whilst Fig. 5.35b shows the results when fitting the background-enriched samples alone.

Comparing these fit results shows that the background-enriched samples put weaker constraints on all the parameters than the signal-enriched samples. Comparing these to Fig. 5.32 demonstrates that fitting both sample together gives the smallest uncertainties on the fit parameters, even when taking into account the larger systematic uncertainties on the background-enriched samples.

As a final validation of the fit, the minimum χ^2 value for each throw of the nominal parameter set is plotted in Fig. 5.36 and fitted with a χ^2 distribution. The fit gives 7.96 degrees of freedom (DoF) for the distribution. This agrees well with the 8 DoF expected, since there are 12 samples and four free parameters, giving more confidence that the χ^2 function used in this fit is minimising correctly.

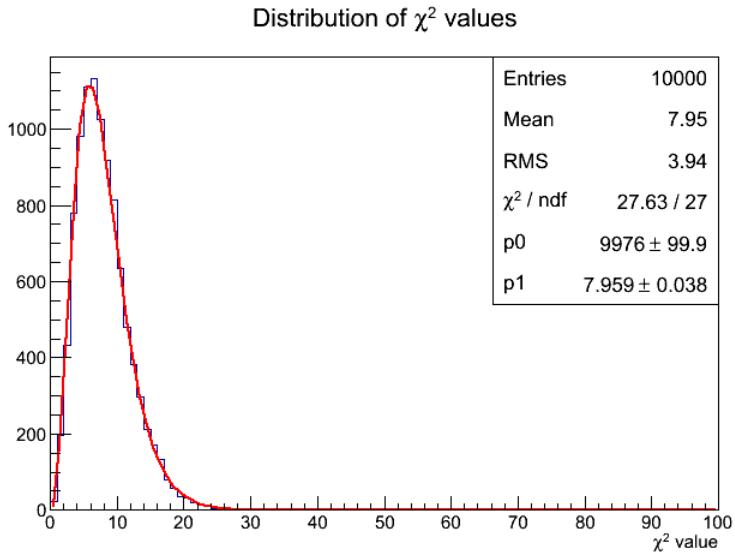


Figure 5.36.: The minimum χ^2 value from 10,000 fits to fake data generated using parameter set 1. Parameter 1 (p1) corresponds to the fitted number of degrees of freedom of the χ^2 distribution.

5.2.4. Fit applied to real data

The final event selection in Chap. 4 was applied to 2.665×10^{20} POT of neutrino beam data collected from T2K Run II and Run III, giving the samples of ECal events shown in Table 5.3. The FGD event sample was taken from T2K Run I and Run II data, corresponding to 1.077×10^{20} POT, and is also shown in Table 5.3. For comparison, the number of events selected in the MC sample is shown, after scaling to match the data POT.

The fit described in this chapter was then performed on the selected data, and the fitted parameter values it returned are listed in Table 5.4. The fit finds a minimum χ^2 value of 49.11, which for a χ^2 distribution with 8 DoF corresponds to a *p*-value of 6×10^{-8} .

The nominal MC is shown plotted against the measured data in Fig. 5.37 and a comparison of the two after the χ^2 fit is in Fig. 5.38.

Once the minimum had been found, all the normalisation parameters were fixed at their minimum value except R^{Pb} , which was varied to map out the χ^2 minimum shown in Fig. 5.39a. This was also done for the R^{C} and R^{External} parameters, shown in Fig. 5.39b

Sample	No. Selected - Data	No. Selected - MC
Top right ECal	83809	91870
Top left ECal	91493	97261
Left ECal	196814	200949
Bottom right ECal	172940	160323
Bottom left ECal	205481	187389
Top right ECal - Background	37908	34751
Top left ECal - Background	42541	37967
Left ECal - Background	79383	72630
Bottom right ECal - Background	80508	61678
Bottom left ECal - Background	89233	68441
DsECal	79158	86411
FGD	4485	4707

Table 5.3.: The total number of events selected in the different sub-detector samples for both data and MC.

Normalisation parameter	Fitted value
Lead	$-0.11^{+0.13}_{-0.13}$
Carbon	$0.50^{+0.07}_{-0.08}$
External	$1.74^{+0.04}_{-0.05}$
Other	$2.78^{+0.28}_{-0.27}$

Table 5.4.: The fitted values of the four normalisation parameters with their errors from the fitting.

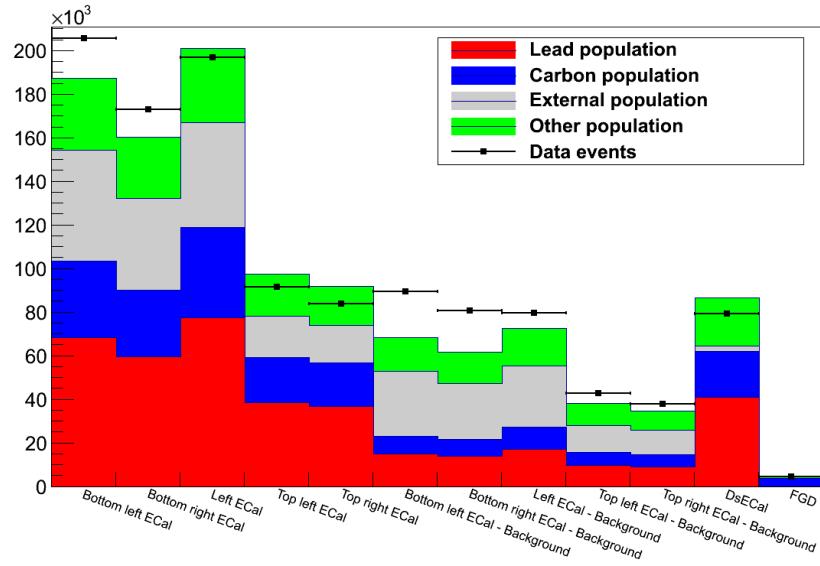


Figure 5.37.: The number of events selected in data (points) compared to the number selected in the nominal MC (histogram) across each of the sub-detector samples. The MC has been scaled to match the data POT and broken down to show the contributions from the separate fit populations.

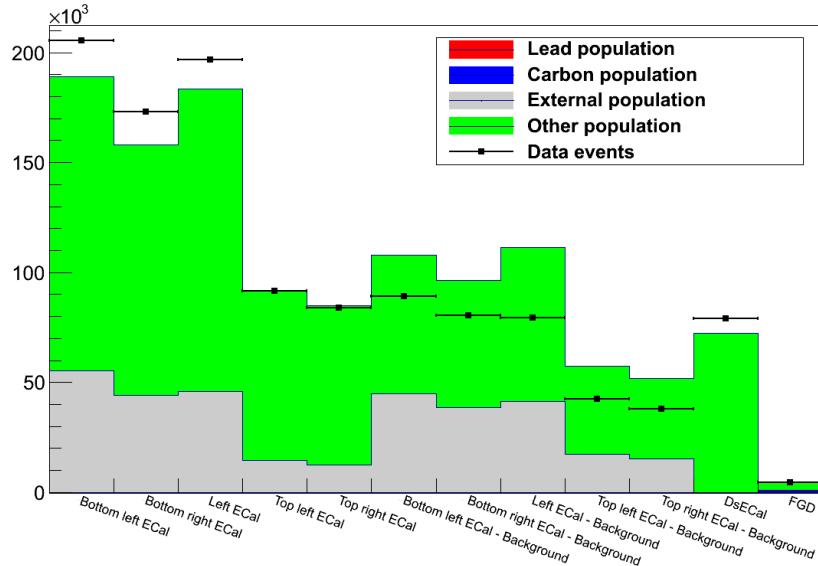


Figure 5.38.: The number of events selected in data (points) compared to the number selected in the fitted MC (histogram) across each of the sub-detector samples. The MC has been scaled to match the data POT and incorporate the fitted normalisation parameters.

and Fig. 5.39c respectively. All three graphs show a smoothly changing χ^2 value where the minimum has been correctly found by the fit.

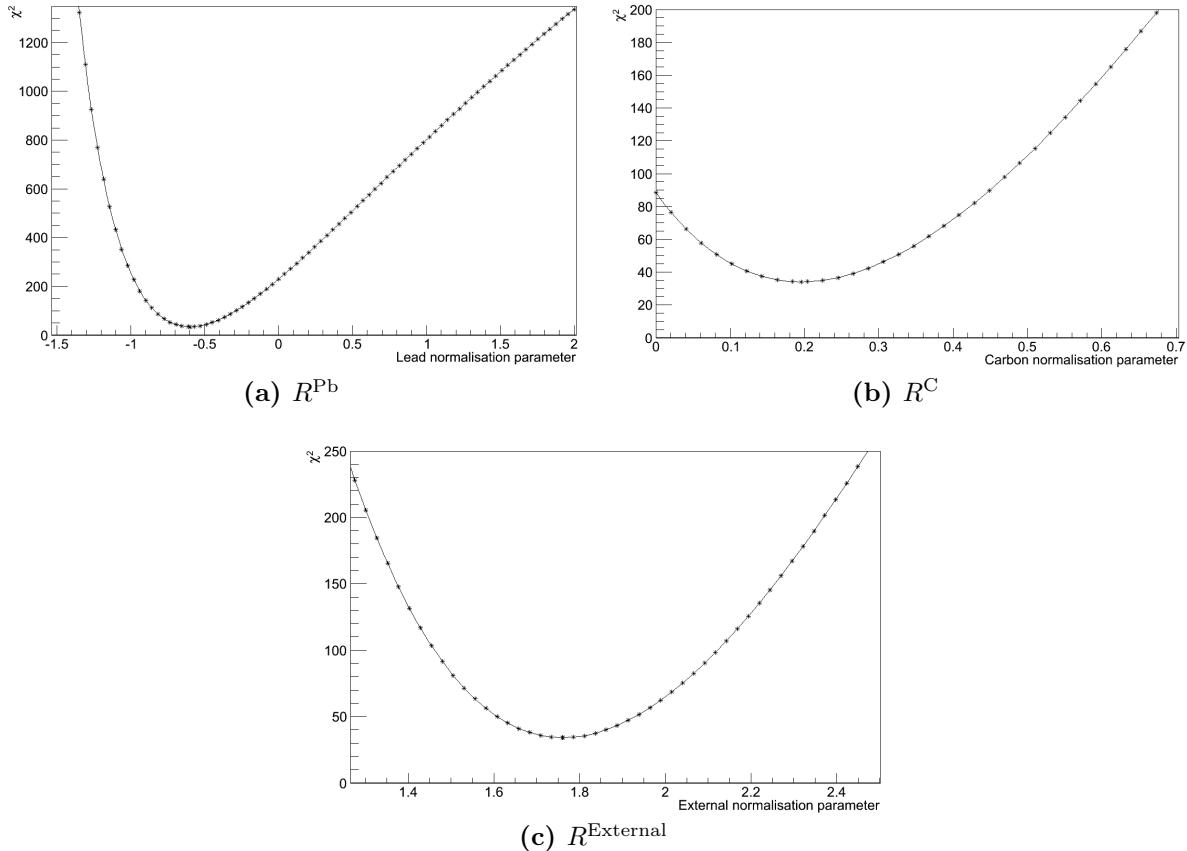


Figure 5.39.: The value of the χ^2 function as a function of (a) R^{Pb} , (b) R^C and (c) R^{External} whilst all the other normalisation parameters are held at their fitted values.

The results in Table 5.4 are clearly unphysical, since the fit chooses a negative number of neutrino interactions on lead nuclei and a greatly increased number of interactions from other sources. The final p -value of 6×10^{-8} for the fit, and the fit results in Fig. 5.38, show that the fitted MC is not a good match to the data. Both before and after the fit there is an excess of data in the bottom signal-enriched samples compared to the MC whilst before the fit there is an excess of MC compared to data in the top ECal module samples. The shape-only covariance matrix in Fig. 5.30a shows that in both the signal and background-enriched selections the uncertainty on the number of events in the top modules is almost the same as the uncertainty on the number of events in the bottom and side modules, and that they are correlated with one another. This gives the χ^2 fit less freedom to change the shape of the MC to fit the observed data.

Barrel module	R^{Pb}	R^{C}	R^{External}
Top right	$0.75^{+0.12}_{-0.12}$	$0.90^{+0.04}_{-0.04}$	$1.17^{+0.13}_{-0.12}$
Top left	$0.70^{+0.16}_{-0.17}$	$0.86^{+0.06}_{-0.07}$	$1.31^{+0.25}_{-0.22}$
Left side	$0.72^{+0.10}_{-0.10}$	$0.87^{+0.04}_{-0.04}$	$1.28^{+0.09}_{-0.10}$
Bottom right	$0.54^{+0.13}_{-0.12}$	$0.77^{+0.05}_{-0.05}$	$1.67^{+0.13}_{-0.13}$
Bottom left	$0.52^{+0.09}_{-0.08}$	$0.76^{+0.04}_{-0.04}$	$1.69^{+0.07}_{-0.07}$

Table 5.5.: The fitted values of the three normalisation parameters found when fitting each of the barrel ECal modules separately with the DsECal and FGD samples.

The tension between the data and MC comes from attempting to fit the bottom modules at the same time as the other barrel ECal samples. To remove this tension, the χ^2 was modified in order to fit four sub-detector samples at once: the DsECal and FGD samples along with the signal-enriched and background-enriched sample from a single barrel ECal module. With only four samples in the fit, it is not possible to use them to constrain four parameters. Since the fit requires that the overall normalisation of the MC can be varied, the ‘Other’ population was combined with the ‘External’ population, so N_i^{exp} from Eqn. 5.3 now becomes

$$N_i^{\text{exp}} = (R^{\text{Pb}} n_i^{\text{Pb}} + R^{\text{C}} n_i^{\text{C}} + R^{\text{Pb}} n_i^{\text{Pb,other}} + R^{\text{C}} n_i^{\text{C,other}} + R^{\text{External}} (n_i^{\text{External}} + n_i^{\text{Other}})). \quad (5.13)$$

Each barrel ECal module was included in the fit sequentially alongside the FGD and DsECal samples, with the fitting covariance matrix being reduced each time to include only those samples present in the fit. The final fitted parameters from this process are shown in Table 5.5.

As an example, Fig. 5.40 compares the fitted MC to the data when the fit is performed on the left barrel ECal module. This shows very good agreement between the data and MC in all of the samples. To confirm that the fit minimum is correct, Fig. 5.41 shows the χ^2 minimum across a range of values of the R^{Pb} and R^{C} parameters. This was created by performing a grid-scan across the R^{Pb} and R^{C} normalisation values and minimising the χ^2 at each point, again fitting only the left barrel ECal alongside the FGD and DsECal.

Fig. 5.41 also displays the χ^2 minimum and 1σ contour that were found by the fit, which agree with the measured χ^2 surface.

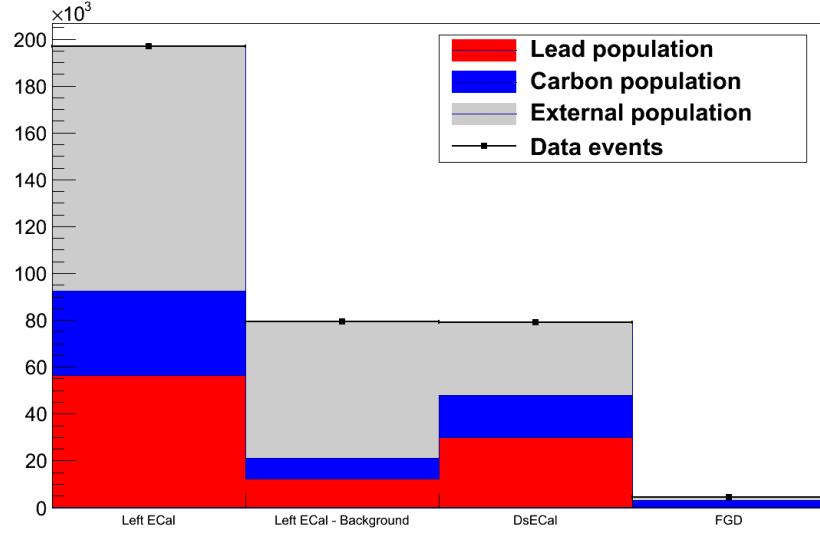


Figure 5.40.: The number of events selected in data (points) compared to the number selected in the fitted MC (histogram) across each of the sub-detector samples when only the left ECal module is included in the fit. The MC has been scaled to match the data POT and incorporate the fitted normalisation parameters.

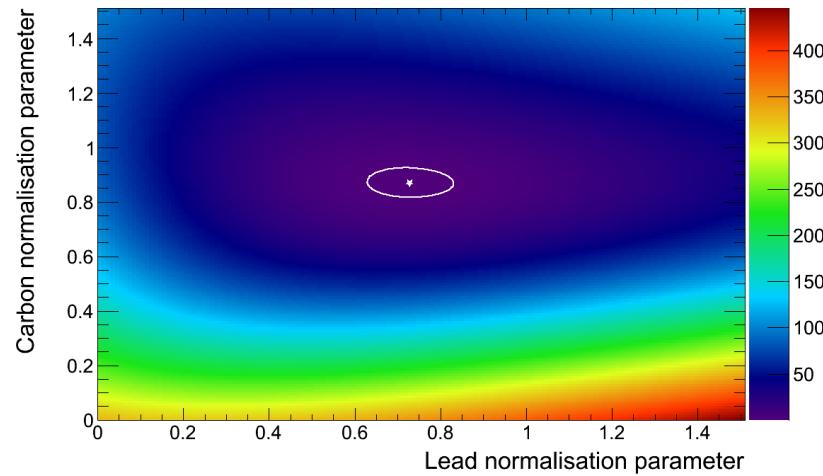


Figure 5.41.: The minimum χ^2 value across a range of values of the R^{Pb} and R^C parameters when fitting the left barrel ECal. The fitted minimum is shown by the black star, with the 1σ contour of the fit shown with a black line.

The measurements of R^{Pb} in Table 5.5 agree with one another within errors, as do the results for R^C , but the external normalisation parameter R^{External} results do not. These findings show that the bottom barrel ECals, where there is a data excess, prefer

a greatly reduced lead normalisation and a greatly increased number of events from external sources. On the other hand, the top and left ECal modules all produce a lead normalisation of about 0.7 and a slightly increased external normalisation. These results indicate the presence of an external, beam related background that is both not modelled in the MC and not covered by the existing systematic uncertainties.

Combining these measurements from Table 5.5 to give a single value of R^{Pb} is not trivial, since the fitting errors will be correlated with each other in some way. By taking the highest and lowest values of R^{Pb} , including their error, a 1σ error envelope can be constructed, and the error-weighted mean of the five measurements can then be used as an estimate of the true value of the lead normalisation. This gives

$$R^{\text{Pb}} = 0.63_{-0.21}^{+0.24}, \quad (5.14)$$

as the measured value of the lead population normalisation.

5.2.5. Extracting the ν_μ -Pb charged current inclusive cross-section

The R^{Pb} parameter can be used to calculate the ν_μ -Pb charged current inclusive cross-section according to

$$\sigma_{\text{Pb}} = \frac{R^{\text{Pb}} N_{\text{MC}}^{\text{Pb}} \eta^{\text{Pb}}}{T^{\text{Pb}} \Phi_{\text{MC}} \epsilon^{\text{Pb}}}, \quad (5.15)$$

where $N_{\text{MC}}^{\text{Pb}}$ is the number of events selected in the DsECal in the MC when scaled to match the data POT total, η^{Pb} is the purity of the DsECal selection, ϵ^{Pb} the efficiency of the selection, Φ_{MC} is the total flux used to create the MC and T^{Pb} is the number of lead nuclei present in the DsECal.

In this analysis, systematic errors on both the number of lead nuclei and the total flux were included in the fit to find R^{Pb} . The error from the fit therefore provides the error from the shape systematic uncertainties. Fig. 5.30b shows the effect of including the normalisation uncertainties on the sub-detector samples, with the DsECal sample having a variance of 0.03. This can be combined in quadrature with the fitting uncertainty to give

$$R^{\text{Pb}} = 0.63_{-0.27}^{+0.29}, \quad (5.16)$$

where the quoted error is almost exclusively from the systematic uncertainties, since the statistical uncertainties are negligible.

As described in Chap 2, the DsECal contains 33 sheets of lead with a total mass of 4949.33 kg. The atomic mass of naturally occurring lead is 207.2 ± 0.1 atomic units [64], giving 1.44×10^{28} lead nuclei in the DsECal in the MC. In reality, the lead sheets contained $2.0 \pm 0.2\%$ antimony, but the effect of this can be ignored given the uncertainty on the normalisation parameter.

From Chap. 4 the purity and efficiency of the DsECal selection for interactions on both carbon and lead nuclei are 59.7% and 55.1% respectively. When considering lead nuclei alone this translates to a selection purity of 41.5% and efficiency of 55.8%, with 86411 events selected in total.

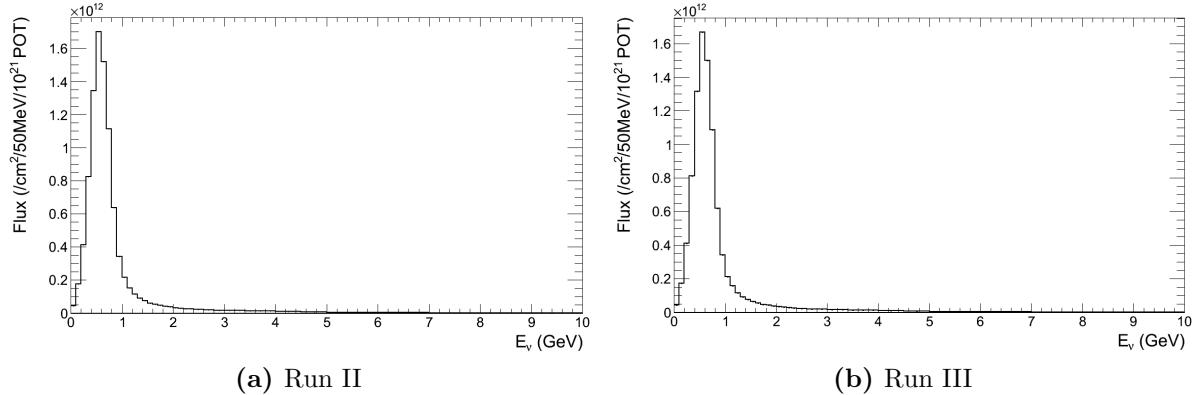


Figure 5.42.: The simulated neutrino flux for T2K Run II, Fig. 5.42a, and Run III, Fig. 5.42b, as a function of neutrino energy.

This analysis ran over 1.015×10^{19} POT of data from T2K Run II and 1.573×10^{20} POT of data from Run III, after data quality cuts. Fig. 5.42 shows the neutrino flux per 500 MeV and per 10^{21} POT for both Run II and Run III. Integrating these histograms and scaling by the relevant POT total, the analysis ran on an integrated neutrino flux of $9.423 \times 10^{11} \nu \text{ cm}^{-2}$ from Run II and $1.441 \times 10^{12} \nu \text{ cm}^{-2}$. Using these figures and Eqn. 5.15, the ν_μ -Pb charged current cross-section is measured to be

$$\sigma_{\text{Pb}}^{\text{CC}} = 5.76_{-2.46}^{+2.65} \times 10^{-39} \text{ cm}^2 \text{ nucleon}^{-1}. \quad (5.17)$$

This can be compared to the ν_μ CC inclusive cross-section reported by Ref. [59]

$$\sigma_{\text{C}}^{\text{CC}} = 6.91 \pm 0.13^{\text{stat}} \pm 0.84^{\text{syst}} \times 10^{-39} \text{ cm}^2 \text{ nucleon}^{-1}, \quad (5.18)$$

showing that, in agreement with the findings from the MINER ν A collaboration [29], the muon neutrino charged current inclusive cross-section per nucleon is the same for both carbon and lead nuclei.

Chapter 6.

Discussion and conclusions

6.1. Discussion of cross-section result

The results shown in Chap. 5 provide some insight into the current T2K simulation. Firstly, Fig. 5.37 shows that there are more events in the bottom ECal modules in the data than appear in the MC. It also shows an excess of data events in each of the background enriched samples, with a greater excess again appearing in the bottom modules.

As mentioned in Chap. 4, there are two types of MC used in this analysis, beam MC and sand MC. The beam MC has been well tested in the tracker region of the ND280, but has not been examined so closely in the surrounding magnet. The sand MC has only recently been developed, and was tuned to match data collected from the tracker region. The distribution of the simulated neutrinos across the face of the ND280 is one aspect of the MC that has had relatively little study, and these results highlight a need for this in the future. A distorted neutrino beam shape could be accounted for by adding an additional parameter to the fit described in Chap. 5. This additional parameter could increase the number of MC events in the bottom ECal modules whilst reducing the number of events in the top modules, changing the shape of the neutrino beam. Though crude, this would help ascertain whether a mis-modelling of the beam shape is the cause of the observed differences and would give a measure of the effect this mis-modelling has.

Table 5.5 shows that when each module is included separately in the fit, the normalisation parameters obtained are all physical and fall into two groups. In the first group, the top and left ECal modules give a lead normalisation of about 0.73, compared to a carbon normalisation of 0.88 and an external normalisation of around 1.25. On the other

hand, the bottom modules find a lead normalisation of 0.53, a carbon normalisation of 0.77 and an external normalisation of 1.68. All the modules indicate that the ratio of external events to internal events is higher in the data than the MC, particularly so in the bottom modules, suggesting that it is the external population of the MC that is not well modelled. This could be tested in future versions of this analysis by using cuts on the BDT classifier distribution to try and remove this poorly understood background. This would then allow the entire barrel ECal to be used in the fit and would reduce the uncertainty on R^{Pb} from the fit itself to around 0.1 rather than the current value of 0.21.

The measured cross-section presented in this thesis has a large error associated with it, preventing it from providing much information about the underlying cross-section physics. This analysis has however shown that a better simulation of the neutrino beam is required to correctly describe events outside the ND280 tracker region, and in particular that it is in the regions outside the barrel ECals where the difference between data and MC is most noticeable. It has also shown that the fitting technique described in Chap. 5 can be successfully applied to the ND280 data samples, and when the neutrino beam is better understood it will provide a competitive measurement of the ν_μ -Pb CC inclusive cross-section.

6.2. Conclusions

The field of neutrino physics has moved on rapidly since the discovery of neutrino oscillations and is now entering the precision era. The current and planned neutrino beam experiments aim to measure both the neutrino mass hierarchy and the CP-violating phase δ , and in order to do this they require a better understanding of neutrino interactions.

This thesis has presented the current state of neutrino physics and discussed some of the challenges that are being faced. It has given an overview of the T2K experiment, which is working to overcome these challenges, and described improvements to the ND280 event reconstruction algorithms that will help the experiment achieve its goals.

The main body of this work was an investigation into neutrino interactions within the ND280 ECals, leading to a measurement of the ν_μ -Pb CC inclusive cross-section. This involved the creation of a BDT multi-variate analysis to select neutrino interactions within the ECals, and the construction of a χ^2 fit to extract the cross-section from the selected data.

The event selection described in Chap. 4 showed that it is possible to select neutrino interactions with high purity and efficiency in both the DsECal and barrel ECals. This event selection could also become a useful tool for other analyses at the ND280, allowing them to remove backgrounds originating from interactions in the ECals. The data-MC comparisons in Chap. 4 highlighted the discrepancies that exist between the ND280 simulation and the data, leading to a robust selection that still provided separation between the signal interactions and backgrounds.

This selection was applied to the T2K Run II and Run III data, selecting approximately 500,000 neutrino interactions in the six ECal modules used in the analysis, described in Chap. 5. The effects of neutrino interaction uncertainties and the neutrino flux uncertainties were examined in the regions outside the ND280 tracker, and the first low-level detector systematic studies for the ECals were performed. Comparing the MC samples to the selected data showed that there was a data excess present in both the signal and background enriched samples, particularly in the bottom ECal modules.

These differences did not allow the cross-section fit to produce a good match between the data and MC, as it returned an unphysical value for the lead normalisation parameter. This tension was removed by fitting the barrel ECal modules sequentially, producing a value of R^{Pb} for each individual module. The range of values seen represents the additional uncertainty arising from the discrepancy between the data and MC, which was incorporated into the normalisation parameter uncertainty. This gave the final measured ν_μ -Pb CC inclusive cross-section as

$$\sigma_{\text{Pb}}^{\text{CC}} = 5.76_{-2.46}^{+2.65} \times 10^{-39} \text{ cm}^2 \text{ nucleon}^{-1}. \quad (6.1)$$

Comparing this result to the ν_μ -carbon CC inclusive cross-section measured by T2K shows that they are equal within errors, agreeing with the recent measurements of the lead-carbon cross-section ratio reported by the MINER ν A collaboration.

6.3. Outlook

The analysis presented here was the first to look in detail at neutrino interactions both outside the main ND280 tracker region and within the ND280 ECals. As such there are many areas that can be improved upon in future iterations.

As mentioned in Chap. 3, work on the Radon transform to include a shower classification, improved 3D matching and allow the fitting of curved tracks will provide better vertex identification and could allow future analyses to split the ECal interaction sample according to the number of final state particles produced. This would improve both the ability to select neutrino interactions and give more information on the type of interaction that occurred.

Ongoing work on understanding the ECal charge response and modelling this correctly in the ND280 MC will provide better agreement between the data and MC for the low level ECal variables, allowing more of them to be used in the BDT, improving its separating power. This would in turn provide a stronger constraint on the lead cross-section. More use could also be made of the BDT information when extracting the cross-section. The current fit splits the BDT distribution to provide signal and background enriched samples which are then used in the χ^2 fit. Instead of performing a two bin fit in this way the BDT distribution itself could be fitted for each of the ECal modules, which should give smaller uncertainties on the normalisation parameters.

The greatest uncertainty on the measured cross-section arose from the discrepancies observed in the distribution of events across the ECal modules between data and MC. There are a number of things that could account for this, such as structural sections of the ND280 pit that are not included in the sand MC or a mis-modelling of the neutrino beam flux, and a number of groups within the T2K collaboration are now focussed on resolving this issue. Once this has been fully understood it will be possible to fit the ECal modules simultaneously, reducing the systematic uncertainty on the cross-section measurement by a factor of 2.

There is also much scope for extending this analysis beyond what has been presented here. In particular, muon ranging and calorimetric measurements of the neutrino interaction could be used to extract the ν_μ -Pb cross-section as a function of neutrino energy, or improved event identification could enable the measurement of exclusive CC interaction modes.

Bibliography

- [1] ATLAS Collaboration, G. Aad *et al.*, Physics Letters B **716**, 1 (2012).
- [2] CMS Collaboration, S. Chatrchyan *et al.*, Physics Letters B **716**, 30 (2012).
- [3] R. Davis, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).
- [4] K. S. Hirata *et al.*, Phys. Rev. Lett. **65**, 1297 (1990).
- [5] A. I. Abazov *et al.*, Phys. Rev. Lett. **67**, 3332 (1991).
- [6] P. Anselmann *et al.*, Physics Letters B **285**, 376 (1992).
- [7] K. S. Hirata *et al.*, Physics Letters B **205**, 416 (1988).
- [8] D. Casper *et al.*, Phys. Rev. Lett. **66**, 2561 (1991).
- [9] W. Allison *et al.*, Physics Letters B **391**, 491 (1997).
- [10] Y. Fukuda *et al.*, Physics Letters B **436**, 33 (1998).
- [11] SNO Collaboration, G. A. McGregor, (2002), nucl-ex/0205006.
- [12] KamLAND Collaboration, T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005).
- [13] Particle Data Group, J. Beringer *et al.*, Phys. Rev. D **86**, 010001 (2012).
- [14] T2K Collaboration, K. Abe *et al.*, Phys. Rev. Lett. **107**, 041801 (2011).
- [15] DAYA-BAY Collaboration, F. An *et al.*, Phys.Rev.Lett. **108**, 171803 (2012), 1203.1669.
- [16] RENO collaboration, J. Ahn *et al.*, Phys.Rev.Lett. **108**, 191802 (2012), 1204.0626.
- [17] J. A. Formaggio and G. P. Zeller, Rev. Mod. Phys. **84**, 1307 (2012).
- [18] MiniBooNE Collaboration, A. A. Aguilar-Arevalo *et al.*, Phys. Rev. D **81**, 092005 (2010).
- [19] NOMAD Collaboration, V. Lyubushkin *et al.*, Eur.Phys.J. **C63**, 355 (2009), 0812.4543.
- [20] K2K Collaboration, R. Gran *et al.*, Phys. Rev. D **74**, 052002 (2006).
- [21] M. Dorman and M. Collaboration, Preliminary Results for CCQE Scattering with the MINOS Near Detector, 2009, 6th International Workshop on Neutrino-Nucleus

Interactions in the Few-GeV Region, Sitges, SPAIN, MAY 18-22, 2009.

- [22] L. Fields, (2012), 1212.0060.
- [23] R. Smith and E. Moniz, Nuclear Physics B **43**, 605 (1972).
- [24] E. J. Moniz *et al.*, Phys. Rev. Lett. **26**, 445 (1971).
- [25] A. Butkevich and S. Mikheyev, Phys.Rev. **C72**, 025501 (2005), hep-ph/0505008.
- [26] Y. Hayato, Acta Physica Polonica B **40**, 2477 (2009).
- [27] C. Andreopoulos *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **614**, 87 (2010).
- [28] CHORUS Collaboration, Eur. Phys. J. C **30**, 159 (2003).
- [29] K. Hurtado, CC Inclusive and DIS Measurements at MINERvA, 2012, NuInt.
- [30] Y. Hayato *et al.*, Letter of Intent Neutrino Oscillation Experiment at JHF, 2003.
- [31] T2K Collaboration, K. Abe *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **659**, 106 (2011).
- [32] F. Retire *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **610**, 378 (2009), New Developments In Photodetection NDIP08.
- [33] A. Vacheret, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **623**, 201 (2010), 1st International Conference on Technology and Instrumentation in Particle Physics.
- [34] A. Vacheret *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **656**, 69 (2011).
- [35] F. Retire, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **623**, 495 (2010), 1st International Conference on Technology and Instrumentation in Particle Physics.
- [36] Super-Kamiokande Collaboration, Y. Ashie *et al.*, Phys. Rev. D **71**, 112005 (2005).
- [37] P. A. S. Ritt and K. Olchanski, MIDAS (Maximum Integration Data Acquisition System), 2001.
- [38] A. Cervera-Villanueva, J. Gmez-Cadenas, and J. Hernando, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **534**, 180 (2004), Proceedings of the IXth International Workshop on Advanced Computing and Analysis Techniques in Physics Research.
- [39] A. Hillairet *et al.*, ND280 Reconstruction, 2011, T2K internal technical note TN 72.

- [40] N. Abgrall *et al.*, Neutrino Flux Prediction for the 2010a Analysis, 2011, T2K internal technical note TN 38.
- [41] FLUKA version 2008.3c, 2008.
- [42] CERN Application Software Group, GEANT3, 1993, A detector description and simulation tool.
- [43] NA61/SHINE Collaboration, N. Abgrall *et al.*, Phys. Rev. C **84**, 034604 (2011).
- [44] NA61/SHINE Collaboration, N. Abgrall *et al.*, Phys. Rev. C **85**, 035210 (2012).
- [45] R. Brun and F. Rademakers, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **389**, 81 (1997), New Computing Techniques in Physics Research V.
- [46] S. Agostinelli *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506**, 250 (2003).
- [47] A. Hoecker *et al.*, PoS **ACAT**, 040 (2007), physics/0703039.
- [48] T2KUK Group, Y. Uchida *et al.*, ND280 DsECAL Status Report for 2010a Analysis, 2010, T2K internal technical note TN 18.
- [49] R. Mankel, Reports on Progress in Physics **67**, 553 (2004).
- [50] T2K Collaboration, K. Abe *et al.*, Phys. Rev. D **87**, 012001 (2013).
- [51] M. Bass *et al.*, Data Quality at the near detectors for Run2, 2011, T2K internal technical note TN 50.
- [52] B. P. Roe *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **543**, 577 (2005).
- [53] Y. Freund and R. E. Schapire, Journal of Computer and System Sciences **55**, 119 (1997).
- [54] A. Carver, Particle Identification in the ND280 Electromagnetic Calorimeter, 2009, T2K internal technical note TN 2.
- [55] G. Barker *et al.*, Implementation of the Second Generation PID for the ND280 Tracker ECals, 2012, T2K internal technical note TN 111.
- [56] P. Guzowski, Reconstruction of neutrino induced neutral current neutral pion events with the T2K ND280 Tracker and ECAL, 2011, Ph.D. Thesis, Imperial College, UK.
- [57] K2K Collaboration, A. Rodriguez *et al.*, Phys. Rev. D **78**, 032003 (2008).
- [58] SciBooNE Collaboration, Y. Nakajima *et al.*, Phys. Rev. D **83**, 012005 (2011).
- [59] T2K Collaboration, K. Abe *et al.*, (2013), 1302.4908.

- [60] P. de Perio *et al.*, Implementation of the NIWG Cross Section Parametrization, 2012, T2K internal technical note TN 113.
- [61] C. Andreopoulos *et al.*, Handling Neutrino Interaction Uncertainties using Event Reweighting, 2009, T2K internal technical note TN 7.
- [62] P. Masliah and A. Vacheret, Characterisation of the MPPC linearity response with TRIPt electronics and Y-11 fibre, 2011, T2K internal technical note TN 47.
- [63] C. Bojeckho *et al.*, CCQE-like and CC-non-QE-like event selections in the ND280 tracker using Run 1+2 data, 2012, T2K internal technical note TN 93.
- [64] J. R. de Laeter *et al.*, Pure Appl. Chem. **75**, 683 (2003).

Appendix A.

Boosted decision tree input variables

Plots showing the distribution of signal, through-going background and signal-like background events in the DsECal and barrel ECals for each of the BDT input variables described in Chapter 4. In all cases the blue histogram contains signal events, the yellow filled histogram holds through-going background events and the red histogram shows signal-like background events. The combined area of the background histograms equals that of the signal histogram.

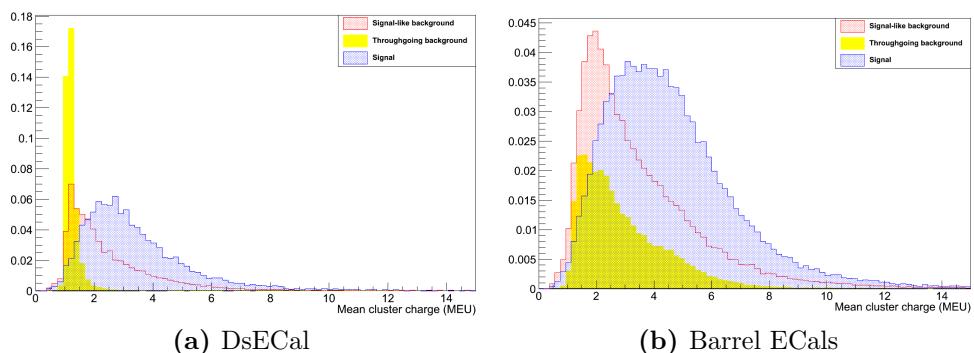


Figure A.1.: The mean charge of all hits in the cluster.

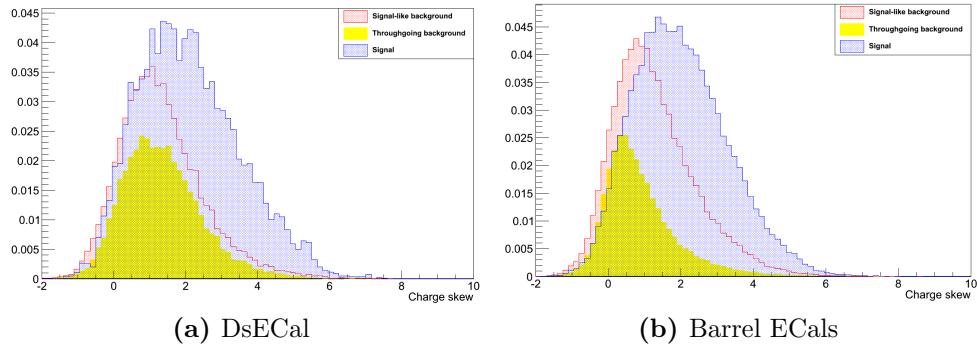


Figure A.2.: The charge skew of hits in the cluster.

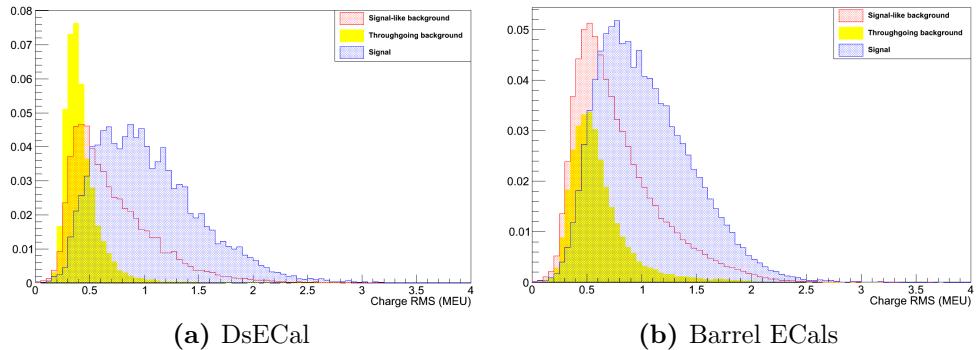


Figure A.3.: The RMS of the charge of hits in the cluster.

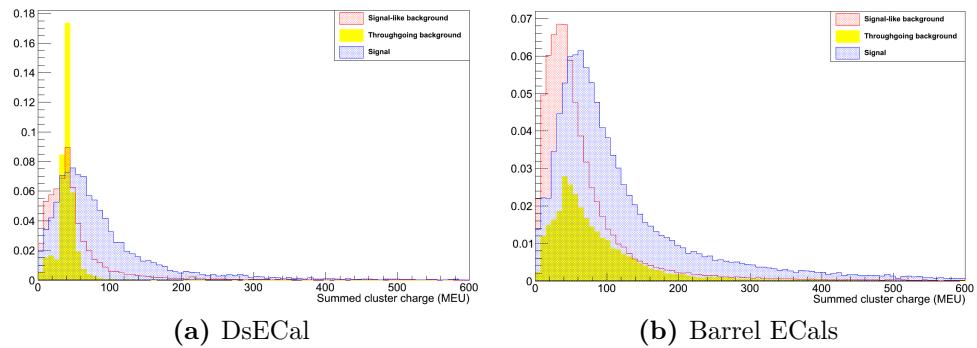


Figure A.4.: The summed charge of all hits in the cluster.

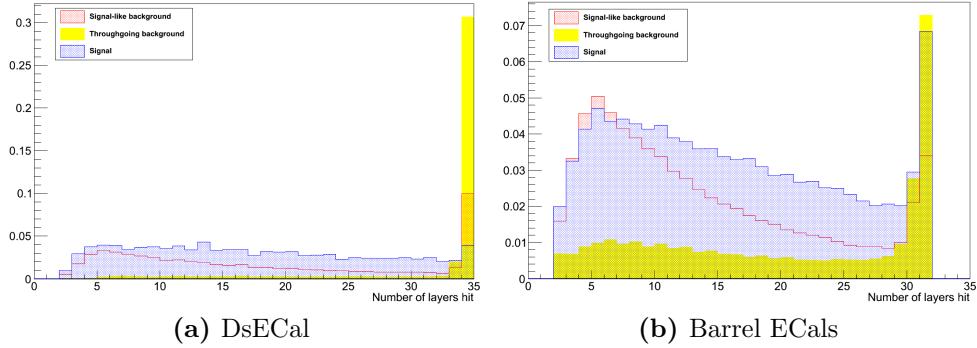


Figure A.5.: The number of ECal layers containing a hit from the cluster.

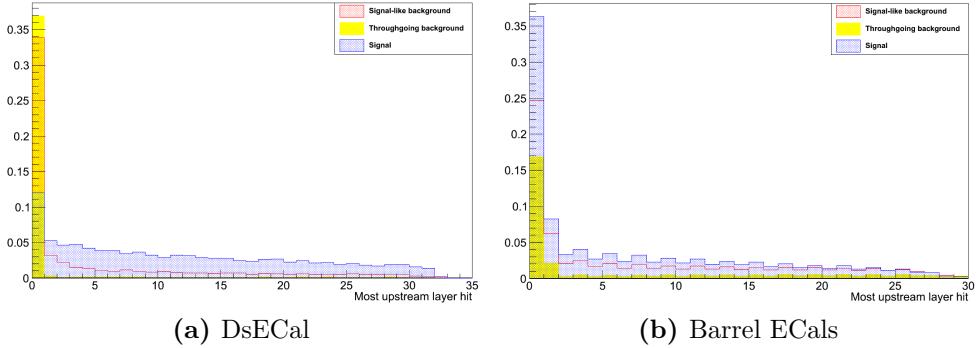


Figure A.6.: The layer closest to the centre of the ND280 that contains a hit from the cluster.

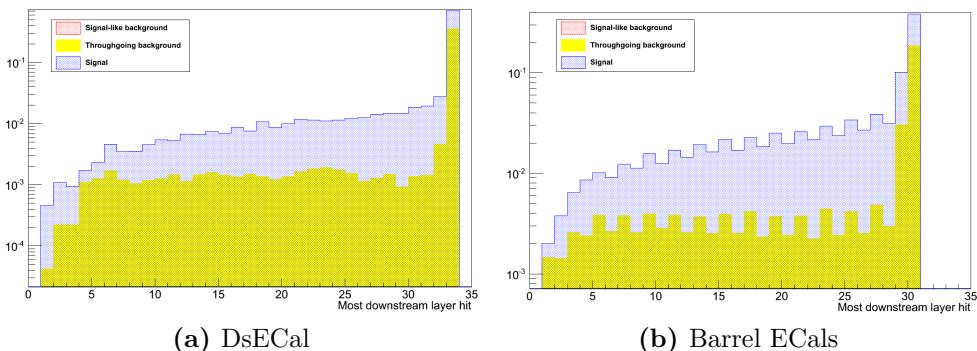


Figure A.7.: The layer furthest from the centre of the ND280 that contains a hit from the cluster.

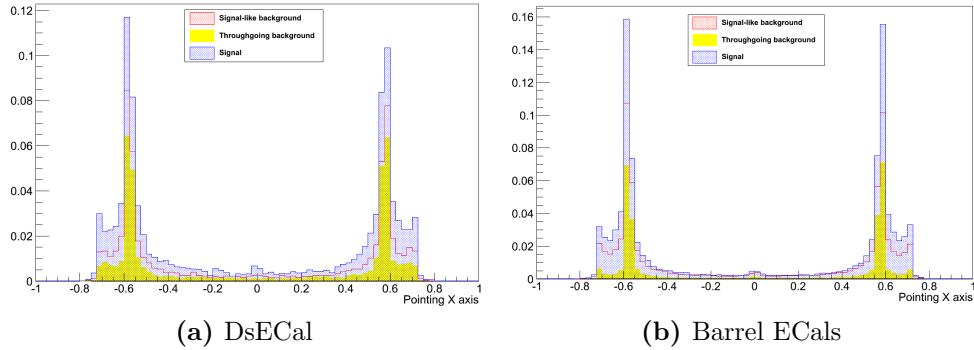


Figure A.8.: Pointing X axis.

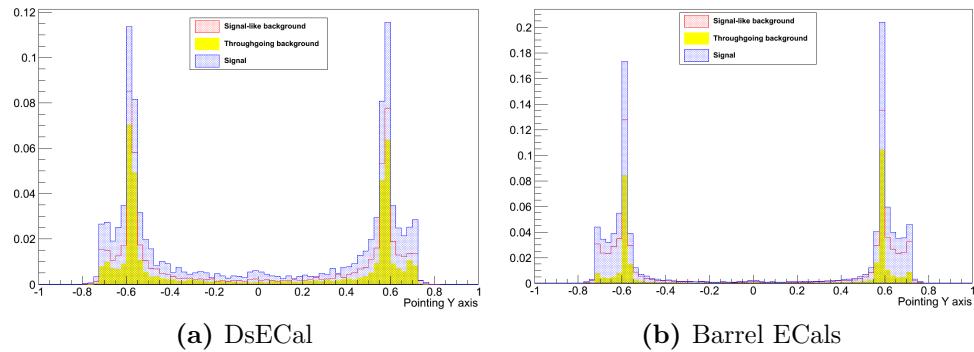


Figure A.9.: Pointing Y axis.

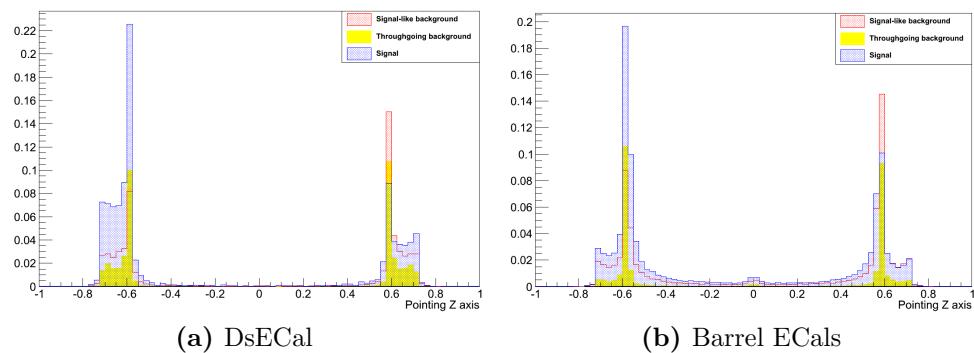


Figure A.10.: Pointing Z axis.

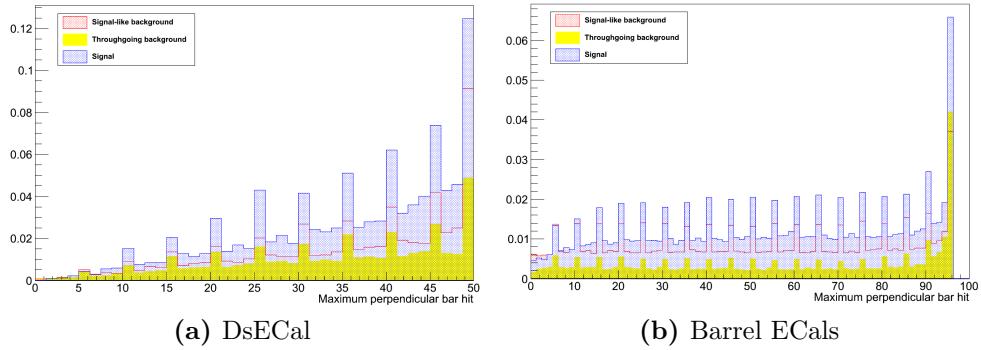


Figure A.11.: Maximum perpendicular bar hit.

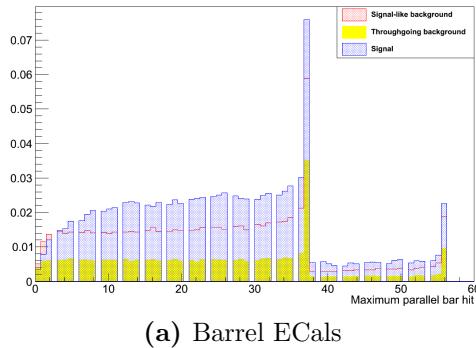


Figure A.12.: Maximum parallel bar hit.

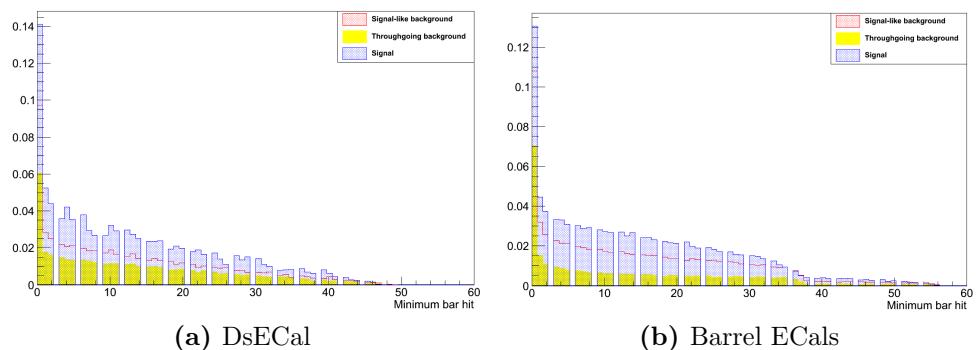


Figure A.13.: Minimum bar hit.

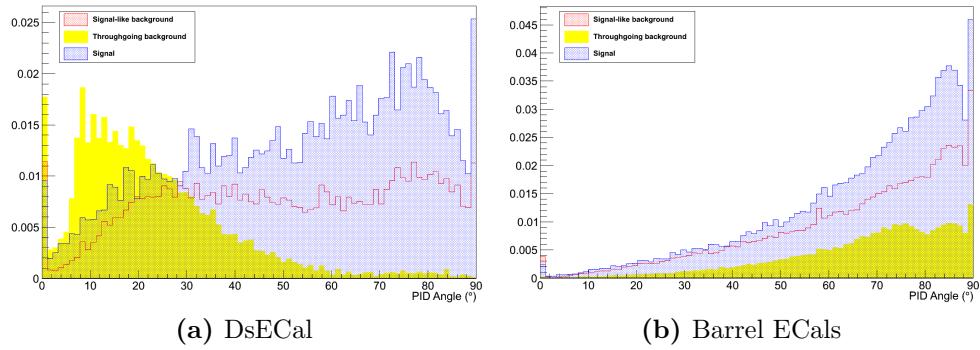


Figure A.14.: PID Angle.

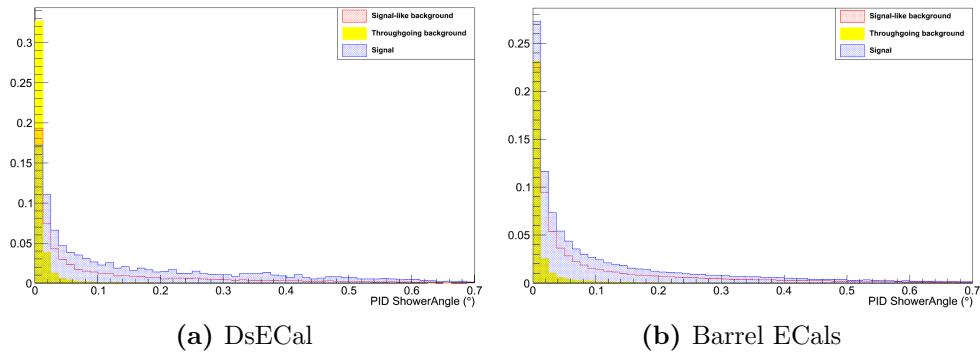


Figure A.15.: PID ShowerAngle.

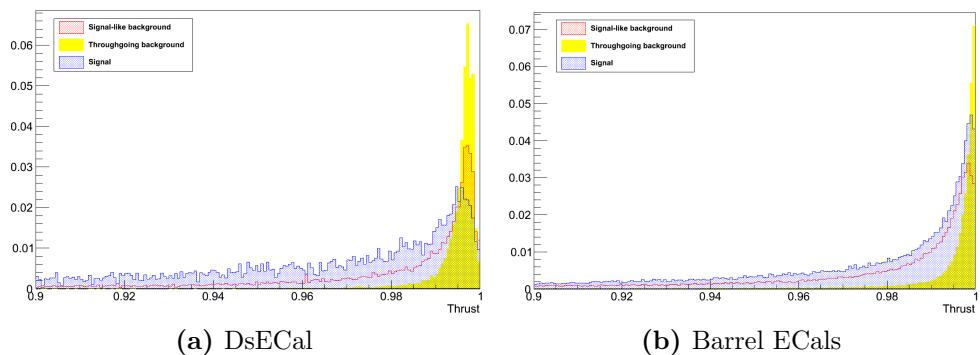


Figure A.16.: Thrust value.

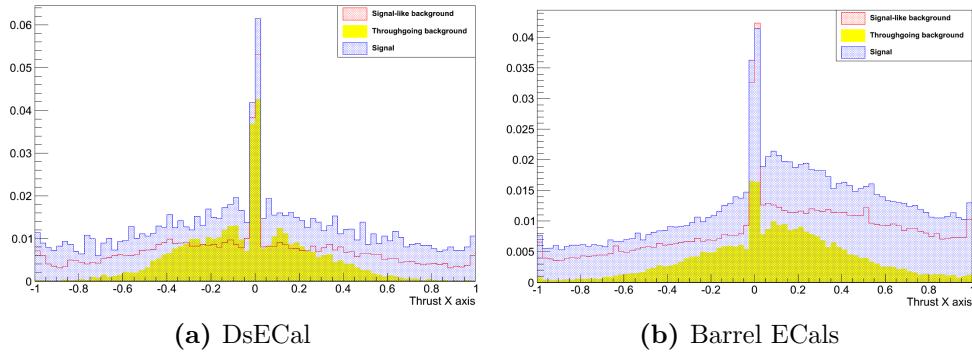


Figure A.17.: Thrust X axis.

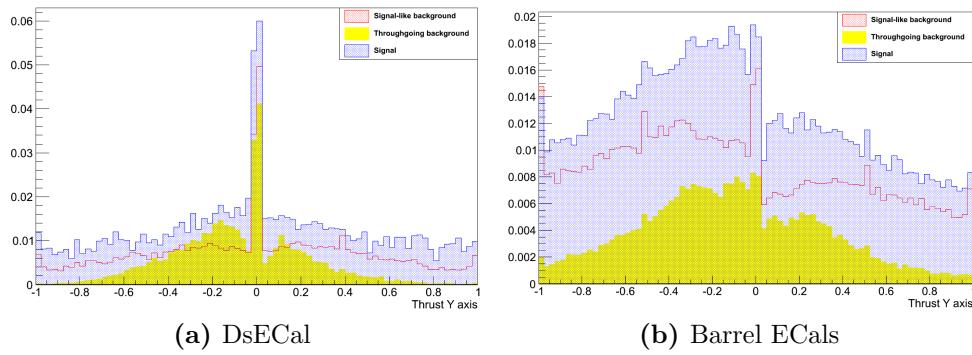


Figure A.18.: Thrust Y axis.

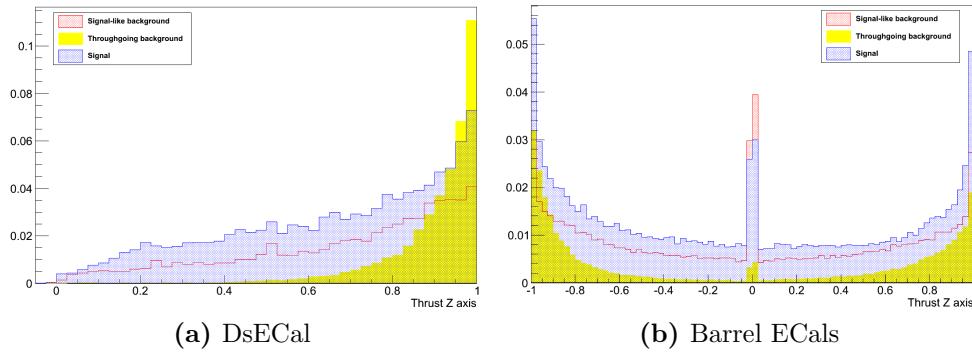


Figure A.19.: Thrust Z axis.

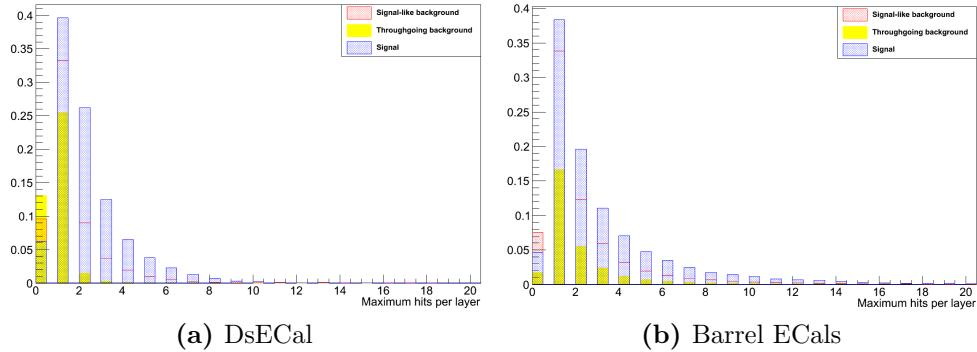


Figure A.20.: Maximum hits per layer.

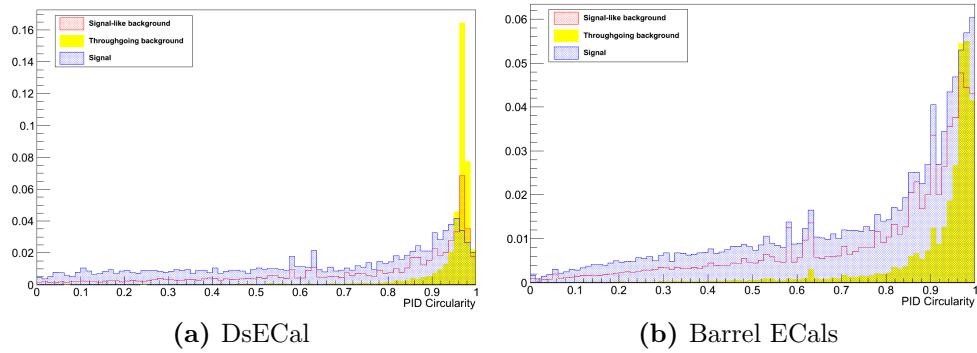


Figure A.21.: PID Circularity.

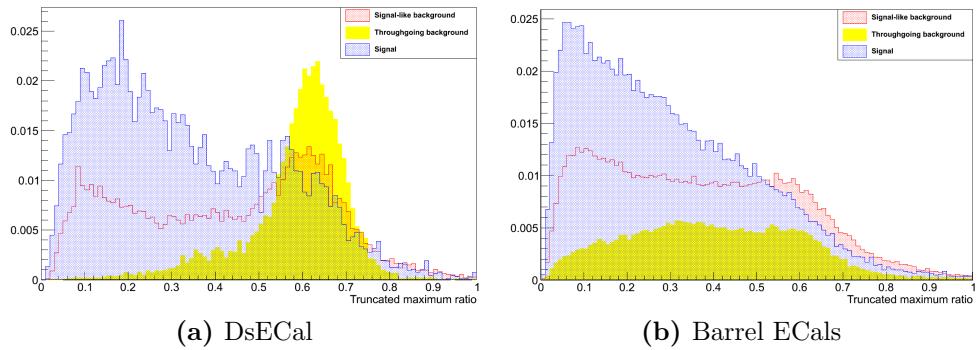


Figure A.22.: Truncated maximum ratio.

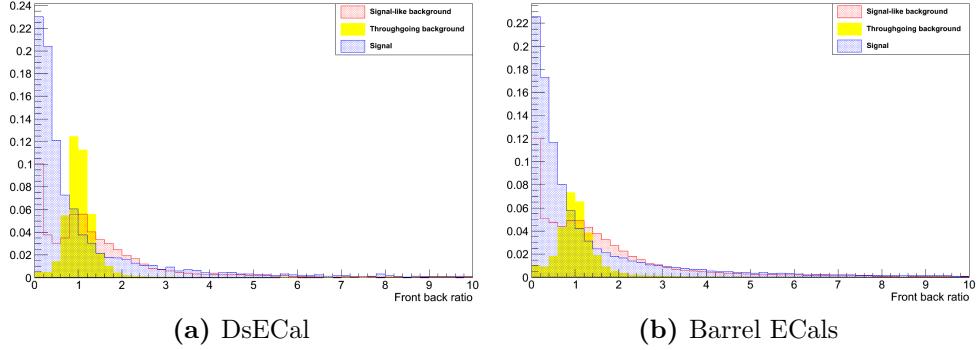


Figure A.23.: Front back ratio.

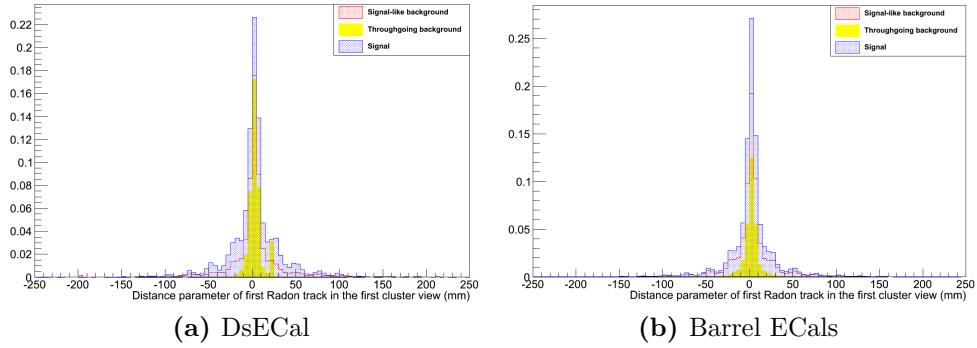


Figure A.24.: Distance parameter of first Radon line in first view of cluster.

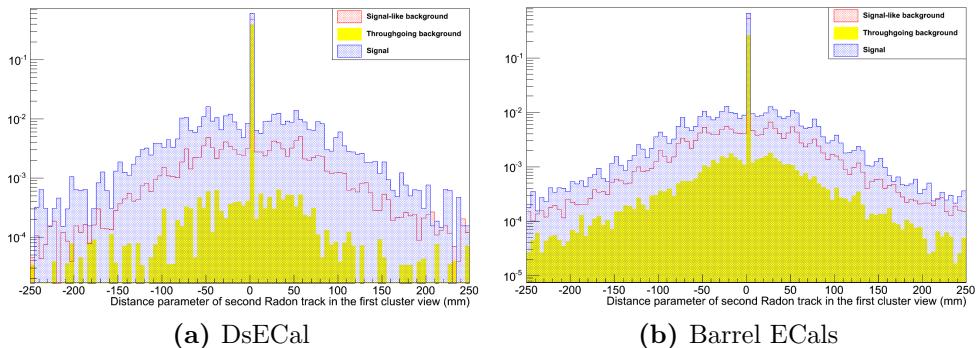


Figure A.25.: Distance parameter of second Radon line in first view of cluster.

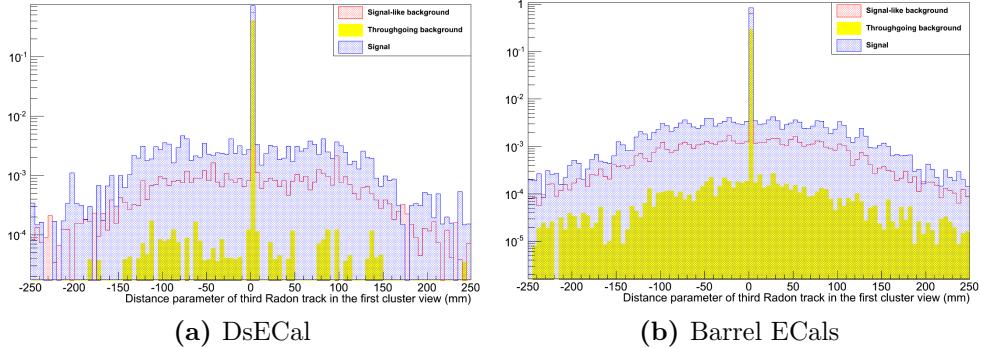


Figure A.26.: Distance parameter of third Radon line in first view of cluster.

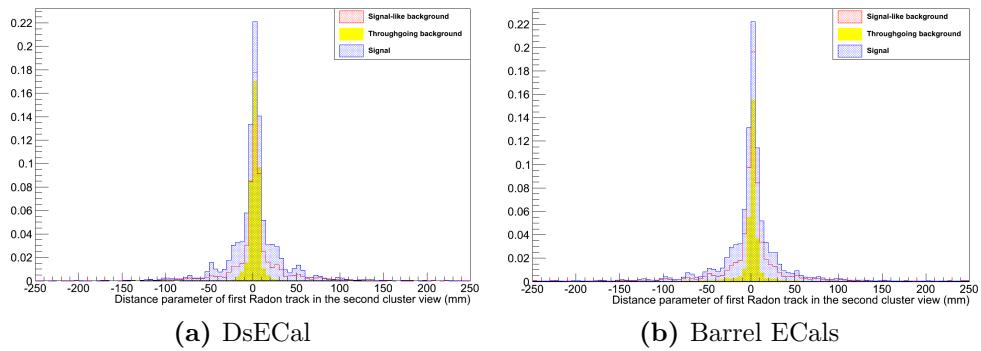


Figure A.27.: Distance parameter of first Radon line in second view of cluster.

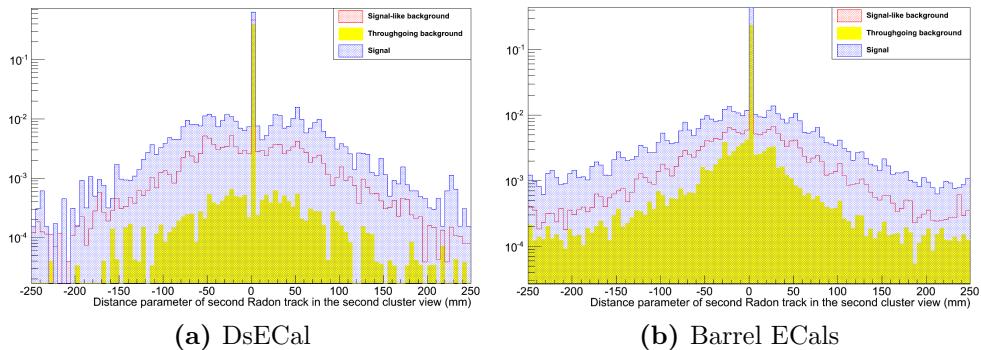


Figure A.28.: Distance parameter of second Radon line in second view of cluster.

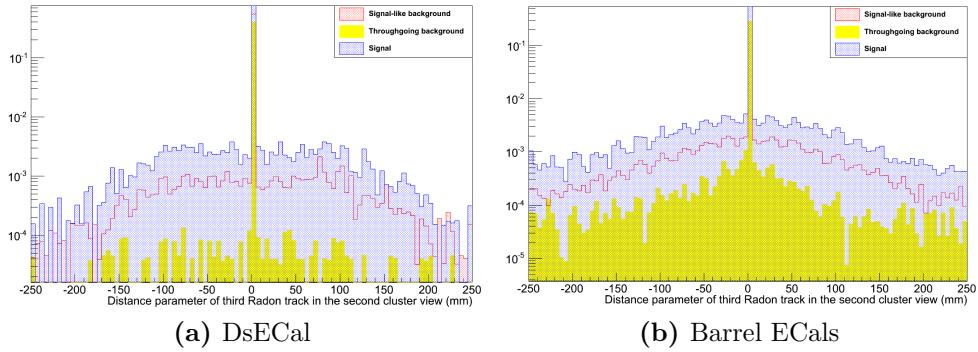


Figure A.29.: Distance parameter of third Radon line in second view of cluster.

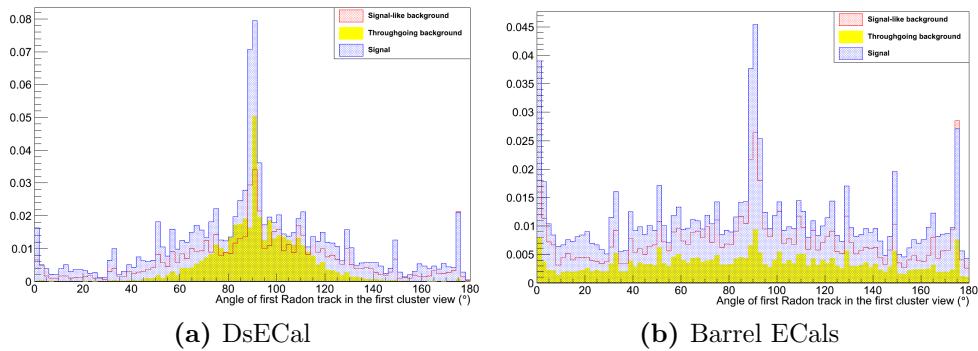


Figure A.30.: Angle of first Radon line in first view of cluster.

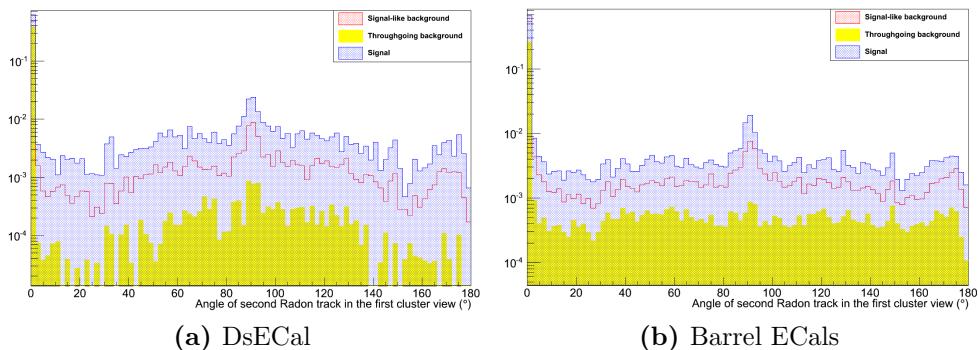


Figure A.31.: Angle of second Radon line in first view of cluster.

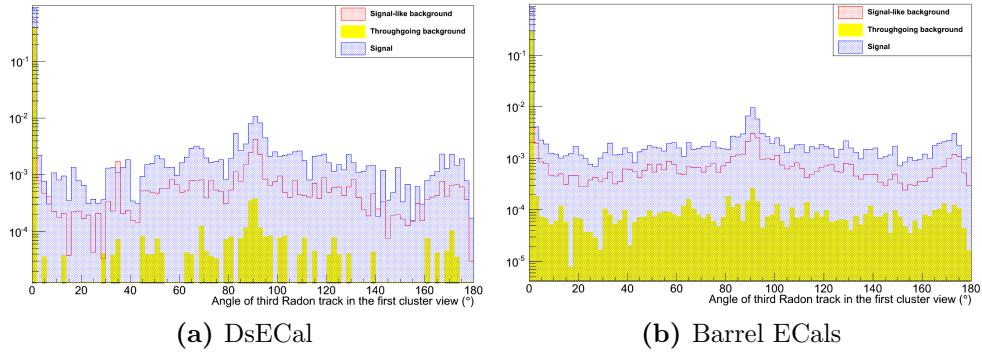


Figure A.32.: Angle of third Radon line in first view of cluster.

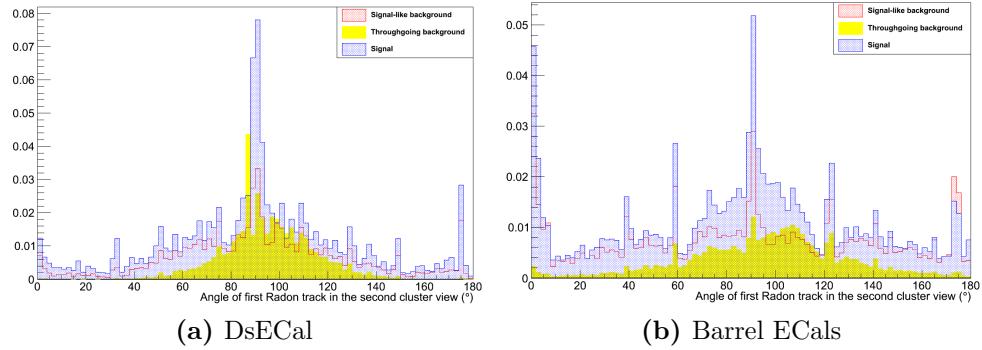


Figure A.33.: Angle of first Radon line in second view of cluster.

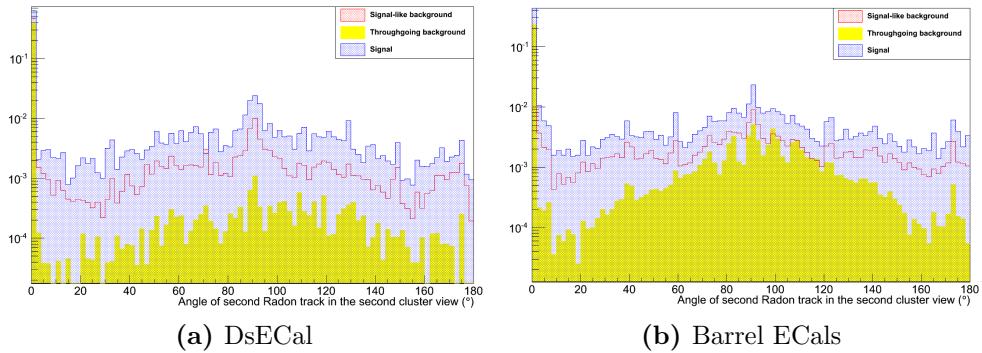


Figure A.34.: Angle of second Radon line in second view of cluster.

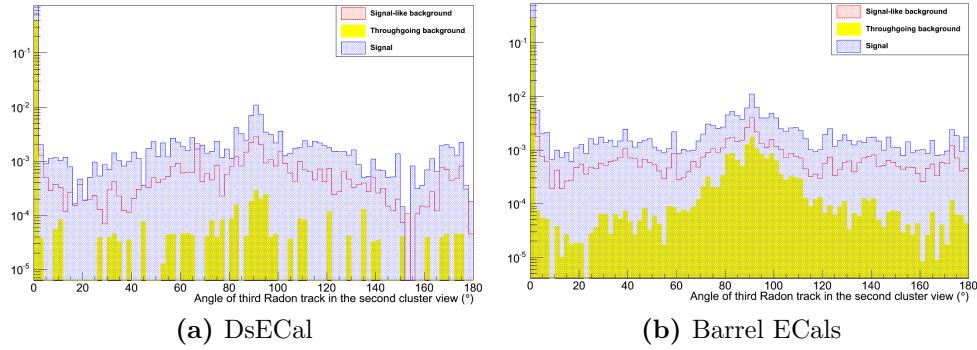


Figure A.35.: Angle of third Radon line in second view of cluster.

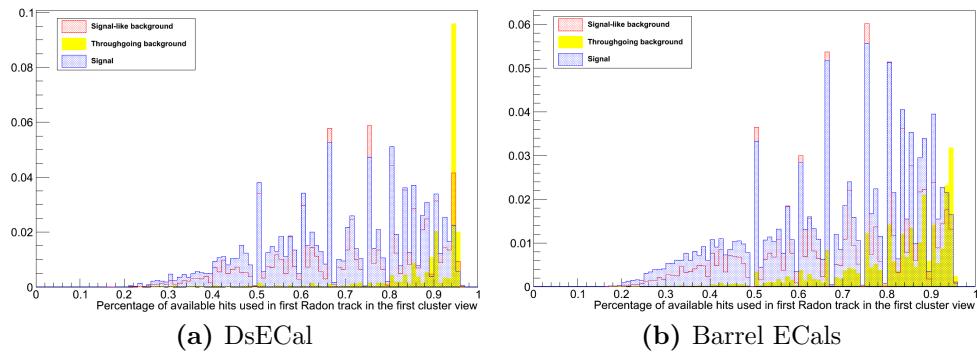


Figure A.36.: Percentage of available hits used in first Radon line in first view of cluster.

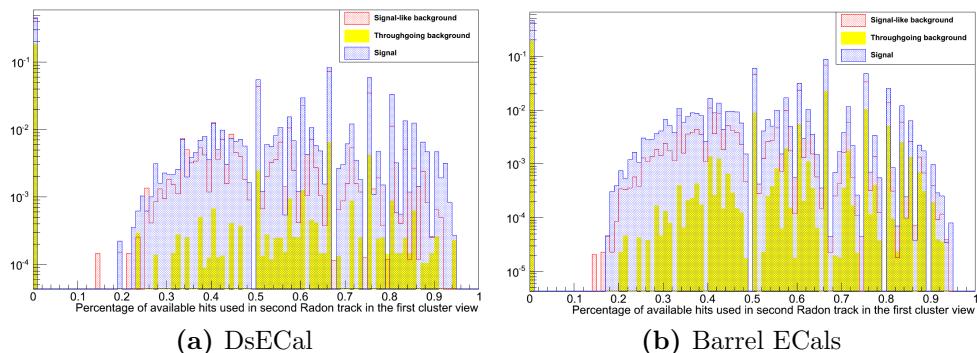


Figure A.37.: Percentage of available hits used in second Radon line in first view of cluster.

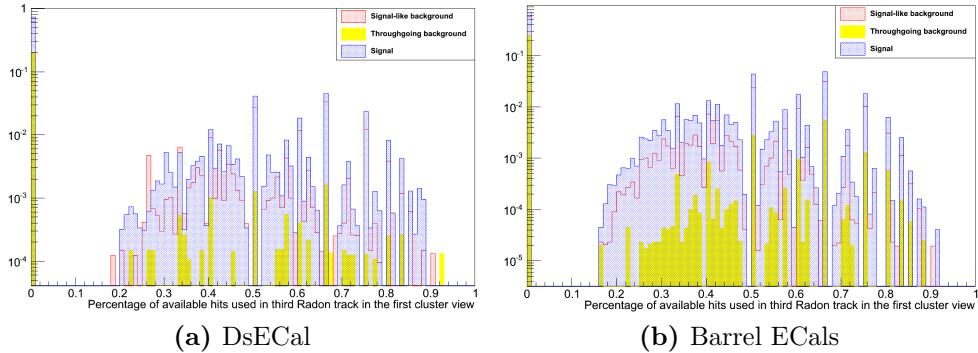


Figure A.38.: Percentage of available hits used in third Radon line in first view of cluster.

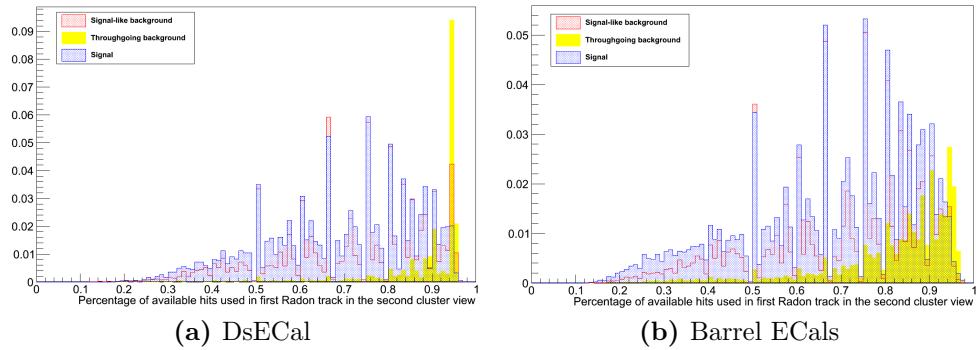


Figure A.39.: Percentage of available hits used in first Radon line in second view of cluster.

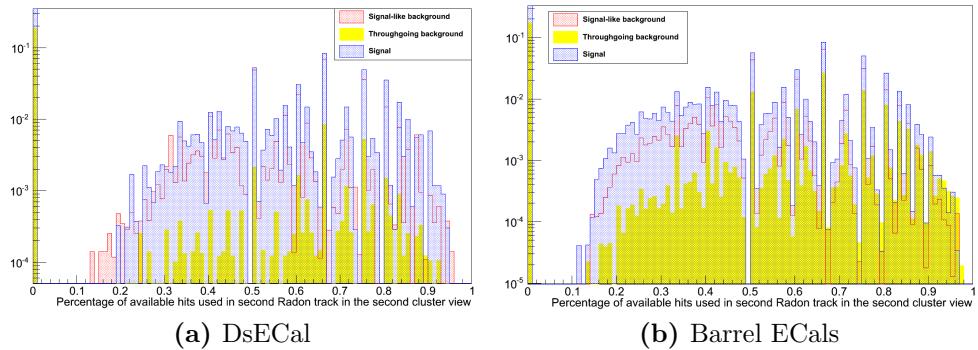


Figure A.40.: Percentage of available hits used in second Radon line in second view of cluster.

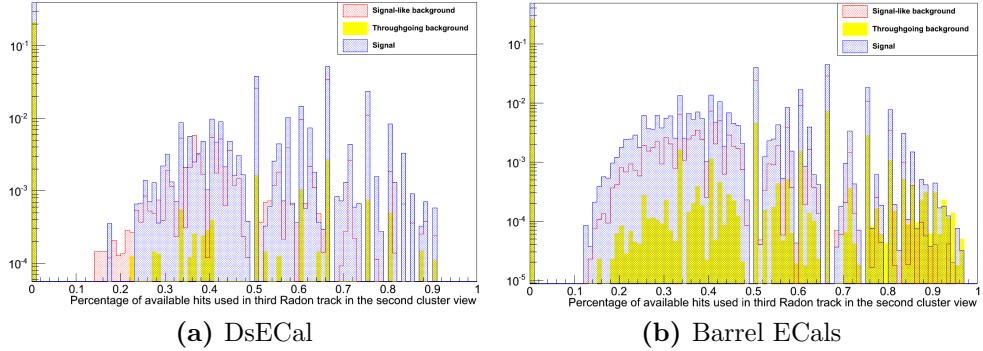


Figure A.41.: Percentage of available hits used in third Radon line in second view of cluster.

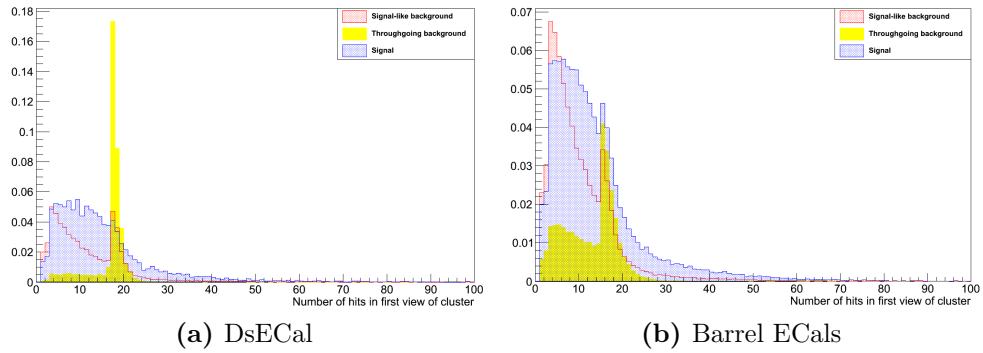


Figure A.42.: Number of hits in the first view of the cluster.

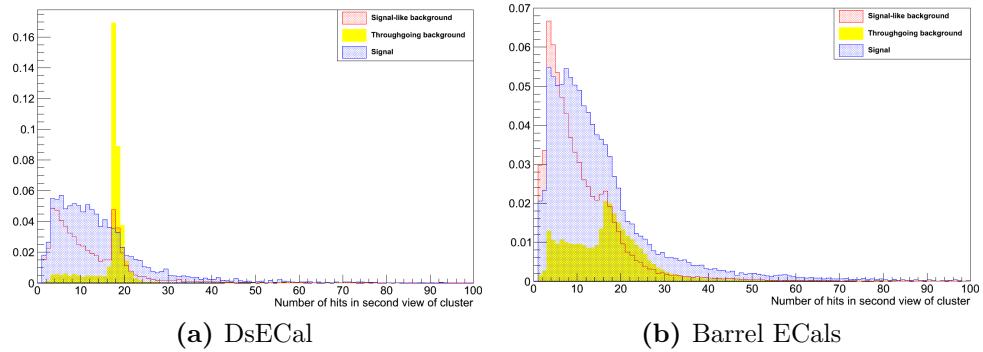


Figure A.43.: Number of hits in the second view of the cluster.

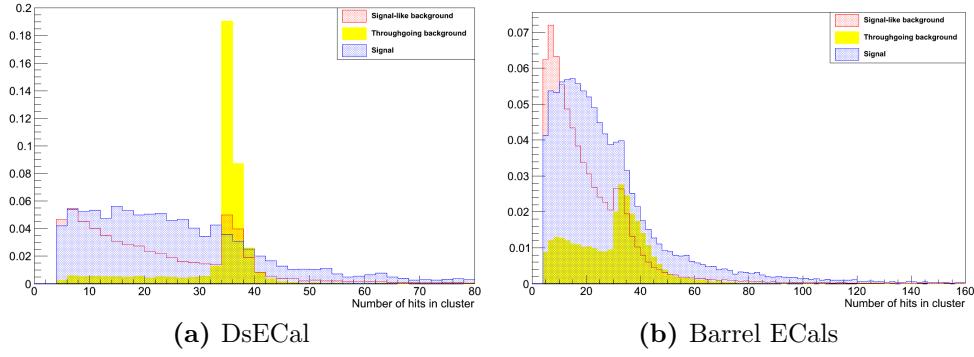


Figure A.44.: The total number of hits in the cluster.

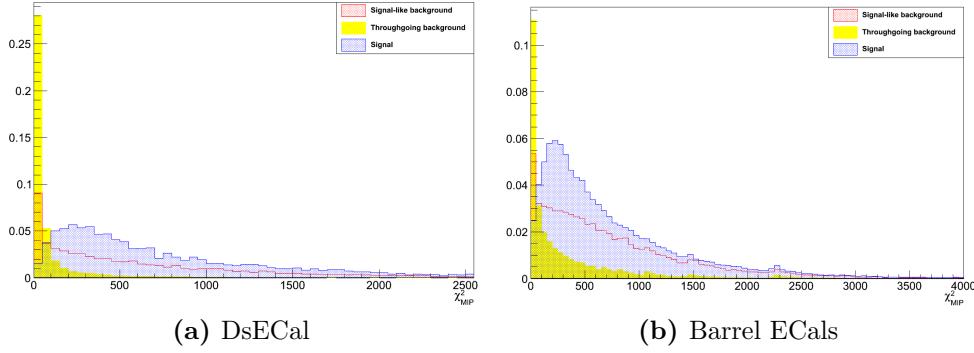


Figure A.45.: χ^2_{MIP} .

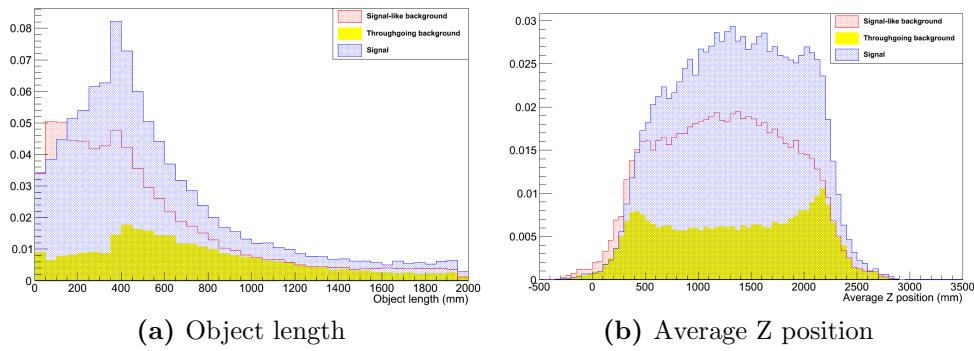


Figure A.46.: The average Z position and length of clusters in the barrel ECal modules.

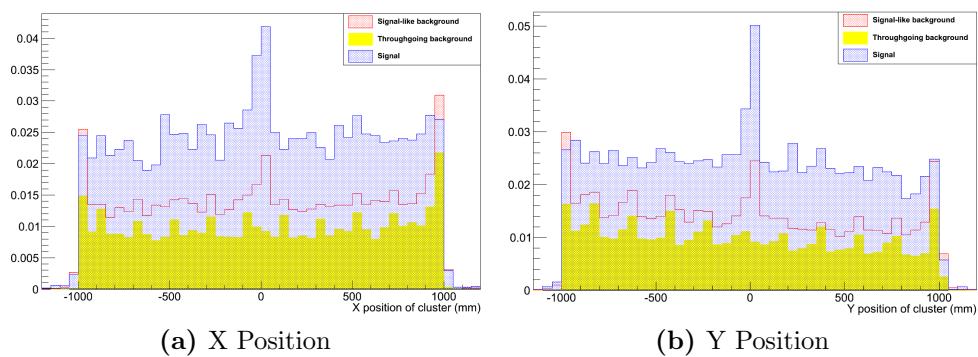


Figure A.47.: The X-Y position of clusters in the DsECal.

Appendix B.

‘External’ population in barrel ECal samples

Plots showing the energy of neutrino interactions in the ‘External’ population of each of the signal enriched and background enriched sub-detector samples. The events have been categorised by the type of interaction that occurred.

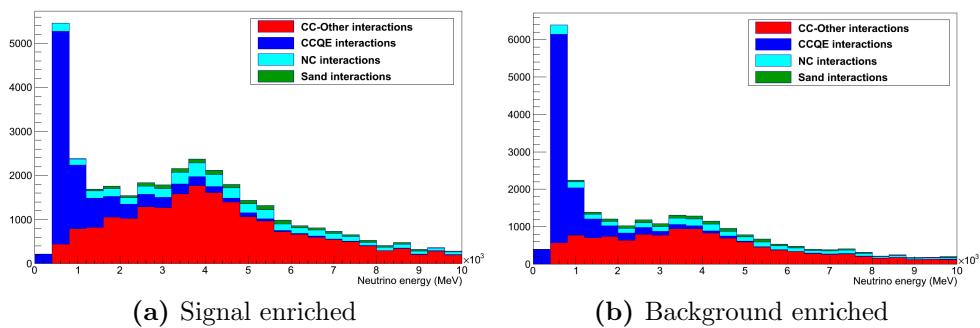


Figure B.1.: Top right barrel ECal module ‘External’ population.

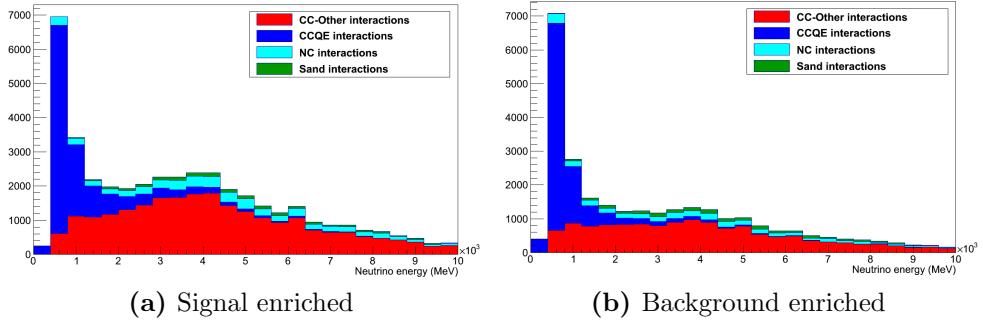


Figure B.2.: Top left barrel ECal module 'External' population.

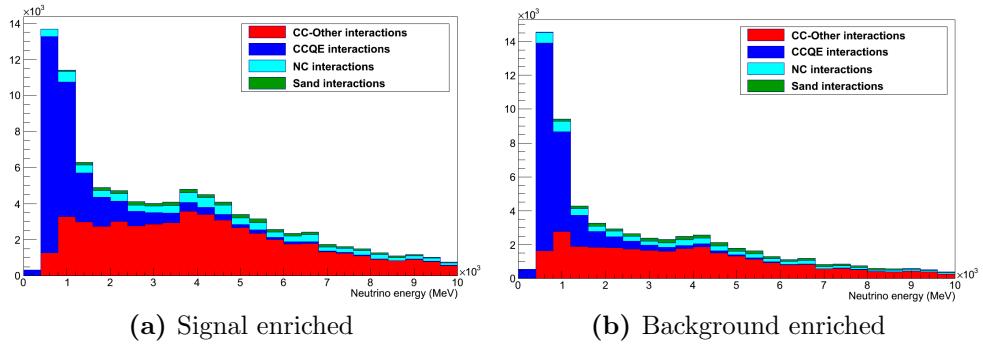


Figure B.3.: Left side barrel ECal module 'External' population.

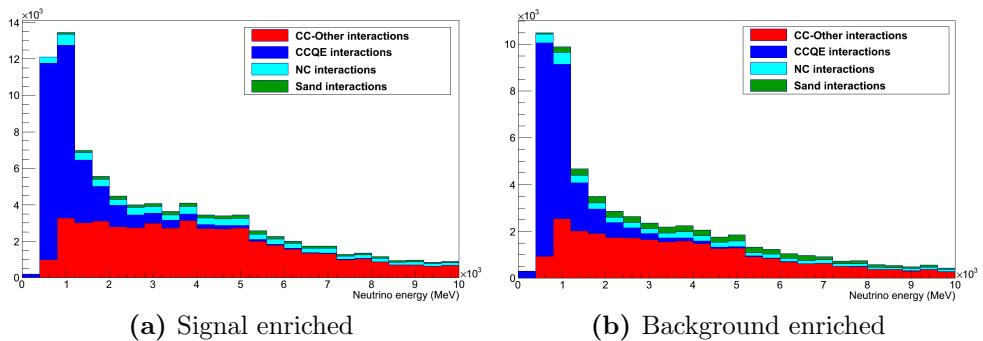


Figure B.4.: Bottom right barrel ECal module 'External' population.

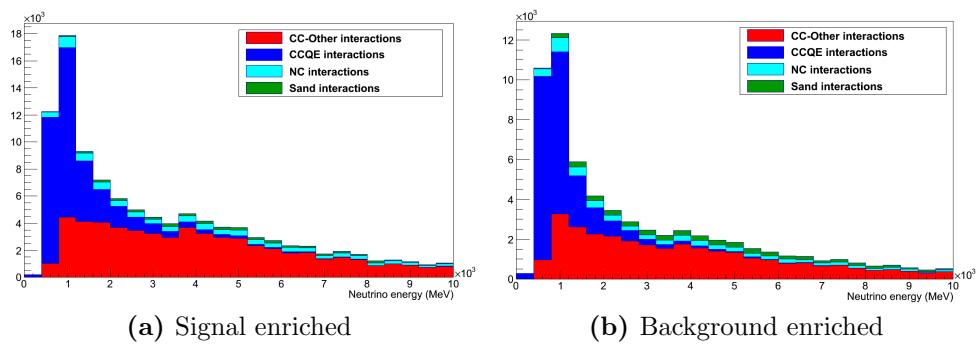


Figure B.5.: Bottom left barrel ECal module 'External' population.