

Identifying Muon Neutrino Charged-Current Interactions in the MicroBooNE Detector



Joris Jan de Vries

Department of Physics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Corpus Christi College

August 2019

Declaration

I hereby declare that this dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. In addition, it is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. Lastly, this dissertation contains fewer than 60,000 words, including summary/abstract, tables, footnotes and appendices, but excluding table of contents, photographs, diagrams, figure captions, list of figures/diagrams, list of abbreviations/acronyms, bibliography and acknowledgements, as outlined by the Degree Committee for Physics and Chemistry.

Joris Jan de Vries

August 2019

Identifying Muon Neutrino Charged-Current Interactions in the MicroBooNE Detector

Joris Jan de Vries

The MicroBooNE experiment is a short-baseline neutrino experiment at the Fermi National Accelerator Laboratory (FNAL) in the US, receiving a highly pure muon neutrino beam from the Booster Neutrino Beam (BNB) and taking data since October 2015. The main physics goal of MicroBooNE is to clarify the nature of the low-energy excess of electron-like events observed by the MiniBooNE Cherenkov detector, which due to its detector technology is unable to resolve whether the observed excess shower-like events are due to electrons or photons. Instead, MicroBooNE employs cutting-edge liquid argon time projection chamber (LArTPC) technology, which offers excellent spatial and calorimetric resolution, which makes it possible to reconstruct complex neutrino interactions and to efficiently distinguish different final state particle types. This thesis presents a fully-automated event selection of interactions of the types $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ developed using the Pandora multi-algorithm approach to pattern recognition for LArTPC experiments. The analysis in this work is performed on a subset of the MicroBooNE Run 1 dataset corresponding to 4.52×10^{19} protons on target acquired by the MicroBooNE detector between February and October 2016. The analysis quantifies the main signal and background interaction channels that produce a reconstructed neutrino interaction with one and two particles in the final state, and quantifies the extent to which different aspects of the Pandora reconstruction affect the selection performance. MicroBooNE is a surface detector, and therefore cosmic-ray tracks are the main observed background. A direction fitting procedure has been implemented in this analysis to formulate directional probabilities, using the Bethe equation and the high-resolution calorimetric information of MicroBooNE, to reduce this cosmic-ray background. This functionality is used to enhance the efficiency of the Pandora cosmic ray tagging logic, and the impact of this procedure on the reconstruction and event selection performance is analysed. The result of this study is two fully-automated selections of $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ (57.91% pure) and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ (78.30% pure) with efficiencies of 64.35% and 62.65%, respectively.

Acknowledgements

I would like to start by thanking the Cambridge Trust for generously awarding me a European Scholarship and enabling me to pursue this degree. In addition, I would particularly like to thank my parents, Thea and Alexander, and my girlfriend, Amy, without whose support this thesis would never have been completed. I will be eternally grateful for your endless encouragement: I could not have done this without you. I am also grateful to Mark Thomson and John Marshall, for their teaching and guidance. A special mention goes out to Lorena Escudero-Sanchez, who always made time to offer words of wisdom and support, and who has done endless proof-reading. Thanks also to Leigh Whitehead, for supplying advice when it was most needed. Lastly, I was fortunate to share my PhD with excellent co-students: thank you Jack, Steven, Bono, Andy and Stefano for making this experience infinitely more enjoyable.

Table of contents

List of figures	xiii
List of tables	xxvii
1 Introduction	1
2 Neutrino Physics	5
2.1 Discovery of the Neutrino	5
2.2 Neutrino Oscillations	8
2.2.1 Oscillation Experiments	11
2.3 Neutrino Masses in the Standard Model	15
2.4 Neutrino Interactions	18
2.4.1 Quasi-Elastic Interactions	18
2.4.2 Deep Inelastic Scattering	23
2.5 The LSND and MiniBooNE Anomalies	24
2.5.1 MiniBooNE	25
2.5.2 Interpretations Beyond the Standard Model	27
2.5.3 MicroBooNE	30
3 The MicroBooNE Experiment	31
3.1 Introduction	31
3.2 The Booster Neutrino Beam	33

3.3	The MicroBooNE LArTPC	35
3.4	Light Collection System	36
3.5	The Cryogenic System	38
3.5.1	Purification System and Purity Monitors	41
3.6	Signal Processing	42
3.6.1	Hit Finding	44
3.7	Simulation & Cosmic Ray Background	46
3.7.1	Space Charge Effects	47
3.8	Detector Operations	51
4	The Pandora Reconstruction Framework	53
4.1	Introduction	53
4.2	The Pandora SDK	53
4.3	Pandora at MicroBooNE	55
4.4	LArTPC Reconstruction Algorithms	57
4.4.1	Cluster and Hit Creation	59
4.4.2	Three-dimensional track reconstruction	59
4.4.3	Vertex Reconstruction	64
4.4.4	Shower reconstruction	66
4.4.5	Particle Refinement & Hierarchy	68
4.5	Performance	70
5	Direction Finding	77
5.1	The Bethe Equation	77
5.2	The Fitting Procedure	81
5.2.1	Formulating Chi Squared Values	85
5.3	Hit Filtering	90
5.3.1	Main Hit Filter	91

5.3.2	Bragg Peak Filter	94
5.4	Single Particle Hypothesis	100
5.5	Particle Splitting	103
5.6	Forwards Probability	107
5.6.1	Applications	111
6	Reconstruction Improvements	115
6.1	Vertexing Improvements	115
6.2	Cosmic Probability	121
6.2.1	Formulating Downwards and Upwards Hypotheses	121
6.2.2	Parametrising Probability Curve Fits	127
6.3	Directional Cosmic Ray Tagging	130
6.3.1	Formulating Cuts	133
6.3.2	Applying Directionality	135
6.3.3	Impact on the Reconstruction	138
7	Selecting Charged-Current Muon Neutrino Final States	141
7.1	Definitions and Pre-Selection Cuts	142
7.2	Reconstruction Studies	145
7.2.1	Characterising Reconstruction Effects	149
7.3	Multivariate Analysis for Event Selection	154
7.3.1	Topological Variables	154
7.3.2	Directional Information and Cosmic Ray Background	157
7.3.3	Calorimetric Information and Particle ID	160
7.3.4	Proton PID	163
7.4	Event Selection	166
7.4.1	Impact of Correlated Variables and Directional Information on BDT Performance	168

7.4.2	Impact of the Event Selection on Simulated Data	171
7.5	Simulated and Real Data Comparison	173
7.6	Examples of Selected Events	180
7.7	Discussion	181
8	Conclusion	185
	References	189

List of figures

2.1	The neutrino mass hierarchy possibilities. Figure taken from [30].	11
2.2	$P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ as a function of distance for $E_{\bar{\nu}_e} = 3$ MeV, assuming that $\theta_{12} = 30^\circ$, $\theta_{23} = 45^\circ$, $\theta_{13} = 10^\circ$, $\Delta m_{12}^2 = 8 \times 10^{-5} \text{eV}^2$ and $\Delta m_{32}^2 = 2.5 \times 10^{-3} \text{eV}^2$. Figure from [30].	12
2.3	Ratio of the background and geo-neutrino-subtracted $\bar{\nu}_e$ spectrum to the expectation based on oscillation parameters as a function of $L_0/E_{\bar{\nu}}$. L_0 is the effective baseline taken as a flux-weighted average ($L_0 \approx 180$ km). Figure from [37].	14
2.4	The (anti)neutrino cross section per unit of energy σ/E_ν , as a function of E_ν . Contributions for quasi-elastic scattering (dashed), resonance production (dot-dashed) and deep inelastic scattering (dotted) are indicated. Cross section predictions as given by the NUANCE generator [46] are indicated by the solid line. Quasi-elastic scattering data and predictions have been averaged over neutron and proton targets and hence have been divided by a factor of two for the purposes of this plot. [47]	19
2.5	Feynman diagrams that illustrate charged-current (CC) quasi-elastic scattering (2.5a) and neutral-current (NC) quasi-elastic scattering (2.5b) processes.	21
2.6	Hadrons produced by the neutrino interactions can subsequently interact with the surrounding nucleons. This image gives a schematic overview of different pion FSI modes. Figure from [57].	23
2.7	Feynman diagrams that illustrate CC resonant (2.7a) and coherent pion (2.7b) production processes.	23
2.8	Feynman diagram for the deep inelastic process $\nu_\mu + p \rightarrow \mu^- + X$, where X denotes the hadronic shower object.	25

- 2.9 The LSND excess, indicating possible $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance is shown in 2.9a. The red and green histograms indicate the background from intrinsic $\bar{\nu}_e$ in the beam and from other sources, respectively. The blue histogram is an example of a potential $\bar{\nu}_\mu \rightarrow \nu_e$ signal that, when added to the predicted backgrounds, is representative of the data. The allowed and excluded regions in the $(\sin^2 2\theta, \Delta m^2)$ parameter space are shown in 2.9b. The inner and outer regions correspond to 90% and 99% CL allowed regions, while the curves are 90% CL limits from other neutrino experiments. Figures from [58]. 26
- 2.10 MiniBooNE results, showing the neutrino E_ν^{QE} distributions for data gathered from 12.84×10^{20} protons-on-target (2.10a). Data points with statistical errors are shown, as well as stacked distributions of the background contributions, with systematic errors. Figure 2.10b shows the appearance probability at MiniBooNE and LSND as a function of L/E , with good agreement. Error bars show statistical uncertainties only, and the shaded band is the MiniBooNE allowed region, assuming two-neutrino oscillations. 28
- 2.11 MiniBooNE allowed regions for a combined neutrino mode (12.84×10^{20} POT) and antineutrino mode (11.27×10^{20}) data sets for events with $200 \leq E_\nu^{QE} \leq 3000$ MeV given a two-neutrino oscillation model. The shaded areas indicate the 90% and 99% C.L. LSND allowed regions. Also shown are 90% C.L. limits from the KARMEN [60] and OPERA [61] experiments. The black point indicates the MiniBooNE best fit at $(\sin^2 2\theta, \Delta m^2) = (0.96, 0.041 \text{ eV}^2)$. Figure from [59]. 29
- 2.12 The neutrino mass hierarchy with the addition of a sterile neutrino flavour. . 30
- 3.1 Aerial diagram showing location of MicroBooNE along the BNB beam line (the orange dashed line) at Fermilab. Other detectors are also shown [62]. . 33
- 3.2 A schematic overview of the Fermilab accelerator complex. Figure from [71]. 34
- 3.3 Schematic of different components that make up the BNB beamline. 34
- 3.4 Neutrino flux prediction at MicroBooNE. Figure from [73]. 35
- 3.5 The coordinate system convention employed by the MicroBooNE collaboration. This figure schematically depicts the physical boundaries of the TPC, the coordinate system axes and the three wire planes on the right. Figure from [74]. 36

3.6	Diagram of the photomultiplier optical unit (left) and an image of the optical units mounted in MicroBooNE, immediately prior to LArTPC installation (right). Image from [62].	38
3.7	The MicroBooNE light collection system consists of a primary system of 32 optical units and a secondary optical system of four lightguide paddles, which are present for R&D purposes. These are mounted behind the anode wire planes such that the view is not obscured by the structural cross bars of the LArTPC [62].	39
3.8	A rendering showing the MicroBooNE cryostat and cryogenic system, and the platform for the electronics racks as installed in the LAr Test Facility (LArTF) [62].	40
3.9	MicroBooNE LArTPC and PMT signal processing and readout stages [62].	42
3.10	Example induction pulses on a W plane wire for a simulated event, with and without Gaussian fits. The x axis corresponds to units of time, and the y axis has units of ADC (analog to digital counts). Image generated using LARSOFT [7, 8].	45
3.11	A neutrino event in the MicroBooNE detector, recorded on October 23rd, 2015 (part of the Run 1 dataset). The direction of the beam line runs along the horizontal axis, and the neutrino interaction vertex (manually labelled <i>V</i>) is clearly present as the point of confluence of the final state particles. Several tracks are present in the event, as well as an electromagnetic shower (labelled <i>T</i> and <i>S</i> , respectively). In the bottom left corner and on the right-hand side, three cosmic ray tracks are also visible (labelled <i>CR</i>). Figure modified from [76].	45
3.12	Illustration of the two types of cosmic-ray background events. Figure 3.12a is an event in which no neutrino interaction is present, but PMT trigger is generated by a cosmic-ray track. In Figure 3.12b a neutrino interaction induces a PMT trigger, but a cosmic ray track may still be selected as the most likely neutrino candidate. The blue outline is a schematic representation of the cryostat. Figure from [81].	47
3.13	Rendering showing the LArTPC muon tagger panel groups (topside, side, and underside) and a simulation of cosmic rays crossing the CRT, the brown lines represent possible cosmic ray trajectories. This design achieves a coverage of 85% of cosmic ray muons passing the TPC [79].	48

3.14	The reconstructed start and end positions of cosmic ray tracks that were tagged by an external muon counter in the x - y plane. In the absence of space charge effects, these points would be distributed along the detector boundaries, which are indicated by dashed lines. The anode is located at $x = 0$ cm and the cathode is located at $x = 256$ cm. Figure from [82].	49
3.15	Illustration of the simulated effects of space charge on the distortions in reconstructed track positions in the MicroBooNE TPC in the x - y plane. Corrections are shown for the x , y and z coordinates for a slice of the detector at $z = 10$ cm. The colour scale in each plot indicates the reconstructed minus the true coordinate $\text{pos}_{\text{reco}} - \text{pos}_{\text{true}}$ in centimeters.	50
3.16	Delivered and recorded protons-on-target from the Booster Neutrino Beam since the start of the MicroBooNE data taking period. Weekly fluxes are shown in blue (with units on the left axis) and the integrated delivered flux is shown in red (with units on the right axis).	51
3.17	The variation of Q_A/Q_C for a period of 56 days during the data taking period (16 February 2016 to 21 April 2016). The black filled circles indicate points with space charge corrections applied and blue triangles indicate points without space charge corrections. After correction, all points have values ≤ 1 within statistical errors.	52
4.1	The hierarchy of the Pandora output data products, which are persisted in the LARSOF Event Data Model. Navigation along PFO hierarchies is indicated by dashed lines. Navigation from each object to their associated objects is indicated by solid arrows. Figure adapted from [85].	56
4.2	An overview of the two Pandora reconstruction chains in MicroBooNE. The neutrino ID and slicing stages are not depicted in this figure. Inputs and outputs are indicated by orange borders, reconstruction stages with black borders, and the cosmic ray tagging logic by a dashed border. Figure adapted from [85].	58

- 4.3 A comparison of the high-purity clusters created by the initial stage of the track cluster creation logic (the `TrackClusterCreation` algorithm), which creates a new cluster at every ambiguous point (Figure 4.3a) and the final output merged clusters, which have both high purity and high completeness (Figure 4.3b). Each pair of horizontal grey lines corresponds to an inactive detector region, and each coloured track corresponds to a separately reconstructed cluster of hits. The x coordinate is in the horizontal direction, the z coordinate in the vertical direction. Modified from [85]. 60
- 4.4 Example topologies that are correctly resolved into unambiguous groupings of 2D clusters by the 3D track reconstruction. Four matching situations are shown, containing: (a) one cluster per view, (b) two small δ -ray clusters (circled), (c) a projection of two particles that appears to be one cluster in the U view and (d) a cluster erroneously split in the V view. The x coordinate is in the horizontal direction, the z coordinate in the vertical direction. Modified from [85]. 62
- 4.5 An example W view projection of an event containing three particles, each indicated by a separate colour, in which five vertex candidates are shown. The table in the top right of the figure shows the S , $S_{energykick}$, $S_{asymmetry}$, $S_{beamdeweight}$ scores for each of the candidates. Candidate E is eventually selected, since it has the highest overall score S . Details of the vertex scoring procedure are given in the text. Figure from [85]. 67
- 4.6 An example projection in the W view of an electromagnetic shower. The red track-like clusters are not affected by the shower reconstruction. The long blue clusters that point towards the interaction vertex are identified as shower spines, after which the green, smaller clusters are added onto the spines iteratively. This procedure yields a shower-like cluster of high completeness. Figure from [85]. 69
- 4.7 An example particle hierarchy in the W view associated to a CC ν_μ event with a muon, proton and pion in the final state. The neutrino particle has three track-like daughter particles. The first is a long proton (blue), which recoils off a nucleus which results in a track-like secondary daughter particle (grey). The second is a π^+ (magenta), which decays into a μ^+ , which rapidly decays into a visible e^+ , which is reconstructed as a shower-like particle (purple). The third neutrino daughter is a muon track (red). Figure from [85]. 71

4.8	Example reconstruction of a simulated 500-MeV CC ν_μ quasi-elastic event. The event final state is $\mu^- + p$. Simulated unresponsive channels are indicated by the horizontal grey lines, which cause the break in the proton track. Figure from [85].	74
5.1	$\langle -dE/dx \rangle$ for μ^+ traversing copper, as a function of $\beta\gamma = p/Mc$. The corresponding momentum scale is also indicated. The two rightmost blue bands indicate the Bethe regime, in which the Bethe formula is an excellent approximation of $\langle \frac{-dE}{dx} \rangle$, which is the region with which this discussion concerns itself [88].	79
5.2	Mean energy loss rate in liquid (bubble chamber) hydrogen, gaseous helium, carbon, aluminium, iron, tin, and lead. Radiative effects, relevant for muons and pions, are not included. These become significant for muons in iron for $\beta\gamma \geq 1000$, and at lower momenta for muons in higher-Z absorbers [88]. . .	80
5.3	Simulated particle kinetic energy T (5.3a) and $\langle -dE/dx \rangle$ (5.3b) as a function of CSDA range R for a muon in LAr with $T_i = 700$ MeV. This plot was generated by iteratively sampling the particle energy as given by the Bethe equation, with a step size dR of 0.1 cm. For low particle energies the Bragg peak becomes apparent as a faster energy loss.	82
5.4	The sequential mapping procedure. A hypothesis for the kinetic energy at the end of a track T_e is mapped to a range value R_i using the mapping \mathcal{M} . A hypothesis for the total range ΔR maps R_e to R_i , which in turn is mapped to T_i . Any value $R_i \leq R_k \leq R_e$ now also has an associated value T_k through the mapping \mathcal{M}	83
5.5	A direction fit to a simulated forwards-going contained muon. The values of \tilde{Q} are indicated by black points, and the fit values \tilde{Q}_{fit} resulting from the forwards direction fit are indicated by magenta points. While the values of \tilde{Q} are affected by secondary interactions such as δ -rays, the values of \tilde{Q}_{fit} are more robust to such effects.	86
5.6	The distribution of $Q_{fit} \equiv \tilde{Q}_{fit} \cdot w$ using a fine binning of 0.03 MeV (5.6a) and a coarse binning of 0.1 MeV (5.6b). The binning used in Figure 5.6b is used when plotting the distributions of the residuals Δ	87

- 5.7 Distributions of the residuals Δ for four equidistant bins in Q_{fit} . The inner 70, 80 and 90% of the distributions are shaded red, beige and blue, respectively. The Gaussian fit to the inner 70% (red) area is shown by the blue curve. The red curve is the extrapolation of the blue curve outside of the 70% region. 88
- 5.8 Three fits to the values of Δ_{RMSE} (5.8a) and σ_{Gauss} (5.8b) resulting from the inner 70% distributions of Figure 5.7. The red fit uses a functional form $a \cdot \sqrt{Q_{fit}}$, the blue fit uses $a \cdot Q_{fit}$ and the magenta fit uses the superimposed form $\sqrt{a \cdot Q_{fit}^2 + b \cdot Q_{fit}}$ 89
- 5.9 The hits deposited by a 779 MeV muon, before applying the filtering procedure (5.9a) and after applying the filtering procedure, with the best direction fit indicated by magenta points (5.9b). 92
- 5.10 The distributions of the fractional position p_f for pure and impure hits. A cut is applied where the distributions of p_f intersect, at $p_f = 0.72$. The region retained by the cut is shaded blue. 93
- 5.11 Plots indicating the performance of the filtering method. The definitions of purity P and completeness C are given in the text. The long tails in these plots are due to events in which many hits are split into multiple fragments, resulting in very low \tilde{Q} values. These fragments are removed and can constitute a significant fraction of the total number of hits. 94
- 5.12 The effect of the Bragg peak filter on an example muon track in the W view. Figure 5.12a shows a muon track for which \tilde{Q} values are inflated near the interaction vertex due to overlapping hit pulses, the inclusion of proton hits. A split hit in the Bragg peak is also present. Figure 5.12b shows the effect of applying the Bragg peak filter on the distribution of \tilde{Q} . Impure hits are coloured red, and are hits for which $\geq 15\%$ of the hit charge has been contributed by a secondary primary Monte Carlo particle. 96
- 5.13 Distributions of $\max(\frac{\tilde{Q}_i}{\tilde{Q}_{i\pm 3}})$ for Bragg peak and non-Bragg peak hits. A cut is applied at $\max(\frac{\tilde{Q}_i}{\tilde{Q}_{i\pm 3}}) = 1.725$, which retains 94.80% of Bragg peak entries and 89.28% of non-Bragg peak entries. 97
- 5.14 A visualisation of the change in χ_{min}^2/N due to applying the track end filter for true forwards going muons (triangles) and backwards going muons (squares). A distinction is made between correct removals, for which the direction hypothesis changes for the better (green), and bad removals (red). 98

5.15	Figures 5.15a and 5.15b show scatter plots of good and bad Bragg peak filter entries for tracks that satisfy $50 \leq N < 100$ and $150 \leq N < 200$, respectively. Figures 5.15c and 5.15d show the corresponding purity P_T and completeness C_T values for 100 different cuts on T in the range $0 \leq T \leq 10$	99
5.16	A definition of optimal cut values T on the observed value of $\delta\chi_{BA}^2/N \equiv \chi_{B,min}^2/N - \chi_{A,min}^2/N$ is given by the exponential fit in 5.16a and the effect of applying these cut values to the scatter plot of Figure 5.14 is given in 5.16b.	101
5.17	An overview of the steps in the direction fitting procedure. Information flow and algorithm type have been coloured according to whether they constitute information resulting from the direction fitting procedure, or otherwise. The word ‘cut’ indicates decision logic associated to the single and merged particle hypotheses, as described in the text. A legend is given in the bottom left of the figure.	101
5.18	Distribution of χ^2/N for true forwards and backwards muons are given in 5.18a and 5.18b. Colour-coded binned scatter plots of χ^2/N with N are given in 5.18c and 5.18d. Based on these distributions, a good fit has been determined to be a fit for which $\chi^2/N \leq 1.5$	102
5.19	An example single cluster in the W view that contains hits from both a muon and a proton is given in 5.19a. The result of applying the splitting procedure described in this section on this hit collection is given in 5.19b. The best fits associated to the two separate particle hypotheses are given as magenta points.	104
5.20	Distributions of the magnitude of the best jump J_B for good and bad particle splits are drawn in good and red, respectively.	106
5.21	Scatter plots for correct and incorrect splits for the entire event sample is given in 5.21a and the same scatter plot after applying the cut values defined by the exponential fit of Figure 5.22.	106
5.22	Splitting cut values for bins in N with exponential fit of the form $T = a_2 + \exp(b_2 + c_2 \lfloor N/50 \rfloor)$	107
5.23	Several examples of successfully split clusters, where the proton and the muon that make up the original cluster are clearly visible to have distinct \tilde{Q} values.	108

- 5.24 Distribution of $\Delta\chi_{FB}^2/N$ for true contained forwards muons (5.24a) and true contained backwards muons (5.24b). The same distributions with extended tails are shown in 5.24c and 5.24d, obtained by setting tail entries to the average value in the tail region (as described in the text). The distributions with the extended tails can be used to calculate a bin-by-bin forwards probability P_F using equation 5.24. 110
- 5.25 Two examples of true backwards events that have been incorrectly reconstructed. 111
- 5.26 Probability curves showing P_F as a function of $\Delta\chi_{FB}^2/N$. The probability P_F has been calculated using equation 5.24 and the distributions of Figures 5.24c and 5.24d. A comparison is made between the values of P_F with both the splitting and the Bragg peak filter functionality disabled (red), with only the Bragg peak filter enabled (blue) and with the Bragg peak filter and particle splitting enabled (green). 112
- 6.1 An illustration showing how crossing vertex candidates are constructed. Sliding fits to tracks (black lines) are extrapolated a fixed amount (indicated by the dashed blue lines), and intersections of these extrapolations (indicated by red points) are matched between views to create new vertex candidates. . . 117
- 6.2 The impact of introducing crossing vertex candidates on the vertex ΔR distributions in the CCQEL μ^- and CCQEL $\mu^- + p$ channels. The distribution due to the current reconstruction is indicated in red, and the distribution after introducing crossing candidates is indicated in blue. The impact of these candidates on the CCQEL μ^- channel is minimal because the vast majority of events in this category are reconstructed as single tracks. In contrast, CCQEL $\mu^- + p$ channel typically have a multiplicity of 2, and so crossing candidates have a large impact. 117
- 6.3 The impact of introducing crossing vertex candidates on the vertex ΔR distributions in the CCQEL μ^- and CCQEL $\mu^- + p$ channels. The distribution due to the current reconstruction is indicated in red, and the distribution after introducing crossing candidates is indicated in blue. The impact of these candidates on the CCQEL μ^- is largest, because the absence of secondary particles introduces information sparsity, such that the relative impact of introducing additional information in the form of new candidates is largest. . . 119

- 6.4 Example event in which z candidates create an improvement. The event contains one track in each view. The track curvature affects sliding fit interpolation, affecting the standard vertex candidate creation logic. Inactive detector regions have obscured the low- z ends of the tracks. Note the different scales in z 119
- 6.5 The effect of turning on z candidates in an example event. Vertex candidates that pass the on hit/in gap quality cut are displayed in red. After turning on the z candidates, new vertex candidates are produced at the low- z track endpoint, near the true vertex position. 120
- 6.6 Distributions of $\Delta\chi_{DU}^2/N$ (6.6a) and χ_{min}^2/N (6.6b) for simulated cosmic ray muons. These plots demonstrate that $\Delta\chi_{DU}^2/N$ entries are predominantly negative, indicating downwards-going particles and that χ_{min}^2/N peaks near 1, indicating high-quality fits. The peak near $\Delta\chi_{DU}^2/N \approx 0$ is predominantly caused tracks for which the Bragg peak is not apparent, due to their orientation in the detector. These figures were created using an older version of the simulation, but they are representative. 123
- 6.7 Distributions of $\Delta\chi_{DU}^2/N$ for true cosmic ray muons. Figure 6.7a shows the effect of demanding that the reconstructed low- y endpoint is fiducial. Figure 6.7b shows the effect of applying the Bragg peak finder (BPF) to the distribution in 6.7a, which demands that the direction fit parameter $T_{end} \leq 25$ MeV. 124
- 6.8 A distribution of P_c , as a function of $-\Delta\chi_{DU}^2/N$. Point errors are binomial. The points on this probability curve are constructed using the prescription in equation 6.2 on the distribution of Figure 6.6a. The fit function in 6.8b is given in equation 6.3. 126
- 6.9 The variation of the downwards probability P_c with N , χ_{min}^2/N , ϕ and θ . . . 127
- 6.10 The variation of P_c with N while keeping χ_{min}^2/N fixed. These plots show that the variation of P_c with N is not solely due to the correlation between N and χ_{min}^2/N . The choice of binning for χ_{min}^2/N corresponds to the choice of binning for χ_{min}^2/N in Figure 6.9. The last bin $4.0 \leq \chi_{min}^2/N \leq \infty$ is excluded, because it has very low statistics. 128
- 6.11 The fit values for the parameters P_{max} , α and β for $75 \leq N \leq 225$ in bins of 75 and $0.25 \leq \chi_{min}^2/N \leq 2.25$ in bins of 0.5. 129

6.12	The fit values for the parameters P_{max} and α as shown in Figure 6.11, but with a fitted 2D surface of the form $p_n = a_n + b_n N + c_n \chi_{min}^2/N$	130
6.13	Underlying P_c distributions taken from the range $0 \leq N \leq 400$ in bins of 100 and 200 for the last bin, to increase the underlying number of events, and $0 \leq \chi_{min}^2/N \leq 4$ in bins of 1 and 2 for the last bin.	131
6.14	Distributions of y_{max} for neutrino-induced and cosmic-ray muons.	132
6.15	Distributions of θ , Q_{PFO} and $p_{v,z}$ for neutrino-induced and cosmic-ray muons in the $y_{max} \geq 100$ cm, $L \geq 25$ cm untagged sample.	134
6.16	A comparison of the normalised distributions of $\Delta\chi_{DU}^2$ and $\Delta\chi_{DU}^2/N$ for cosmic rays, $\nu \downarrow$ and $\nu \uparrow$. The event sample is the untagged sample with $y_{max} \geq 100$ cm and all cuts in 6.2 applied in order. A Bragg peak finder has been applied, as described in the text.	137
6.17	The curves in this figure quantify what fraction of cosmic muon tracks are removed and what fraction of $\nu \uparrow$ and $\nu \downarrow$ are removed, as a function of cut values applied to directional variables. Figure 6.17a compares cuts on $\Delta\chi_{DU}^2$ and $\Delta\chi_{DU}^2/N$, and Figure 6.17b compares cuts on P_c and $\Delta\chi_{DU}^2/N$	138
7.1	The fiducial volume, which avoids regions with a significant number of unresponsive channels. Taken from [96].	145
7.2	Plots that characterise the NO RECONSTRUCTABLE interaction type. Figure 7.2a indicates that events of this type are not just low energy, and Figure 7.2b demonstrates that such events contain few true neutrino-induced hits, of which only a subset are incorporated into PFOs. This is due to events containing neutrons that induce sparse clouds of hits, as discussed in the text.	149
7.3	An example event in which a neutron induces a diffuse cloud of hits, leading to the event being classified as a ν NO RECONSTRUCTABLE event. The cosmic ray muon indicated in blue is selected as ν_{reco}	150
7.4	Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$	156
7.5	Stacked plots of the distributions of the PFO y_{max} and z_{max} variables.	156
7.6	Stacked plots of the distributions of the $N=2$ shortest PFO summed charge and length variables.	157

7.7	Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$	158
7.8	Stacked distributions for the $N=1$ $\Delta\chi_{FB}^2/N$ and $\Delta\chi_{DU}^2/N$ variables.	160
7.9	An example CCQEL μ^- event, in which the muon observed charge deposits \tilde{Q}_{obs} and the magenta fitted charges \tilde{Q}_{fit} are shown. The empty region for $15 \lesssim L_{2D} \lesssim 50$ cm is due to an inactive detector region in the collection plane. The spike in \tilde{Q}_{obs} near $L_{2D} = 0$ cm is due to overlapping hits near the interaction vertex.	162
7.10	Stacked plots the two most important simple PID variables, both for the shortest PFO in $N = 2$ events.	162
7.11	The charge calibration values for the x coordinate (7.11a) and the yz plane (7.11c). A fit to the distribution of Figure 7.11a of the form $y = a \cdot x + b \cdot \sin(x)$ is shown in 7.11b. Images 7.11a and 7.11c from [97].	165
7.12	The stacked distribution of χ_p^2 for $N = 2$. The OTHER channel collection has been added to the stack to demonstrate the separating power of this variable when distinguishing between signal and background.	167
7.13	ROC curves for several methods for the $N=1$ and $N=2$ contained event sample.	168
7.14	Correlation matrices showing the correlation coefficients ρ (defined in equation 7.6) for the complete set of $N = 2$ variables (7.14a) and the $N = 2$ variables after removing highly-correlated variables (7.14b). The correlation coefficients are calculated for $N = 2$ signal events. The variable $\langle Q_{fit} \rangle$ is indicated as Q_{av} in 7.14a.	169
7.15	The data-driven pre-selection cuts, excluding $Q_L < 100$ MeV (7.15a) for $N = 1$ and retaining $0.1 \leq \zeta \leq 3.04$ rad for $N = 2$	174
7.16	Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$ with data points overlaid before applying the event selection. A data-driven pre-selection cut excluding $Q_L < 100$ MeV is applied in these distributions.	176
7.17	Stacked distributions for the $N=1$ $\Delta\chi_{FB}^2/N$ and $\Delta\chi_{DU}^2/N$ variables with data points overlaid before applying the event selection. A data-driven pre-selection cut excluding $Q_L < 100$ MeV is applied in these distributions.	177

- 7.18 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=2$ with data points overlaid before applying the event selection. A data-driven pre-selection cut retaining $0.1 \leq \zeta \leq 3.04$ rad is applied in these distributions. 177
- 7.19 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$ after applying the event selection, with data points overlaid. 178
- 7.20 Stacked distributions for the $N=1$ $\Delta\chi_{FB}^2/N$ and $\Delta\chi_{DU}^2/N$ variables with data points overlaid, after applying the event selection. 178
- 7.21 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=2$ after applying the event selection, with data points overlaid. 179
- 7.22 A short $N = 1$ track chosen by the event selection (7.22a). Cosmic ray hierarchies are indicated in red, and v_{reco} is indicated in blue. Figures 7.22b, 7.22c and 7.22d show the projections of v_{reco} in the U, V and W views. . . 182
- 7.23 A long $N = 1$ track chosen by the event selection (7.23a). Cosmic ray hierarchies are indicated in red, and v_{reco} is indicated in blue. Figures 7.23b, 7.23c and 7.23d show the projections of v_{reco} in the U, V and W views. . . 183
- 7.24 An $N = 2$ interaction chosen by the event selection (7.24a). Cosmic ray hierarchies are indicated in red, and v_{reco} is indicated in blue. Figures 7.24b, 7.24c and 7.24d show the projections of v_{reco} in the U, V and W views. . . 184

List of tables

3.1	MicroBooNE LArTPC design parameters and nominal operating conditions. Adapted from [62].	37
4.1	Reconstruction performance for charged-current channels resulting in μ^- and $\mu^- + p$ final states. The table displays the number of reconstructed particles matched to each primary Monte Carlo particle per interaction channel. A channel is only deemed correct if all primary Monte Carlo particles have the correct number of matches. When a cosmic-ray particle has been matched to a target Monte Carlo particle, this ends up in the ‘0 Matches’ category. . . .	75
6.1	The impact of the the new vertex candidates on the Pandora reconstruction performance. The indicated fractions are the fraction of correct events, using the Pandora correctness criterion, and the Total column shows the total number of events per interaction channel.	120
6.2	Cut values retaining as close to 90% of the cosmic ray entries as possible. These cut values are obtained by applying each cut sequentially to the untagged $y_{max} \geq 100$ cm $L \geq 25$ cm sample, such that each cut affects the subsequent variable distribution.	135
6.3	The number of cosmic ray, ν_{\uparrow} and ν_{\downarrow} primary particles left after each cut is applied.	136
6.4	$\Delta\chi^2_{DU}/N$ cut values that retain as close to 98% of ν_{\uparrow} as possible, which corresponds to removing at most 2700 ν_{\uparrow} . The number of cuts applied is indicated, as well as the number of additional cosmic ray particles, actual ν_{\uparrow} and ν_{\downarrow} removed. The $\Delta\chi^2_{DU}/N$ is that retains closest to the target ν_{\uparrow} is found by scanning with a resolution of 0.01. The cut $L \geq 25$ cm is applied everywhere in this table.	138

6.5	Various $\Delta\chi^2_{DU}/N$ corresponding to certain fractions of ν_\uparrow to be retained. These fractions refer to the fractions of ν_\uparrow remaining after applying the N and y_{max} cuts, so out of 30514 total ν_\uparrow . The impacts of these cuts on the reconstruction performance are compared in Table 6.6.	139
6.6	The impact of the different $\Delta\chi^2_{DU}/N$ requirements on the Pandora reconstruction performance. The columns are arranged by the strictness of the cut (left to right, left being no cut). The indicated fractions are the fraction of correct events, using the Pandora correctness criterion.	140
7.1	All target channels: these are the interaction types of the form $\nu_\mu \rightarrow \mu^- + X$ and $\nu_\mu \rightarrow \mu^- + p + X$, which result in the target $1\mu^-$ and $1\mu^-1p$ final states.	144
7.2	A breakdown of interaction types that result in reconstructed primary particle multiplicities of $N=1$ and $N=2$ in a simulated sample with neutrino interaction generated by GENIE and cosmic rays simulated by CORSIKA. The vertex containment cut, described in Section 7.1, has been applied in these tables. Furthermore, a directional cosmic ray tagging (DCRT) cut has been applied, details of which are outlined in the text. Blank entries have been introduced in the $N=1$ table for row consistency between tables. The event sample contains 520448 events in total.	147
7.3	Comparison of the $N=1$ interaction type tables when cheating several aspects of the reconstruction. From left to right: cheating the ν ID, cheating the ν ID and the cosmic ray tagging, and cheating the ν ID, the cosmic ray tagging and the slicing.	152
7.4	Comparison of the $N=2$ interaction type tables when cheating several aspects of the reconstruction. From left to right: cheating the ν ID, cheating the ν ID and the cosmic ray tagging, and cheating the ν ID, the cosmic ray tagging and the slicing.	153
7.5	A comparison of the values of f_{ROC} when using all variables, the reduced variable set and the reduced variable set without directional information (labelled ‘ND’ for ‘no direction’). The reduced variable set consists of the variables that remain after removing highly-correlated variables and variables with little discriminative power.	172
7.6	The impact of applying the event selection on the $N = 1$ interaction type tables.	172
7.7	The impact of applying the event selection on the $N = 2$ interaction type tables.	173

- 7.8 Event counts after applying the event selection for $N = 1$ and $N = 2$. Event counts are normalised by protons-on-target and trigger counts, as outlined in the text. The ‘Before Cuts’ column indicates the number of events that pass the selection before applying the data-driven pre-selection cuts on \tilde{Q}_L and ζ . The ‘After Cuts’ column indicates event counts with these cuts applied. Within the EXTBNB + MC category, a distinction is made between i) signal events (SIGNAL) ii) neutrino-induced background events (OTHER ν) and iii) all other background (OTHER BACKGROUND). Indented rows in this table indicate subsets, such that e.g. SIGNAL is the subset of EXTBNB + MC that contains signal events. Percentages indicate the fractions of events remaining in the normalised sample after applying the event selection. 180

Chapter 1

Introduction

Neutrino physics has been an active area of research since the 1950s, although the existence of the neutrino was postulated as early as 1930 by Wolfgang Pauli. Investigations into the theory of neutrinos have resulted in a number of surprising experimental results, including the deficit of solar neutrinos measured in the 1960s. This discovery was eventually explained by the theory of neutrino oscillations, the experimental confirmation of which was decorated with the 2015 Nobel prize in physics. Although the resulting incorporation of small neutrino masses into the Standard Model (SM) of particle physics allows for the explanation of this process, the mechanism by which neutrinos obtain their small masses is not understood. In addition, neutrino experiments have made observations that are not entirely in line with the Standard Model description of three light flavours of neutrinos. The Liquid Scintillator Neutrino Detector (LSND) observed an excess of $\bar{\nu}_e$ (electron-like) events in a $\bar{\nu}_\mu$ beam [1] which was later observed by the MiniBooNE experiment, first in the 200-475 MeV region [2] and recently extended to 1250 MeV, using neutrino and antineutrino data [3]. These discoveries have sparked a vigorous interest in possible theoretical explanations, including the postulation of a fourth sterile neutrino flavour, which would not be weakly-interacting but would enhance the $\nu_\mu \rightarrow \nu_e$ transition probability and hence the rate of charged-current quasi-elastic ν_e interactions in MiniBooNE [4]. Alternative explanations include the mis-estimation of single photon backgrounds in MiniBooNE [5], or an entirely new photon production process [6]. A summary of relevant neutrino physics can be found in Chapter 2.

The MiniBooNE low energy excess (LEE) is tantalising but inconclusive, since MiniBooNE measured charged-current quasi-elastic ν_e appearance, but as a Cherenkov detector is incapable of distinguishing between photons and electrons. To conclusively resolve the observed LEE, the MicroBooNE experiment was commissioned, which uses liquid argon

time projection chamber (LArTPC) technology with excellent spatial and calorimetric resolution that will be able to distinguish photons and electrons by identifying the initial track-like segment in electron-induced showers, as well as the distance between the neutrino interaction vertex and photon starting point, which is different for photons due to their conversion length. The MicroBooNE experiment also uses the same neutrino beam as MiniBooNE, the Booster Neutrino Beam (BNB) at the Fermi National Accelerator Laboratory (Fermilab). MicroBooNE will be able to perform a suite of high-precision cross section measurements and is capable of detecting neutrinos induced by supernovae. Furthermore, MicroBooNE is a crucial step in research and development (R&D) programme, from which future liquid argon neutrino experiments will benefit. MicroBooNE has had a very successful data taking period, taking its first neutrino beam data in October 2015, and having exceeded the projected delivered protons-on-target (POT). Chapter 3 will outline the physics goals of the MicroBooNE experiment in more detail, and will describe the operating principles behind LArTPC detectors, and will provide detail about detector subsystems and the Fermilab Booster Neutrino Beam (BNB).

The large quantity of high-resolution data produced at LArTPC detectors requires novel techniques in order to fully exploit the quality of the images in a fully-automated reconstruction. The reconstruction of different final state particles in the detector, a mixture of track-like and shower-like objects, due to a variety of neutrino interaction types, is a challenge for pattern recognition logic, unlikely to be overcome by a single clustering algorithm. Instead, the Pandora software development kit (SDK) is used, which provides a general multi-algorithm framework within which pattern recognition problems can be formulated. Chapter 4 gives an overview of the Pandora SDK, as well as the many bespoke algorithms that have been developed in Pandora for use in the reconstruction of MicroBooNE events. Furthermore, the high-resolution calorimetric information that can be resolved by LArTPC detectors can be exploited in reconstruction algorithms, and Chapter 5 outlines a novel tool that uses predicted values of $-\langle dE/dx \rangle$ from the Bethe equation to create a direction fit. The addition of a directional probability has direct application within an automated reconstruction process, and can help resolve difficult neutrino interactions, such as those where two tracks appear back-to-back in the detector.

In addition, directional information can be used to reduce the cosmic-ray background inherent in any surface detector such as MicroBooNE. Such background events are difficult to resolve, since neutrinos interact only weakly with matter and only few beam spills result in a neutrino interaction, despite the use of high-intensity neutrino beams. In contrast, cosmic ray particles are ubiquitous. Many steps have been taken by the MicroBooNE collaboration to reduce this cosmic-ray background, including the installation of a cosmic ray tagger (CRT)

system and leveraging the scintillation light produced in the detector. Chapter 6 will outline how directional information can be used to formulate a probability that a particle is travelling in the downwards direction, which can be utilised as a novel approach to cosmic ray removal. It will be shown in this Chapter that incorporating directional information can enhance the existing cosmic ray tagging logic within Pandora, and that doing so has a significant beneficial impact on the reconstruction performance.

MicroBooNE will perform a suite of high-precision cross section measurements, which inherently rely on the accurate identification of particular neutrino interaction final states. This thesis will develop an event selection for events of the forms $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ using the Pandora reconstruction. This event selection incorporates the reconstruction improvements and directional tools developed in the rest of this work. A study will quantify the contributions of the dominant signal and background channels in these selections, and will investigate the extent to which different aspects of the reconstruction affect the selection performance. The model employed to perform the event selection is a Boosted Decision Tree (BDT), which is trained on a number of bespoke topological, calorimetric and directional variables. It will be shown that the event selection is highly effective and results in a pure selection. Furthermore, it will be shown that the addition of directional information has a significant impact on the performance of the event selection. The event selection is developed on a simulated dataset produced using LARSOF [7, 8] that corresponds to 5.37×10^{20} protons-on-target. The event selection is then applied to a Run 1 dataset gathered by the MicroBooNE detector between February and October 2016 corresponding to 4.52×10^{19} protons-on-target. Finally, a comparison of the distributions of input variables before and after the selection and selection studies are presented in Chapter 7 and conclusions are given in Chapter 8.

Chapter 2

Neutrino Physics

This chapter will give an overview of neutrino physics and starts by giving an overview of the history behind the discovery of the neutrino in Section 2.1. Section 2.2 describes the phenomenon of neutrino oscillations, the formalism used to describe oscillation probabilities, and important physics experiments that have investigated this aspect of neutrino physics. A brief discussion is included in Section 2.3 regarding how neutrinos could get their light masses, and the seesaw mechanism is introduced. Subsequently, Section 2.4 discusses how neutrinos interact with matter in the neutrino energy regime in which MicroBooNE operates, outlining the dominant interaction types and giving an overview of historical neutrino cross section measurements in this energy regime. Lastly, Section 2.5 discusses the low-energy excess of $\bar{\nu}_e$ events discovered by the LSND experiment and also observed by the MiniBooNE experiment. Possible interpretations of this experimental result are given in the context of physics beyond the Standard Model, and the motivation behind the MicroBooNE experiment is outlined.

2.1 Discovery of the Neutrino

The existence of the neutrino was first postulated by Wolfgang Pauli in 1930, motivated by the puzzling observation that β -decays did not seem to conserve energy, momentum and angular momentum. At the time, β -decay was thought to be the process

$${}^A_ZX \rightarrow {}^A_{Z+1}X' + e^-. \quad (2.1)$$

If this were the case, then for a specific β -decay one would expect the electron to have a discrete energy equal to the difference in nuclear mass. However, the energies of β -decay electrons were observed to follow a continuous energy spectrum. To solve this problem, Wolfgang Pauli postulated that another low mass, neutral and unobserved particle was present in the β -decay, which he called a ‘neutron’, continuing the naming scheme of the already discovered electron and proton. After James Chadwick discovered the particle we now know as the neutron, the Italian physicist Edoardo Amaldi once jokingly referred to Pauli’s neutron as a ‘neutrino’ (an Italian diminutive meaning ‘little neutral one’). Fermi used the term neutrino at a conference in Paris in July 1932, and the term was adopted into scientific vocabulary from then on.

In 1942, Wang Ganchang proposed using beta capture as a method for detecting neutrinos via inverse β -decay, where the predicted particle $\bar{\nu}_e$ interacts with a proton p to create a neutron n and a positron e^+ [9] such that

$$\bar{\nu}_e + p \rightarrow n + e^+. \quad (2.2)$$

This interaction has a unique signature since the positron annihilates with an electron from the environment, creating two gamma rays, and the neutron can be captured by a nucleus, also producing a gamma ray. This experimental setup was realised by Clyde L. Cowan and Frederick Reines in 1956 in what is now known as the Cowan-Reines neutrino experiment [10]. In this experiment, the antineutrino source was a nearby nuclear reactor, the detector medium was water, and cadmium chloride was added to the tank to enable neutron capture. The photons produced by neutron capture were detected in tanks filled with liquid scintillator situated on either side of the water tank. After months of data collection, they collected neutrinos at a rate of about three per hour in their detector, and observed a difference in the neutrino rate when the reactor was off. Their conclusive results for the existence of an antineutrino were published in July 1956 [10] and their work was awarded the 1995 Nobel prize in physics [11].

After the existence of the neutrino had been experimentally established, a number of important results on their properties followed. It was established in 1958 by Goldhaber *et al.* that the neutrino has negative helicity [12]. In 1959, Davis *et al.* hypothesised that neutrinos and antineutrinos react differently with nucleons [13]. In 1962, an experiment was performed by Danby *et al.*, in which the ν_μ was discovered through the observation that the neutrinos in pion decays produce μ leptons, but no electrons, and were therefore distinct from ν_e [14]. In 1975, the τ lepton was discovered [15], confirming the existence of three generations of quarks and leptons. It was known that each generation consisted of particle and anti-particle

versions of two quarks, a charged lepton and an associated neutrino. The discovery of a third generation of lepton therefore also implied the existence of a third associated neutrino: the ν_τ . The DONUT collaboration discovered the ν_τ in 2000 [16]. There is strong evidence, gathered by the LEP experiments at CERN, that there are only three light active neutrino flavours. This was deduced from observations of the width of the Z boson resonance [17], which is governed by a hadronic, leptonic and neutrino component. This result strongly implies that if other neutrino flavours do exist, they are either extremely massive, or do not interact through the weak force.

A pivotal moment in the history of the neutrino was the *solar neutrino problem*. Nuclear fusion in the sun proceeds through a number of distinct processes, each of which has several stages. All of these processes produce electron neutrinos, causing the sun to be a pure ν_e source. The exploration of solar neutrinos began in the mid-1960s when Ray Davis and his team carried out a solar neutrino detection experiment based in the Homestake Mine in South Dakota, USA [18]. The Homestake experiment used a radiochemical technique to measure the flux of solar neutrinos and consisted of a tank of 615 tonnes of dry-cleaning fluid, C_2Cl_4 . The solar neutrino flux was measured by counting the number of ^{37}Ar atoms produced in the inverse β -decay process



The ^{37}Ar atoms were extracted from the tank and counted through their radioactive decays. Roughly 1.7 neutrino interactions per day were expected, but the observed rate was only 0.48 ± 0.04 neutrino interactions per day [19]. This apparent deficit of solar neutrinos became known as the solar neutrino problem. The Homestake experiment was sensitive to the relatively high-energy ^8B neutrinos produced in the sun. Subsequently, the SAGE and GALLEX radiochemical experiments used gallium as a target, and were sensitive to the low-energy neutrinos from the first step of the solar proton-proton chain. These experiments also observed a deficit of solar neutrinos [20][21].

Radiochemical experiments such as the Homestake, SAGE and GALLEX experiments played an important role in demonstrating the existence of the solar neutrino deficit. However, large water Cherenkov detectors established the origin of the deficit. A notable experiment in this category is the Super-Kamiokande experiment [22], for which the detector is a large vessel of water surrounded by photo-multiplier tubes (PMTs) that are capable of detecting single photons. In this experiment, solar neutrinos are detected by the elastic scattering process $\nu_e + e^- \rightarrow \nu_e + e^-$ and the neutrino energy and electron direction can be reconstructed from the number of detected photons and their orientation, respectively. Super-Kamiokande

observed clear evidence that solar neutrino oscillate, since the flux of electron neutrinos was measured to be about half the expected flux [23].

The Super-Kamiokande experiment and others established clear evidence that the observed rate of solar electron neutrinos was lower than the predicted rate. The Sudbury Neutrino Observatory (SNO) experiment in Canada was designed to measure both the ν_e and total solar neutrino flux [24]. The target medium in the SNO experiment was 1000 tonnes of heavy water, D_2O , and like Super-Kamiokande was a Cherenkov detector. This experiment measured the electron and non-electron component of the solar neutrino spectrum by comparing the Charged Current (CC), Neutral Current (NC) and Elastic Scattering (ES) neutrino reactions on deuterium. The use of heavy water as a detector medium is critical, since it is transparent to Cherenkov light and only ν_e can interact in the CC interaction $\nu + d \rightarrow p + n + \nu_e$, since solar neutrinos have energies smaller than the mass of muons and τ leptons. In addition, in the NC interaction, a neutrino dissociates the deuteron, breaking it into its constituent neutron and proton. The subsequent neutron capture produces a signature 6.25 MeV photon. All three neutrino flavours are equally likely to participate in the NC interaction. Lastly, electron elastic scattering has a higher interaction cross section for electron neutrinos, since exchange of both Z and W bosons is possible. The SNO data demonstrated that the total flux of neutrinos from the sun is consistent with the theoretical expectation, but rather than consisting of only ν_e , there is a large ν_μ and/or ν_τ component [25]. Since ν_μ and ν_τ cannot be produced in the fusion processes in the sun, SNO provides clear evidence of neutrino flavour transformations over large distances.

2.2 Neutrino Oscillations

The neutrino flavour transformations observed by SNO and other experiments can be explained by the phenomenon of neutrino oscillations. Neutrino oscillations result from the fact that there is a difference between neutrino mass and flavour eigenstates. The flavour eigenstates are labelled by lepton flavour e , μ , τ , and govern the interactions of neutrinos with matter. The mass eigenstates diagonalise the free Hamiltonian and describe the evolution of neutrinos in time and space. They are usually labelled 1, 2 and 3. The flavour and mass eigenbases are related by the PMNS matrix U [26–28], such that

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (2.4)$$

The PMNS matrix U_{PMNS} is given by

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (2.5)$$

where c_{ij} is shorthand for $\cos(\theta_{ij})$ and s is shorthand for sine. The complex Dirac phase δ quantifies the extent of leptonic CP violation, which is absent if $\delta = 0$. Two further phases α_1, α_2 are required if the neutrino is a Majorana particle, i.e. it is its own antiparticle. However, these phases affect only the diagonal matrix elements and are unobservable in neutrino oscillations. Therefore the Majorana phases are neglected here. The matrix U is represented here with dimensions 3×3 according to the scope of our current knowledge, but if experiments show that the matrix U is not unitary this points towards new physics, such as sterile neutrinos, and a different form for U , with different dimensions.

When, at a certain time $t = 0$, a neutrino of a particular flavour is produced, it exists as a coherent linear superposition of the three mass eigenstates. Consider for instance a ν_α state

$$|\psi(0)\rangle = |\nu_\alpha\rangle = U_{\alpha 1} |\nu_1\rangle + U_{\alpha 2} |\nu_2\rangle + U_{\alpha 3} |\nu_3\rangle. \quad (2.6)$$

Since $|\nu_i\rangle$ are mass eigenstates, their propagation can be described by plane wave solutions with position four-vector x_i and four-momentum p_i , such that

$$|\nu_i(x)\rangle = |\nu_i\rangle e^{-ip_i \cdot x_i} \equiv |\nu_i\rangle e^{-i\phi_i}. \quad (2.7)$$

Here the parameter ϕ is the phase of the plane wave representing the mass eigenstate, where

$$\phi_i \equiv p_i \cdot x = E_i T - p_i L. \quad (2.8)$$

It will shortly become apparent that the differences in phase are what drive neutrino oscillations. The wavefunction at position \mathbf{x} and time t can thus be written as

$$|\psi(\mathbf{x}, t)\rangle = |\nu_\alpha\rangle = U_{\alpha 1} |\nu_1\rangle e^{-i\phi_1} + U_{\alpha 2} |\nu_2\rangle e^{-i\phi_2} + U_{\alpha 3} |\nu_3\rangle e^{-i\phi_3}. \quad (2.9)$$

Expressing the mass eigenstates in terms of the flavour eigenstates in equation 2.9 and regrouping the terms allows equation 2.9 to be expressed as

$$|\psi(\mathbf{x}, t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_{i=1,2,3} e^{-i\phi_i} U_{\alpha i} U_{\beta i}^* \right) |\nu_\beta\rangle. \quad (2.10)$$

The probability that a neutrino in flavour state α oscillates into a flavour β is then given by

$$P(\nu_\alpha \rightarrow \nu_\beta) = |\langle \nu_\beta(t) | \nu_\alpha \rangle|^2 = \sum_{i,j} U_{\alpha i}^* U_{\beta i}^* U_{\alpha j} U_{\beta j} e^{-i(\phi_i - \phi_j)}. \quad (2.11)$$

In addition, the phase difference $\phi_i - \phi_j \equiv \Delta\phi_{ij} = (E_i - E_j)T - (p_i - p_j)L$ can be simplified by assuming that the momenta of the two states are equal and using the fact that $m \ll E$, such that

$$\Delta\phi_{ij} = (E_i - E_j)T = \left(p \left(1 + \frac{m_i^2}{p^2} \right)^{\frac{1}{2}} - p \left(1 + \frac{m_j^2}{p^2} \right)^{\frac{1}{2}} \right) T \approx \frac{m_i^2 - m_j^2}{2p} L, \quad (2.12)$$

where L is the propagation distance and using $T \approx L$ in natural units. Although this is a simplification, a proper wave packet treatment yields the same result [29]. Equation 2.12 shows that, if the neutrino masses are different, the different phases will evolve at different rates as the neutrino propagates through space. In particular, this phase difference only exists if at least two of the mass splittings Δm_{ij}^2 are non-zero, which is only possible if two of the neutrino flavours have mass. Here a conventional shorthand is used to indicate the mass splittings

$$\Delta m_{ij}^2 \equiv m_i^2 - m_j^2. \quad (2.13)$$

Hence, a neutrino in a particular flavour eigenstate will develop components belonging to other flavour eigenstates, meaning that after propagating some distance it may be detected as a different flavour eigenstate. Crucially, because neutrino oscillations are an empirical fact, neutrinos must have mass. Using equation 2.12, the two-flavour oscillation probability can be approximated by

$$P(\nu_\alpha \rightarrow \nu_\beta) \approx \sum_{i,j} U_{\alpha i}^* U_{\beta i}^* U_{\alpha j} U_{\beta j} e^{-i \frac{\Delta m_{ij}^2 L}{2p}}. \quad (2.14)$$

The fact that neutrino flavour states oscillate as they propagate is evidence of the fact that neutrinos have non-zero masses. However, it is clear from equation 2.14 that neutrino oscillations depend only on the differences between neutrino masses Δm_{ij}^2 , and so from oscillation experiments the individual neutrino masses remain unknown. While the absolute masses are yet unknown, the mass-squared splittings Δm_{ij}^2 have been measured in neutrino oscillation experiments. This introduces two possible ordering configurations of the neutrino masses: the ‘normal’ hierarchy in which $m_1 < m_2 < m_3$ and the ‘inverted hierarchy’ $m_3 < m_1 < m_2$. These two mass hierarchy hypotheses are visually represented in Figure 2.1.




Fig. 2.1 The neutrino mass hierarchy possibilities. Figure taken from [30].

2.2.1 Oscillation Experiments

When considering neutrino oscillation experiments, a distinction can generally be made between reactor and accelerator experiments and between short-baseline and long-baseline experiments. It will be shown that each of these four combinations are sensitive to different mixing angles θ_{ij} and neutrino mass differences Δm_{ij} .

Reactor Experiments

Nuclear fission reactors produce a large flux of $\bar{\nu}_e$ from the β -decays of various radioisotopes such as ^{235}U , ^{238}U , ^{239}Pu and ^{241}Pu , which are produced in nuclear fission. The mean energy of the reactor antineutrinos is about 3 MeV and the flux is known precisely from the power produced by the reactor, which is closely monitored. The $\bar{\nu}_e$ can be detected through the inverse β -decay process

$$\bar{\nu}_e + p \rightarrow e^+ + n. \quad (2.15)$$

If the $\bar{\nu}_e$ oscillate to other neutrino flavours, they will not be detected since the neutrino energy is well below threshold to produce a muon or tau-lepton in the final state, and hence the survival probability of $\bar{\nu}_e$ is measured. Inverse β -decay events can be detected because the e^+ annihilates and creates two photons. Subsequently, the neutron that is produced is captured by nuclei in the detector medium (a scintillator liquid), a process that produces another photon approximately 100 μs later. These photons produce Compton scattered electrons, which ionise the liquid scintillator and produce scintillation light. The signature for a $\bar{\nu}_e$ interaction is therefore the coincidence of a prompt pulse of scintillation light from the annihilation and a delayed pulse from the neutron capture 10-100 μs later. The observed amount of prompt light provides a measure of the neutrino energy. Using equation 2.14, the oscillation probability $P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ can be expressed as

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = 1 - \cos^4(\theta_{13}) \sin^2(2\theta_{12}) \sin^2\left(\frac{\Delta m_{21}^2 L}{4E_{\bar{\nu}}}\right) - \sin^2(2\theta_{13}) \sin^2\left(\frac{\Delta m_{32}^2 L}{4E_{\bar{\nu}}}\right). \quad (2.16)$$

Therefore, the oscillations are governed by the mass splittings Δm_{21}^2 and Δm_{32}^2 . The short wavelength component, which depends on Δm_{32}^2 , oscillates with an amplitude of $\sin^2(2\theta_{13})$ about the longer wavelength component, with wavelength determined by Δm_{21}^2 . For illustration, Figure 2.2 shows $P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ as a function of L for $E_{\bar{\nu}_e} = 3$ MeV, fixing the values of the mixing angles and mass splittings to sensible values (detailed in the caption). The oscillations occur on two different length scales.




Fig. 2.2 $P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ as a function of distance for $E_{\bar{\nu}_e} = 3$ MeV, assuming that $\theta_{12} = 30^\circ$, $\theta_{23} = 45^\circ$, $\theta_{13} = 10^\circ$, $\Delta m_{12}^2 = 8 \times 10^{-5} \text{eV}^2$ and $\Delta m_{32}^2 = 2.5 \times 10^{-3} \text{eV}^2$. Figure from [30].

In short-baseline reactor experiments, the long wavelength contribution is negligible, and $P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ can be approximated to be

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) \approx 1 - \sin^2(2\theta_{13}) \sin^2\left(\frac{\Delta m_{32}^2 L}{4E_{\bar{\nu}}}\right). \quad (2.17)$$

And as such, short-baseline experiments can impose constraints on θ_{13} and m_{32}^2 . An example of a short-baseline experiment is the Daya Bay experiment in China [31]. The Daya Bay experiments detect antineutrinos from six reactor cores, three at $L \approx 0.5$ km and three at $L \approx 1.65$ km. Each detector consists of a large vessel containing 20 tonnes of liquid scintillator loaded with gadolinium. By comparing the data recorded in the three far detectors at 1.65 km from the reactors, with the data from the three near detectors,

uncertainties cancel. The $\bar{\nu}_e$ were detected through inverse β -decay, and a difference in detected events between the near and far detectors of $N_{far}/N_{near} = 0.940 \pm 0.012$ was observed, which is consistent with $\sin^2(2\theta_{13}) = 0.0841 \pm 0.0027$ [32]. Recent results from the RENO reactor experiment [33] in South Korea, which is similar in design to the Daya Bay experiment, also show a deficit of $\bar{\nu}_e$, compatible with $\sin^2(2\theta_{13}) = 0.086 \pm 0.006$ [34]. Similarly, the Double Chooz experiment in France measured $\sin^2(2\theta_{13}) = 0.086 \pm 0.071$ [35].

An example of a long-baseline reactor experiment is the KamLAND experiment [36], which detected $\bar{\nu}_e$ from a number of reactors located 130-240 km from the detector through the inverse β -decay process. With this setup, the L/E dependence of the rapid oscillations due to the Δm_{32}^2 term in equation 2.16 is not resolved because the neutrino sources (the reactors) are not at a single distance L and also because the energy resolution is insufficient to resolve the rapid neutrino energy dependence. Consequently, only the average value of $\langle \sin^2 \Delta_{32} \rangle = \frac{1}{2}$ matters. Therefore, $P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ can be approximated as

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) \approx 1 - \cos^4(\theta_{13}) \left[1 - \sin^2(2\theta_{12}) \sin^2 \left(\frac{\Delta m_{21}^2 L}{4E_{\bar{\nu}}} \right) \right]. \quad (2.18)$$

Hence, long-baseline reactor experiments like KamLAND are sensitive to the values of θ_{12} and m_{21}^2 . KamLAND observed 1609 reactor $\bar{\nu}_e$ events where 2179 ± 89 were expected in the absence of oscillations. Because the amount of prompt scintillation light provides a measure of the neutrino energy $E_{\bar{\nu}}$, the survival probability at KamLAND can be plotted as a function of $L/E_{\bar{\nu}}$, which is shown in Figure 2.3. This figure shows a clear oscillation signal, where the survival probability depends on $L/E_{\bar{\nu}}$. The location of the minimum at $L/E_{\bar{\nu}} \approx 50$ km is consistent with $\Delta m_{21}^2 = 7.58 \pm {}^{+0.14}_{-0.13}(\text{stat}) \pm {}^{+0.15}_{-0.15}(\text{syst}) \times 10^{-5} \text{eV}^2$ and $\tan^2(\theta_{12}) = 0.56 \pm {}^{+0.10}_{-0.07}(\text{stat}) \pm {}^{+0.10}_{-0.06}(\text{syst})$ [37].

Accelerator Experiments

Many higher-energy neutrino experiments utilise neutrinos produced at accelerators and in the atmosphere. In both instances, neutrinos are produced predominantly through meson decays. High-energy cosmic ray particles interact with nuclei in the atmosphere to produce a shower of mesons, which subsequently decay into neutrinos. In accelerator experiments, protons are accelerated to high energies and impinged on a target, often consisting of beryllium or carbon, which then produces a large number of mesons. A good example of an accelerator neutrino experiment is the MINOS long-baseline oscillation experiment [38], which used the NuMI neutrino beam produced at the Fermilab accelerator complex (see Section 3.2). The




Fig. 2.3 Ratio of the background and geo-neutrino-subtracted $\bar{\nu}_e$ spectrum to the expectation based on oscillation parameters as a function of $L_0/E_{\bar{\nu}}$. L_0 is the effective baseline taken as a flux-weighted average ($L_0 \approx 180$ km). Figure from [37].

neutrino energy spectrum of this beam is concentrated in the range 1-5 GeV and peaks at 3 GeV. The 1000 ton MINOS near detector was located 1 km from the source and the 5400 ton MINOS far detector was located in a mine in Northern Minnesota, 735 km from the source. The detectors were relatively simple, consisting of planes of iron interleaved with planes of 4 cm wide strips of plastic scintillator. When a charged particle traverses the scintillator, scintillation light is produced, which is transmitted to small PMTs using optical fibres that are embedded in the scintillator. The detector was magnetised to enable the measurement of the momentum of muons produced in $\nu_\mu N \rightarrow \mu^- + X$ interactions from their curvature. The amount of scintillation light gives a measure of the energy of the hadronic final state X produced in the interaction. Hence, on an event-by-event basis, the neutrino energy is reconstructed: $E_\nu = E_\mu + E_X$. MINOS studied the neutrino oscillations of an almost pure ν_μ beam. Because θ_{13} is relatively small, $\nu_\mu \rightarrow \nu_\tau$ oscillations dominate, and most of the oscillated ν_τ are below threshold for producing a tau-lepton. MINOS therefore performs a disappearance measurement of $|m_{32}^2|$ and θ_{32} . The MINOS results are consistent with $|m_{32}^2| = (2.43 \pm 0.13) \times 10^{-3} \text{eV}^2$ (68 % CL) and $\theta_{32} \gtrsim 0.90$ (90 % CL) [39].

2.3 Neutrino Masses in the Standard Model

In Section 2.2, it was outlined that $m_2^2 - m_1^2 = (7.6 \pm 0.2) \times 10^{-5} \text{eV}^2$ and $|m_3^2 - m_2^2| = (2.3 \pm 0.1) \times 10^{-3} \text{eV}^2$. Oscillation experiments, being a type of interference experiment, are not sensitive to the absolute neutrino masses, but only their differences, or ‘mass splittings’. However, assuming that one of the neutrino masses is exactly zero, and given that the largest mass splitting is $\Delta_{32}^2 = (2.3 \pm 0.1) \times 10^{-3} \text{eV}^2$, would lead to a neutrino mass of $\sqrt{\Delta m^2} \approx 0.05 \text{eV}$. Furthermore, measurements of the ν_μ and ν_τ masses can be performed by measuring the masses of π and τ decay final states, and the ν_e mass affects the energy distribution of electrons emitted in β -decays, such as will be performed by the Katrin experiment [40]. Such measurements have also placed small upper bounds on the neutrino masses. Naturally, theoretical investigations have been carried out regarding the potential origin of these extremely small masses. One compelling explanation is the ‘seesaw mechanism’ [41, 42], which is described here.

From experiment, it is known that neutrinos interact only through the weak force, and that the weak charged current due to the exchange of W^\pm bosons is a vector minus axial vector (V-A) interaction of the form $\gamma^\mu - \gamma^\mu \gamma^5$, with a vertex factor of the form

$$\frac{-ig_W}{\sqrt{2}} \frac{1}{2} \gamma^\mu (1 - \gamma^5). \quad (2.19)$$

Here, g_W is the weak coupling constant, the strength of which is the same for all three generations of leptons. This is called lepton universality of the weak charged current. The expression of the vertex factor in equation 2.19 includes the left-handed chiral projection operator $P_L = \frac{1}{2}(1 - \gamma^5)$. This means that when both spinors are right-handed chiral states, the current reduces to zero, since $(1 - \gamma^5)u_R(p) = 0$, such that

$$j_{RR}^\mu = \frac{g_W}{\sqrt{2}} \bar{u}_R(p') \gamma^\mu (1 - \gamma^5) u_R(p) = 0. \quad (2.20)$$

Hence the only non-zero contributions to the interaction matrix element are those that involve only left-handed chiral particle states. For antiparticle spinors, P_L projects out right-handed chiral states, and so only right-handed chiral antiparticle states participate in the charged current weak interaction. The weak interaction is therefore said to be ‘left-handed’.

With no right-handed partner, the neutrino can have no Dirac mass term in the Lagrangian. To see this, note that the free-particle Lagrangian for a massive, spin 1/2 particle is

$$\mathcal{L} = i\bar{\psi}\gamma_\mu\partial^\mu\psi - m\bar{\psi}\psi. \quad (2.21)$$

But in this expression, $\bar{\psi}\psi = \bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L$. Hence a Dirac mass term mixes left-handed and right-handed neutrino states and vanishes. Because the weak interaction is left-handed, neutrinos cannot have a Dirac mass term in the Standard Model. For this reason, a standard Dirac mass term for the neutrino will require the existence of both right-handed neutrino and left-handed antineutrino states. However, because neutrinos carry no strong or electric charge, it is also possible that they are Majorana particles, i.e. they are their own antiparticles, such that $\psi^c = \psi$, where the superscript c indicates the charge conjugate. The operators which appear in the Lagrangian for the neutrino interaction in this case are the set $(\psi_L, \psi_R, \psi_L^c, \psi_R^c)$ and $(\bar{\psi}_L, \bar{\psi}_R, \bar{\psi}_L^c, \bar{\psi}_R^c)$. This still allows for Dirac mass terms of the form $m_D(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L)$, but terms of the form

$$\frac{1}{2}M_L(\bar{\psi}_L^c\psi_L + \bar{\psi}_L\psi_L^c) + \frac{1}{2}M_R(\bar{\psi}_R^c\psi_R + \bar{\psi}_R\psi_R^c) \quad (2.22)$$

also appear. These are the Majorana mass terms, which mix pairs of charge-conjugated states of the fermion. One implication of this concerns the conservation of lepton number. If neutrinos are Dirac particles, the neutrino has lepton number $+1$ and the antineutrino has lepton number -1 . Taking the muon family as an example, neutrinos must produce μ^- , as they both have lepton number $+1$, and antineutrinos must produce μ^+ . Such considerations disappear for Majorana neutrinos, allowing for reactions that violate the conservation of lepton number.

The existence of Majorana neutrinos would be experimentally detectable, since certain rare isotopes are forbidden from decaying through standard β -decay and would instead decay through neutrino-less double β -decay, a nuclear process whereby the nuclear charge changes by 2 units while the atomic mass is left unchanged, such that

$$(Z, A) \rightarrow (Z + 2, A) + 2e^-. \quad (2.23)$$

Such a process is forbidden in the Standard Model of particle physics, since it is a lepton number violating process. Therefore, observing neutrino-less double β -decay would confirm the Majorana nature of the neutrino [43, 44]. Various experiments, such as Cuoricino [45] have searched for neutrino-less double β -decays, but no conclusive evidence of their existence has been found to date.

Excluding neither the possibility of the existence of Dirac masses nor of Majorana masses, the most general renormalisable Lagrangian for the neutrino masses includes both the Dirac

and Majorana mass terms, which is of the form

$$\begin{aligned}\mathcal{L}_{DM} &= \frac{1}{2} (m_D \psi_L \bar{\psi}_R + m_D \bar{\psi}_L \psi_R + m_L \bar{\psi}_L^c \psi_L + m_R \bar{\psi}_R^c \psi_R) + h.c. \\ &= (\bar{\psi}_L^c \bar{\psi}_R) \begin{pmatrix} M_L & m_D \\ m_D & M_R \end{pmatrix} \begin{pmatrix} \psi_L \\ \psi_R^c \end{pmatrix} + h.c.\end{aligned}\quad (2.24)$$

In equation 2.24, the physical masses associated to the particles can be found by diagonalising the mass matrix M , where M is given by

$$M = \begin{pmatrix} m_L & m \\ m & m_R \end{pmatrix} \rightarrow \begin{pmatrix} m_1 & 0 \\ 0 & m_2 \end{pmatrix}. \quad (2.25)$$

In other words, the masses of the neutrino states are given by the eigenvalues m_1 and m_2 of the 2×2 matrix M , given by

$$m_{1,2} = \frac{1}{2} \left((m_L + m_R) \pm \sqrt{(m_L - m_R)^2 + 4m_D^2} \right). \quad (2.26)$$

Using equation 2.26 with different values for m_L , m_R and m_D will result in different values for the physical masses m_1 and m_2 . However, the most interesting behaviour occurs if one chooses $m_L = 0$ and $m_R \gg m_D$, in which case m_1 and m_2 evaluate to

$$m_1 = \frac{m_D^2}{m_R}, \quad (2.27)$$

$$m_2 = m_R \left(1 + \frac{m_D^2}{M_R^2} \right) \approx m_R. \quad (2.28)$$

In short, this procedure yields one very light neutrino mass $m_1 = m_D^2/m_R$ and a very large neutrino mass $m_2 \approx m_R$. If the Dirac mass term is of the same order of magnitude as the observed lepton masses $\mathcal{O}(1.0)$ GeV, then m_R is made sufficiently large such that $m_1 \sim 0.01$ eV. This implies that m_R must be very large: $m_R \sim 10^{11}$ GeV. This is the origin of the name of the seesaw mechanism: the smaller m_ν becomes, the larger m_R must be. The seesaw mechanism provides a natural explanation for the relative smallness of the neutrino masses compared to the other leptons. If the neutrino is a Majorana particle, for each neutrino flavour, the left-handed and right-handed chiral mass eigenstates would be almost entirely the light left-handed neutrino and the massive right-handed neutrino, respectively. Such a heavy right-handed neutrino with $m_R \sim 10^{11}$ would have a mass much too large to be created in experiments, and hence would remain unobserved.

2.4 Neutrino Interactions

The study of neutrino interactions with matter, and hence neutrino cross sections, is a central part of furthering neutrino physics. The mechanisms through which neutrinos interact with matter are dependent on the neutrino energy E_ν . The energy regime in which MicroBooNE operates is sufficiently high to introduce complexity into the description of neutrino scattering, as several distinct neutrino scattering mechanisms start to play a role. The scattering possibilities fall into three main categories: (quasi-)elastic (QE) scattering, resonance production (RES) and deep inelastic scattering (DIS). As a result of these various mechanisms, neutrino interaction final states within this energy regime can take myriad forms, ranging from nucleon emission to final states containing multiple mesons, and hadronic showers at higher energies. This energy regime is often referred to as the ‘transition region’, since increasing neutrino energy increases the spatial resolution of the interaction, allowing the neutrino to transition from interacting with the entire nucleon to the constituent partons within the nucleus.

Figure 2.4 gives an overview of existing cross sections measurements for neutrino interactions ($\nu_\mu + N \rightarrow \mu^- + X$) and antineutrino interactions ($\bar{\nu}_\mu + N \rightarrow \mu^+ + X$) across this energy range, gathered by different experiments over several decades. This figure demonstrates that the total cross section approaches a linear dependence on neutrino energy, which is predicted by the quark parton model: when neutrino energies increase, point-like scattering off quarks dominates the scattering mechanism. The different neutrino interaction mechanisms that contribute to the cross sections plotted in Figure 2.4 are discussed in this section.

2.4.1 Quasi-Elastic Interactions

The majority of neutrino-hadron interactions for $E_\nu \lesssim 2$ GeV are quasi-elastic, and as such this channel provides a large fraction of the signal events for many neutrino experiments, including MicroBooNE. In a quasi-elastic interaction, the neutrino scatters off of an entire nucleon, rather than its constituent partons. By convention, quasi-elastic scatters are taken to be CC elastic scatters, and NC elastic scatters are referred to simply as ‘elastic’. Quasi-elastic (anti)neutrino-hadron interactions are of the form

$$\begin{aligned} \nu_l n &\rightarrow l^- p \\ \bar{\nu}_l p &\rightarrow l^+ n. \end{aligned} \tag{2.29}$$

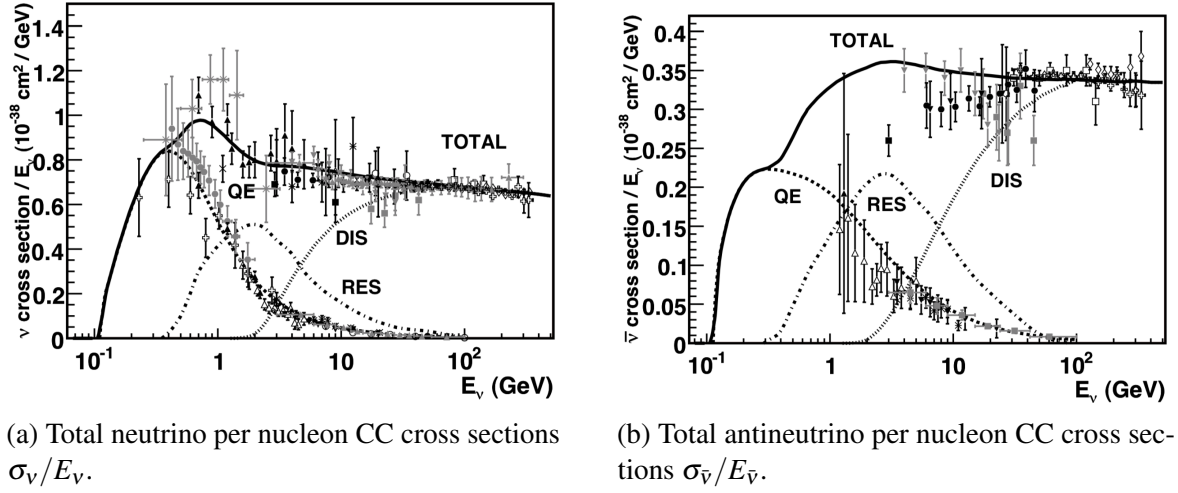


Fig. 2.4 The (anti)neutrino cross section per unit of energy σ/E_ν , as a function of E_ν . Contributions for quasi-elastic scattering (dashed), resonance production (dot-dashed) and deep inelastic scattering (dotted) are indicated. Cross section predictions as given by the NUANCE generator [46] are indicated by the solid line. Quasi-elastic scattering data and predictions have been averaged over neutron and proton targets and hence have been divided by a factor of two for the purposes of this plot. [47]

The Feynman diagrams for these processes are given in Figure 2.5. Much of the formalism regarding quasi-elastic neutrino interactions relies on the Llewellyn-Smith parametrisation [48], which expresses the quasi-elastic differential cross section in terms of functions of the nuclear form factors A , B and C , such that

$$\frac{d\sigma}{dQ^2} = \frac{G_F M^2 |V_{ud}|^2}{8\pi E_\nu^2} \left(A \pm \frac{(s-u)}{M^2} B + \frac{(s-u)^2}{M^4} C \right). \quad (2.30)$$

where the $-$ and $+$ refer to (anti)neutrino scattering, G_F is the Fermi coupling constant, $Q^2 \equiv -q^2$ is the squared four-momentum transfer, M is the nucleon mass, m is the lepton mass, E_ν is the neutrino energy and $(s-u) = 4ME_\nu - Q^2 - m^2$. The factors A , B and C are

functions of the form factors of the nucleon and are defined as follows:

$$\begin{aligned}
A &= \frac{(m^2 + Q^2)}{M^2} [(1 + \eta)F_A^2 - (1 - \eta)F_1^2 \\
&\quad + \eta(1 - \eta)F_2^2 + 4\eta F_1 F_2 \\
&\quad - \frac{m^2}{4M^2} ((F_1 + F_2)^2 + (F_A + 2F_B)^2) - \left(\frac{Q^2}{M^2} + 4 \right) F + P^2] \\
B &= \frac{Q^2}{M^2} F_A (F_1 + F_2) \\
C &= \frac{1}{4} (F_A^2 + F_1^2 + \eta F_2^2).
\end{aligned} \tag{2.31}$$

The functions F_1 , F_2 , F_A and F_P are the vector, axial-vector and pseudoscalar form factors, respectively, which can naively be interpreted as the Fourier transform of the internal charge distribution of the nucleus. Modern neutrino experiments often use complicated nuclei as their neutrino targets, such that nuclear effects become important and can produce significant corrections to the quasi-elastic cross section of equation 2.30. For quasi-elastic interactions, the nucleus is often modelled in terms of individual quasi-free nucleons that are part of the scattering process. This approximation is often called the impulse approximation (IA) [49], and in it the neutrino interacts with only a single nucleon, but this nucleon can have short-range correlations with the other nucleons in the nucleus. Many neutrino experiments use a Relativistic Fermi Gas approach [50] to model the nucleus and simulate quasi-elastic neutrino interactions, although more sophisticated treatments have emerged in recent years, such as in [51].

The hadrons that are produced in neutrino-nucleus interactions and exit the nucleus may re-scatter while propagating through the nuclear medium. Such re-scatters are often called Final State Interactions (FSI), and can change the charge and multiplicity of the outgoing hadrons, as well as altering their final state kinematics. It is also possible that multiple nucleons transition to an excited state in a quasi-elastic scatter, mainly through the meson exchange current (MEC). The meson exchange current interaction channel is characterised by the emission of two nucleons, a topology which is difficult to distinguish from quasi-elastic interactions. In this process, a weak boson from the leptonic current is exchanged by a pair of nucleons (2-body current). The existence of this process was first discussed in the context of inelastic electron scattering, for instance in [52], and was modelled in [53] in order to explain the charged-current quasi-elastic inclusive cross section measurements performed by MiniBooNE. The inclusion of MEC in neutrino interaction generators [54] has lead

several groups to successfully reproduce the MiniBooNE cross sections results, as is done for instance in [55].

Neutrinos can also scatter elastically off the nucleus through an NC interaction, which typically ejects a nucleon from the nucleus. The two interaction channels are:

$$\begin{aligned} \nu_\mu p &\rightarrow \nu_\mu p \\ \nu_\mu n &\rightarrow \nu_\mu n \end{aligned} \quad (2.32)$$

Calculations involving NC elastic scatters are very similar to those of CC quasi-elastic scatters, and as such they are usually treated using similar formalism.




Fig. 2.5 Feynman diagrams that illustrate charged-current (CC) quasi-elastic scattering (2.5a) and neutral-current (NC) quasi-elastic scattering (2.5b) processes.

Resonant and Coherent Interactions

In addition to (quasi-)elastic interactions, inelastic collisions are also possible, in which the target nucleon enters an excited state after the interaction. The neutrino interaction produces a baryon resonance N^* , which has a short lifetime and often decays to a nucleon and a single pion final state, according to the interaction

$$\begin{aligned} \nu_\mu N &\rightarrow \mu^- N^* \\ N^* &\rightarrow \pi N' \end{aligned} \quad (2.33)$$

where N denotes a neutron or a proton. While higher multiplicity decay modes are also possible, the focus here will lie on single pion production. In the intermediate neutrino

energy regime, this mechanism is the most common mechanism through which single pions are produced. In scattering off of free nucleons, there are seven possible resonant single pion reaction channels. Three are charged current:

$$\begin{aligned} \nu_\mu p &\rightarrow \mu^- p \pi^+ & \bar{\nu}_\mu p &\rightarrow \mu^+ p \pi^- \\ \nu_\mu n &\rightarrow \mu^- p \pi^0 & \bar{\nu}_\mu p &\rightarrow \mu^+ n \pi^0 \\ \nu_\mu n &\rightarrow \mu^- n \pi^+ & \bar{\nu}_\mu n &\rightarrow \mu^+ n \pi^-. \end{aligned} \quad (2.34)$$

And four are neutral current:

$$\begin{aligned} \nu_\mu p &\rightarrow \nu_\mu p \pi^0 & \bar{\nu}_\mu p &\rightarrow \bar{\nu}_\mu p \pi^0 \\ \nu_\mu p &\rightarrow \nu_\mu n \pi^+ & \bar{\nu}_\mu n &\rightarrow \bar{\nu}_\mu n \pi^0 \\ \nu_\mu n &\rightarrow \nu_\mu n \pi^0 & \bar{\nu}_\mu n &\rightarrow \bar{\nu}_\mu n \pi^0 \\ \nu_\mu n &\rightarrow \nu_\mu n \pi^- & \bar{\nu}_\mu p &\rightarrow \bar{\nu}_\mu p \pi^-. \end{aligned} \quad (2.35)$$

To describe such resonant single pion production processes, experiments commonly employ the Rein-Sehgal model [56] with the additional inclusion of lepton mass terms. This model gives predictions for both CC and NC resonance production and a prescription for handling interferences between overlapping resonances. The cross sections for the production of different resonances can be modelled, and at lower energy the production of the $\Delta(1232)$ resonance dominates. Furthermore, the pion can subsequently interact with the nucleons in the nucleus, which can cause it to be scattered elastically, absorbed by the nucleus, produce new pions or to exchange charge with the nucleus. These FSI processes are depicted in Figure 2.6.

In addition to creating a resonance, neutrinos can also interact coherently with the entire nucleus, transferring negligible energy to the nuclear target A . Such coherent interactions can produce single pion final states. In these low- Q^2 interactions, there is no nuclear recoil and a strongly forward-scattered pion is produced. Both NC and CC coherent pion production mechanisms exist:

$$\begin{aligned} \nu_\mu N &\rightarrow \nu_\mu N' \pi^0 & \bar{\nu}_\mu N &\rightarrow \bar{\nu}_\mu N' \pi^0 \\ \nu_\mu N &\rightarrow \mu^- N' \pi^+ & \bar{\nu}_\mu N &\rightarrow \mu^+ N' \pi^-. \end{aligned} \quad (2.36)$$

Figure 2.7 shows the Feynman diagrams for charged-current (CC) resonant π^+ production and coherent π^+ production as illustrations of both these processes. Finally, neutrino interactions can also create strange quarks in the final state, and are therefore capable of producing kaons. In the regime $E_\nu \lesssim 2$ GeV, kaons are predominantly produced through the






Fig. 2.6 Hadrons produced by the neutrino interactions can subsequently interact with the surrounding nucleons. This image gives a schematic overview of different pion FSI modes. Figure from [57].

Cabibbo-suppressed interaction $\nu_\mu N \rightarrow \mu^- K^+ N$. Kaon interactions are less relevant in the context of MicroBooNE physics, and will not be discussed here.



(a) Feynman diagram for charged-current (CC) resonant π^+ production.



(b) Feynman diagram for coherent π^+ production.

Fig. 2.7 Feynman diagrams that illustrate CC resonant (2.7a) and coherent pion (2.7b) production processes.

2.4.2 Deep Inelastic Scattering

Given sufficient energy, the neutrino can start to resolve the internal structure of the target nucleon and can scatter off of individual quarks within the nucleon, via the exchange of

a virtual W or Z boson producing a lepton and a hadronic system in the final state. This process is called deep inelastic scattering (DIS) and has been employed to test the Standard Model, measure its parameters, and to probe the structure of the nucleon. Charged and neutral current DIS processes are possible:

$$\begin{aligned} \nu_\mu N &\rightarrow \mu^- X & \bar{\nu}_\mu N &\rightarrow \mu^+ X \\ \nu_\mu N &\rightarrow \nu_\mu X & \bar{\nu}_\mu N &\rightarrow \bar{\nu}_\mu X. \end{aligned} \quad (2.37)$$

The interaction $\nu_\mu N \rightarrow \mu^- X$ is illustrated in Figure 2.8. The inclusive cross section for DIS scattering of neutrinos and antineutrinos can then be written as:

$$\frac{d^2\sigma^{\nu,\bar{\nu}}}{dx dy} = \frac{G_F^2 M_N E_\nu}{\pi(1 + Q^2/M_{W,Z}^2)} \left[\frac{y^2}{2} 2xF_1 + \left(1 - y - \frac{M_N xy}{E_\nu} \right) F_2 \right] \pm y(1 - y/2)xF_3, \quad (2.38)$$

where x , y and Q^2 are dimensionless kinematic variables:

$$\begin{aligned} x &\equiv \frac{Q^2}{2p_2 \cdot q} = \frac{Q^2}{2M_N E_\nu y} \\ y &\equiv \frac{p_2 \cdot q}{p_2 \cdot p_1} = \frac{E_X}{E_\nu} \\ Q^2 &\equiv -q^2 = -m_\mu^2 + 2E_\nu(E_\mu - p_\mu \cos \theta_\mu). \end{aligned} \quad (2.39)$$

Here, E_ν is the incident neutrino energy, M_N is the nucleon mass, E_X is the energy of the hadronic system X and E_μ , p_μ and θ_μ are the energy, momentum, and scattering angle of the outgoing muon in the laboratory frame. The DIS neutrino cross section depends on $F_i(x, Q^2)$, the nucleon structure functions. The structure functions depend on the parton distribution functions (PDFs), which describe the distribution of quark momenta within the proton and as such, measurements of DIS cross sections helps reveal the structure of the nucleon. Many neutrino experiments report measurements of F_i as part of their output.

2.5 The LSND and MiniBooNE Anomalies

The Liquid Scintillator Neutrino Detector (LSND) experiment [1] ran at the LAMPF accelerator at Los Alamos National Laboratory between 1993 and 1998. The LSND detector was a short-baseline Cherenkov detector, studying the appearance of $\bar{\nu}_e$ in a beam consisting predominantly of $\bar{\nu}_\mu$ through inverse β -decay. The neutrino beam was produced by firing an 800 MeV proton beam onto a target, which produces charged pions. The π^+ component




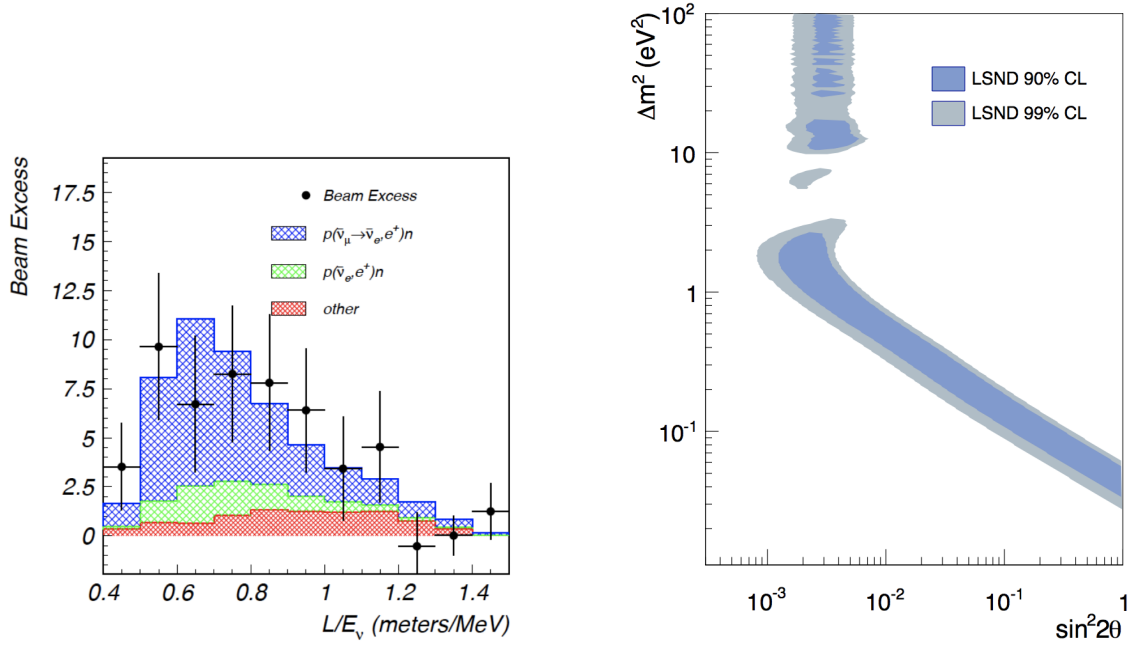
Fig. 2.8 Feynman diagram for the deep inelastic process $\nu_\mu + p \rightarrow \mu^- + X$, where X denotes the hadronic shower object.

was used to produce $\bar{\nu}_\mu$ through the subsequent decays $\pi^+ \rightarrow \mu^+ \nu_\mu$ and $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$. Furthermore, a distinction was made between neutrinos resulting from decays at rest (DAR) and decays in flight (DIF), by selecting events with $E_\nu \geq 60$ MeV for the DAR beam, since the maximum energy of a $\bar{\nu}_\mu$ produced by a stopping muon is $m_\mu/2 = 52.8$ MeV. The detector consisted of a vessel filled with 167 tonnes of mineral oil (CH_2), which was doped with the organic scintillation material butyl-PBD.

The LSND experiment observed an excess of $87.9 \pm 22.4 \pm 6.0$ $\bar{\nu}_e$ events, a 3.8σ excess in the DAR $\bar{\nu}_\mu$ beam [58]. A plot of the data, shown with a stacked plot of background (green and red) and fitted signal (blue), is shown in Figure 2.9a. This observation could be explained if the $\bar{\nu}_\mu$ oscillate into $\bar{\nu}_e$, but given that $L/E \approx 0.75$ m/MeV for the LSND detector, this would imply a mass splitting value of $\Delta m^2 \approx 1$ eV², which is an order of magnitude larger than any of the other mass splittings measured by neutrino experiments. Allowed and excluded regions in the $(\sin^2 2\theta, \Delta m^2)$ parameter space are shown in Figure 2.9b.

2.5.1 MiniBooNE

The main purpose behind the MiniBooNE experiment [59] was to test the oscillation signal observed by LSND. To do this, the MiniBooNE experiment used a very different experimental setup from LSND, while keeping the L/E ratio at approximately the same value. For MiniBooNE, $\langle L/E \rangle \approx 500/600$ m/MeV and for LSND $\langle L/E \rangle \approx 35/40$ m/MeV. The MiniBooNE experiment used the BNB at Fermilab, Illinois, and was capable of running in both neutrino and antineutrino mode, by charge-selecting the beam mesons. This allowed the



(a) Amount of measured $\bar{\nu}_e$ events as a function of L/E_ν .

(b) Allowed regions in the $(\sin^2 2\theta, \Delta m^2)$ parameter space.

Fig. 2.9 The LSND excess, indicating possible $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance is shown in 2.9a. The red and green histograms indicate the background from intrinsic $\bar{\nu}_e$ in the beam and from other sources, respectively. The blue histogram is an example of a potential $\bar{\nu}_\mu \rightarrow \nu_e$ signal that, when added to the predicted backgrounds, is representative of the data. The allowed and excluded regions in the $(\sin^2 2\theta, \Delta m^2)$ parameter space are shown in 2.9b. The inner and outer regions correspond to 90% and 99% CL allowed regions, while the curves are 90% CL limits from other neutrino experiments. Figures from [58].

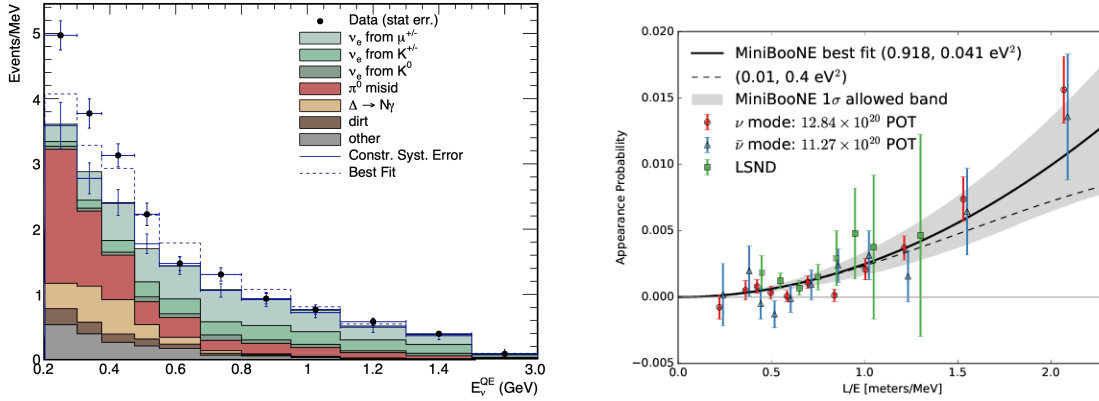
experiment to search for both $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations. The experiment also utilised information gathered by the SciBooNE experiment, located at around 100 m from the same beam source. By using both a near and a far detector, the neutrino flux and energy distribution can be measured in the near SciBooNE detector before oscillations occur, allowing for a more precise oscillation search.

The MiniBooNE detector is a spherical, 950,000 litre mineral oil Cherenkov detector. Events were detected using mainly Cherenkov light associated with the outgoing particles, since the oil contained no scintillator. The Cherenkov light formed rings on the 1280 photomultiplier tubes around the periphery of the spherical detector. The amount of light and Cherenkov ring geometry were used to measure the particle energies and direction as well as the particle type: muon, pion, or electron. Heavier particles such as muons produce sharp rings of Cherenkov light, whereas lighter particles create fuzzier rings. A $\pi^0 \rightarrow \gamma\gamma$ decay results in two potentially overlapping rings. Since electrons and photons induce the same type of Cherenkov ring, it is very difficult to distinguish between them.

The latest MiniBooNE results indicate a 4.7σ excess of 460 ± 99.0 events in the combined ν_e and $\bar{\nu}_e$ analysis, using data corresponding to 12.84×10^{20} and 11.27×10^{20} POT, respectively. The MiniBooNE results are shown in Figure 2.10, showing the neutrino E_ν^{QE} distributions for data (2.10a) and the appearance probability at MiniBooNE and LSND as a function of L/E (2.10b). The best oscillation fit in a two-neutrino oscillation model $\nu_\mu \rightarrow \nu_e$ has a probability of 21.1%, whereas the background-only fit has a χ^2 probability of 6×10^{-7} relative to the best fit. The results gathered by MiniBooNE are consistent in energy and magnitude with the excess of events measured by the LSND collaboration, and the significance of the combined LSND and MiniBooNE excesses is 6.0σ [3]. The MiniBooNE allowed regions for the combined neutrino and antineutrino modes for events with $200 \leq E_\nu^{QE} \leq 3000$ MeV, given a two-neutrino oscillation model, are shown in Figure 2.11.

2.5.2 Interpretations Beyond the Standard Model

If one accepts that the LSND and SciBooNE/MiniBooNE anomalous results are due to neutrino oscillations, new physics has to be introduced to explain these observations. The simplest addition to the three-flavour oscillation model is to add an additional neutrino mass



(a) The MiniBooNE neutrino mode low-energy excess for all CC quasi-elastic data.

(b) The appearance probability at LSND and MiniBooNE as a function of L/E .

Fig. 2.10 MiniBooNE results, showing the neutrino E_v^{QE} distributions for data gathered from 12.84×10^{20} protons-on-target (2.10a). Data points with statistical errors are shown, as well as stacked distributions of the background contributions, with systematic errors. Figure 2.10b shows the appearance probability at MiniBooNE and LSND as a function of L/E , with good agreement. Error bars show statistical uncertainties only, and the shaded band is the MiniBooNE allowed region, assuming two-neutrino oscillations.

state, so that the unitary mixing matrix becomes

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} & U_{\mu 4} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} & U_{\tau 4} \\ U_{s1} & U_{s2} & U_{s3} & U_{s4} \end{pmatrix}, \quad (2.40)$$

where the new flavour is denoted by the label s . The new mass hierarchy that results from the addition of this fourth neutrino flavour is shown in Figure 2.12. Adding this additional neutrino flavour has important consequences. First, the precise measurements of the Z^0 width made at LEP and SLD, determined that there are only three families of light-mass, weakly interacting neutrinos [17]. Hence, the additional neutrino, if it is lighter than half the mass of the Z , must not interact via the exchange of W or Z bosons, and is, therefore, sterile. The addition of a single sterile neutrino flavour in addition to the three existing ‘active’ flavours is often referred to as the 3+1 model.

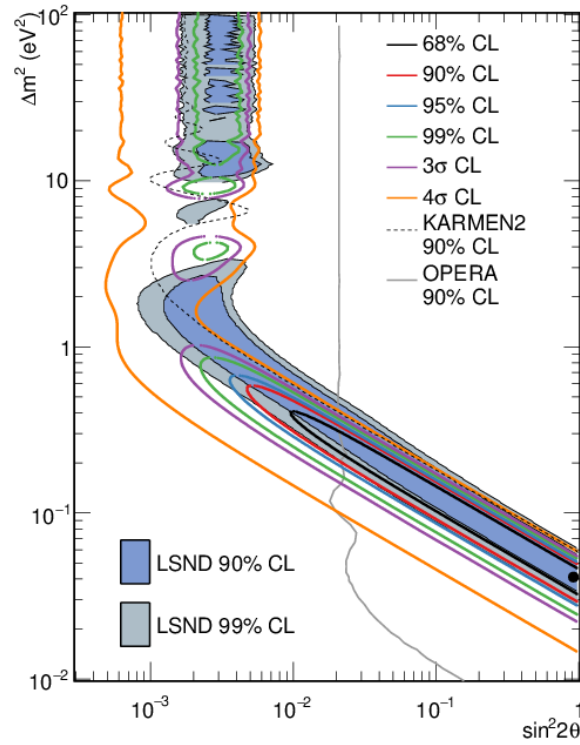


Fig. 2.11 MiniBooNE allowed regions for a combined neutrino mode (12.84×10^{20} POT) and antineutrino mode (11.27×10^{20}) data sets for events with $200 \leq E_{\nu}^{QE} \leq 3000$ MeV given a two-neutrino oscillation model. The shaded areas indicate the 90% and 99% C.L. LSND allowed regions. Also shown are 90% C.L. limits from the KARMEN [60] and OPERA [61] experiments. The black point indicates the MiniBooNE best fit at $(\sin^2 2\theta, \Delta m^2) = (0.96, 0.041 \text{ eV}^2)$. Figure from [59].




Fig. 2.12 The neutrino mass hierarchy with the addition of a sterile neutrino flavour.

2.5.3 MicroBooNE

The MiniBooNE experiment was subject to a number of background channels. The predominant two background channels are misidentified π^0 and $\Delta \rightarrow N\gamma$ events. Of these, the former represents the largest component of background events. The two photons that result from π^0 decay appear as a single ring in the Cherenkov detector when they are very forward-boosted or when one photon is below-threshold, resulting in a final state that mirrors signal events. The contribution from this background can be reduced significantly by reconstructing the invariant π^0 mass of the event and filtering events on this basis. Similarly, a neutral current resonant interaction can produce a Δ baryon, which can decay into a nucleon and a photon: $\Delta \rightarrow N + \gamma$ with a small branching ratio. Both these backgrounds are due to the fact that Cherenkov detectors cannot distinguish between electrons and photons. However, liquid argon time projection chamber (LArTPC) technology provides excellent spatial and calorimetric resolution, and will allow for electrons and photons to be distinguished from one another. For this reason, the MicroBooNE experiment is taking data, which has the primary physics goal of understanding the MiniBooNE and LSND low-energy excess. The following chapter will discuss the MicroBooNE experiment and detector in detail.

Chapter 3

The MicroBooNE Experiment

This chapter will provide an overview of the MicroBooNE experiment. Section 3.1 will define the physics goals underlying the MicroBooNE experiment, and will discuss the data it has taken so far. MicroBooNE receives a high-intensity neutrino beam from the FNAL Booster Neutrino Beam (BNB), which is discussed in Section 3.2. Section 3.3 discusses the MicroBooNE detector and how its liquid argon time projection chamber (LArTPC) has been implemented. Subsequently, sections 3.4 and 3.5 will outline the light collection and cryogenic subsystems. Section 3.6 will discuss how LArTPC and light collection signals are processed at MicroBooNE, and Section 3.7 discusses the simulation packages that are used to simulate neutrino events in the MicroBooNE detector. Finally, section 3.8 gives a brief overview of detector operations, including protons-on-target (POT) received by the detector and the purity of the liquid argon in the LArTPC.

3.1 Introduction

The MicroBooNE (Micro Booster Neutrino Experiment) [62] is a neutrino experiment using a liquid argon time projection chamber (LArTPC) detector of approximately 100 tonnes located at the Fermi National Accelerator Laboratory (FNAL) in Batavia, IL, USA. The MicroBooNE experiment is part of a sequence of experiments to utilise a LArTPC technology to study neutrino interactions. The potential of using cryogenic noble liquids as a detector medium in particle physics was first recognised in the 1970s and LArTPC technology was first employed in the ICARUS T600 detector [63], which started operation in 2010 and ran until 2014. In 2017, the ICARUS detector was moved to Fermilab, and it is expected to once again start taking data in late 2019. On a smaller scale, the ArgoNeuT experiment

operated a LArTPC with a volume of approximately 0.25 tonnes from 2009-2010 in the NuMI beam at Fermilab [64], which performed a series of detailed studies on the interaction of medium-energy neutrinos [65]. Next generation LArTPCs for the Short Baseline Neutrino Detector (SBND) experiment and the Deep Underground Neutrino Experiment (DUNE) [66–68] are now being designed and constructed, including a proof-of-principle detector called ProtoDUNE, which has been constructed at CERN [69] and operated on a test beam in autumn 2018 and is now taking cosmic ray data.

The primary physics goal for the MicroBooNE experiment is to resolve the low-energy excess observed by the LSND and MiniBooNE experiments. In addition, the state-of-the-art MicroBooNE detector will allow for other valuable physics analyses to be performed. For instance, MicroBooNE is capable of performing a suite of high-statistics cross section measurements, which will be of great interest to the neutrino physics community. In addition, MicroBooNE is capable of detecting supernovae through a sudden sharp burst in neutrino-induced events. If a supernova were to occur in our galaxy, MicroBooNE would detect more events than were observed in total from supernova (SN) 1987a, the only nearby supernova event to occur when neutrino experiments were operational. Data gathered by the MicroBooNE detector can be combined with data gathered elsewhere, furthering our potential for understanding supernova events. MicroBooNE employs a separate supernova data stream, which continually monitors for signals indicative of supernova events. Lastly, the development of the MicroBooNE detector has been an important step in the research and development (R&D) of LArTPC detectors. Many of the lessons learned regarding the construction and stable operation of LArTPC detectors can be applied to the construction of other detectors, such as DUNE.

MicroBooNE began operations in late 2015, and is currently still taking data. MicroBooNE is part of the short baseline neutrino program (SBN), for which the SBND detector is currently being constructed and the ICARUS detector is being refurbished [70]. Figure 3.1 shows an aerial diagram showing the location of MicroBooNE, the BNB beam line, and other experiments at FNAL, including the SBN near and far detectors. As can be seen in this figure, the MicroBooNE detector is located along the BNB beamline, 470 m from the target. The detector itself is a 60 metric ton fiducial mass (170 metric ton total mass) LArTPC detector housed within a cylindrical cryostat. It contains many auxiliary systems, such as the light collection, DAQ and CRT systems. Details of the Booster Neutrino Beam and the components of the detector are given in the following sections.




Fig. 3.1 Aerial diagram showing location of MicroBooNE along the BNB beam line (the orange dashed line) at Fermilab. Other detectors are also shown [62].

3.2 The Booster Neutrino Beam

The neutrino beam used in the MicroBooNE experiment is produced at the Fermilab accelerator complex. This complex consists of four separate accelerators: the Main Injector, the linear accelerator (Linac), the Booster and the Recycler. These four accelerators produce neutrino beams that serve several experiments, and they are depicted in Figure 3.2. The accelerator complex produces a low energy proton beam at the Booster ($E_p \approx 8$ GeV) and a high-energy beam at the Main Injector ($E_p \approx 120$ GeV). The primary neutrino beam used by the MicroBooNE experiment is the Booster neutrino beam (BNB). This beam starts as a beam of H^- ions, which are accelerated to a kinetic energy of roughly 400 MeV in a linear accelerator. Then, the ions are converted to bare protons by stripping electrons from the H^- ions using a carbon foil. The protons are accelerated in the 15 Hz Booster synchrotron up to 8.89 GeV momentum, and are released in ‘beam spills’ that occur every $1.6\mu\text{s}$ and contain approximately 4×10^{12} protons.

The protons in the beam spill are impinged on a beryllium target, and the number of protons-on-target is measured by two toroids upstream of the target, with an uncertainty of about 2% [72]. Other information such as beam position on target, intensity, timing and direction are also recorded, and is used to monitor the beam. When the protons hit the beryllium target, many secondary particles are produced, including pions and kaons which form the primary source of neutrinos and antineutrinos through their decay modes. The secondary particles are focused by a toroidal electromagnet, also called the horn, which has a neutrino and an antineutrino mode. In the neutrino mode, π^+ and K^+ , which decay into ν_μ are focused while π^- and K^- , which produce $\bar{\nu}_\mu$ are dispersed, resulting on a focused beam

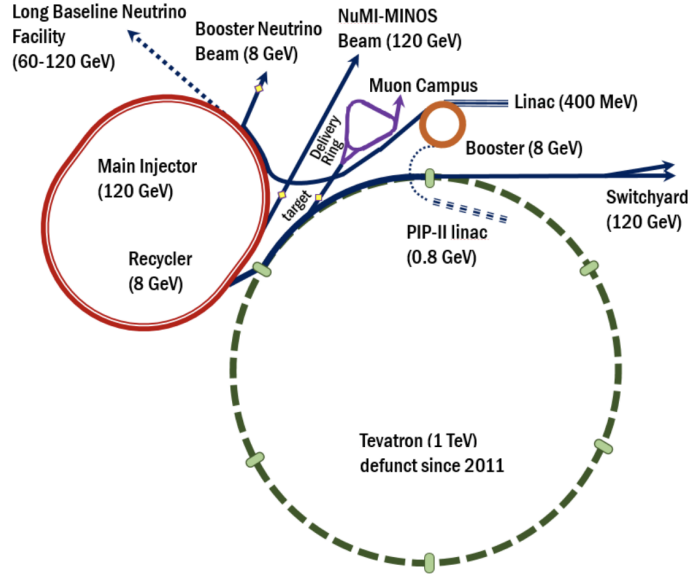


Fig. 3.2 A schematic overview of the Fermilab accelerator complex. Figure from [71].

of ν_μ , and vice versa for the antineutrino mode. After the secondary particles have been focused by the horn, they travel through a 50 m decay region, where these particles decay into the (anti)neutrinos that eventually reach the detector. To filter any charged particles which have not decayed, a concrete absorber has been placed in the beam line, as well as 450 m of dirt. After passing through these two types of absorbers, the highly pure muon neutrino beam impinges on the MicroBooNE detector. The BNB beamline is depicted schematically in Figure 3.3.




Fig. 3.3 Schematic of different components that make up the BNB beamline.

The analysis in this work uses a dataset collected when the beam was in the neutrino mode, such that the beam consists primarily of ν_μ . In this mode, the secondary particles are predominantly π^+ , with small contributions of K^+ and μ^- . The branching ratio of the decay $\pi^+ \rightarrow \mu^+ + \nu_\mu$ is very large, resulting in a very pure beam of ν_μ . However, some

contamination arises from μ^+ decay in the interaction $\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$ and the decay of μ^- ($\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu$) which are very energetic and hence not easily deflected by the focusing horn. Kaon decays also contribute to the neutrino flux, but in lesser measure due to the smaller kaon production rate. Decays of K^\pm , K^0 , K_L^0 result in higher-energy neutrinos, such that the ν_μ flux below roughly 2.5 GeV is due predominantly to pion decays, and the flux above this threshold is due almost exclusively to kaon decays. Kaons also contribute a $\bar{\nu}_e$ flux through decays such as $K^- \rightarrow \pi^0 e^- \bar{\nu}_e$. Through all these various decay modes, the Booster Neutrino Beam consists of 93.57% ν_μ , 5.86% $\bar{\nu}_\mu$, 0.52% ν_e and 0.05% $\bar{\nu}_e$. The simulated neutrino flux as a function of E_ν is shown in Figure 3.4.




Fig. 3.4 Neutrino flux prediction at MicroBooNE. Figure from [73].

3.3 The MicroBooNE LArTPC

The physical structure of the MicroBooNE LArTPC consists of three major structures: the cathode, the field cage, and the anode. These three components together measure roughly $2.6\text{m} \times 2.3\text{m} \times 10.4\text{m}$ in length, and a strong electric field (500 V/cm) is created between the cathode and anode planes located on the long edges of the detector. The cathode plane consists of 9 stainless steel sheets that are fixed to a frame and is kept at a nominal voltage of -128 kV produced by a dedicated high voltage system. A field cage encloses the active detector volume between the cathode and anode planes, and is designed to create a uniform electric field in this region. The coordinate system used in the MicroBooNE detector is shown in Figure 3.5.




Fig. 3.5 The coordinate system convention employed by the MicroBooNE collaboration. This figure schematically depicts the physical boundaries of the TPC, the coordinate system axes and the three wire planes on the right. Figure from [74].

The field cage consists of 64 tubular stainless loops spaced 4.0 cm apart that are held in place by a support structure. Each field cage loop is connected to the adjacent loops by a resistor divider chain such that the difference in electric potential between adjacent loops is 2kV. An ever-present danger is that of electrical discharge between the high voltage field cage loops and the inner cryostat surface and in order to prevent this the electric field strength is minimised at the corner and edges of the field cage. The anode plane consists of a frame that holds 8256 stainless steel wires distributed across three wire planes, and is situated opposite the cathode. Both of the induction planes, labelled U and V, are made up of 2400 wires and are at an angle of $\pm 60^\circ$ to the vertical, respectively. The collection plane, labelled W, consists of 3456 wires oriented vertically. The wires in each plane are separated by a 3 mm wire pitch, and the distance between each plane is also 3 mm. The wires are connected to application-specific integrated circuits (ASICs), which include a preamplifier, shaper and signal driver that propagate the signals induced on the wires to dedicated cold electronics. The ASICs are designed to operate at 83K, the temperature of liquid Argon. The MicroBooNE design parameters are summarised in Table 3.1.

3.4 Light Collection System

An additional valuable source of information is provided by the scintillation light that is produced inside of the detector, which is important both for reconstruction and event selection. The light collection system at MicroBooNE is designed to meet two physics goals: firstly, the

<i>MicroBooNE LArTPC Design Parameters</i>	
Property	Value
Number of anode planes	3
Wire pitch & plane spacing	3 mm
Number of wires & W	8256, 2400, 2400, 2456 (total, U, V & W)
Wire orientation w.r.t. vertical	+60 deg, -60 deg, 0 deg (U,V,W)
Cathode voltage (nominal)	-128 kV
Bias voltage (U,V,Y)	-200 V, 0 V, +440 V (U,V,W)
Nominal drift electric field	500 V/cm
Maximal drift time cathode to U plane (at nominal electric field)	1.6 ms
Number of field cage steps	64
Ring-to-ring voltage step	2.0 kV

Table 3.1 MicroBooNE LArTPC design parameters and nominal operating conditions. Adapted from [62].

light collection system should enable triggering on low-energy protons in accelerator-induced events and secondly, for non beam-induced events, the light collection system is designed to be able to efficiently observe electrons induced by supernovae with energies of 5 to 10 MeV.

Light produced in liquid argon arises from two processes: scintillation and Cherenkov radiation. Liquid argon is a bright scintillator, and scintillation light is produced by the de-excitation of excited argon dimers, and is emitted isotropically. Liquid argon produces a large amount of scintillation light at about 24,000 photons per MeV, at 500 V/cm drift field, and is transparent to its own scintillation light. The scintillation light has a prompt and slow component with decay times of about 6 ns and 1.6 μ s, respectively. Both components consist of vacuum ultraviolet (VUV) photons with a wavelength of 128 nm.

The light collection system allows for the formation of a beam event trigger when a pulse of light is observed to coincide with the beam spill. Since only about 1 in 600 beam spills produce a neutrino event, this significantly reduces the total amount of data produced. Furthermore, for accelerator-induced events, the time of the beam spill (1.6 μ s) functions as a t_0 for the event, but for non-beam physics studies, the interaction time t_0 can only be determined by the light collection system. Because the maximum drift time of electrons is about 1.6 ms at the nominal electric field strength of 500 V/cm, over 99% of accelerator-induced events includes one or more cosmic ray muon tracks. The distribution of hits

across the photodetectors can aid in rejecting cosmic rays and separating them from the neutrino-induced interactions.




Fig. 3.6 Diagram of the photomultiplier optical unit (left) and an image of the optical units mounted in MicroBooNE, immediately prior to LArTPC installation (right). Image from [62].

The light collection system consists of a primary system, which collects light for use in triggering and reconstruction, and a secondary sub-system, which was included mostly for R&D purposes. The primary light collection system is made up of optical units, each one consisting of a photomultiplier tube (PMT) located behind a wavelength-shifting plate, which shifts the 128 nm VUV liquid argon scintillation photons to the visible spectrum at roughly 425 nm, through its use of the organic molecule Tetraphenyl Butadiene (TPB). In total, 32 optical units were installed, yielding 0.9% photocathode coverage. The MicroBooNE PMT unit is depicted schematically in Figure 3.6, as well as a photograph of their placement within the TPC. The secondary system consists of four light guide paddles. These paddles were introduced for R&D studies for future LArTPCs, and are placed near the primary optical units to allow a comparison of their performances. The light collection system is located behind the three anode wire planes, and the combined transparency of these three planes is about 86% for light at normal incidence, assuming 100% absorption by the wires themselves. The optical unit and paddle placement in the detector is shown in Figure 3.7.

3.5 The Cryogenic System

The operation of the MicroBooNE LArTPC relies on the use of large quantities of highly-purified liquid argon. To cool the argon, maintain its purity, and to control pressure and temperature gradients, an intricate cryogenic system has been developed that can maintain stable operations with minimal downtime. The MicroBooNE cryogenic system is depicted




Fig. 3.7 The MicroBooNE light collection system consists of a primary system of 32 optical units and a secondary optical system of four lightguide paddles, which are present for R&D purposes. These are mounted behind the anode wire planes such that the view is not obscured by the structural cross bars of the LArTPC [62].

schematically in Figure 3.8. The main component of the cryogenic system is the cryostat, which houses the entire LArTPC as well as the light-collection subsystem. The cryostat is supported by three additional systems: the argon purification system, the nitrogen refrigeration system, and the controls and monitoring system. The cryostat and the subsystems will now be discussed in turn.

Three major components make up the MicroBooNE cryostat: a stainless steel vessel, front and rear supports, and foam insulation covering the cryostat outer surfaces. The vessel is cylindrical with domed caps at each end. The cryostat is 12.2 m in overall length, with an inner diameter of 3.81 m, and a wall thickness of 11.1 mm. The main design consideration for the cryostat is that the temperature gradients within the liquid argon must be kept to a minimum, so as not to introduce non-constant electron drift velocities. The electron drift velocity is $v_d = 1600$ m/s at an electric field of 500 V/cm, with a liquid argon temperature dependence $\Delta v_d/v_d = -0.019\Delta T$. This velocity must remain as constant as possible in magnitude and direction to prevent incorrect mapping from drift time to position along the drift direction. This requirement limits the allowable temperature variations of the liquid argon to less than 0.1 K and the laminar and turbulent flow rate of liquid argon to less than 1 m/s. These requirements limit fractional errors in velocity, and therefore in the drift coordinate determination, to be less than 0.1%. In addition to temperature requirements, the cryostat is designed to minimise outgassing (desorption), which the introduction of air into the system.




Fig. 3.8 A rendering showing the MicroBooNE cryostat and cryogenic system, and the platform for the electronics racks as installed in the LAr Test Facility (LArTF) [62].

An essential part of the cryogenic system is the liquid argon purification subsystem. Ionisation electrons within the LArTPC must not be significantly attenuated by impurities in the liquid argon, such as oxygen, nitrogen and water. The maximum drift distance of 2.5 m dictates that the argon be kept free of electronegative contaminants to the level of 100 parts-per-trillion (ppt) O_2 -equivalent. In addition to oxygen, nitrogen levels must be kept at less than 2 ppm (parts per million), as nitrogen detrimentally affects the quenching and attenuation of the scintillation photons. The argon purification system consists of liquid argon pumps and filters that circulate the argon and remove the aforementioned impurities. Unlike O_2 and water, the purification system does not have the ability to filter N_2 , and therefore the level of nitrogen contamination is determined by the quality of the delivered argon. There are two sets of pumps and filters in the system, arranged in parallel in order to allow for continuous recirculation while parts are being serviced.

The liquid argon detector medium is cooled by a liquid nitrogen refrigeration system, which consists of two condensers that are arranged in parallel. One of these condensers functions as a backup in case the first condenser fails. Each condenser contains two liquid nitrogen coils, and typically only one coil runs at any given time, but the second coil can be switched on when the heat load increases for any reason. Furthermore, to prevent ambient heating increasing the temperature within the cryostat and potentially boiling the liquid argon,

41 cm of spray-on, closed cell, polyurethane insulation was applied to the exterior of the cryostat upon installation at the LArTF.

3.5.1 Purification System and Purity Monitors

The temperature and purity of the liquid argon are constantly monitored to ensure they remain within the set bounds. MicroBooNE utilises resistive thermal devices (RTDs) to measure temperature throughout the experimental infrastructure. Twelve RTDs are arranged along the inside of the cryostat, and an additional ten RTDs are mounted inside screws attached to the structure of the LArTPC. Furthermore, each of the filter vessels in the purification system contains nine RTDs. The purity of the liquid argon is measured using double-gridded ion chambers, which can measure O_2 equivalent contamination of 50 to 300 ppt. The electron lifetime inside the detector active volume is estimated by measuring the fraction of electrons generated at the purity monitor cathode that subsequently arrives at the purity monitor anode (Q_A/Q_C) after a drift time $t - t_0$. The ratio of (Q_A/Q_C) is related to the electron lifetime τ :

$$Q_A/Q_C = e^{-(t-t_0)/\tau} \quad (3.1)$$

This makes it possible to estimate the electron lifetime based on the value of Q_A/Q_C . Three purity monitors are used, which are of various lengths and are positioned at different points inside the detector. One purity monitor, with a drift distance of 50 cm, is located just downstream of the filters and is used to monitor filter effectiveness. Two purity monitors, one with a drift distance of 19 cm and the other with a drift distance of 50 cm, sit within the primary vessel at both ends of the LArTPC. They are installed at different heights to allow purity measurements at different depths of the argon.

The MicroBooNE cryogenic system was designed to allow the cryostat to go from containing air to containing high purity liquid argon without evacuating the cryostat, which is essential from an R&D perspective for future multi-kiloton experiments where the cost of an evacuable cryostat would be prohibitive. MicroBooNE was the first fully-instrumented detector to attain this goal, which was attained in three stages. The first stage is the so-called ‘piston purge’, where the detector is filled with gaseous argon at its lowest point, which then pushes the atmosphere out of the detector. The second stage is a recirculation stage where the gaseous argon loop is closed, and water is removed from it using a water filter. Finally, the gaseous argon is cooled down to the point where injection of liquid argon does not damage the detector. Two major technical challenges in the cooling down phase were to prevent sense wires from breaking, since the massive TPC frame cools and shrinks more slowly than the

wires themselves. Secondly, the temperature difference between the top and bottom of the TPC should be less than 20 K during the entire procedure to avoid warping the frame.

3.6 Signal Processing

The operations of a LArTPC produce a large number of analog signals, which must be amplified, digitised and written to disk for further analysis. The extremely low temperatures of a liquid argon environment pose a design challenge, and dedicated low-noise cold electronics have been developed that are responsible for amplifying the signal and transmitting it from the LArTPC electronic channels and the PMTs to a readout system. The readout system then digitises the information, before sending it to the data acquisition system (DAQ) that stores it on disk. The stages of signal processing are illustrated in Figure 3.9 and are described in what follows.

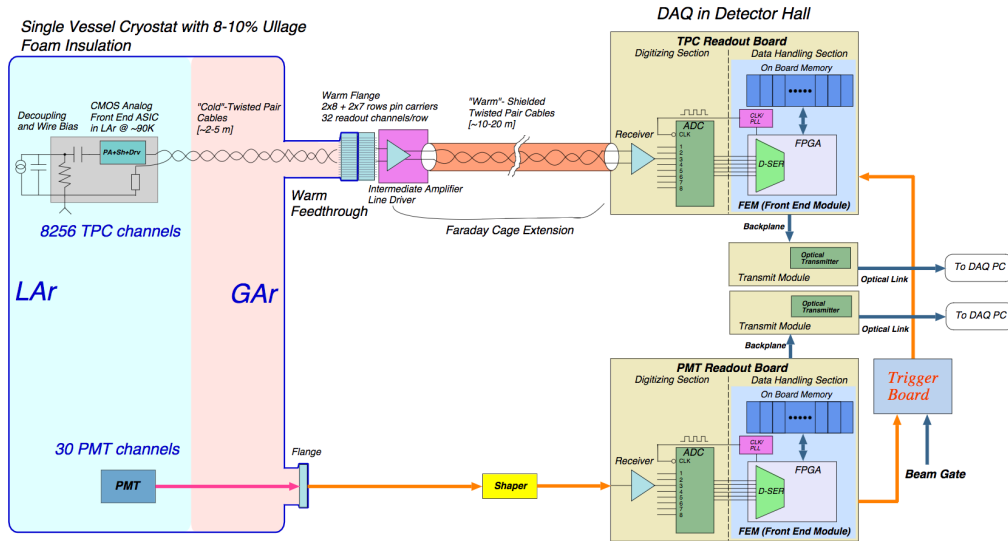


Fig. 3.9 MicroBooNE LArTPC and PMT signal processing and readout stages [62].

MicroBooNE utilises cryogenic low-noise electronics for the readout of the LArTPC. The cold electronics consist of three main parts: the ASICs (application-specific integrated circuits), cold motherboards and cold cables. To reduce noise, the distance between the LArTPC wires and the preamplifier is made as short as possible, by placing the ASICs inside the cryostat. A total of 512 ASICs are required for the readout of the entire detector. The ASICs are housed on a number of cold motherboards, which provide functionality for the calibration and the application of bias voltage to the wire planes, which is crucial to

allow them to function optimally. A specially designed type of cold cable is used both for transmitting the detector signals and also to distribute power to the ASICs. The MicroBooNE cold electronics provide a high-precision charge calibration system, which supplies charge calibration signals through the cold cables and the cold motherboards. The calibration signals enter the cryostat through a feed-through and reach the preamplifiers. The calibration system allows for gain calibration of each individual channel, which determines how sensitive a wire is to charge and allows for measurements of wire integrity and wire noise levels.

The LArTPC readout electronics process the signals produced on the 8,256 wires in the MicroBooNE detector after they have been transmitted from the cold electronics to the warm electronics outside the cryostat. The signal is transmitted to nine readout crates, which digitise and process the analog signal and transmit them to dedicated DAQ machines (one per crate). Data is also transmitted to the slow monitor, which is software designed for the monitoring of all detector subsystems. A data transmitter module is responsible for sending two separate data streams to the DAQ machines. The first data stream is called the NU (neutrino) data stream, which contains losslessly compressed LArTPC data associated with event triggers received by the LArTPC readout crates, such as the BNB trigger. The other data stream is the continuous SN (supernova) data stream, which is compressed with some information loss. The PMT readout electronics process the signals from the 32 PMT units in the detector and identify light signals that coincide with the BNB and NuMI beam spills. Such coincidences generate PMT triggers that can be used in conjunction with other triggers by the trigger board. PMT data is transmitted to DAQ machines in NU and SN streams, similarly to the LArTPC data.

The MicroBooNE trigger board (TB), which is placed within the PMT readout crate, generates ‘level-1’ triggers that flag frames that are of interest and must be treated differently. In the LArTPC and PMT readouts, this trigger flags up four 1.6 ms frames from the NU datastream that are destined for readout. The trigger board has multiple inputs, including BNB and NuMI triggers, and a PMT trigger. The PMT front-end module (FEM) can generate two different PMT triggers: a cosmic and a beam gate PMT trigger. The latter PMT trigger is designed to help identify neutrino interactions, and requires that the PMT multiplicity (the number of PMTs with measurable signal) is above 1, and places a lower bound on the total summed PMT pulse height within the beam spill duration. The PMT trigger is mixed with the beam triggers and other triggers in the trigger board, and when a trigger is generated by the TB it is distributed to all LArTPC and PMT crate controllers. The trigger is then propagated to each FEM, and when an FEM receives a trigger it inhibits its SN datastream and directs its data to the NU readout stream.

The MicroBooNE DAQ reads in the data from the readout electronics and writes it to local disk before it is transferred to long-term storage. It also configures and controls the readout electronics when data is taken in the detector, and is responsible for monitoring data flow and detector conditions to ensure stable operations. The information from each readout crate is sent to a dedicated server called the sub-event buffer (SEB), where it is assembled into a sub-event fragment. The collection of sub-event fragments is sent to a single event-building machine (EVB) over an internal network, which assembles them into full event files. A large collection of data processing systems, as well as the database which holds and monitors the state of the data flow is known collectively as the Python/Postgres for MicroBooNE Scripting system (PUBS). As data flows through the system using PUBS, the progress of each stage is monitored and can be viewed via a GUI. PUBS can also be used to assess the state of the SN stream data.

3.6.1 Hit Finding

As charged particles traverse the detector medium, they induce ionisation electrons, which are drifted onto the wire planes by the electric field. A bias voltage is applied to the wire planes to ensure full transparency of the first two wire planes, such that signals are produced on all three planes. As electrons propagate past the wire planes, they induce a bipolar signal on the induction planes and a unipolar signal on the collection plane. Bipolar signals are converted into a unipolar signal before further processing. The area under this curve, in units of analog-to-digital (ADC) counts, is a measure of hit energy which can be calibrated to units of MeV. For large stretches of time, no electrons are near the wires and there is only noise, such that the signal is first subdivided in regions of interest (ROI), which correspond to time periods with higher-than-noise activity. An ROI waveform is shown in Figure 3.10a for a single collection plane wire in a simulated event. After an ROI has been identified, the Gaussian hit finder algorithm attempts to find individual hits in the waveform by fitting Gaussian peaks [75], as shown in Figure 3.10b. Three sets of 2D hits can be reconstructed in this way, one for each wire plane. The x coordinate of each hit is defined as $x = (t - t_0)/v_d$, where t_0 is the time since the start of the beam spill t_0 and $v_d = 1600$ m/s (at an electric field of 500 V/cm) is the electron drift velocity. Information regarding the z coordinate is given by the collection plane wire positions, and because the U and V sense wires intersect the wires of the collection plane, y information can be retrieved by combining information from the different views. Through this process, 3D hits can be reconstructed (discussed further in Section 4.4).




Fig. 3.10 Example induction pulses on a W plane wire for a simulated event, with and without Gaussian fits. The x axis corresponds to units of time, and the y axis has units of ADC (analog to digital counts). Image generated using LARSOF [7, 8].

The reconstructed hit collections can be used to reconstruct particles, and to visualise events. An example of an event in the MicroBooNE detector, recorded 23 October 2015, is shown in Figure 3.11. This example image illustrates the principal components of LArTPC event topologies: long, straight tracks, two electromagnetic showers and cosmic-ray background tracks.




Fig. 3.11 A neutrino event in the MicroBooNE detector, recorded on October 23rd, 2015 (part of the Run 1 dataset). The direction of the beam line runs along the horizontal axis, and the neutrino interaction vertex (manually labelled *V*) is clearly present as the point of confluence of the final state particles. Several tracks are present in the event, as well as an electromagnetic shower (labelled *T* and *S*, respectively). In the bottom left corner and on the right-hand side, three cosmic ray tracks are also visible (labelled *CR*). Figure modified from [76].

3.7 Simulation & Cosmic Ray Background

Sophisticated simulations of the MicroBooNE beam and detector are available and represent an estimate of neutrino production and interactions. Simulated events serve as a basis for the development of reconstruction logic and analyses. Simulations rely on random sampling of estimated parameter spaces and are therefore classed as Monte Carlo (MC) methods. The MicroBooNE flux simulation is based on the flux simulation developed at MiniBooNE [72] and neutrino interactions are simulated using the GENIE event generator [77], which simulates the primary interaction within the argon nucleus, the production of final states through hadronisation and the propagation and scatters of final state particles. Various configurations are available within GENIE that affect simulation aspects such as the nuclear model and how different interactions are modelled. In this work, the baseline MicroBooNE configuration is used. The simulation of the MicroBooNE detector is done through the use of the GEANT4 package [78], which models particle propagation, drift of ionisation electrons to the wire planes, as well as propagation of scintillation light to the PMTs.

Another important aspect to the simulation is modelling cosmic-ray tracks. MicroBooNE is a surface detector that is subject to a large amount of cosmic rays at a rate of about 5 kHz [62]. This results in an average of 24 cosmic-ray tracks and showers that are included in each 4.8 ms DAQ window [79]. Since only 1 in 600 beam spills on average result in a neutrino interaction, the cosmic-ray background leads to two types of background: one in which there is no neutrino interaction in the $1.6 \mu\text{s}$ beam window, but a cosmic ray induces a PMT trigger and its topology mimics that of a neutrino interaction; and a second type in which a neutrino interaction is present, inducing a PMT trigger, but a cosmic ray track is selected as the most likely neutrino candidate. These two types of events are illustrated in Figure 3.12. The first type of background can be estimated from data gathered when the beam is off (also called off-beam or EXTBNB data), but the second type of background must either be simulated. Alternatively, simulated neutrino events can be combined with cosmic interactions from data. The CORSIKA [80] event generator is used to produce the simulated cosmic ray interactions. A variety of configurable parameters are available in the CORSIKA generator, but only the MicroBooNE default settings are used in this study.

To assist in the removal of cosmic-ray background, the MicroBooNE cosmic ray tagger (CRT) has been built [79]. MicroBooNE began CRT data-taking in March 2017. The CRT system in MicroBooNE is an additional component situated above and to the sides of the cryostat. An illustration of the CRT, showing its different components as well as a simulation of cosmic rays crossing the CRT, is shown in Figure 3.13. The CRT is composed of four panels of scintillators, the top panel, two side panels and the underside panel, providing

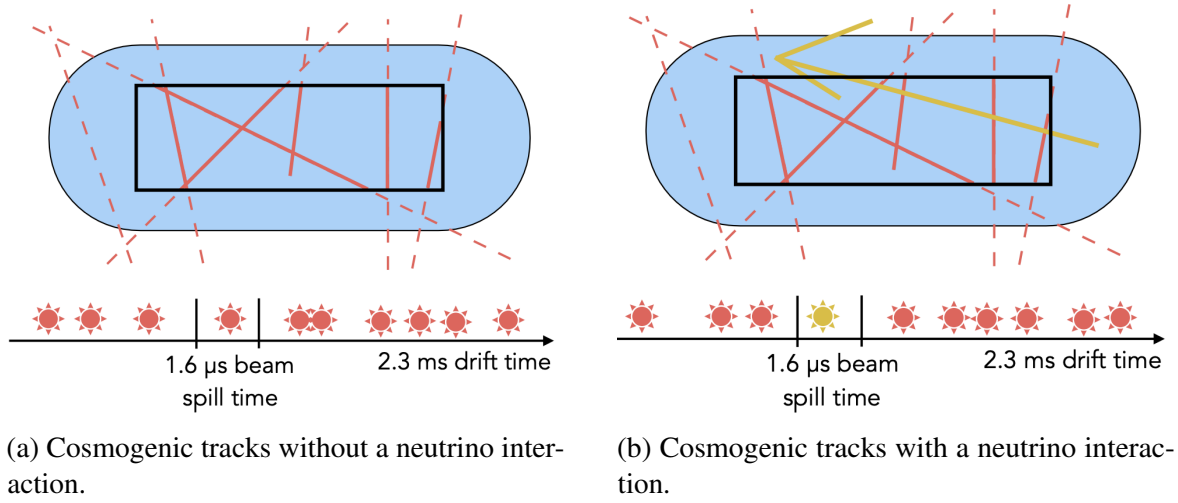


Fig. 3.12 Illustration of the two types of cosmic-ray background events. Figure 3.12a is an event in which no neutrino interaction is present, but PMT trigger is generated by a cosmic-ray track. In Figure 3.12b a neutrino interaction induces a PMT trigger, but a cosmic ray track may still be selected as the most likely neutrino candidate. The blue outline is a schematic representation of the cryostat. Figure from [81].

roughly 85% coverage of the LArTPC volume. The design of the CRT has been motivated by attempting to attain maximum coverage, given the physical space constraints imposed by the CRT module being within the liquid argon test facility (LArTF), where the MicroBooNE detector is housed. Each plane consists of several scintillator modules, with a total of 73 modules divided across the four planes, and the dimensions of the panels range from $4\text{ m} \times 9\text{ m}$ to $14\text{ m} \times 8\text{ m}$. As a cosmic ray particle crosses a scintillator module, photons are created within each individual scintillator strip, which represents the entry position of the cosmic ray particle in the CRT. The cosmic ray particle subsequently traverses the LArTPC volume, if it has non-negligible momentum, and with a large probability intersects either the side or bottom panels after leaving the LArTPC. The CRT has therefore provided an entry and exit position, which can be compared to the reconstructed positions of tracks within the LArTPC, which greatly aids in their removal. The MicroBooNE CRT was installed in two phases, with phase one starting in July-September 2016 to install the bottom, feedthrough side and pipe side planes and phase two starting in March 2017 to install the top pane.

3.7.1 Space Charge Effects

Another complication introduced by the cosmic-ray background is that the constant interactions within the detector create slow-moving positive ions that distort the electric field within

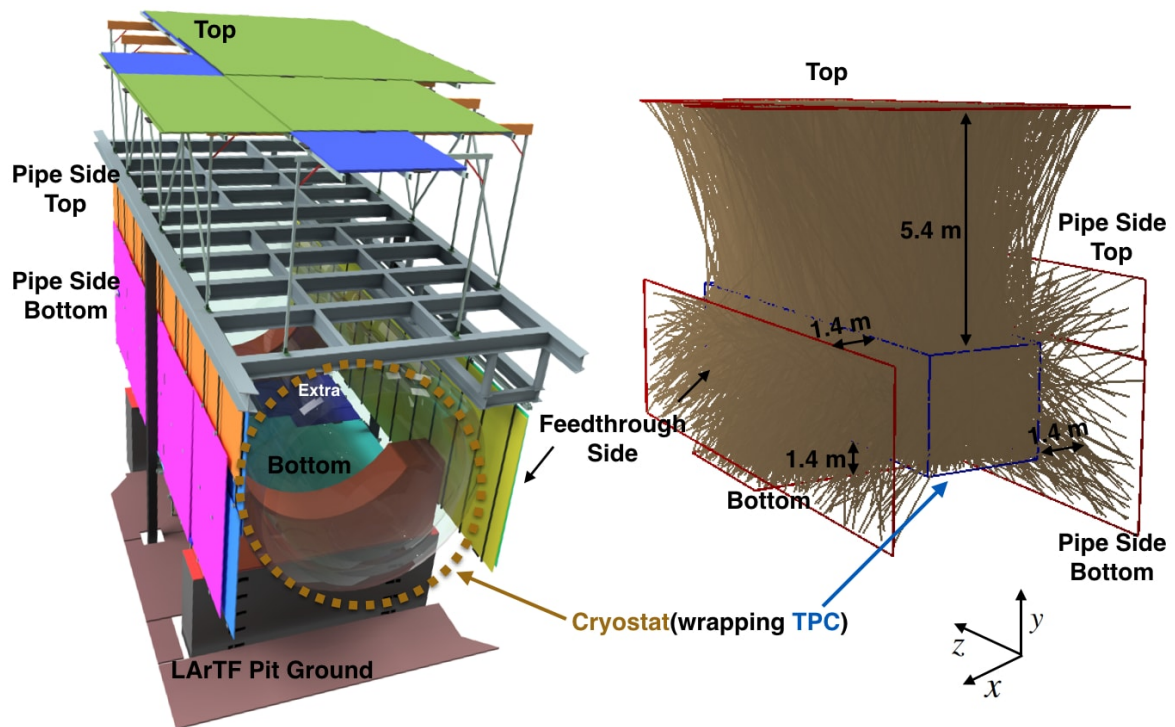


Fig. 3.13 Rendering showing the LArTPC muon tagger panel groups (topside, side, and underside) and a simulation of cosmic rays crossing the CRT, the brown lines represent possible cosmic ray trajectories. This design achieves a coverage of 85% of cosmic ray muons passing the TPC [79].

the detector. The drift velocity of such ions under nominal electric field in MicroBooNE is drastically lower than that of electrons, at approximately 0.8 cm/s, and while these positive ions are neutralised when they reach the cathode, an equilibrium is reached within the detector, such that positive ions are constantly present. The distortion of the electric field due to these positive ions in turn distorts the reconstructed trajectories of tracks within the detector by up to 10 cm. These effects are referred to as 'space charge effects' (SCE) [82]. An example of the effect of space charge effects on the reconstructed endpoints of cosmic ray tracks is shown in Figure 3.14.




Fig. 3.14 The reconstructed start and end positions of cosmic ray tracks that were tagged by an external muon counter in the x - y plane. In the absence of space charge effects, these points would be distributed along the detector boundaries, which are indicated by dashed lines. The anode is located at $x = 0$ cm and the cathode is located at $x = 256$ cm. Figure from [82].

Software has been developed to simulate the impact of space charge on the electric field. The simulation assumes a constant charge deposition rate of $2 \times 10^{-10} \text{ C/m}^3/\text{s}$, which can be calculated from the rate of cosmic rays, their average track length, and assuming that they are minimally ionising muons that have an energy loss of 2.1 MeV/cm. A recombination factor is applied to this rate, which reduces it by roughly 25%. Since the charge deposition rate is constant, to a good approximation higher-order effects can be ignored, such that the space charge density is linear with respect to the distance from the anode plane. The positive ions

accumulate near the cathode, and as such the distortions to the electric field are largest for large x positions. In addition, in the centre of the detector electrons experience an attractive force due to the positive ions isotropically, such that tracks remain undistorted. At the edges of the detector, however, drift electrons are pulled back inwards into the detector because of the accumulation of positive charge. Figure 3.15 shows the simulated effects of space charge on track x , y and z coordinates. These values are used within this thesis to correct for space charge effects within the Pandora framework.

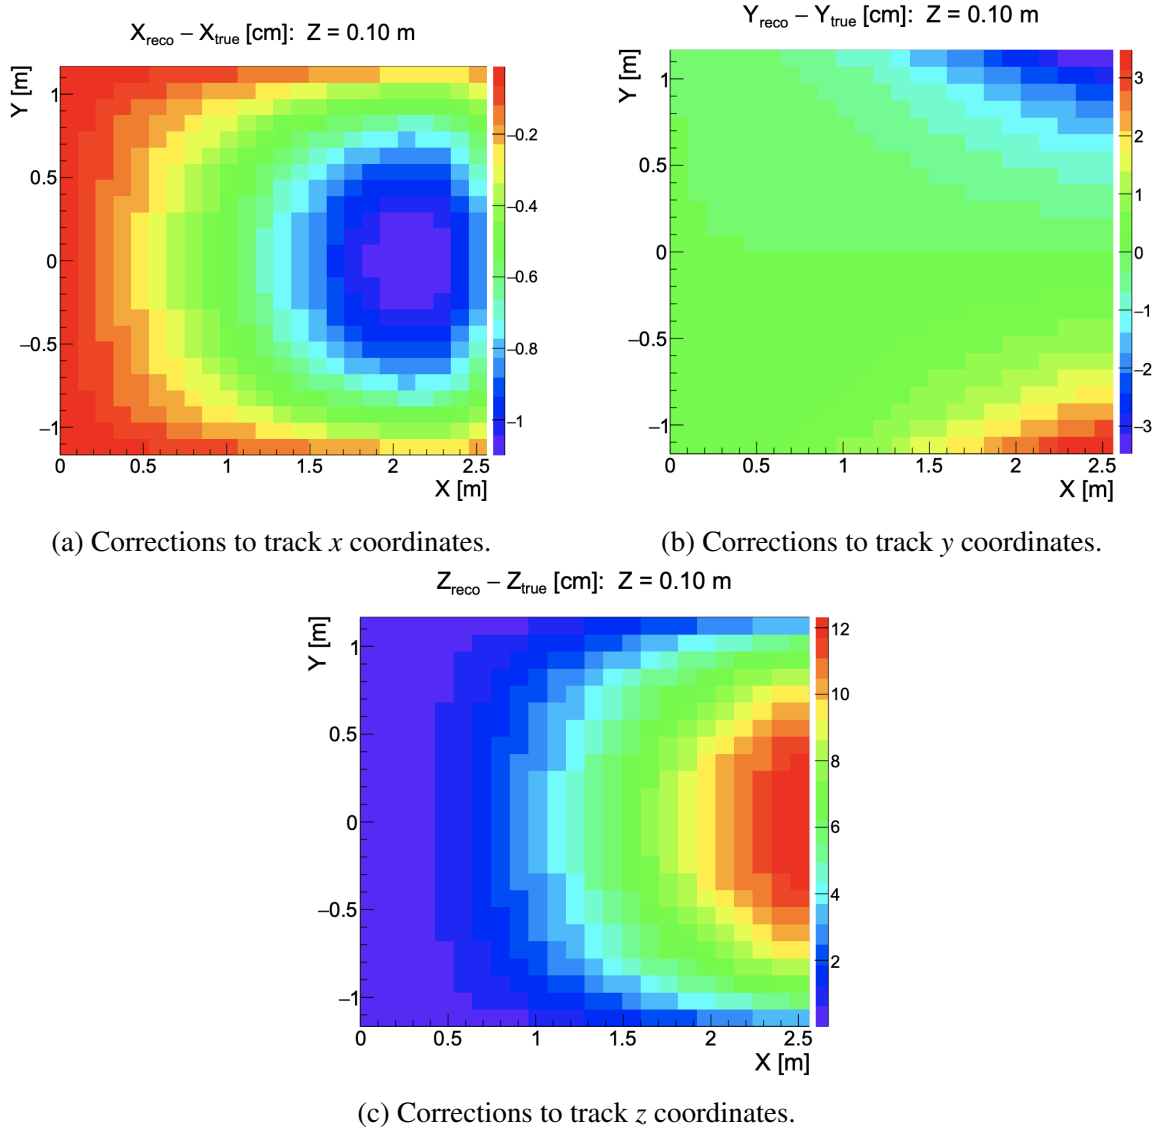


Fig. 3.15 Illustration of the simulated effects of space charge on the distortions in reconstructed track positions in the MicroBooNE TPC in the x - y plane. Corrections are shown for the x , y and z coordinates for a slice of the detector at $z = 10 \text{ cm}$. The colour scale in each plot indicates the reconstructed minus the true coordinate $\text{pos}_{\text{reco}} - \text{pos}_{\text{true}}$ in centimeters.

3.8 Detector Operations

The MicroBooNE detector has been taking data from the BNB and NuMI neutrino beams since the autumn of 2015. The MicroBooNE detector was designed to receive 6.6×10^{20} protons-on-target from the BNB, but as of August 2019, the equivalent of approximately 1.3×10^{21} protons-on-target has been written to tape. The weekly fluxes and delivered and recorded protons-on-target since data taking are shown in Figure 3.16. The data used in this work is Run 1 data, corresponding to the data taking period from February to October 2016. This dataset consists of 1.6×10^{20} protons-on-target in total. At this time, the cosmic ray tagger system had not yet been installed, and as such information from this module is not available in this dataset. In addition to delivering an integrated flux that is well above expectation, MicroBooNE has had a very stable and succesful run. For instance, the argon purity has consistently been very high at $Q_A/Q_C \approx 1$ during the entire data taking period. Figure 3.17 shows the variation of Q_A/Q_C , defined in equation 3.1, for a period of 56 days during the data taking period 16 February 2016 to 21 April 2016, a subset of the Run 1 period.




Fig. 3.16 Delivered and recorded protons-on-target from the Booster Neutrino Beam since the start of the MicroBooNE data taking period. Weekly fluxes are shown in blue (with units on the left axis) and the integrated delivered flux is shown in red (with units on the right axis).




Fig. 3.17 The variation of Q_A/Q_C for a period of 56 days during the data taking period (16 February 2016 to 21 April 2016). The black filled circles indicate points with space charge corrections applied and blue triangles indicate points without space charge corrections. After correction, all points have values ≤ 1 within statistical errors.

Chapter 4

The Pandora Reconstruction Framework

4.1 Introduction

This chapter will discuss the Pandora reconstruction framework. Section 4.2 will discuss the Pandora Software Development Kit (SDK), which describes many of the experiment-agnostic features of the Pandora toolkit, including the event data model (EDM) and the philosophy behind the Pandora reconstruction framework. Section 4.3 will outline how Pandora can be used for reconstructing neutrino-induced and cosmic interactions in the MicroBooNE detector, and Section 4.4 describes many of the algorithms that form the core of the MicroBooNE Pandora reconstruction chain. Finally, Section 4.5 will discuss the performance metrics that are used by the Pandora development team to continually improve the quality of the Pandora reconstruction. These performance metrics are important when discussing how reconstruction improvements outlined in this thesis, such as those in Chapter 6 and 7, impact the reconstruction.

4.2 The Pandora SDK

The problem of reconstructing particle interactions in the context of high energy physics is essentially a pattern recognition problem. The particle interactions within a detector can be described by a limited number of interaction channels, and each of these channels will have a distinct signature. The Pandora SDK provides a framework in which to develop and run pattern recognition algorithms. It consists of a dependency-free C++ library, APIs, and provides an Event Data Model for managing pattern recognition problems. The Application

Programming Interfaces (APIs) define many of the core operations within the Pandora SDK, and allow the user to define the input building blocks that define a pattern recognition problem, implement pattern recognition algorithms, and perform additional operations, such as accessing and modifying the created building blocks.

The Pandora EDM provides a hierarchical framework with which to describe the input and output objects that describe the pattern recognition problem. The Pandora EDM is designed to be self-describing, such that each object provides all the information that is required to interact with other objects in the EDM. A distinction is made within the Pandora EDM between input and algorithm objects. Input objects are immutable representations of the input data, whereas algorithm objects are hierarchical containers of input objects that represent (partial) solutions to the pattern recognition problem. The Pandora SDK monitors the usage of input objects so that each may only appear in a single algorithm object, so that they are not double-counted. The Pandora input objects are outlined below:

- *CaloHit*: designed to represent a calorimetric hit, which is typically the lowest-level input in high energy physics pattern recognition problems. A CaloHit defines a position and extent in space and time, together with an associated intensity or energy measurement;
- *Track*: tracks are continuous trajectories of hits, and store information about particle positions and momenta (track states). Tracks can have parent-daughter and sibling relationships that preserve the hierarchy of the underlying particle interaction.
- *Monte Carlo particle*: Monte Carlo particle objects are primarily used for development, and represent the true state of the simulated particle given to the Pandora instance. As such, they represent hierarchical target solutions for the target pattern recognition problem, and algorithm performance can be assessed using this information.

The Pandora algorithm objects are outlined below:

- *Cluster*: a cluster is a coherent grouping of CaloHit objects, and often represents a candidate particle, or a part or projection thereof. The cluster object stores information about the combined properties of the CaloHit objects that constitute it, such as summed energy estimators, and results of linear fits to the collection of spatial positions.
- *Vertex*: the identification and classification of a specific point in space, vertices are typically used to flag positions of particle creation or decay.

- *ParticleFlowObject*: a Particle Flow Object (PFO) is a hierarchical container consisting of several different types of algorithm objects. Typically, it contains Clusters, Tracks and Vertices, together with metadata describing the particle type and four-momentum. PFOs can also store parent-daughter links in order to describe particle decay hierarchies. The PFO represents the ultimate output of the pattern recognition, grouping the input objects into structures that completely define the solution.

Pandora algorithms form the core of the pattern recognition reconstruction in the Pandora framework. They are designed to be highly modular, and represent a set of pattern-finding instructions. The algorithms use the Pandora APIs to access the Pandora objects. The reconstruction chain is defined by a number of XML settings files, in which algorithms are represented by top-level tags. Within these settings files, the order in which algorithms run can be modified, and their inputs and outputs can be modified. The Pandora framework is general enough to be framework-independent, and has been used to create automated reconstruction solutions for the Compact Linear Collider (CLIC) [83], a high-energy e^+e^- collider, as well as LArTPC experiments such as MicroBooNE and ProtoDUNE [69]. The Pandora SDK is discussed in rich detail in [84].

4.3 Pandora at MicroBooNE

The functionality provided by the Pandora SDK makes it extremely suitable for the reconstruction of particle interactions in the context of LArTPCs, fully utilising the excellent spatial and calorimetric resolution provided by such detectors. In excess of a hundred algorithms and tools have been developed inside Pandora for the reconstruction of cosmic ray muon and neutrino events within LArTPC detectors. A dedicated liquid argon reconstruction framework has been developed at Fermilab, called LARSOF [7, 8]. The Pandora reconstruction is integrated into the LARSOF framework via the `larpandora` translation module, which converts the input data from the LARSOF EDM into the Pandora EDM. In the context of the MicroBooNE reconstruction, the inputs to Pandora are 2D calorimetric hits. The translation module is responsible for controlling the Pandora reconstruction, and converts the output objects from the Pandora EDM back into the LARSOF EDM, after which they can be used by analysers. It also provides detector information in the form of an XML or PNDR file, which includes information about the physical dimensions of the detector, as well as wire pitches, wire angles and potential inactive detector regions. It is also possible to specify multiple drift volumes, which has applications in experiments like DUNE. On a per-event basis, the translation module converts input hits from the LARSOF EDM to the

Pandora EDM, and stores a record of the Monte Carlo particles in the event if the input data is simulated data.

The pattern recognition reconstruction output is typically provided in the form of particle flow objects (PFOs), which represent distinct particle hypotheses, and can be either tracks or showers. PFOs persist links to the 2D clusters of which they consist, which in turn consist of CaloHit objects. Furthermore, they also contain a list of 3D positions, called 3D hits or SpacePoints, and a 3D reconstructed vertex position, representing its interaction point or initial energy deposit in the detector. The PFOs are placed in a hierarchy, which identifies parent-daughter relationships and describes the particle flow in the observed interactions. An illustration of a particle hierarchy is provided in Figure 4.1. In addition, a neutrino particle is created as part of the hierarchy and can form the primary parent particle for a neutrino interaction. The type of each particle is not currently reconstructed, but they are instead identified as track-like or shower-like. Track and shower objects carry additional metadata, such as position and momentum information for tracks or principal-axis information for showers.

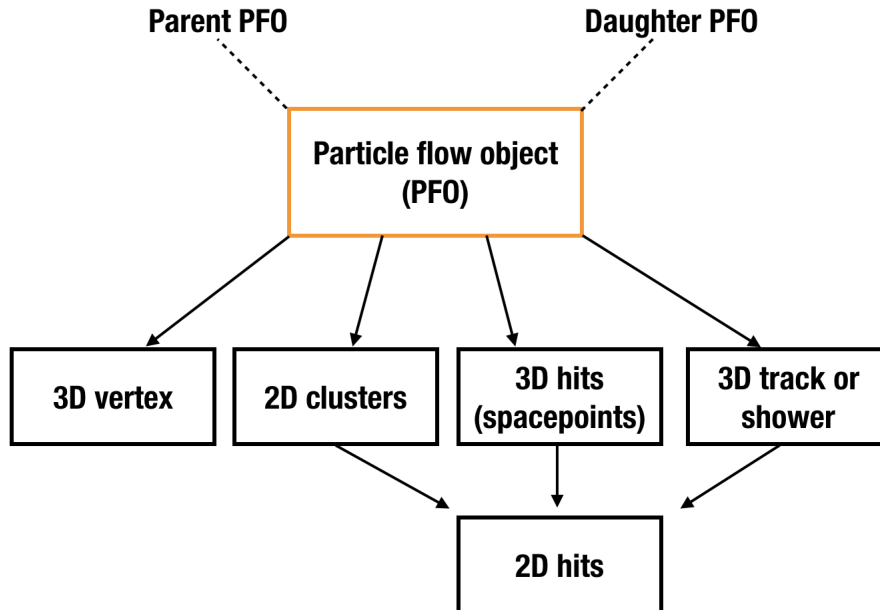


Fig. 4.1 The hierarchy of the Pandora output data products, which are persisted in the LARSOFT Event Data Model. Navigation along PFO hierarchies is indicated by dashed lines. Navigation from each object to their associated objects is indicated by solid arrows. Figure adapted from [85].

For reconstructing MicroBooNE data, the Pandora reconstruction has been separated into two reconstruction paths. Both Pandora realisations are tailored to reconstruct the two main

aspects of a MicroBooNE event: neutrino-induced and cosmic-induced particles. The first reconstruction path is denoted PandoraCosmic, and it is optimised to reconstruct cosmic ray muons and their associated δ -rays. The second path, PandoraNu, is designed to reconstruct neutrino interactions. In contrast to the PandoraCosmic pass, where the vertex of each particle is taken to be its highest-y point, several algorithms are dedicated to reconstructing the neutrino interaction vertex. In addition, neutrino interactions can induce shower-like particles, and in the PandoraNu pass a pre-trained support vector machine attempts to distinguish between the two particle types. Lower-level algorithms are re-used between the two reconstruction paths, and there are also a large number of path-specific algorithms.

The PandoraCosmic and PandoraNu reconstruction chains are applied to MicroBooNE in sequence, where the PandoraCosmic reconstruction is run first on all the hits present in a readout window, providing a list of candidate cosmic ray particles. This list of cosmic ray candidates is examined by a cosmic ray tagging module, which identifies unambiguous cosmic rays and flags the particles being such. This flagging is done based on the entry and exit points of the track in the detector, and whether they are in-time with the beam trigger. After this stage, a new cosmic-removed hit collection is created by subtracting all the cosmic ray hits from the total hit collection. This cosmic-removed hit collection forms the input to PandoraNu. After running part of the PandoraCosmic reconstruction, the input hits are separated into ‘slices’, with the purpose of separating neutrino-induced hits from cosmic ray hits in separate subsets of the detector space. The hit collection in each slice is processed by the remainder of the PandoraNu reconstruction chain, which returns a list of candidate neutrino particles, one for each slice. This double reconstruction chain is illustrated in Figure 4.2.

4.4 LArTPC Reconstruction Algorithms

This section will outline many of the detector-agnostic reconstruction algorithms, such as the 2D cluster creation process, the 3D track reconstruction, 3D hit reconstruction, neutrino interaction vertex reconstruction, shower reconstruction and construction of particle hierarchies within particle flow objects.

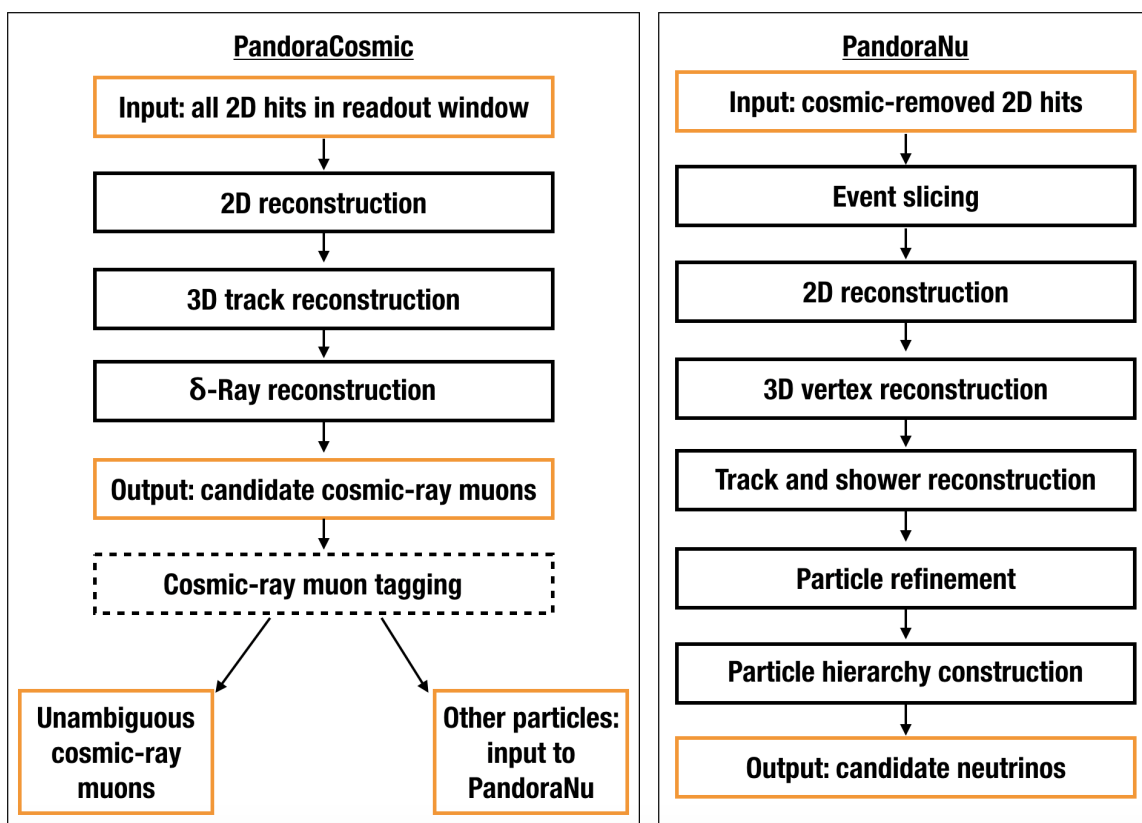


Fig. 4.2 An overview of the two Pandora reconstruction chains in MicroBooNE. The neutrino ID and slicing stages are not depicted in this figure. Inputs and outputs are indicated by orange borders, reconstruction stages with black borders, and the cosmic ray tagging logic by a dashed border. Figure adapted from [85].

4.4.1 Cluster and Hit Creation


The entry point of the reconstruction chain is the reconstruction of 2D clusters. The `EventPreparation` algorithm separates the input hits into three lists, one for each view. For each wire plane, the `TrackClusterCreation` algorithm produces a list of 2D clusters that represent continuous, unambiguous lines of hits. The `TrackClusterCreation` algorithm aims to identify coherent groupings of hits, and a new cluster is created any time it encounters a bifurcation or ambiguous feature. This process results in a large amount of small clusters with high ‘purity’, i.e. clusters for which most hits belong to the same underlying particle and low ‘completeness’: most clusters will contain only a small fraction of the hits deposited by each distinct particle.

The next step is for cluster-merging algorithms to identify and merge clusters in the 2D views that are likely related. Cluster associations are identified by spatial proximity, and by finding clusters for which the directions of their principal axes align significantly. Figure 4.3 shows the two stages of cluster creation in an example event containing simulated cosmic ray muons. The final 2D clusters that result from the application of a series of topological algorithms form the input to the next stage of the reconstruction, which aims to match cluster features between readout planes, in order to construct 3D particles.

4.4.2 Three-dimensional track reconstruction

The aim of the 3D track reconstruction is to identify 2D clusters from separate readout planes that belong to individual track-like particles. The 3D track reconstruction is performed primarily by the `ThreeDTransverseTracks` algorithm, which quantifies the suitability of all cross-plane combinations of clusters and stores the results in a 3D tensor. The tensor indices are the clusters in the U, V and W views and a custom `TransverseOverlapResult` object is stored for each combination, which contains information about the matching process. The information in the tensor represents all possible cluster combinations, and the individual tensor entries are used to identify appropriate and inappropriate cluster matches. In addition, 2D clusters can be split or combined iteratively, until the tensor for the resulting 2D clusters contains the smallest degree of matching ambiguity.

The construction of `TransverseOverlapResult` objects relies on the shared x (drift time) coordinate between views. First, a number of sampling points are defined in x that is common to all three clusters. Then, 2D sliding fit objects are constructed for each cluster:



(a) High-purity clusters, which are truncated at each ambiguous point, such as the threshold of an inactive region.

(b) Refined and merged clusters formed by the series of topological algorithms, combined to form one cluster for each cosmic ray muon track, indicated in red and purple.


Fig. 4.3 A comparison of the high-purity clusters created by the initial stage of the track cluster creation logic (the `TrackClusterCreation` algorithm), which creates a new cluster at every ambiguous point (Figure 4.3a) and the final output merged clusters, which have both high purity and high completeness (Figure 4.3b). Each pair of horizontal grey lines corresponds to an inactive detector region, and each coloured track corresponds to a separately reconstructed cluster of hits. The x coordinate is in the horizontal direction, the z coordinate in the vertical direction. Modified from [85].

- *Pandora sliding fits*: a Pandora sliding fit object is created from a Pandora cluster by sorting the cluster hits by their z coordinate, and defining a sliding fit window. The sliding fit window (in cm) is defined in the z direction, and partitions the cluster into local regions. For each sliding fit window, a linear least-squares fit is created to the hits in the local cluster region, and each linear fit shares its endpoints with its neighbours. The sliding fit object is the total collection of linear fits. Pandora 2D and 3D sliding fits are capable of capturing track curvature, and are able to provide accurate longitudinal coordinates along the fit for each hit.


Sliding fit positions at a sampling point in x in two views can then be projected into the third view, and the predicted position can be compared to the actual position at the sampling position in this view. By comparing all combinations ($U, V \rightarrow W$; $V, W \rightarrow U$; $U, W \rightarrow V$), a χ^2 -like quantity can be calculated for each tensor entry. This χ^2 -like value, together with the common x -overlap span, the number of sampling points and the number of consistent sampling points, is stored in the `TransverseOverlapResult` in the tensor.

The matching tensor can be ‘traversed’, starting from a given cluster, to obtain a complete set of connected clusters. In addition, the tensor is queried by a number of algorithm tools. These algorithm tools can create new particles or modify the 2D pattern recognition based on the information contained in the `TransverseOverlapResult` objects. These algorithm tools are run in a fixed order, and if any of them make a change, the entire series is run again from the beginning. This iterative process is repeated until none of the algorithm tools make any changes to the existing reconstruction. The tools, in the order in which they are executed, are listed below:


- `ClearTracks` tool: this tool creates particles from three clusters (one in each view) that represent a ‘clear’ track, which in this context means that the cluster grouping contains only one cluster from each view. An example of three clusters that constitute a clear track is given in Figure 4.4a;
- `LongTracks` tool: this tool is designed to resolve obvious matching ambiguities. An example is given in Figure 4.4b, where the presence of two δ -ray clusters introduces multiple possible pairings, though one combination is the obvious best match;
- `OvershootTracks` tool: the topology this tool aims to address is a cluster matching of the form 1:2:2, in which a cluster grouping contains one cluster in one view, and two clusters in the other two views. In other words, there is a single long cluster that should be split into two new clusters. An example is given in 4.4c;




(a) Matching of the type 1:1:1 (U:V:W).



(b) Matching of the type 1:2:2 (U:V:W).



(c) Matching of the type 1:2:2 (U:V:W).



(d) Matching of the type 1:2:1 (U:V:W).

Fig. 4.4 Example topologies that are correctly resolved into unambiguous groupings of 2D clusters by the 3D track reconstruction. Four matching situations are shown, containing: (a) one cluster per view, (b) two small δ -ray clusters (circled), (c) a projection of two particles that appears to be one cluster in the U view and (d) a cluster erroneously split in the V view. The x coordinate is in the horizontal direction, the z coordinate in the vertical direction. Modified from [85].

- `UndershootTracks` tool: this tool does the opposite of the `OvershootTracks` tool, and identifies two clusters that have likely been erroneously split. As such, it addresses topologies of the form 1:2:1, such as in 4.4d;
- `MissingTracks` tool: this tool accounts for the fact that particle features may be obscured in one view due to the 3D particles being projected into 2D, such that a single cluster represents multiple overlapping particles. If this tool identifies appropriate cluster overlap, using the cluster-relationship information available from the tensor, the tool can create particles using only two clusters.
- Finally, the `TrackSplitting` and `MissingTrackSegment` tools split tracks if necessary and can add hits to tracks to extend them, respectively. As a final step, the `LongTracks` tool is run again with a reduced matching threshold, in an attempt to match all clusters in the event.

After the 3D track reconstruction, the `PandoraCosmic` pass dissolves any cluster that has not been associated to a particle, since such clusters likely represent fragments of δ -ray showers. These hits are reclustered, and the `DeltaRayMatching` algorithm is responsible for matching δ -ray clusters between views to create shower-like δ -ray particles. These δ -ray particles are then tagged as daughters to an appropriate parent track. Having processed all the clusters in the event, 3D hits can be reconstructed from each 3D cluster. The construction of 3D hits proceeds differently depending on the cluster topology. The two most important particle types for 3D hit reconstruction are transverse tracks, with significant extent in the x coordinate, and shower-like objects:

- When there is one transverse track in each view, 2D sliding fit objects are constructed for each particle, and for each 2D hit in each view, the sliding linear fit positions are evaluated in the other two views (using the shared x coordinate). An analytic χ^2 minimisation procedure is then applied, which extracts the optimal y and z coordinates at the given x coordinate;
- In the shower-like case, an input 2D hit is selected in one of the views, and all combinations of 2D hits are considered in a narrow band in x in the other two views. For a given combination of hit U , V and W values, the most appropriate y and z coordinates are again selected by a χ^2 minimisation procedure. An upper bound is applied to the χ^2 value, to ensure that the 3D hits that result from this procedure are maintain spatial coherence.

4.4.3 Vertex Reconstruction

After the 3D hits have been created, the PandoraCosmic reconstruction is completed by placing vertices at the high- y endpoints of each cosmic ray muon particle. Vertices are also created for δ -ray particles, and are placed at the 3D point of closest approach between the parent cosmic ray muon and daughter δ -ray. In neutrino-induced interactions, correctly reconstructing the neutrino interaction vertex is a crucial part of the reconstruction, as the vertex position is used to split or merge clusters downstream in the reconstruction chain. The creation of a 3D vertex proceeds by two stages: the first step is the creation of a multitude of candidate vertex positions, and the second step is the subsequent scoring of the vertex candidates, after which a pre-trained support vector machine (SVM) selects a single candidate as the neutrino interaction vertex.

The vertex candidates are created by the `CandidateVertexCreation` algorithm, which compares 2D clusters from different readout planes that have an above-threshold overlap in the x coordinate. Two-dimensional sliding fits are created for every 2D cluster in the event, and for every cluster endpoint the position corresponding to the same x coordinate as the endpoint is identified in another view, allowing for some extrapolation of the sliding fits beyond the cluster boundaries. The two cluster positions are then projected into three dimensions, yielding a candidate vertex position. This process is repeated for every cluster endpoint in every view, such that four candidate vertices are created for each cluster pairing. The total amount of vertex candidates is reduced by applying a quality cut to each candidate, which requires that candidates are on or near a hit, or in an inactive detector region, in all three views.

Once the vertex candidates have been created, a collection of vertex selection algorithms aim to select the vertex candidate that is most likely to represent the neutrino interaction vertex. The selection process is based on a number of distinct scores that are formulated for each vertex candidate, which are enumerated below:

1. $S_{energykick}$, the energy kick score: this component of the score attempts to incorporate and quantify the observation that primary neutrino daughters point back towards the true interaction vertex and have higher energies than secondary daughters, such as particle decays or recoils. The score is formulated by projecting each 3D vertex candidate into the U, V and W views. A parameter $E_{ij}^{T'}$ is then calculated to assess whether the candidate is consistent with observed cluster j in view i . This parameter is closely related to the transverse energy, but has additional degrees of freedom that

introduce a dependence on the displacement between the cluster and vertex projection. The score S and the quantity $E_{ij}^{T'}$ are defined as

$$S_{energykick} = \exp \left\{ - \sum_{view\ i} \sum_{cluster\ j} \frac{E_{ij}^{T'}}{\varepsilon} \right\} \quad (4.1)$$

$$E_{ij}^{T'} = \frac{E_j \times (x_{ij} + \delta_x)}{d_{ij} + \delta_d}, \quad (4.2)$$

where x_{ij} is the transverse impact parameter between the vertex and a linear fit to cluster j in view i , d_{ij} is the closest distance between the vertex and cluster and E_j is the cluster energy. The parameters ε , δ_d and δ_x are tunable constants: ε determines the relative importance of the energy kick score, δ_d is necessary because d_{ij} may be zero and controls weighting as a function of d_{ij} , and δ_x controls weighting as a function of x_{ij} .

2. $S_{asymmetry}$, the asymmetry score: this component of the score suppresses candidates that are incorrectly placed along single, straight clusters. This score is calculated by counting the numbers of hits upstream $N_i \uparrow$ and downstream $N_i \downarrow$ of the candidate position, which are identified by constructing a 2D principal axis in each view. For the true vertex a large asymmetry is expected, since the hits emerge from a single point in the detector. A fractional asymmetry, denoted A_i for view i , is then calculated for each vertex candidate:

$$S_{asymmetry} = \exp \left\{ \sum_{view\ i} \frac{A_i}{\alpha} \right\} \quad (4.3)$$

$$A_i = \frac{|N_i \uparrow - N_i \downarrow|}{N_i \uparrow + N_i \downarrow}, \quad (4.4)$$

where α is a constant that determines the relative importance of the asymmetry score.

3. $S_{beamdeweight}$, the beam deweighting score: this score incorporates information about the beam direction into the vertex score. Since most neutrino interaction vertices ought to be near the most upstream (low- z) hits in the event, vertex candidates can be deweighted when they have high z values. The beam deweighting contribution to the score is defined as

$$S_{beamdeweight} = \exp \{ -Z/\zeta \}, \quad (4.5)$$

where $Z \equiv \frac{z - z_{min}}{z_{max} - z_{min}}$ and ζ is a tunable parameter that determines the relative importance of the beam deweighting score.

Figure 4.5 shows a number of candidate vertex projections in the W view and their assigned scores, including information about each of the components of the score S . This figure demonstrates that the score associated to the selected vertex candidate, labelled E , has high contributions from $S_{energykick}$, $S_{asymmetry}$ and $S_{beamdeweight}$. These three scores are used to train a support vector machine (SVM), which performs two vertex selection tasks:

- *Identifying the correct vertex region*: the vertex candidate creation process produces a large number of vertex candidates, which are clustered together near regions of interest within the event topology. The SVM is trained using aggregate information from the vertex candidate clusters, such as their scores $\{S\}$ and topological information, to identify the cluster of vertex candidates that is most likely to be near the true neutrino interaction vertex;
- *Identifying the correct vertex candidate*: once the best vertex region has been selected, the next selection task performed by the SVM is to select the best vertex candidate from the selected candidate cluster. This training task once again utilises the vertex scores, as well as hit-level topological information.

The trained SVM is saved to an XML file, and is used by the `VertexSelection` algorithm to select the most probable neutrino interaction vertex.

4.4.4 Shower reconstruction

In neutrino-induced interactions, electromagnetic showers are present in some event topologies, induced by electrons and photons. The PandoraNu pass, therefore, also attempts to reconstruct electromagnetic showers. An example of an event topology containing an electromagnetic shower is shown in Figure 4.6. The 2D shower reconstruction starts by adding branches to any long clusters, which are labelled as ‘shower spines’. The 2D reconstruction of electromagnetic showers proceeds along the following steps:

- The 2D clusters are categorised as being either track-like or shower-like by a pre-trained support vector machine, based on a number of variables such as cluster length, variations in sliding fit direction along and the extent of the cluster transverse to the primary linear fit axis;




Fig. 4.5 An example W view projection of an event containing three particles, each indicated by a separate colour, in which five vertex candidates are shown. The table in the top right of the figure shows the S , $S_{energykick}$, $S_{asymmetry}$, $S_{beamdweight}$ scores for each of the candidates. Candidate E is eventually selected, since it has the highest overall score S . Details of the vertex scoring procedure are given in the text. Figure from [85].

- If an existing track particle is re-labelled as being shower-like, it is dissolved into its constituent clusters, so that a shower-like particle may be constructed from them;
- Among these newly-dissolved clusters, long 2D clusters are identified as potential shower spines. In addition to cluster length, the fact that shower spines typically point back towards the projected neutrino interaction vertex is used to identify such clusters;
- When a shower spine has been identified, shorter clusters are added onto the spine as branches by the `ShowerGrowing` algorithm. This algorithm operates recursively, first finding the most likely primary branches to add onto the shower spine, and afterwards adding branches onto other branches. For every branch, the algorithm records the ‘strength of association’ to each spine, which is a measure of how likely it is to represent a correct shower branch. Branch addition decisions are then made in the context of the overall event topology.

After the 2D shower reconstruction, the 2D shower-like clusters are matched between views to construct 3D shower particles, using logic similar to the 3D track reconstruction described in Section 4.4.2. A rank-three tensor is constructed by the `ThreeDShowers` algorithm, which stores cluster-overlap and relationship information in `ShowerOverlap` objects. The tensor is ‘diagonalised’ by multiple sequential algorithms, thereby minimising matching ambiguities. The 2D `ShowerOverlapResult` are formulated by constructing 2D envelopes from the edges of shower-like clusters in two views and used to predict the envelope in the third view. The fraction of hits in the third cluster contained within the envelope is then stored, alongside details of the common cluster x -overlap.

A number of algorithms run sequentially and query the tensor to construct combinations of 2D clusters, operating along similar lines as the 3D track reconstruction. After the 3D showers have been refined, the `ParticleCharacterisation` algorithm classifies each particle as being either track-like or shower-like. This algorithm uses an SVM-based approach and uses various topological and calorimetric features to estimate the probability that a particle is track-like or shower-like.

4.4.5 Particle Refinement & Hierarchy

A further particle refinement stage is included in the reconstruction chain, which finalises the assignment of hits to particles after identifying and addressing potential problems with constructed particles. The primary issue to address is often the completeness of sparse




Fig. 4.6 An example projection in the W view of an electromagnetic shower. The red track-like clusters are not affected by the shower reconstruction. The long blue clusters that point towards the interaction vertex are identified as shower spines, after which the green, smaller clusters are added onto the spines iteratively. This procedure yields a shower-like cluster of high completeness. Figure from [85].

showers. Single sparse showers are often erroneously represented as multiple, distinct shower-like particles. A number of algorithms have been designed to address this specific failure mode:

- The `ClusterMopUp` algorithms aim to grow existing 2D shower-like particles by incorporating remaining, unassociated 2D clusters. It does so by considering the 2D clusters that have been included in shower-like particles, using parameterisations of the 2D cluster extents, including cone fits and sliding linear fits to the edges of the showers. If unassociated clusters are bounded by the assigned clusters, or are in close proximity to them, they are incorporated into the appropriate shower-like particle;
- The `SlidingConeParticleMopUp` algorithm constructs 3D cones for 3D shower-like particles and incorporates particle fragments that fall within the cone into the appropriate existing particle. The local 3D cone axes and apices are constructed from particles topologies, and cone opening angles can be specified as algorithm parameters, or derived from the topology of the 3D shower hits. The cones are then extrapolated in order to assess which particle fragments are encompassed by them;

- The `SlidingConeClusterMopUp` algorithm projects the 3D cones created by the `SlidingConeParticleMopUp` algorithm down into each of the three views, and uses the resulting 2D cones to incorporate unassociated 2D clusters into existing particles;
- The `IsolatedClusterMopUp` algorithm dissolves any remaining unassociated 2D clusters and looks to add their hits to nearby shower-like particles.

Having reconstructed track-like and shower-like particles in the event, a particle hierarchy is constructed for each detector slice. The particle hierarchy is a tree-like object with a neutrino candidate ν_{reco} at its root, which initially only contains the reconstructed neutrino vertex. Subsequently, each particle is flagged either as a primary daughter of the reconstructed neutrino, or as a daughter of another particle based on the proximity of their 3D hits to the interaction vertex. A number of algorithm tools seek to add unassigned particles as daughters to the primary daughters of the neutrino, and depending on whether the largest primary daughter is track-like or shower-like, the neutrino particle is labelled as being a ν_μ or ν_e . Finally, the points of closest approach for each particle are added as 3D vertex positions in the neutrino hierarchy. An illustration of a particle hierarchy is given in Figure 4.7.

4.5 Performance

Improvements to the Pandora reconstruction are motivated by a set of well-defined performance metrics. These metrics are based on the matching of hits between the underlying Monte Carlo truth particles and the reconstructed particles created by Pandora. To construct these metrics, primary final states in the neutrino interaction (e^\pm , μ^\pm , γ , K^\pm , π^\pm or p) are identified as targets for the reconstruction and each 2D hit is matched to the Monte Carlo particle that deposited the majority of its charge. The totality of these hits are referred to as the ‘true hits’ associated to the Monte Carlo particle and hits associated to downstream (daughter) Monte Carlo particles in the Monte Carlo particle hierarchy are associated to the relevant primary Monte Carlo particle. For instance, the hits associated to the Monte Carlo particle resulting from a particle decay are flagged as belonging to the true hit collection of the parent particle.

Not all Monte Carlo particles are considered as reconstructable to prevent obfuscating the performance metrics. Examples include Monte Carlo particles that do not deposit sufficient true hits, or Monte Carlo particles that deposit diffuse hits scattered throughout the detector. As such, hits are neglected in the performance evaluation if the hierarchy shows they are




Fig. 4.7 An example particle hierarchy in the W view associated to a CC ν_μ event with a muon, proton and pion in the final state. The neutrino particle has three track-like daughter particles. The first is a long proton (blue), which recoils off a nucleus which results in a track-like secondary daughter particle (grey). The second is a π^+ (magenta), which decays into a μ^+ , which rapidly decays into a visible e^+ , which is reconstructed as a shower-like particle (purple). The third neutrino daughter is a muon track (red). Figure from [85].

associated to Monte Carlo particles downstream of a far-travelling neutron, or if the primary Monte Carlo particle is a far-travelling photon. Furthermore, target Monte Carlo particles are only considered reconstructable if they are associated to at least 15 hits passing the selection, including at least five hits in at least two views. When counting hits associated with a target Monte Carlo particle, this particle must be responsible for at least 90% of the true energy deposition recorded for the hit. This selection corresponds to true momentum thresholds of approximately 60 MeV/c for muons and 250 MeV/c for protons in the MicroBooNE simulation.

After determining the true hit collection for each Monte Carlo particle, the reconstructed particles are matched to the Monte Carlo particles. The number of hits shared between each target Monte Carlo particle and each reconstructed particle is stored in an association matrix. Analogously to what was done in the case of Monte Carlo particles, all hits belonging to daughters of reconstructed particles are associated to the hit collection of the visible primary parent particles. Now that hit collections have been defined for both Monte Carlo and reconstructed particles, the following quantities represent useful performance metrics on a per-particle basis. For a pairing between a Monte Carlo particle with true hits $\{T\}$ and a reconstructed particle with reconstructed hits $\{R\}$:

- Efficiency η : for a type of target Monte Carlo particle, the efficiency is the fraction of such target Monte Carlo particles with at least one matched reconstructed particle;
- Completeness $C = |T \cap R| / |T|$: for a given pairing of reconstructed particle and target Monte Carlo particle, is the fraction of the Monte Carlo particle true hits that are shared with the reconstructed particle;
- Purity $P = |T \cap R| / |R|$: for a given pairing of reconstructed particle and target Monte Carlo particle, is the fraction of hits in the reconstructed particle that are shared with the target Monte Carlo particle.

The matching procedure below is used to provide a final, human interpretation of the reconstruction output:

- First, identify the match between an available Monte Carlo particle and an available reconstructed particle that contains the largest number of shared hits. After a match has been made, flag both the Monte Carlo particle and reconstructed particle as unavailable, meaning they can not participate in another match;

- This procedure is repeated until no further matches are possible;
- Any reconstructed particles that have not been matched are matched to the target Monte Carlo particle with which they share most hits, even if that Monte Carlo particle already has reconstructed particles associated to it. Consequently, the number of reconstructed particles matched to a target Monte Carlo particle can become larger than one if it already has a matched particle, but can never increase if the number of matches is zero when the iterative matching procedure has concluded.

This procedure yields a list of primary target Monte Carlo particles with their associated reconstructed particles. An event is deemed to be correctly reconstructed if there is exactly one reconstructed particle for each target Monte Carlo particle at the end of this procedure. The fraction of correct events within an event sample is a useful and very sensitive metric through which the pattern-recognition performance may be assessed. The PandoraNu reconstruction performance is characterised separately for each neutrino interaction type and a selection of possible final states in the MicroBooNE detector simulation. Furthermore, only neutrino interactions inside the fiducial volume of the LArTPC are considered, which is defined to be the active volume excluding the regions within 10 cm of the detector edges in x and z , and within 20 cm of the edges in y .

An example $\nu_\mu + \text{Ar} \rightarrow \mu^- + p$ is shown in Figure 4.8. There are two primary Monte Carlo particles in this event, and two reconstructed particles are output by Pandora. Both Monte Carlo particles have exactly one reconstructed particle associated to them, and as such this event is deemed to be correct. The number of matches for each primary Monte Carlo particle can be quantified for each event in a simulated sample to quantify reconstruction performance. The main analysis sample that is used in this work contains neutrino interactions simulated using the MicroBooNE default GENIE configuration and contains cosmic rays simulated by CORSIKA. Details of the simulation packages are given in Section 3.7. The notation used to indicate interaction types such as CCQE μ^- is supplied in Section 7.1. The number of matched particles and the fraction of correct events for charged-current interactions channels producing $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ final states are given in Table 4.1. These performance metrics form the reconstruction baseline that is improved upon by the new methods outlined in chapters 6. Table 4.1 shows that the fraction of correct events for most charged current $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ final state channels lie between 60% and 70%. Matches of 0 particles to a Monte Carlo particle are mainly due to cosmic background events being identified as the reconstructed neutrino ν_{reco} , but can also be due to particles being difficult to resolve, e.g. in the case of

short protons. Events in which the reconstructed particle multiplicity is larger than the true multiplicity are caused by tracks being broken up by inactive detector regions, or by daughter particles of primaries being flagged as primaries themselves.




Fig. 4.8 Example reconstruction of a simulated 500-MeV CC ν_μ quasi-elastic event. The event final state is $\mu^- + p$. Simulated unresponsive channels are indicated by the horizontal grey lines, which cause the break in the proton track. Figure from [85].

<i>Number of Matches and Fraction Correct Per Channel</i>						
Interaction Type	Total	Number Correct	0 Matches	1 Match	2 Matches	3+ Matches
CCQEL_MU	25892	17071 (65.93%)				
μ^-			8093 (31.26%)	17092 (66.01%)	686 (2.65%)	15 (0.06%)
CCQEL_MU_P	29720	19228 (64.70%)				
μ^-			7226 (24.31%)	22010 (74.06%)	464 (1.56%)	8 (0.03%)
p			9046 (30.44%)	19761 (66.49%)	860 (2.89%)	41 (0.14%)
CCRES_MU	1048	703 (67.08%)				
μ^-			317 (30.25%)	704 (67.18%)	26 (2.48%)	2 (0.19%)
CCRES_MU_P	3162	2056 (65.02%)				
μ^-			663 (20.87%)	2450 (77.48%)	60 (1.90%)	0 (0.00%)
p			917 (29.00%)	2140 (67.68%)	103 (3.26%)	4 (0.13%)
CCMEC_MU	5697	3800 (66.70%)				
μ^-			1685 (29.58%)	3804 (66.77%)	204 (3.58%)	2 (0.35%)
CCMEC_MU_P	13270	9133 (68.82%)				
μ^-			2655 (20.01%)	10424 (78.55%)	190 (1.43%)	4 (0.03%)
p			3618 (27.26%)	9365 (70.57%)	283 (2.13%)	7 (0.05%)

Table 4.1 Reconstruction performance for charged-current channels resulting in μ^- and $\mu^- + p$ final states. The table displays the number of reconstructed particles matched to each primary Monte Carlo particle per interaction channel. A channel is only deemed correct if all primary Monte Carlo particles have the correct number of matches. When a cosmic-ray particle has been matched to a target Monte Carlo particle, this ends up in the ‘0 Matches’ category.

Chapter 5

Direction Finding

This chapter will outline a method for finding the direction of tracks in LArTPC detectors using the Pandora reconstruction output. The fitting method works by finding a $\langle -dE/dx \rangle$ curve that best describes the observed hit charge distribution of a track. The value of mean particle energy loss $\langle -dE/dx \rangle$ is accurately described by the Bethe equation, which is discussed in Section 5.1. Section 5.2 discusses the fitting procedure, and describes how a χ^2 value is formulated as a measure of goodness of fit. To improve the performance of the direction fitting procedure, a hit filtering procedure is applied before fitting, which is described in Section 5.3, and the impact of this filtering procedure on the quality of the direction fit is quantified. Sections 5.4 and 5.5 analyse the fitting hypotheses that a cluster consists of a single particle, or two distinct particles, respectively. In particular, Section 5.5 describes a method of identifying two separate particles that have been merged into a single cluster based on improving the fit χ^2 value. Finally, Section 5.6 explains how the distribution of χ^2 differences between forwards and backwards fits, $\Delta\chi_{FB}^2/N$, can be used to formulate a probability that a track-like particle is travelling in the forwards direction along the z axis. Section 5.6.1 briefly mentions potential applications of the direction fitting method. The study in this chapter is based on a simulated sample of quasi-elastic $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ events, where X is the final-state nucleus.

5.1 The Bethe Equation

Observed hit charge values Q in LArTPC detectors are quantified response signals to drift electrons collected on the wire planes of the detector. As such, these hit charge values are manifestations of the mean energy loss of the ionising particle. The Bethe equation [86]

provides an accurate prediction for the mean energy loss (also called mean stopping power) of a particle $\langle -dE/dx \rangle$ as it traverses a medium, as a function of particle momentum

$$\left\langle -\frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} - \frac{C}{Z} \right], \quad (5.1)$$

where β and γ have their usual relativistic definitions. For a particle with mass M and momentum $M\beta\gamma$, the quantity W_{max} is the maximum kinetic energy which can be imparted to a free electron in a single collision, and is given by

$$W_{max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2}. \quad (5.2)$$

The mean energy loss of a particle traversing matter changes with its momentum $\beta\gamma$. As an example of energy loss in a dense medium, the different momentum regimes and the behaviour of $\langle -dE/dx \rangle$ in copper is shown in Figure 5.1. This figure describes energy loss in copper and not liquid argon, but the behaviour of $\langle -dE/dx \rangle$ as a function of $\beta\gamma$ is similar in different media. Equation 5.1 accurately describes the true mean energy loss in the Bethe region $10^{-1} \lesssim \beta\gamma \lesssim 10^3$, where the mean energy loss is a function of β alone. The Bethe region is indicated in Figure 5.1 by the second and third blue bands from the left. The mass dependence and behaviour of $\langle -dE/dx \rangle$ is more complex in other momentum regions, the description of which falls outside of the scope of this discussion. It is assumed that in the context of LArTPC reconstruction, the particles are always in the Bethe region. The real feature of interest in Figure 5.1 is the rise in $\langle -dE/dx \rangle$ as $\beta\gamma \rightarrow 10^{-1}$, which is called the Bragg peak. Since the Bragg peak occurs at low particle energies, at the end of their trajectories in the detector, it can be used to determine the track direction for a charged particle that is contained within the LArTPC. As such, identifying the Bragg peak forms the foundation of the fitting procedure discussed in Section 5.2.

Besides depending on the particle momentum $\beta\gamma$, equation 5.1 also exhibits a dependence on Z/A , where Z is the atomic number of the absorber and A is the atomic mass of the absorber in g/mol. Example $\langle -dE/dx \rangle$ curves are shown for different absorbers in Figure 5.2. Typically, Z increases more slowly than A across the periodic table, resulting in a suppression of $\langle -dE/dx \rangle$ for heavier elements. In Figure 5.2, the $\langle -dE/dx \rangle$ curve for argon would be situated in between the curves for iron ($Z = 26$) and aluminium ($Z = 13$), and would have a minimum energy loss value of 2.1 MeV/cm [87]. The correction terms δ and C in equation 5.1 are the density and shell corrections, respectively. The shell correction is important for low momenta outside of the Bethe region, and will henceforth be neglected by setting $C = 0$. As the particle momentum $\beta\gamma$ increases, the density effect correction $\delta(\beta\gamma)$




Fig. 5.1 $\langle -dE/dx \rangle$ for μ^+ traversing copper, as a function of $\beta\gamma = p/Mc$. The corresponding momentum scale is also indicated. The two rightmost blue bands indicate the Bethe regime, in which the Bethe formula is an excellent approximation of $\langle -dE/dx \rangle$, which is the region with which this discussion concerns itself [88].

becomes relevant. As the particle energy increases, its electric field flattens and extends, so that the distant-collision contribution increases as $\ln(\beta\gamma)$. However, the electric field also tends to polarise atoms along its path. Due to this polarisation, electrons far from the incoming particle are effectively shielded from the full electric field and collisions with these shielded electrons will contribute less to the total energy loss than predicted by the Bethe equation. The density effect increases with higher density, since the polarisation is larger in condensed materials. This polarisation effect limits the field extension and effectively truncates the part of the logarithmic rise of $\langle -dE/dx \rangle$ as $\beta\gamma \rightarrow 10^3$. The density effect correction is usually computed using Sternheimer's parameterisation [89]:

$$\delta(\beta\gamma) \Rightarrow \begin{cases} 2\log_{10}(x) - \bar{C} & \text{if } x \geq x_1 \\ 2\log_{10}(x) - \bar{C} + a(x_1 - x)^k & \text{if } x_0 \leq x < x_1 \\ 0 & \text{if } x < x_0 \text{ (nonconductors)} \\ \delta_0 10^{2(x-x_0)} & \text{if } x < x_0 \text{ (conductors)}, \end{cases}$$

where $x \equiv \log_{10}(p/Mc)$ and x_0 is the value of x which corresponds to the momentum below which $\delta = 0$. The quantity x_1 corresponds to the momentum above which the relation between δ and x can be considered to be linear. In equation 5.3, liquid argon is treated as a




Fig. 5.2 Mean energy loss rate in liquid (bubble chamber) hydrogen, gaseous helium, carbon, aluminium, iron, tin, and lead. Radiative effects, relevant for muons and pions, are not included. These become significant for muons in iron for $\beta\gamma \geq 1000$, and at lower momenta for muons in higher-Z absorbers [88].

conductor. The values of x_0 , x_1 , k , \bar{C} and δ_0 are medium-dependent and are taken from [87]. Figure 5.1 shows the effect of including the density correction, and it can be seen that $\delta(\beta\gamma)$ has a non-negligible effect in the Bethe region, and is therefore included in this treatment.

5.2 The Fitting Procedure

Equation 5.1 describes the relation between particle momentum p and mean energy loss $\langle -dE/dx \rangle$. Integrating this equation over all possible particle energies yields the ‘continuous slowing-down approximation’ (CSDA) range R for a particle which loses energy only through ionisation and atomic excitation, such that

$$\Delta R = \int_{E_i}^{E_e} -\frac{dx}{dE} dE. \quad (5.4)$$


In this approximation, the rate of energy loss at every point along the track is assumed to be equal to the total stopping power, and fluctuations in the energy loss are neglected: $-dE/dx \approx \langle -dE/dx \rangle$. The CSDA range R is a useful quantity that relates the distance traversed by the particle in the medium to its momentum p and hence to its kinetic energy T , through the relationship

$$p = \sqrt{T^2 + 2TM}, \quad (5.5)$$


where M is the particle mass. Conversely, it is also possible to integrate $\langle -dE/dx \rangle$ to obtain the kinetic energy loss of the particle, under the same assumptions:

$$\Delta E = \int_{R_i}^{R_e} -\frac{dE}{dx} dx. \quad (5.6)$$

Equation 5.6 can also be used to numerically integrate the energy loss of a particle with a certain starting energy T_i , which results in a curve of T as a function of R . An example is shown in Figure 5.3a, in which the energy loss of a muon in LAr with $T_i = 700$ MeV was numerically integrated over R , with a range step size $dR = 0.1$ cm. The Bragg peak is evident in Figure 5.3a as accelerated energy loss for lower particle energies, which manifests as visible curvature in the region where $T \leq 20$ MeV. Figure 5.3b plots the magnitude of $\langle -dE/dx \rangle$ for the range $0 \leq T \leq 700$ MeV, which increases asymptotically as the particle kinetic energy approaches zero.



(a) Simulated particle kinetic energy T with CSDA range R .



(b) Simulated particle mean energy loss $\langle -dE/dx \rangle$ with CSDA range R .

Fig. 5.3 Simulated particle kinetic energy T (5.3a) and $\langle -dE/dx \rangle$ (5.3b) as a function of CSDA range R for a muon in LAr with $T_i = 700$ MeV. This plot was generated by iteratively sampling the particle energy as given by the Bethe equation, with a step size dR of 0.1 cm. For low particle energies the Bragg peak becomes apparent as a faster energy loss.

A simulated energy loss curve like the one in Figure 5.3a can be constructed for any starting kinetic energy T_i . If the step size dR is set to a small enough value, this will result in a highly accurate mapping \mathcal{M} between T and R , which is treated like an operator, such that $\mathcal{M}(T) = R$ and $\mathcal{M}(R) = T$. Consider a track observed in the MicroBooNE detector with a certain amount of kinetic energy at its reconstructed endpoint T_e , where $T_e \geq 0$ because the track may not be contained, or because its true endpoint may be obscured by an inactive detector region. The mapping \mathcal{M} will then uniquely identify a value $R_e = \mathcal{M}(T_e)$, the value of R at the track endpoint. If the total distance traversed by the particle ΔR is known, this uniquely identifies the initial range value $R_i \equiv R_e - \Delta R$. The mapping \mathcal{M} then also determines $T_i = \mathcal{M}(R_i)$, which defines the total kinetic energy deposited by the particle in the detector medium $\Delta T \equiv T_i - T_e$. This sequential mapping procedure is illustrated in Figure 5.4. In short, for an observed track, the following two parameters uniquely define ΔT , ΔR and the kinetic energy of the particle at any point along the track T_k :

- T_e : the kinetic energy of the particle at its reconstructed endpoint;
- ΔR : the 3D distance traversed by the particle between its observed endpoints, under the CSDA approximation.

The outlined mapping procedure can be applied to simulated or real MicroBooNE tracks by constructing a Pandora sliding linear fit (defined in Chapter 4) for the relevant clusters. The sliding fit object will provide a longitudinal coordinate L_k for every hit in a track cluster,

in order of ascending z . The mapping \mathcal{M} makes it possible to map every longitudinal position L_k to a kinetic energy value T_k , given T_e and ΔR . Equations 5.1 and 5.5 then also define $\langle -dE/dx \rangle_k$, the mean particle energy loss at every point L_k .




Fig. 5.4 The sequential mapping procedure. A hypothesis for the kinetic energy at the end of a track T_e is mapped to a range value R_i using the mapping \mathcal{M} . A hypothesis for the total range ΔR maps R_e to R_i , which in turn is mapped to T_i . Any value $R_i \leq R_k \leq R_e$ now also has an associated value T_k through the mapping \mathcal{M} .

Each hit in the cluster will have an associated hit charge value Q_{ADC} in integrated ADC counts, which can be rescaled to units of MeV, such that

$$Q \equiv \frac{Q_{ADC} \cdot F \cdot W}{\mathcal{R}} \equiv \Omega Q_{ADC}, \quad (5.7)$$

where the constant TPC gain factor $F = 145.12$ scales integrated ADC counts to a number of electrons in the collection plane. This value has been obtained by extensive calibration work performed by the MicroBooNE collaboration [90]. The factor $W = 23.6$ eV is the ionisation energy of liquid argon, and $\mathcal{R} = 0.62$ is a recombination factor that accounts for the fact that ionisation electrons, after liberation, can recombine with positive argon ions in their vicinity [91]. In what follows, the decision has been made to normalise Q by the observed hit width w (cm), such that

$$\tilde{Q} \equiv \frac{Q}{w}, \quad (5.8)$$

where the shorthand \tilde{Q} has been introduced, which is employed throughout this chapter. This additional normalisation is applied to correct for the case in which the waveforms associated

to multiple hits overlap and are reconstructed by the hit finding procedure as a single hit. Such a merged hit would have a large associated hit width.

A MINUIT fitting procedure can be implemented to construct a minimal- χ^2 value of $\langle -dE/dx \rangle_k$ for each observed hit charge \tilde{Q} . MINUIT [92] is an optimised numerical optimisation package, which has been implemented in ROOT. This fitting procedure can be performed using only collection plane information, by making the assumption that the observed 2D track length L_{2D} is related to the 3D quantity ΔR by an unknown scaling factor λ , which is defined as

$$\lambda \equiv \frac{\Delta R}{L_{2D}}. \quad (5.9)$$

This fixes one of the two necessary unknowns as $\Delta R = \lambda L_{2D}$, and λ can be made a MINUIT fit parameter. Furthermore, performing the direction fit using 2D hit information has two distinct advantages. Firstly, the high-fidelity calorimetric information from the W plane can be used in isolation. Secondly, by relying only on 2D information, directional information can be incorporated earlier in the reconstruction chain, before the 3D hit reconstruction.

The only other unknown quantity is T_e , the kinetic energy of the track at its endpoint, which is made the second MINUIT fit parameter. In addition to the unknown quantities λ and T_e , a third parameter S is added, which is an overall scaling factor that aids the fitting procedure in finding a global minimum in the parameter space. However, the observed values of \tilde{Q} cannot be directly compared to the fit output $\langle -dE/dx \rangle_k$, since $\sum \tilde{Q} \neq \Delta T$ and $\sum w \neq \Delta R$. It is therefore necessary to rescale the total track hit width to ΔR and the total hit charge to ΔE , by defining two scaling parameters

$$\begin{aligned} \alpha &\equiv \Delta T / \sum_i \tilde{Q}_i, \\ \beta &\equiv \Delta R / \sum_i w_i. \end{aligned} \quad (5.10)$$

The scaling factor α makes the fitting procedure agnostic to the absolute scale of the charge values \tilde{Q} , since $\sum_i \alpha \tilde{Q}_i = \Delta T$ by construction. A fit charge value $\tilde{Q}_{fit,k}$ can now be constructed for every mean energy loss value returned by the fit $\langle -\frac{dE}{dx} \rangle_{fit,k}$, resulting in the expression

$$\tilde{Q}_{fit,k} \equiv \frac{\beta S}{\alpha} \left\langle -\frac{dE}{dx} \right\rangle_{fit,k}. \quad (5.11)$$

5.2.1 Formulating Chi Squared Values

The MINUIT optimisation works by selecting values for the fit parameters that minimise a χ^2 value, which quantifies goodness of fit. The value of χ^2 is defined in terms of the fit residuals Δ_i and the error on the target points σ_i , and is given by

$$\chi^2 = \sum_i^N \left(\frac{\tilde{Q}_i - \tilde{Q}_{fit,i}}{\sigma_i} \right)^2 \equiv \sum_i^N \left(\frac{\Delta_i}{\sigma_i} \right)^2, \quad (5.12)$$

where $i = 1, \dots, N$ indexes the hits in the cluster. This describes a χ^2 distribution with N degrees of freedom. The formalism of the previous section, combined with equation 5.12, is sufficient to construct direction fits.

Figure 5.5 shows an example direction fit for a forwards contained muon. A direction fit can be created for the hypothesis that the particle is travelling in the forwards direction in z and for the hypothesis that the particle is travelling backwards in z . For the forwards hypothesis, the longitudinal coordinates L_i for the hits are given by the Pandora 2D sliding fit, since the sliding fit beginpoint is the lowest- z hit. For the backwards hypothesis, the hit longitudinal coordinates are defined to be $L_{2D} - L_i$. The direction fit values $\tilde{Q}_{fit,i}$ corresponding to the forwards hypothesis are shown in figure 5.5 in magenta. Each direction hypothesis will result in a χ^2 value, which will be annotated by $_F$ for true forwards particles and by $_B$ for true backwards particles. The smallest χ^2 value of both hypotheses is labelled χ_{min}^2 , and is defined as

$$\chi_{min}^2 \equiv \min(\chi_F^2, \chi_B^2). \quad (5.13)$$

The values of χ_F^2 and χ_B^2 are compared in order to select a final direction hypothesis, which is discussed in Sections 5.4 and 5.5. However, the errors σ_i in equation 5.12 remain undefined, affecting the resulting values of χ^2 .

In equation 5.12, the residuals $\Delta \equiv \tilde{Q} - \tilde{Q}_{fit}$ are expected to increase with \tilde{Q} , both due to the scale change and due to the associated difficulty of reconstructing hits in the Bragg peak. Because \tilde{Q} is based on Q , which is a measured quantity, each hit charge \tilde{Q}_i will have an associated uncertainty σ_i . In what follows, the simplifying assumption is made that the hit width w is an uncertainty-less value, such that the uncertainty on \tilde{Q} is equal to the uncertainty on Q . By determining the dependence of the uncertainty σ on Q , a value of σ_i can be estimated for each observed hit charge \tilde{Q}_i , such that the values of χ^2/N peak near 1 (behave like a reduced χ^2 statistic) and can be used to inform the direction fitting logic. However, instead of estimating the function $\sigma(Q)$, a functional form of $\sigma(Q_{fit})$ is developed, where $Q_{fit} \equiv \tilde{Q}_{fit} \cdot w$. The reason for this is that various effects, such as the presence of

high-charge hits due to δ -ray interactions (discussed in Section 5.3) can cause a large spread in the observed hit charge values. These effects introduce multiple effects into each bin of Q , causing large biases in the distribution of Δ for high values of Q . The values of Q_{fit} , however, are unaffected by such effects, and the residuals in each bin of Q_{fit} are roughly normally distributed and are unbiased. The more regular behaviour of \tilde{Q}_{fit} compared to \tilde{Q} can be seen in Figure 5.5.




Fig. 5.5 A direction fit to a simulated forwards-going contained muon. The values of \tilde{Q} are indicated by black points, and the fit values \tilde{Q}_{fit} resulting from the forwards direction fit are indicated by magenta points. While the values of \tilde{Q} are affected by secondary interactions such as δ -rays, the values of \tilde{Q}_{fit} are more robust to such effects.

The values of \tilde{Q}_{fit} are recorded for each hit in each true forwards-going, contained muon in the event sample, resulting in the distribution shown in Figure 5.6a. To estimate $\sigma(Q_{fit})$, the following steps are performed:

- A 2D distribution of Δ with Q_{fit} is made (omitted here), which is subsequently binned in Q_{fit} using the binning defined by Figure 5.6b;
- For each bin in Q_{fit} , the distribution of the values of Δ is plotted. A Gaussian fit is performed to the distributions of Δ , resulting in a fitted standard deviation value σ_{Gauss} , as well as the root mean squared error of the residuals, defined as

$$\Delta_{RMSE} \equiv \sqrt{1/N(\Delta_1^2 + \dots + \Delta_N^2)}. \quad (5.14)$$

This gives two measures of the error, which can be compared. It is shown in what follows that these two error measures have a similar dependence on Q_{fit} ;

- The values of σ_{Gauss} and Δ_{RMSE} are plotted as a function of Q_{fit} and an appropriate functional form is determined that captures the observed behaviour of both σ_{Gauss} and Δ_{RMSE} . This functional form can then be used in the direction fit to estimate accurate values of σ_i in equation 5.12, resulting in accurate values of χ^2 .

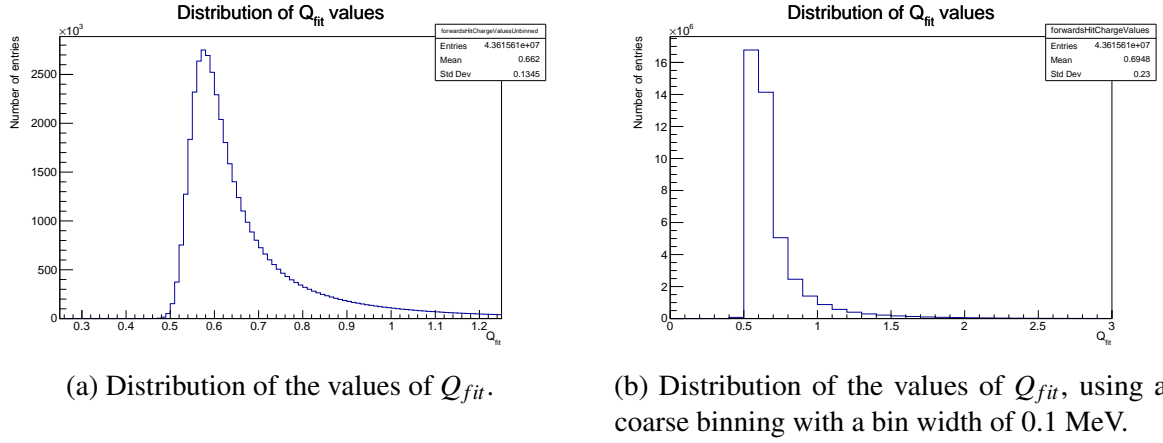




Fig. 5.6 The distribution of $Q_{fit} \equiv \tilde{Q}_{fit} \cdot w$ using a fine binning of 0.03 MeV (5.6a) and a coarse binning of 0.1 MeV (5.6b). The binning used in Figure 5.6b is used when plotting the distributions of the residuals Δ .


A number of example Δ distributions per bin in Q_{fit} are shown in Figure 5.7. This figure clearly shows that as Q_{fit} increases, the distributions of Δ become wider, reflecting a larger uncertainty σ . As the Δ distributions become wider, they also acquire long tails. These observed long tails in the distributions of the residuals are due to a combination of secondary effects. As mentioned, δ -ray hits contribute large values of Q , leading to large values of Δ . In addition, effects such as imperfect clustering contribute similarly large Δ values. These effects are discussed in more detail in Section 5.3. The long tails in the distributions of the residuals significantly increase uncertainty estimates for both σ_{Gauss} and Δ_{RMSE} . To correct for this effect, the Δ distributions have been divided up into three regions: the inner 70, 80 and 90% of the entries, sorted by their distance from the mean (shaded red, beige and blue in Figure 5.7, respectively). The distributions of the inner 70% of entries are found to be robust to the long tails, while retaining the dependence of the uncertainty on Q_{fit} . These distributions are therefore used in what follows. The Gaussian fits performed on the inner 70% of the Δ distributions are shown as blue lines in Figure 5.7, and their extrapolations are shown as red lines.




(a) Δ distribution in the bin with bin centre $Q_{fit} = 0.55$ MeV.



(b) Δ distribution in the bin with bin centre $Q_{fit} = 1.05$ MeV.



(c) Δ distribution in the bin with bin centre $Q_{fit} = 1.55$ MeV.



(d) Δ distribution for in the bin with bin centre $Q_{fit} = 2.05$ MeV.

Fig. 5.7 Distributions of the residuals Δ for four equidistant bins in Q_{fit} . The inner 70, 80 and 90% of the distributions are shaded red, beige and blue, respectively. The Gaussian fit to the inner 70% (red) area is shown by the blue curve. The red curve is the extrapolation of the blue curve outside of the 70% region.



(a) Δ_{RMSE} of the inner 70% of Δ entries.(b) Gaussian fit σ of the inner 70% of Δ entries.

Fig. 5.8 Three fits to the values of Δ_{RMSE} (5.8a) and σ_{Gauss} (5.8b) resulting from the inner 70% distributions of Figure 5.7. The red fit uses a functional form $a \cdot \sqrt{Q_{fit}}$, the blue fit uses $a \cdot Q_{fit}$ and the magenta fit uses the superimposed form $\sqrt{a \cdot Q_{fit}^2 + b \cdot Q_{fit}}$.

The values of Δ_{RMSE} and σ_{Gauss} resulting from the inner 70% of the Δ distributions, as a function of the centre of each Q_{fit} bin, are shown in Figure 5.8. It is clear from Figure 5.8 that the values of Δ_{RMSE} and the σ_{Gauss} behave similarly. The values of σ_{Gauss} tend to rise more quickly for higher values of Q_{fit} , especially in the region $Q_{fit} \geq 3$ MeV, where the underlying Gaussian distribution is still obscured by the tails in the distribution. As such, the uncertainty as estimated by σ_{Gauss} is slightly overestimated in this region. To determine a functional form for the dependence of Δ_{RMSE} and σ_{Gauss} on Q_{fit} , a comparison is made between three hypotheses:

$$\sigma = \begin{cases} a \cdot \sqrt{Q_{fit}} \\ a \cdot Q_{fit} \\ \sqrt{a \cdot Q_{fit}^2 + b \cdot Q_{fit}} \end{cases}, \quad (5.15)$$

where a and b are fit parameters. The fits for each of these three functions on the observed values of $\Delta(Q_{fit})$ are shown as red, blue and magenta lines in Figure 5.8. The combined fit $\sigma = \sqrt{a \cdot Q_{fit}^2 + b \cdot Q_{fit}}$ captures the behaviour of both σ_{Gauss} and Δ_{RMSE} values well, and similar values for a and b are found for both the Δ_{RMSE} and σ_{Gauss} fits. Hence, the values of σ_i in equation 5.12 are estimated as

$$\sigma = \sqrt{a \cdot Q_{fit}^2 + b \cdot Q_{fit}}, \quad (5.16)$$

where $a \approx 4.2 \times 10^{-3}$ and $b \approx 9.7 \times 10^{-3}$. Using equation 5.16, the fitting procedure can accurately fit observed hit charge distributions in real and simulated LArTPC data. However, the resulting fits and $\Delta\chi^2$ values are strongly affected by outlying values of \tilde{Q} , such as those in Figure 5.5. It is therefore useful to filter out such hit charge values before fitting, in order to improve the quality of the direction fit. As such, a hit filtering procedure has been developed, which is the topic of the next section.

5.3 Hit Filtering

The values of $-dE/dx$ follow a skewed Landau distribution, causing a variation in the observed values of \tilde{Q} . Entries in the tail of the equivalent distribution of \tilde{Q} should be removed, since they can negatively affect the fit and the resulting χ^2 values. In addition, muons can induce δ -rays as they propagate, which induce their own ionisation electrons in the liquid argon. These additional ionisation electrons will impinge on the wire planes at the same time as the electrons induced by the muon, leading to reconstructed hits with higher

values of \tilde{Q} . The impact of these effects on the observed values of \tilde{Q} has been visualised in Figure 5.9a, in which hits for which more than one Monte Carlo particle contributes $\geq 15\%$ of the hit charge have been coloured red. These ‘impure’ hits can be seen to have significantly higher \tilde{Q} values than the pure muon hits. Figure 5.9b shows the effect of the filtering procedure that is described in this section, which removes the majority of impure hits. In addition, the direction fit on the filtered hit collection is shown in magenta.


When developing the hit filter, a distinction is made between filtering the ends of the track and the inner part of the track. The reason for this is that particular care must be taken when filtering hits in the endpoints of the track, where the Bragg peak may appear, since the Bragg peak contains the directional information. For this reason, the first and last 5% of the 2D track length are filtered separately by the ‘Bragg peak filter’, discussed in Section 5.3.2. The remaining 90% of the track length is filtered using the main filtering procedure, which is discussed first.

5.3.1 Main Hit Filter


It is clear from Figure 5.9a that one of the primary characteristics of pure hits, compared to impure hits, is that they are coherently grouped together. Impure hits are much less densely clustered in (\tilde{Q}, L_{2D}) space and the main principle behind the hit filter is to quantify the density of hits in this space, and to remove the most isolated hits. The extent of the parameter space along the spatial and charge axes is defined by the variables ΔL and $\Delta\tilde{Q}$, respectively. The extent along the spatial axis is simply the 2D track length: $\Delta L = L_{2D}$, and the extent along the charge axis is given by the difference between its two extremal points: $\Delta\tilde{Q} = \max(\tilde{Q}) - \min(\tilde{Q})$. For every pair of hits indexed by i and j , a distance measure D can then be defined as

$$D_{ij} \equiv \sqrt{|L_i - L_j|^2 + \left(\frac{\Delta L}{\Delta\tilde{Q}}\right)^2 |\tilde{Q}_i - \tilde{Q}_j|^2}. \quad (5.17)$$

In equation 5.17, the ratio $\Delta L/\Delta\tilde{Q}$ is used to scale the charge distance $|\tilde{Q}_i - \tilde{Q}_j|$ to the same scale as the length distance $|L_i - L_j|$. Using the distance measure D_{ij} , each hit in the target cluster can now be ordered according to the average distance $\langle D_k \rangle$ to its k nearest neighbours. In the hit filter implementation, $k = 5$, since the main source of impure hits is low-energy δ -rays, which deposit few hits that are widely spaced apart in (\tilde{Q}, L_{2D}) space. The ‘fractional position’ p_f of $\langle D_5 \rangle$ is used to attain good separation between pure and impure hits. The fractional position p_f is constructed by finding the index p in a sorted vector of $\langle D_5 \rangle$ values in a cluster, and dividing this index by the length of the vector N , such that $p_f \equiv p/N$.



(a) The hit collection due to the muon, with pure hits in black and impure hits in red.



(b) The same hits as in 5.9a after applying the filtering procedure, with direction fit.

Fig. 5.9 The hits deposited by a 779 MeV muon, before applying the filtering procedure (5.9a) and after applying the filtering procedure, with the best direction fit indicated by magenta points (5.9b).

Figure 5.10 shows the distributions of the fractional position p_f of $\langle D_5 \rangle$ for pure and impure hits. To apply the hit filter, a cut is applied where the p_f distributions for pure and impure hits intersect, at $p_f = 0.72$. The region of p_f that is retained by this cut is shaded blue in Figure 5.10.




Fig. 5.10 The distributions of the fractional position p_f for pure and impure hits. A cut is applied where the distributions of p_f intersect, at $p_f = 0.72$. The region retained by the cut is shaded blue.

The performance of the hit filter can be quantified, using the concepts of filter purity, filter completeness and by quantifying the fraction of pure hits retained by the filter. If the target cluster contains a set of impure hits $\{I\}$ and a collection of removed hits $\{F\}$, then the filter purity P is defined as the fraction of hits removed by the filter that were impure, and the filter completeness C is defined as the total fraction of impure hits that has been removed, such that

$$P = \frac{|I \cap F|}{|F|}, \quad (5.18)$$

$$C = \frac{|I \cap F|}{|I|}.$$

Figure 5.11 shows distributions of the filter purity and completeness, as well as distributions of the fraction of hits removed by the filter, and the fraction of pure hits retained by the filter. These distributions indicate that the filter is effective at removing impure hits, and does not remove a significant fraction of pure hits. The distribution of the fraction of pure hits retained by the filter peaks at a value < 1 , since some pure hits have large \tilde{Q} values. Similarly, some impure hits are very close to the mean of the distribution of \tilde{Q} values in the track. The long tails in the distributions of P and C are mainly due to clusters with few hits, for which the

entire track length is not minimally ionising, leading to pure hits with \tilde{Q} values that have larger fluctuations.

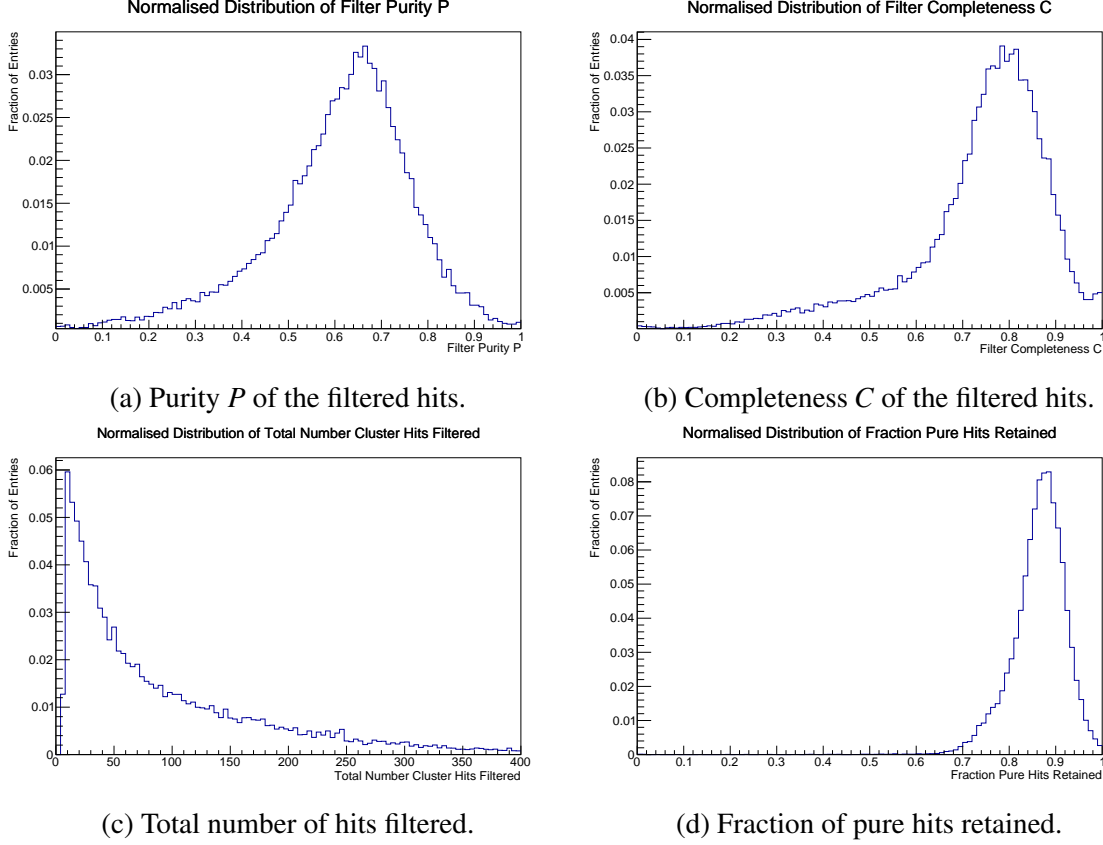


Fig. 5.11 Plots indicating the performance of the filtering method. The definitions of purity P and completeness C are given in the text. The long tails in these plots are due to events in which many hits are split into multiple fragments, resulting in very low \tilde{Q} values. These fragments are removed and can constitute a significant fraction of the total number of hits.

5.3.2 Bragg Peak Filter


Either track endpoint may contain the Bragg peak, if the track is contained and if the Bragg peak is not obscured by inactive detector regions. As such, hits with longitudinal coordinate values L_i for which the distance to either track endpoint is $\leq 5\%$ of the total 2D track length L_{2D} are filtered differently from the rest of the track, in order not to erase the directional information contained in the hits of the Bragg peak. A hit filter in the Bragg peak region is necessary, since impure high-charge hits near the interaction vertex may resemble a fake Bragg peak. This can lead to an incorrect direction hypothesis with a smaller value of χ^2/N

than the correct direction hypothesis. High-charge hits near the track beginpoint are mainly due to the following three effects:


- *Overlapping hits or ‘charge sharing’ near the interaction vertex*: multiple particles often emerge from the neutrino interaction vertex, which initially ionise the same region of detector space. This will result in ionisation electrons liberated by multiple particles impinging on the wire planes simultaneously, leading to merged hits with very large induction pulses, and very high hit charge values;
- *Clustering errors*: the Pandora clustering algorithms use only topological information, and hence it is possible to assign hits induced by other particles to the muon cluster. This is especially noticeable when proton hits are included in the muon cluster, since the proton mass is much larger than the muon mass, leading to much higher $-dE/dx$ values;
- *Split hits*: the hit finding algorithms can reconstruct large induction pulses on the wire planes that belong to a single hit as multiple hits. In many cases, this will result in a value of \tilde{Q} within the expected range, and a small hit ‘fragment’. This is especially likely to happen in the Bragg peak, where individual hits have very high charges.

Figure 5.12a shows the \tilde{Q} values for a true contained forwards muon that illustrates these three effects. In this example there are a large number of hits near the interaction vertex ($L_{2D} \approx 0$) that have high \tilde{Q} values due to overlapping hit pulses near the interaction vertex and due to a few proton hits that are part of the muon cluster. In addition, the last hit in the Bragg peak has been split into a hit with a high value of \tilde{Q} and a low- \tilde{Q} hit fragment. The value of the residuals $\Delta \equiv \tilde{Q} - \tilde{Q}_{fit}$ for the hits near $L_{2D} \approx 0$ are large, leading to large χ^2 contributions from these hits for the correct direction hypothesis, and small contributions for the incorrect direction hypothesis. Figure 5.12b shows the same \tilde{Q} distribution after the Bragg peak filter has been applied, which is described in this section. This figure demonstrates that the Bragg peak filter is effective at removing impure hits near the track endpoints while retaining hits in the Bragg peak.

Hit fragments in the Bragg peak such as in Figure 5.12b can be identified by the large \tilde{Q} difference between such hits and their nearest neighbours. As such, the two hits in the cluster for which L_{2D} is minimal and maximal are removed if their \tilde{Q} values are $\leq 90\%$ of that of their nearest neighbours. This procedure is applied iteratively, until no hits remain that meet this criterion. Unlike split hits, impure hits due to charge sharing near the interaction vertex or clustering errors are often clustered together in (\tilde{Q}, L_{2D}) space, making them more



(a) Example track for which the \tilde{Q} distribution displays both charge sharing near the interaction vertex and a split hit in the Bragg peak.



(b) The track of Figure 5.12a after applying the Bragg peak filter.

Fig. 5.12 The effect of the Bragg peak filter on an example muon track in the W view. Figure 5.12a shows a muon track for which \tilde{Q} values are inflated near the interaction vertex due to overlapping hit pulses, the inclusion of proton hits. A split hit in the Bragg peak is also present. Figure 5.12b shows the effect of applying the Bragg peak filter on the distribution of \tilde{Q} . Impure hits are coloured red, and are hits for which $\geq 15\%$ of the hit charge has been contributed by a secondary primary Monte Carlo particle.

difficult to remove. Despite this, the \tilde{Q} values of such impure hits can be compared to their k -th nearest neighbours, and the ratio $\max(\frac{\tilde{Q}_i}{\tilde{Q}_{i\pm k}})$ is likely to be large. However, this quantity may also be large for hits in the Bragg peak, which can be far-removed from their nearest neighbours. As such, this is a difficult filtering problem, but the removal of hits that can cause the direction estimate to be incorrect is essential to the accuracy of the fitting procedure.

Figure 5.13 shows the normalised distributions of $\max(\frac{\tilde{Q}_i}{\tilde{Q}_{i\pm 3}})$ for hits within the first and last 5% of the track length nearest the interaction vertex ('non-Bragg hits'), and hits furthest from the interaction vertex ('Bragg peak hits'). This fractional distance is chosen as it works well for both long and short tracks. This figure shows that, despite the large overlap in the distributions, there is a larger fraction of non-Bragg hits in the tail of the distribution. The distributions for Bragg peak hits and non-Bragg hits intersect in the bin with bin centre $\max(\frac{\tilde{Q}_i}{\tilde{Q}_{i\pm 3}}) = 1.725$ MeV/cm. By applying a cut at this value, the Bragg peak filter is more likely to remove an impure hit than a hit in the Bragg peak. This cut retains 94.80% of Bragg peak entries and 89.28% of non-Bragg peak entries. The performance of the Bragg peak filter can be quantified by considering the effect that removing impure hits near the track endpoints has on the direction hypothesis, as shown in Figure 5.14. In this figure, green points indicate tracks for which the direction hypothesis was incorrect without applying the Bragg filter, and correct after applying the filter. Red points indicate the reverse, and tracks for which no direction hypothesis change occurred as a result of the Bragg peak filter have been omitted. In this plot, true forwards-going muon tracks are represented by triangles, and true backwards tracks are represented by squares.




Fig. 5.13 Distributions of $\max(\frac{\tilde{Q}_i}{\tilde{Q}_{i\pm 3}})$ for Bragg peak and non-Bragg peak hits. A cut is applied at $\max(\frac{\tilde{Q}_i}{\tilde{Q}_{i\pm 3}}) = 1.725$, which retains 94.80% of Bragg peak entries and 89.28% of non-Bragg peak entries.





Fig. 5.14 A visualisation of the change in χ_{min}^2/N due to applying the track end filter for true forwards going muons (triangles) and backwards going muons (squares). A distinction is made between correct removals, for which the direction hypothesis changes for the better (green), and bad removals (red).


Figure 5.14 shows that the direction hypotheses for many events are corrected by the application of the Bragg peak filter. Furthermore, correct Bragg peak filter applications result in a value of $\chi_{A,min}^2/N \approx 1$, as expected for a χ^2 distribution with N degrees of freedom. In addition, the probability of a removal being correct increases with the change in χ^2/N , denoted $\delta\chi_{BA}^2/N$, which is defined as

$$\delta\chi_{BA}^2/N \equiv \chi_{min,B}^2/N - \chi_{min,A}^2/N, \quad (5.19)$$


where $\chi_{min,B}^2/N$ is the χ_{min}^2/N value before applying the Bragg peak filter, and $\chi_{min,A}^2/N$ is the equivalent after applying the filter. The performance of the Bragg peak filter can be enhanced by cutting on the value of $\delta\chi_{BA}^2/N$ to remove bad entries, such that the Bragg peak filter is only applied for $\delta\chi_{BA}^2/N \geq T$. The quantity $\delta\chi_{BA}^2/N$ depends on the number of hits N , since particles with small numbers of hits N tend to have low energies, are not minimally ionising, contain fewer data points, and are harder to fit. Scatter plots like the one in Figure 5.14 can be drawn for bins in N , which reveals the dependence of the fraction of incorrect entries on N . Figures 5.15a and 5.15b show two examples of such binned scatter plots, for $50 \leq N \leq 100$ and $150 \leq N \leq 200$, which reveal that for smaller N , a larger fraction of entries is incorrect, which implies that a stricter cut value T is necessary for low N .




(a) Scatter plot of good and bad particle splits for $50 \leq N < 100$



(b) Scatter plot of good and bad particle splits for $150 \leq N < 200$



(c) Purity P_T and completeness C_T values for 100 different cuts on T in the range $0 \leq T \leq 10$ for the scatter plot in 5.15a.



(d) Purity P_T and completeness C_T values for 100 different cuts on T in the range $0 \leq T \leq 10$ for the scatter plot in 5.15b.

Fig. 5.15 Figures 5.15a and 5.15b show scatter plots of good and bad Bragg peak filter entries for tracks that satisfy $50 \leq N < 100$ and $150 \leq N < 200$, respectively. Figures 5.15c and 5.15d show the corresponding purity P_T and completeness C_T values for 100 different cuts on T in the range $0 \leq T \leq 10$.


The efficacy of each cut value T on $\delta\chi_{BA}^2/N$ can be quantified using the concepts of cut purity P_T and completeness C_T , such that for a scatter plot with a set of entries removed by the cut $\{R\}$ and a set of bad entries $\{B\}$, $P_T = |B \cap R|/|R|$ and $C_T = |B \cap R|/|B|$. A reasonable cut will aim to retain as many good entries as possible, while also removing a reasonable number of bad entries. An optimal value for T can be found by sampling the range $T \in [0, 10]$ equidistantly for a large number of points, 100 in this example, and selecting the value of T for which the purity and completeness measures P_T and C_T are closest together. Figures 5.15c and 5.15d show the purity and completeness curves generated for the scatter plots in Figures 5.15a and 5.15b. This procedure represents a consistent method to find a cut value T that balances the trade-off between excluding bad entries and including good entries. The optimal values of T for the entire binned range N is shown in Figure 5.16a and a fit to the observed values of T is displayed in red. The fit function is an exponential function with three fit parameters a_1 , b_1 and c_1 , with a functional form given by

$$T = a_1 + \exp(b_1 + c_1 \lfloor N/50 \rfloor), \quad (5.20)$$

where $a_1 = 0.405$, $b_1 = 2.122$ and $c_1 = -0.494$ are fit parameters. The fit values given by the red curve in Figure 5.16a are used to determine the minimum value of $\delta\chi_{BA}^2/N$, as a function of N , that the Bragg peak filter must achieve in order to be applied. The result of applying the cuts defined by Figure 5.16a on the Bragg peak filter scatter plot of Figure 5.14 is given in Figure 5.16b. Applying this cut procedure retains 89.32% of good removals and rejects 52.57% of bad entries. This figure shows that this procedure is effective at preventing incorrect Bragg peak filter removals with small values of $\chi_{min,A}^2/N$. While the cut procedure retains incorrect removals with large values of $\chi_{min,A}^2/N$, this has very little impact on the quality of the direction fit, because such large values of $\chi_{min,B}^2/N$ indicate that the direction fit was not able to resolve the direction.

5.4 Single Particle Hypothesis

The hit filtering procedure ensures that the majority of contained particles have a good best direction hypothesis with small values of χ_{min}^2/N . In other words, there is a large difference in the fit quality between the forwards and backwards direction fit hypotheses, with the best hypothesis being a good fit to the observed values of \tilde{Q} . If the initial fit is sufficiently good in terms of χ_{min}^2/N , the decision can be made to forego the subsequent steps in the direction fitting logic, as depicted in Figure 5.17. This decision should also incorporate information about the number of track hits N , as particles with higher momenta are far less likely to



(a) Optimal cut values T as a function of N and a fit to the observed T values with the function $T = a_1 + \exp(b_1 + c_1 [N/50])$.

(b) The result of applying the cut values T defined by the exponential fit in Figure 5.16a on the Bragg peak filter scatter plot of Figure 5.14.

Fig. 5.16 A definition of optimal cut values T on the observed value of $\delta\chi^2_{BA}/N \equiv \chi^2_{B,\min}/N - \chi^2_{A,\min}/N$ is given by the exponential fit in 5.16a and the effect of applying these cut values to the scatter plot of Figure 5.14 is given in 5.16b.

scatter backwards on argon nuclei. Small values of χ^2_{\min}/N indicate that it is highly likely that a direction fit has been found that accurately describes the observed values of \tilde{Q} .

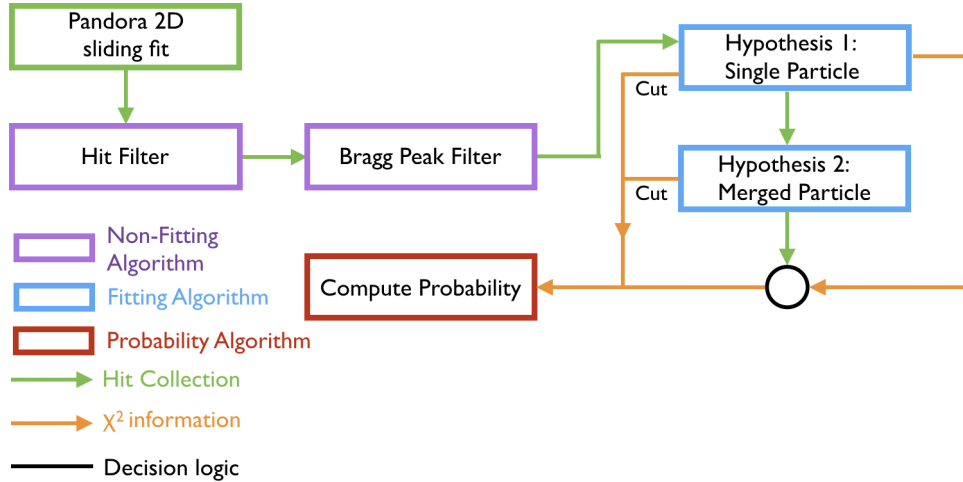


Fig. 5.17 An overview of the steps in the direction fitting procedure. Information flow and algorithm type have been coloured according to whether they constitute information resulting from the direction fitting procedure, or otherwise. The word 'cut' indicates decision logic associated to the single and merged particle hypotheses, as described in the text. A legend is given in the bottom left of the figure.

Figures 5.18a and 5.18b show the distributions of χ^2/N for true forwards and backwards muons, with both χ^2/N distributions peaking at $\chi^2/N \approx 1$. Figures 5.18c and 5.18d show colour-coded binned scatter plots of χ^2/N with N for true forwards and backwards muon

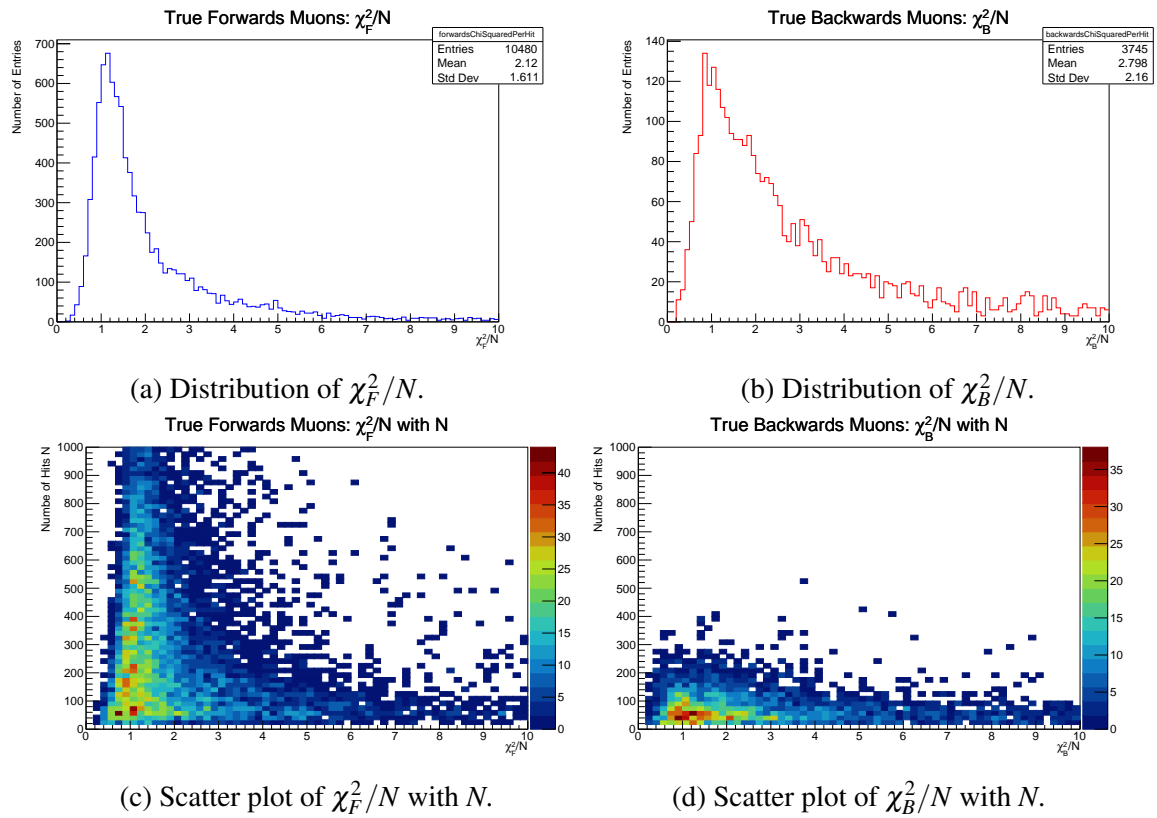


Fig. 5.18 Distribution of χ^2/N for true forwards and backwards muons are given in 5.18a and 5.18b. Colour-coded binned scatter plots of χ^2/N with N are given in 5.18c and 5.18d. Based on these distributions, a good fit has been determined to be a fit for which $\chi^2/N \leq 1.5$.

tracks, respectively. Based on the distributions of Figure 5.18, a good direction fit is defined to be a fit for which $\chi_{min}^2/N \leq 1.5$. Within this range, the χ^2/N value is near enough to 1 that the set of points \tilde{Q}_{fit} accurately describes the data \tilde{Q} , and the probability of secondary effects occurring is likely to be low. Figures 5.18c and 5.18d also show that when $N \geq 400$, the particle is almost certainly travelling in the forwards direction. The following criteria are applied for forwards and backwards particles to determine whether the single particle hypothesis is employed, causing the direction fitting logic to skip straight to the probability calculation of Section 5.6:

- *Single forwards particle:* $\chi_{min}^2 \leq 1.5$ and $\chi_F^2/N \leq \chi_B^2/N$, or if $N \geq 400$;
- *Single backwards particle:* $\chi_{min}^2 \leq 1.5$ and $\chi_F^2/N > \chi_B^2/N$.


The above criteria constitute the first ‘cut’ entry in Figure 5.17. As can be seen in Figure 5.17, if the observed values of χ_{min}^2 and N do not meet these criteria, the hypothesis is explored that the cluster consists of two merged back-to-back particles, which is discussed in the next section.

5.5 Particle Splitting


Due to recoils on argon nuclei, tracks can emerge back-to-back from the interaction vertex, and may be merged into the same cluster by the reconstruction logic. An example is given in Figure 5.19a, which shows the \tilde{Q} distribution of a single W plane cluster that contains hits from both a proton and a muon. If the opening angle θ between these two tracks is sufficiently close to π rad, such a merged cluster is topologically indistinguishable from a single particle track. However, the merge is very obvious in the calorimetric information of Figure 5.19a, since the higher values of $\beta\gamma$ due to the large value of the proton mass $m_p \approx 938.27$ MeV/c, compared to $m_\mu \approx 105.66$ MeV/c [93], manifests as much higher values of $-\langle dE/dx \rangle$ and hence \tilde{Q} . Such sudden changes in \tilde{Q} , or ‘jumps’, can be identified by comparing the \tilde{Q} values of each hit (indexed by i) to their k nearest neighbours, and taking the largest difference:

$$J_i \equiv \max \left(\frac{\tilde{Q}_i}{\tilde{Q}_{i \pm k}} \right) \quad (5.21)$$

By scanning the entire cluster hit collection $i = 1 \cdots N$, a set of potential split positions $\{S_i\}$ can be identified, which are the longitudinal positions L_i associated to jumps J_i . If the \tilde{Q} distribution in Figure 5.19a were fitted using the fit procedure of Section 5.2, this would



(a) A single W cluster that contains hits from both a muon and a proton. The proton has a much higher mass, which results in much higher values of $\langle -dE/dx \rangle$.



(b) The same cluster as in 5.19a, after identifying a split position $S \approx 88\text{cm}$ and fitting the two hit collections separately.

Fig. 5.19 An example single cluster in the W view that contains hits from both a muon and a proton is given in 5.19a. The result of applying the splitting procedure described in this section on this hit collection is given in 5.19b. The best fits associated to the two separate particle hypotheses are given as magenta points.

result in a very large value of χ_{min}^2/N , since both the forwards and backwards hypotheses would only be able to fit one of the two Bragg peaks in the cluster. However, the value of χ_{min}^2/N can be reduced significantly by positing a split position at $L_{2D} \approx 90$ cm and fitting both the resulting particle hypotheses (labelled 1 and 2) separately, such that

$$\chi_{min,split}^2/N = \chi_{min,1}^2/N + \chi_{min,2}^2/N. \quad (5.22)$$

The constraint is imposed that the direction hypotheses D for both particles satisfy $D_1 \neq D_2$, as the particles are assumed to emerge from the same point. For each potential split, the best split S_B is taken to be the split position that maximises $\delta\chi_{split}^2/N \equiv \chi_{min}^2/N - \chi_{min,split}^2/N$. Furthermore, each optimal split position S_B can be classified as either correct or incorrect based on the underlying Monte Carlo information. A correct split is defined to be a split for which either of these conditions are true:

- For at least 10 hits in the cluster, the majority of the charge is deposited by a secondary Monte Carlo primary particle (a simulated particle without any daughters). Hits deposited by daughter particles, such as δ -rays, are considered to be deposited by the parent particle and do not contribute to this criterion;
- Either of the resulting best direction fits contains a majority of hits for which $\geq 50\%$ of the charge is deposited by a secondary Monte Carlo primary particle. This condition accounts for clusters with < 10 hits deposited by a secondary particle, which are still correctly removed by the splitting procedure.

The jump value J_B associated to S_B is highly correlated to the probability that S_B is a correct split, and a significant fraction of incorrect splits can be removed with an appropriate cut on J_B . Figure 5.20 shows the distribution of J_B for correct and incorrect split positions. In many tracks, no significant jumps can be found, resulting in the peaks at $J_B = 0$. In many tracks where the cluster consists of a single particle, jumps with a magnitude $J_B \approx 0.2$ MeV/cm can still be found, either because of inflated values of \tilde{Q} values that have not been removed by the hit filter, or large differences in \tilde{Q} between hits in the Bragg peak. The distribution of J_B for good splits has a much longer tail, associated with large, correct jumps such as the example in Figure 5.19. Based on Figure 5.20, a cut is applied at $J_B = 0.39$ MeV/cm, which is the point at which the distributions in this figure intersect. The values of $\chi_{min,B}^2/N$ and $\chi_{min,A}^2/N$, the values of χ_{min}^2 before and after splitting can be visualised in a scatter plot, as in Figure 5.21a. In this plot, a distinction has also been made between true forwards- and backwards-going particles, which are denoted by triangular and square markers, respectively. The cut on J_B has been applied in this figure.


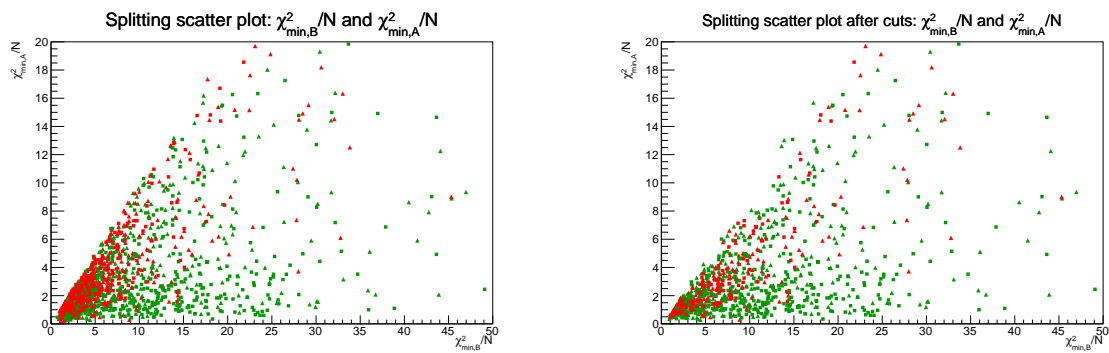


Fig. 5.20 Distributions of the magnitude of the best jump J_B for good and bad particle splits are drawn in good and red, respectively.



(a) Scatter plot for correct and incorrect particle splits.

(b) Scatter plot for correct and incorrect particle splits after applying the cut values defined by the exponential fit of Figure 5.22.

Fig. 5.21 Scatter plots for correct and incorrect splits for the entire event sample is given in 5.21a and the same scatter plot after applying the cut values defined by the exponential fit of Figure 5.22.

Figure 5.21 demonstrates that there is a relatively large number of correct splits for which $\chi_{\min,A}^2/N \ll \chi_{\min,B}^2/N$ and $\chi_{\min,A}^2/N \approx 1$. A procedure identical to the one outlined in Section 5.3.2 can now be applied to this scatter plot, finding an optimal cut T on the value of $\chi_{\min,B}^2/N - \chi_{\min,A}^2/N$ as a function of N . The cut value T as a function of bins in N is given in Figure 5.22, together with a fit to an exponential function of the form $T = a_2 + \exp(b_2 + c_2 \lfloor N/50 \rfloor)$, where $a_2 = 0.369$, $b_2 = 1.158$ and $c_2 = -0.799$ are the values of the fit parameters. Figure 5.21b shows the result of applying the cut values resulting from this exponential fit to the scatter plot of Figure 5.21a, which removes 72.85% of bad splits and retains 48.44% of good splits. Furthermore, Figure 5.23 gives a number of examples of successful splits, their underlying \tilde{Q} distributions, and the direction fits for particles 1 and 2. This figure demonstrates that the splitting procedure is effective at identifying both small ($N \lesssim 10$) and large ($N \gtrsim 10$) secondary particles, which in this event sample are short and long protons. Incorrect splits in 5.21 predominantly occur when Bragg peak hits appear separated from other hits in the track, for instance by an inactive detector region. Such a topology is calorimetrically similar to a proton and an uncontained muon that have been merged into the same cluster.






Fig. 5.22 Splitting cut values for bins in N with exponential fit of the form $T = a_2 + \exp(b_2 + c_2 \lfloor N/50 \rfloor)$.

5.6 Forwards Probability


The distributions of χ_F^2/N and χ_B^2/N can be used to formulate a pseudo-probability P_F that a particle is travelling forwards in z . In particular, the difference in χ^2/N between the two direction hypotheses $\Delta\chi_{FB}^2/N$ quantifies which direction hypothesis fits the observed \tilde{Q}




(a) A correct split with $\chi^2_{min,B}/N = 10.1236$ and $\chi^2_{min,A}/N = 0.84273$



(b) A correct split with $\chi^2_{min,B}/N = 10.1236$ and $\chi^2_{min,A}/N = 0.84273$



(c) A correct split with $\chi^2_{min,B}/N = 15.8661$ and $\chi^2_{min,A}/N = 1.06321$



(d) A correct split with $\chi^2_{min,B}/N = 24.5591$ and $\chi^2_{min,A}/N = 2.04642$

Fig. 5.23 Several examples of successfully split clusters, where the proton and the muon that make up the original cluster are clearly visible to have distinct \tilde{Q} values.

distribution best, and is defined as

$$\Delta\chi_{FB}^2/N \equiv \chi_F^2/N - \chi_B^2/N. \quad (5.23)$$


Hence, $\Delta\chi_{FB}^2/N \leq 0$ indicates that the forwards hypothesis is best, and $\Delta\chi_{FB}^2/N > 0$ denotes that the backwards hypothesis is the best fit. Figures 5.24a and 5.24b show the distributions of $\Delta\chi_{FB}^2/N$ for true forwards and true backwards contained muon tracks, respectively. For true forwards muon tracks, 91.31% of entries satisfy $\Delta\chi_{FB}^2/N \leq 0$, and for true backwards muons, 84.94% of entries satisfy $\Delta\chi_{FB}^2/N > 0$. Clusters for which the sign of $\Delta\chi_{FB}^2/N$ is incorrect are typically events in which the \tilde{Q} distribution appears unstructured, which may happen when the opening angle of a track with the x axis of the detector is small, such that the ionisation electrons associated to the particles are collected on a small number of collection plane wires. Examples of two such events are given in Figure 5.25.

By taking the ratio between the bins in the normalised $\Delta\chi_{FB}^2/N$ distributions, P_F can be calculated for an observed range (bin) in $\Delta\chi_{FB}^2/N$. If, for a given bin B in $\Delta\chi_{FB}^2/N$, N_F and N_B forwards and backwards entries are present in the normalised χ_F^2/N and χ_B^2/N distributions, the forwards probability P_F in B is defined as


$$P_F = \frac{N_F}{N_F + N_B}. \quad (5.24)$$

The low statistics in the tails of both $\Delta\chi_{FB}^2/N$ distributions can lead to large relative fluctuations in N_F and N_B between adjacent bins, and the possibility exists that $N_F = 0$ or $N_B = 0$, leaving the value of P_F undefined. For this reason, the probability of assigning the wrong direction hypothesis to both forwards and backwards particles is approximated to be constant by setting the bin entries in the tails of both distributions to a constant value. The tail region is defined by finding the last bin for which $\Delta\chi_{FB}^2/N > 0$ for true forwards muons and $\Delta\chi_{FB}^2/N < 0$ for true backwards muons and for which the bin content $\geq 0.25\%$ of the total number of entries in the distribution. This bin is labelled by the index t . In the true forwards $\Delta\chi_{FB}^2/N$ distribution, the tail region consists of every bin for which the bin index $i > t$, and in the true backwards distribution every bin for which $i < t$. The threshold value of 0.25% has been chosen to ensure that a consistent definition for the tail region exists, which can be readily applied to distributions with different numbers of entries. The bin content C_i for each bin i in the tail regions is set to the average bin entry in the tail region. For forwards tail entries


$$C_{i>t} = \frac{1}{n-t} \sum_{k=t+1}^n C_k, \quad (5.25)$$




(a) Distribution of $\Delta\chi^2_{FB}/N$ for true forwards muons.



(b) Distribution of $\Delta\chi^2_{FB}/N$ for true backwards muons.




(c) Distribution of $\Delta\chi^2_{FB}/N$ for true forwards muons with extended tails: each bin with index $i > t$ is set to the average value in this region, as described in the text.



(d) Distribution of $\Delta\chi^2_{FB}/N$ for true backwards muons with extended tails: each bin with index $i < t$ is set to the average value in this region, as described in the text.

Fig. 5.24 Distribution of $\Delta\chi^2_{FB}/N$ for true contained forwards muons (5.24a) and true contained backwards muons (5.24b). The same distributions with extended tails are shown in 5.24c and 5.24d, obtained by setting tail entries to the average value in the tail region (as described in the text). The distributions with the extended tails can be used to calculate a bin-by-bin forwards probability P_F using equation 5.24.



(a) An example of a true backwards track that is reconstructed as a forwards track because a large fraction of the track lies within gaps, and high- \tilde{Q} values are present in the event due to charge merging.

(b) An example of a true backwards track that is reconstructed as a forwards track because there is little structure in the observed values of \tilde{Q} and there is a high- \tilde{Q} impure value present that is not removed by the Bragg peak filter.

Fig. 5.25 Two examples of true backwards events that have been incorrectly reconstructed.

and for backwards tail entries:

$$C_{i<t} = \frac{1}{t-1} \sum_{k=1}^{t-1} C_k. \quad (5.26)$$

Applying the probability calculation defined in equation 5.24 to the $\Delta\chi_{FB}^2/N$ distributions with extended tails shown in Figures 5.24c and 5.24d yields the probability curves shown in Figure 5.26. A comparison is made between the values of P_F i) with the splitting and the Bragg peak filter functionality disabled (red), ii) with only the Bragg peak filter enabled (blue) and iii) with the splitting and the Bragg peak filter enabled (green). This figure shows that both the Bragg peak filter and the particle splitting method have a beneficial effect on the underlying $\Delta\chi_{FB}^2/N$ distributions, and hence on the values of P_F . In practice, the values of P_F are set to $P_F = 0.5$ in the region where $|\Delta\chi_{FB}^2/N| \geq 40$, since the extended tails start to dominate in this region.

5.6.1 Applications

Now that a probability P_F for the forwards direction in the W plane has been formulated, this probability can be used to improve the reconstruction logic. In particular, the 2D forwards probability can be mapped onto a 3D cluster to get a direction measure for any track in the detector. An immediate application of this technique is the removal of cosmic rays, where




Fig. 5.26 Probability curves showing P_F as a function of $\Delta\chi^2_{FB}/N$. The probability P_F has been calculated using equation 5.24 and the distributions of Figures 5.24c and 5.24d. A comparison is made between the values of P_F with both the splitting and the Bragg peak filter functionality disabled (red), with only the Bragg peak filter enabled (blue) and with the Bragg peak filter and particle splitting enabled (green).

the direction of long tracks can aid tremendously in separating neutrino-induced particles from cosmic-ray tracks. For instance, in the simplest case, no neutrino-induced particle will enter the detector through its top face and have a downwards direction. The applications of the direction finding method are discussed in Chapter 6.

Chapter 6

Reconstruction Improvements

This chapter outlines several improvements that have been made to the reconstruction logic. Section 6.1 gives details of two new types of vertex candidates (as defined in Section 4.4.3) that have been developed, which significantly improve the vertex resolution for several channels. In addition, the direction finding method that was described in Chapter 5 is applied to the removal of cosmic ray muons. Section 6.2 outlines how directional information can be used to formulate a probability measure that an observed track is a cosmic ray muon, which is designed to be used in probabilistic cosmic ray removal approaches. Furthermore, Section 6.3 uses directional information to formulate a directional cosmic ray tagging (DCRT) procedure, which aims to work with the existing Pandora cosmic ray tagging functionality, and targets those cosmic ray particles that were not flagged by the existing cosmic ray tagging tool. This section demonstrates that it is possible, using directional information, to tag additional cosmic ray particles while removing virtually no neutrino-induced events. The event sample that is used in this chapter is the main analysis sample, introduced in Section 4.5.

6.1 Vertexing Improvements

Section 4.4 in Chapter 4 has outlined that accurate 3D vertex reconstruction is a crucial point in the reconstruction process. Information from the vertex position is used to inform reconstruction decisions, such as whether to split clusters, and improvements to the vertex reconstruction are therefore likely to improve the reconstruction. This section outlines several such improvements, which mainly revolve around creating more accurate vertex candidates for the vertex selection procedure to choose from. Although the vertex selection logic is very sophisticated, the vertex candidate creation logic outlined in Chapter 4 is relatively

simple, relying only on the consistency of the x coordinate between views. While matching 2D track endpoints between views works particularly well for simple event topologies, when using only these regular vertex candidates, there are still events in which there is no vertex candidate near the true position of the interaction vertex. This is often caused by a sparsity of hits near track endpoints, or by the fact that parts can be obscured by inactive detector regions. As such, other vertex candidate creation mechanisms must be developed to compensate for these possibilities. The performance of the vertex reconstruction is assessed through the quantity ΔR , the 3D distance between the selected vertex candidate and the true interaction vertex after applying space charge correction, for the existing vertex reconstruction logic. The vertex ΔR distributions in this section are made using 1×10^5 events from the simulated sample, which contain sufficient statistics to resolve how ΔR changes with the outlined reconstruction changes.

In addition to using the correspondence of the x coordinate between views, an additional fact that may be leveraged to improve the vertex candidate creation is that in each view, all particles point back to their point of origin. The 2D sliding linear fits (introduced in Chapter 4) to the particles in the event can be extrapolated in each view, as is done in the standard vertex candidate creation logic, but instead of matching track endpoints between views, it is possible to match the points at which these extrapolations intersect between views. These points are referred to as ‘intersection points’, and if such intersection points can be matched between views, they represent high-quality vertex candidates. The logic used to match intersection points between views is identical to the logic used to match track endpoints in the candidate creation logic described in Section 4.4. This process is illustrated schematically in Figure 6.1, where the sliding fits are displayed in black, the extrapolations are shown in blue, and the crossing points are shown in red. The crossing points are said to match if they are within 3 cm of one another in x . Each pair of matched crossing points are used to create a new vertex candidate, and the logic for doing so is the same as the logic used to create vertex candidates from matched sliding fit endpoints in the standard vertex candidate creation logic.

The result of introducing these new ‘crossing vertex candidates’ on the vertex ΔR can be seen in Figure 6.2. It is clear from the increased height of the ΔR peak in the CCQEL $\mu^- + p$ channel that the introduction of crossing candidates has a beneficial impact on the vertex resolution in the CCQEL $\mu^- + p$ channel. This same beneficial impact is visible in other channels with primary particle multiplicity 2, as is outlined in Table 6.1. The crossing vertex prescription works especially well for events in MicroBooNE, where most of the neutrino energies result in quasi-elastic scatters or resonance production, which result in only a few tracks. However, for deep inelastic scatters, or events in other liquid Argon contexts, such as





Fig. 6.1 An illustration showing how crossing vertex candidates are constructed. Sliding fits to tracks (black lines) are extrapolated a fixed amount (indicated by the dashed blue lines), and intersections of these extrapolations (indicated by red points) are matched between views to create new vertex candidates.

DUNE, the number of tracks can be very large, and investigating all crossing point pairings and evaluating all new vertex candidates can be very computationally expensive.



(a) Vertex ΔR distribution for the CCQEL μ^- channel.

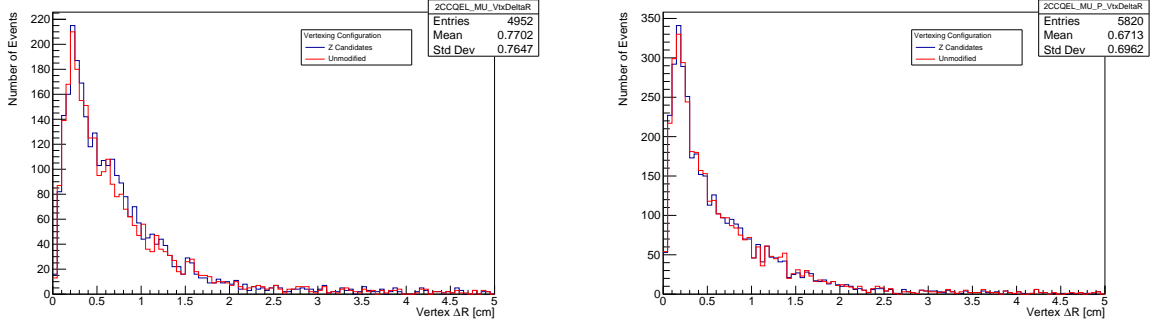
(b) Vertex ΔR distribution for the CCQEL $\mu^- + p$ channel.

Fig. 6.2 The impact of introducing crossing vertex candidates on the vertex ΔR distributions in the CCQEL μ^- and CCQEL $\mu^- + p$ channels. The distribution due to the current reconstruction is indicated in red, and the distribution after introducing crossing candidates is indicated in blue. The impact of these candidates on the CCQEL μ^- channel is minimal because the vast majority of events in this category are reconstructed as single tracks. In contrast, CCQEL $\mu^- + p$ channel typically have a multiplicity of 2, and so crossing candidates have a large impact.

Both the standard candidate vertex creation and the crossing vertex candidate creation assume that there are well-formed track-like particles in at least two views, for which 2D Pandora sliding fit objects can be created and extrapolated. However, this is not always the case. Figure 6.4 shows an event that has a simple topology: it is a CCQEL μ^- event, with one track in each view. The track exhibits significant curvature, but this does not affect the creation of a sliding fit, by construction. What is problematic, however, is that inactive detector regions have left the tracks quite broken up near the interaction vertex, in both the U and W views. As a result, the clusters and sliding fits are truncated near the low- z end of the event, and the standard candidate creation is unlikely to create useful candidates, as the extrapolation is linear and is affected by the curvature. However, in each view there are still hits present in the region near the interaction vertex. This simple hit-based information can be used to construct an additional type of vertex candidate, henceforth denoted as ‘ z candidates’. The logic relies on the same principle of using positions from different views and using these positions to create vertex candidates. In this case, the hit positions in each view are sorted by their z coordinates, and the k hits with the lowest z positions are selected in each view. In what follows, the number of low- z hits per view k that are used is set to $k = 5$. For these $3k$ hit positions across the three views, each combination of pairs of hit positions is used to create a vertex candidate. This procedure results in the creation of $3k^2$ vertex candidates, a number that becomes very large even for moderate k . For the example case $k = 5$, 75 new vertex candidates are produced, which would represent a very large increase in computation time, if each vertex candidate was evaluated. However, most z candidates are removed by the vertex candidate quality cut, as many hit pairings result in candidate positions that are not near 3D hits.

The result of introducing these ‘ z candidates’ on the vertex ΔR distribution is shown in Figure 6.3, which shows that this candidate type provides a clear benefit in the vertex resolution. The impact of these candidates is largest in events with single tracks, since the absence of secondary particles represents a sparsity of information. When more than one track is present in the event, the standard vertex candidate creation algorithm is able to create a large number of good candidates, and the relative impact of introducing z candidates is less. As such, there is a clear benefit to using both new candidate types simultaneously, since the two types of vertex candidates are effective in a complementary set of interaction channels.

The improved vertex resolution that results from introducing these two new types of vertex candidates also has a beneficial effect on the reconstruction performance as a whole. Table 6.1 shows the impact of the new vertex candidates on the Pandora reconstruction performance. In this table, the first column corresponds to the performance outlined in Section 4.5. The new vertex candidates have a general beneficial effect on the fraction of correct events, increasing



(a) Vertex ΔR distribution for the CCQEL μ^- channel.

(b) Vertex ΔR distribution for the CCQEL $\mu^- + p$ channel.

Fig. 6.3 The impact of introducing crossing vertex candidates on the vertex ΔR distributions in the CCQEL μ^- and CCQEL $\mu^- + p$ channels. The distribution due to the current reconstruction is indicated in red, and the distribution after introducing crossing candidates is indicated in blue. The impact of these candidates on the CCQEL μ^- is largest, because the absence of secondary particles introduces information sparsity, such that the relative impact of introducing additional information in the form of new candidates is largest.




Fig. 6.4 Example event in which z candidates create an improvement. The event contains one track in each view. The track curvature affects sliding fit interpolation, affecting the standard vertex candidate creation logic. Inactive detector regions have obscured the low- z ends of the tracks. Note the different scales in z .




Fig. 6.5 The effect of turning on z candidates in an example event. Vertex candidates that pass the on hit/in gap quality cut are displayed in red. After turning on the z candidates, new vertex candidates are produced at the low- z track endpoint, near the true vertex position.

it by roughly 1% in the CCQEL μ^- and CCMEC μ^- channels. Small reductions in the fraction of correct events are generally due to the complex impact of the reconstruction improvements on the Pandora matching logic described in Section 4.5. Furthermore, as was concluded from the vertex ΔR distributions for the different vertex candidate configurations, introducing both crossing- and z -candidates has the largest beneficial impact. As such, these vertex candidates are switched on by default in all plots and proceedings henceforth.

<i>Fractions Correct Per Vertex Configuration</i>					
Interaction Type	Unmodified	Crossing Candidates	Z Candidates	Both Candidates	Total
CCQEL μ^-	17071 (65.93%)	17052 (65.86%)	17264 (66.68%)	17315 (66.87%)	25892
CCQEL $\mu^- + p$	19228 (64.70%)	19272 (64.85%)	19207 (64.63%)	19250 (64.77%)	29720
CCRES μ^-	703 (67.08%)	701 (66.89%)	705 (67.27%)	704 (67.18%)	1048
CCRES $\mu^- + p$	2056 (65.02%)	2053 (64.93%)	2062 (65.21%)	2059 (65.12%)	3162
CCMEC μ^-	3800 (66.70%)	3801 (66.72%)	3859 (67.74%)	3854 (67.65%)	5697
CCMEC $\mu^- + p$	9133 (68.82%)	9143 (68.90%)	9142 (68.89%)	9155 (68.99%)	13270

Table 6.1 The impact of the the new vertex candidates on the Pandora reconstruction performance. The indicated fractions are the fraction of correct events, using the Pandora correctness criterion, and the Total column shows the total number of events per interaction channel.

6.2 Cosmic Probability

The MicroBooNE detector, as a surface detector, is subject to a large number of cosmic-ray background events. The muon cosmic-ray rate in the MicroBooNE detector is estimated to be 5.5 kHz, which corresponds to approximately 20 muons per TPC drift time window of 4.8 ms [94]. In addition, only roughly 1 in 600 beam spills is expected to result in a neutrino interaction. As such, removing as many cosmic rays as possible has been a point of focus for the MicroBooNE reconstruction effort. Many avenues have been explored to reduce the cosmic background, including the installation of a cosmic ray tagger (CRT) module [79], matching PMT activity to the beam spill windows [62] and multivariate cosmic ray rejection [95].

Since cosmic-ray tracks originate in the atmosphere, they are almost exclusively travelling in the downwards- y direction, which is information that has not yet been leveraged in cosmic ray removal efforts. This section outlines the application of the directional probability outlined in the previous Chapter to the removal of cosmic ray tracks, by recasting the forwards-backwards fit quantity $\Delta\chi_{FB}^2/N$ to the upwards-downwards equivalent $\Delta\chi_{DU}^2/N$. Subsequently, this new quantity is used to formulate a probabilistic measure P_c that a track is a cosmic ray. Furthermore, the correlation of P_c with hit count, polar and azimuthal angles and fit χ_{min}^2/N is taken into account. The P_c probability curves are then parametrised, so that for any observed track and its resulting direction fit, an appropriate downwards probability may be formulated. Finally, this procedure is assessed on real cosmic-only data from the MicroBooNE detector.

The cosmic analysis in this section uses the simulated cosmic-ray tracks from a sample also containing simulated neutrino interactions. Ignoring neutrino-induced events is a simplification, but the main aim of this treatment is to formulate a probabilistic measure that a track is a cosmic ray, based on observed directional information. Such a probabilistic measure can be applied in other cosmic ray removal treatments, and presents a possible extension to the aforementioned multivariate methods in [95], in which context a probabilistic quantity is especially useful.

6.2.1 Formulating Downwards and Upwards Hypotheses

The direction fitting procedure described in Chapter 5 was developed using a sample of quasi-elastic ν_μ interactions, and a distinction was made between particles travelling in the forwards and backwards directions along the z axis. The quantity P_c is designed to identify

downwards-travelling cosmic track-like particles without removing neutrino-induced tracks. Within this context, a cosmic ray tagging procedure aims to distinguish between cosmic-ray tracks that enter the detector near its top face and neutrino-induced track-like particles that leave the detector. The former category is exclusively downwards-going and may have a fiducial low- y endpoint, with associated directional information contained in the Bragg peak. Cosmic-ray tracks that do not stop in the detector can be removed trivially, since any track that intersects two detector faces must be a cosmic ray. In addition, neutrino-induced tracks that leave the detector do not have contained Bragg peaks, implying an absence of directional information.

Three-dimensional track information from the Pandora reconstruction can be used to associate the high- and low- z endpoints of the 2D direction fit to 3D high- and low- y endpoints. In doing so, one can reformulate the forwards-backwards directional quantity $\Delta\chi_{FB}^2/N$, defined in equation 5.23, to the upwards-downwards equivalent

$$\Delta\chi_{DU}^2/N \equiv (\chi_D^2 - \chi_U^2)/N, \quad (6.1)$$

where the subscripts D and U now represent ‘downwards’ (high- y to low- y) and ‘upwards’ (low- y to high- y), respectively. An example $\Delta\chi_{DU}^2/N$ distribution is given in Figure 6.6a. The distribution in Figure 6.6a is predominantly negative, as expected for downwards-travelling particles. Note that there is a larger peak at $\Delta\chi_{DU}^2/N \approx 0$, compared to the distribution of $\Delta\chi_{FB}^2/N$ in Figure 5.24a in the previous chapter. This is mainly due to the fact that the direction fit is performed using calorimetric information from the W view, and while neutrino-induced events have a net momentum p_v that is directed along the z -axis ($p_v \cdot \hat{z} \approx 1$), this is not the case for cosmic rays. Cosmic-ray tracks can be almost parallel with the wires in the collection plane, making it much harder to reconstruct individual calorimetric hits, causing a loss of directional information.

Figure 6.6a was created using cosmic ray muons with true (Monte Carlo) fiducial low- y endpoints. In the rest of this section, only such tracks are considered. When identifying cosmic rays with Bragg peaks in practice, it is not sufficient to demand that they have contained reconstructed low- y endpoints, since cosmic rays can be out of sync with respect to the beam spill time window. This effectively shifts the track in x , due to the violated assumption that t_0 coincides with the time of the beam spill, making some cosmic rays appear as though they have a fiducial endpoint, but in reality they have left the detector and no Bragg peak is present in the track. Figure 6.7a shows the distribution of $\Delta\chi_{DU}^2/N$ for cosmic ray muons with a reconstructed low- y endpoint. It is apparent from this figure that time-shifted cosmic rays contribute a large peak at $\Delta\chi_{DU}^2/N \approx 0$. To address this issue, a ‘Bragg peak





Fig. 6.6 Distributions of $\Delta\chi^2_{DU}/N$ (6.6a) and χ^2_{min}/N (6.6b) for simulated cosmic ray muons. These plots demonstrate that $\Delta\chi^2_{DU}/N$ entries are predominantly negative, indicating downwards-going particles and that χ^2_{min}/N peaks near 1, indicating high-quality fits. The peak near $\Delta\chi^2_{DU}/N \approx 0$ is predominantly caused tracks for which the Bragg peak is not apparent, due to their orientation in the detector. These figures were created using an older version of the simulation, but they are representative.

finder' (BPF) is implemented using the information contained in the direction fit parameters (described in Section 5.2). The fit parameter T_e quantifies the kinetic energy of the particle near the Bragg peak, and a cut on this quantity retains only those direction fits for which the direction fit has identified a Bragg peak (either correctly or incorrectly). Figure 6.7b shows the effect of applying the cut $T_{end} \leq 25$ MeV to the distribution of 6.7a. The result of the Bragg peak finder is effectively to suppress the peak at $\Delta\chi^2_{DU}/N = 0$ by removing tracks for which the direction fit was unable to resolve directional information. The Bragg peak finder replaces the Monte Carlo low- y containment criterion in the cosmic ray removal study of Section 6.3, but to develop the cosmic probability logic of this section, the Monte Carlo contained endpoint requirement is used.

The $\Delta\chi^2_{DU}/N$ distribution in Figure 6.6a can be used to formulate a probability P_F that a track is travelling in the downwards- y direction, as was done for the forwards/backwards case in Section 5.6. In addition, Figure 6.6b demonstrates that χ^2_{min}/N peaks near $\chi^2_{min}/N \approx 1$, indicating that the fits fit the observed distributions of \tilde{Q} well. One possible approach to formulating a probability distribution is to utilise a method similar to what was developed for the forwards probability in Chapter 5 in Figure 5.26, which is a bin-by-bin frequentist probability based on the underlying distributions of $\Delta\chi^2_{FB}/N$ for true forwards- and backwards-going muons. For cosmic ray particles, however, it is possible to use the additional information that cosmic-ray tracks are exclusively downwards-going, and that the tracks they should be distinguished from are neutrino-induced tracks that leave the detector through the top face.



(a) The distribution of $\Delta\chi^2_{DU}/N$ for all tracks with a reconstructed fiducial low-y endpoint.

(b) The distribution of $\Delta\chi^2_{DU}/N$ for all tracks with a reconstructed fiducial low-y endpoint and with the Bragg peak finder (BPF) applied.

Fig. 6.7 Distributions of $\Delta\chi^2_{DU}/N$ for true cosmic ray muons. Figure 6.7a shows the effect of demanding that the reconstructed low-y endpoint is fiducial. Figure 6.7b shows the effect of applying the Bragg peak finder (BPF) to the distribution in 6.7a, which demands that the direction fit parameter $T_{end} \leq 25$ MeV.

The latter class of tracks do not have a contained endpoint, and hence no Bragg peak, and predominantly contribute values of $\Delta\chi^2_{DU}/N \approx 0$, with tails where the direction fit has found a ‘fake’ Bragg peak in the form of a spurious highly-charged hit or hits. One can assume that there is a fixed probability P_{fake} of reconstructing a false Bragg peak where none exists, such that P_{fake} is defined by the fraction of entries with $\Delta\chi^2_{DU}/N > 0$ in Figure 6.6a. Since such entries and neutrino-induced tracks with fake Bragg peaks for which $|\Delta\chi^2_{FB}/N| \gtrsim 1$ occur through this same mechanism, one can assume that the probability of reconstructing a neutrino-induced track with $|\Delta\chi^2_{DU}/N| \gtrsim 1$ is also P_{fake} . In other words, it is possible to estimate the behaviour of the neutrino-induced $\Delta\chi^2_{DU}/N$ distribution from the cosmic-induced distribution by quantifying P_{fake} . Hence, for a given bin with index $-k$ where $k = 1 \dots N_B$ in Figure 6.6a, dividing the bin content $N(-k)$ of b by $N(k)$ yields the probability P_c that the observed track is downwards-going, or equivalently that it is a cosmic-ray track, where

$$P_c = \frac{N(-k)}{N(-k) + N(k)} \quad \forall k = 1, 2, \dots, N_B. \quad (6.2)$$

This procedure is equivalent to dividing the cosmic-induced distribution by the neutrino-induced distribution, if the direction fit failures are equally likely for both. The advantage of this procedure is that the machinery relating to P_c can be developed using a simulated cosmics-only sample. Note that equation 6.2 is only valid in the regime where $\Delta\chi^2_{FB}/N < 0$. In the region where $\Delta\chi^2_{FB}/N \geq 0$, the tails of the cosmic- and neutrino-induced distributions

are divided by one another, and since a constant failure rate P_{fake} is assumed for both, $P_c = 0.5$ in this regime, indicating a lack of directional information.

An example probability curve is given in Figure 6.8a, constructed from the $\Delta\chi^2_{DU}/N$ distribution of Figure 6.6a using equation 6.2. The probability curve has certain consistent features. For instance, in the region where $\Delta\chi^2/N \approx 0$, the probability P_c that a track is travelling in the downwards-y direction is approximately 0.5. In this region, no Bragg peak has been found by the direction fit, and the fits corresponding to both hypotheses are of comparable quality, such that $\Delta\chi^2/N$ is equally likely to take small positive or small negative values. After this, as $|\Delta\chi^2/N|$ increases, P_c quickly rises to a maximum value P_{max} , corresponding to the peak of the $\Delta\chi^2/N$ distribution, after which the value of P gradually decays, as bins in the tails of the distribution start to be considered. These regular features in Figure 6.8a can be used to posit a functional form for the probability curve, which can be used to create a fit, as shown in Figure 6.8b. A functional form is posited for P_c of the form

$$P_c(x) = \frac{1}{2} + \left(P_0 - \frac{1}{2}\right) (1 - e^{-\alpha x}) e^{-\beta x}, \quad (6.3)$$

where x is shorthand for $\Delta\chi^2_{DU}/N$. The functional form of equation 6.3 has the following desirable properties:

- $P(0) = 1/2$ and $1/2 \leq P(x) \leq 1$
- A growth factor α determines how quickly the value of $P(x)$ goes to P_0
- A decay factor β determines how quickly the value of $P(x)$ decays to $P(x) = 1/2$

The values of P_0 , α and β can be minimised by a package like Minuit [92] to obtain the fit in Figure 6.8b. Note, however, that in equation 6.3 $P_0 \neq P_{max}$, and that P_0 can take values larger than 1. This is undesirable, as it makes the parameters less interpretable. In addition, it has been observed in practice that optimisation packages like Minuit converge faster and more reliably when P_{max} is given as a fit parameter, in addition to α and β . The following modification resolves this problem. First, define the maximum of $P(x)$ through the equation

$$\frac{d}{dx}P(x_0) = 0, \quad (6.4)$$

such that $P(x_0) \equiv P_{max}$. The value of x_0 can be found by differentiating equation 6.3, giving

$$x_0 = \frac{\ln\left(1 + \frac{\beta}{\alpha}\right)}{\alpha}, \quad (6.5)$$


and equation 6.3 can be evaluated at $x = x_0$, which yields

$$P(x_0) = \left(1 - \frac{\beta}{\alpha + \beta}\right) \left(\frac{\alpha + \beta}{\beta}\right)^{-\beta/\alpha}. \quad (6.6)$$

The next step is to ensure that $P(x_0) = P_{max}$, which is true when

$$P_0 = \frac{1}{2} + \frac{1}{2\alpha} \left((2\alpha P_{max} + 2\beta P_{max} - \alpha - \beta) \left(\frac{\alpha + \beta}{\beta}\right)^{\beta/\alpha} \right). \quad (6.7)$$

Using the definition of P_0 given by equation 6.7 together with 6.3 and by making P_{max} , α and β fit parameters, accurate probability fits are obtained, such as the fit shown in Figure 6.8b.



(a) P_c probability curve for all cosmic ray muon tracks.

(b) P_c probability curve for all cosmic ray muon tracks with fit.

Fig. 6.8 A distribution of P_c , as a function of $-\Delta\chi^2_{DU}/N$. Point errors are binomial. The points on this probability curve are constructed using the prescription in equation 6.2 on the distribution of Figure 6.6a. The fit function in 6.8b is given in equation 6.3.

Lastly, the correlation of P_c with important quantities such as χ^2_{min}/N , N , L , ϕ and θ has been examined. Figure 6.9 shows how the probability curves vary for bins in these quantities. A number of things are apparent in Figure 6.9. Firstly, the probability curves change minimally with L , ϕ and θ . Secondly, there is a dependence of P_c on χ^2_{min}/N and N . In particular, the dependence on χ^2_{min}/N is most pronounced, which is due to the fact that χ^2_{min}/N is a direct measure of the goodness of the best fit. Thirdly, there is a dependence of P_c on N , because this variable is itself highly correlated with the particle energy E . Tracks with few hits have empirically been found to be more difficult to fit, since such tracks have low energies and all hits deposited by them are within the Bragg peak energy regime, which can have high variance. Lastly, it must be taken into account that χ^2_{min}/N and N are by construction highly correlated. To resolve this correlation, the dependence of P_c on N is plotted while keeping χ^2_{min}/N ‘fixed’ by constraining this quantity to one of the χ^2_{min}/N bins

in Figure 6.9d. The resulting variation plots are given in Figure 6.10, and it is apparent from these plots that there is a dependence of P_c on N even when χ_{min}^2/N is constrained.

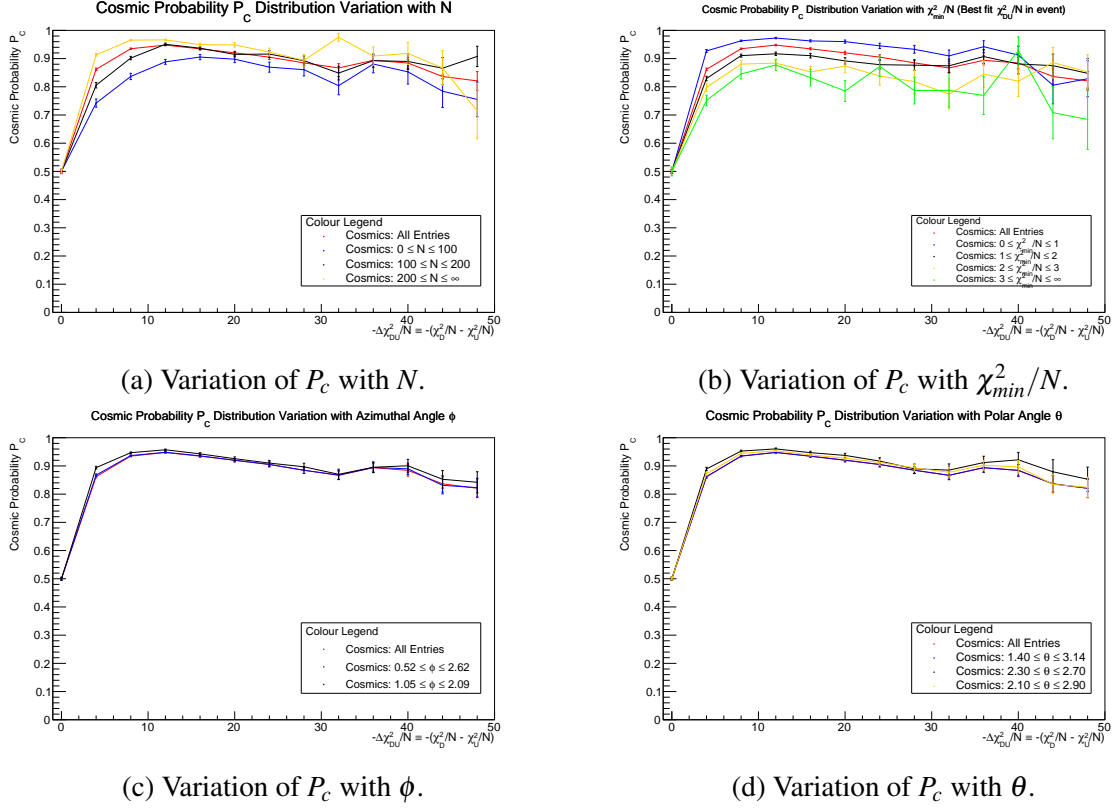


Fig. 6.9 The variation of the downwards probability P_c with N , χ_{min}^2/N , ϕ and θ .

6.2.2 Parametrising Probability Curve Fits

Having established that P_c is a function of N and χ_{min}^2/N and that the probability curves can be fitted accurately using equation 6.3, the fit parameters P_{max} , α and β can be parametrised as functions of N and χ_{min}^2/N . Doing so will make it possible to analytically calculate an estimate of P_c for a value of $\Delta\chi_{DU}^2/N$, N and χ_{min}^2/N without having to perform a lookup in pre-stored probability tables, for instance. The parametrisation of P_{max} , α and β in terms of χ_{min}^2/N and N is achieved by creating a ‘parameter surface’ that shows the values of these fit parameters for a binned grid of χ_{min}^2/N and N values. The parameter surfaces for P_{max} , α and β are shown in Figure 6.11. Note that as N increases, the fits typically become better, and P_c can be determined with higher accuracy. This is reflected in the behaviour of the parameters: the value of P_{max} increases with N , reflecting the fact that the probability curves peak at a higher value. Furthermore, α increases with N , representing the fact that P_c increases

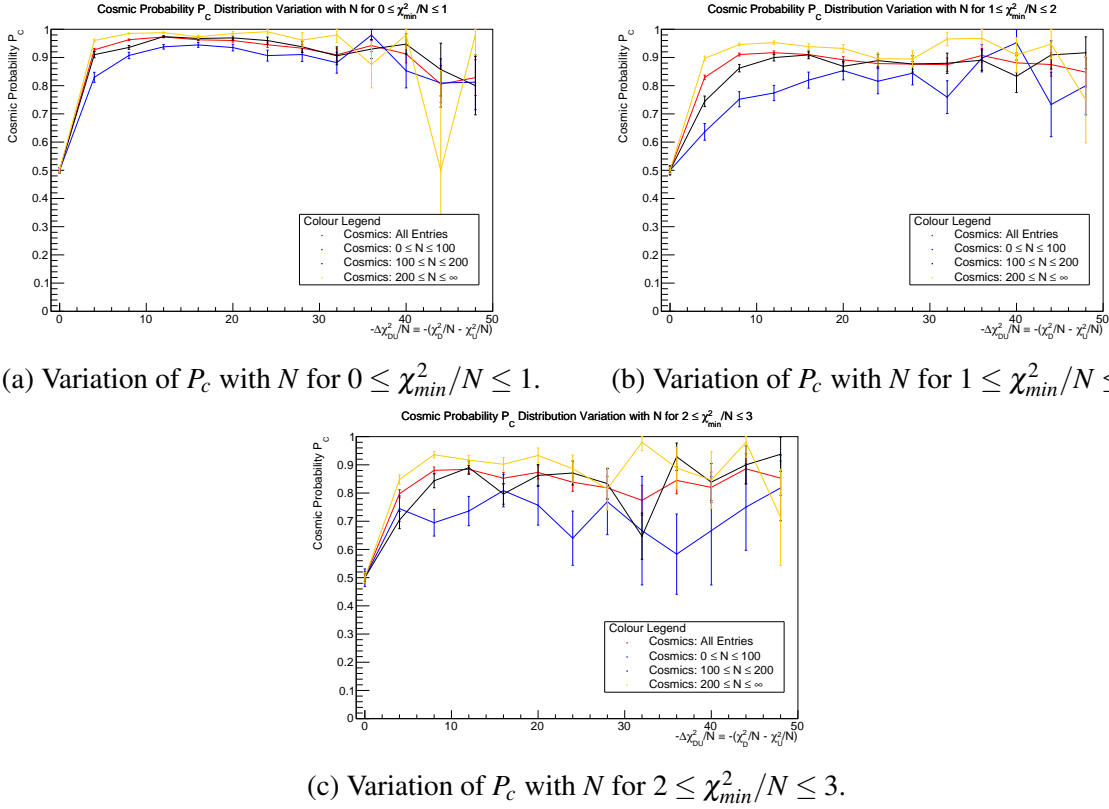


Fig. 6.10 The variation of P_c with N while keeping χ^2_{min}/N fixed. These plots show that the variation of P_c with N is not solely due to the correlation between N and χ^2_{min}/N . The choice of binning for χ^2_{min}/N corresponds to the choice of binning for χ^2_{min}/N in Figure 6.9. The last bin $4.0 \leq \chi^2_{min}/N \leq \infty$ is excluded, because it has very low statistics.

faster with $|\Delta\chi^2/N|$. The reverse relationships hold for the behaviour of these parameters with χ_{min}^2/N , because the underlying fits become better as χ_{min}^2/N decreases. In general, the values of β are small and the errors on the points are very large, which heuristically is due to the fact that the probability curves decay slowly, and variations in this parameter have minimal impact on the probability curve. Since the variation of this parameter is so small, in what follows this parameter is fixed to $\beta = 0.01$ and excluded from the fit.

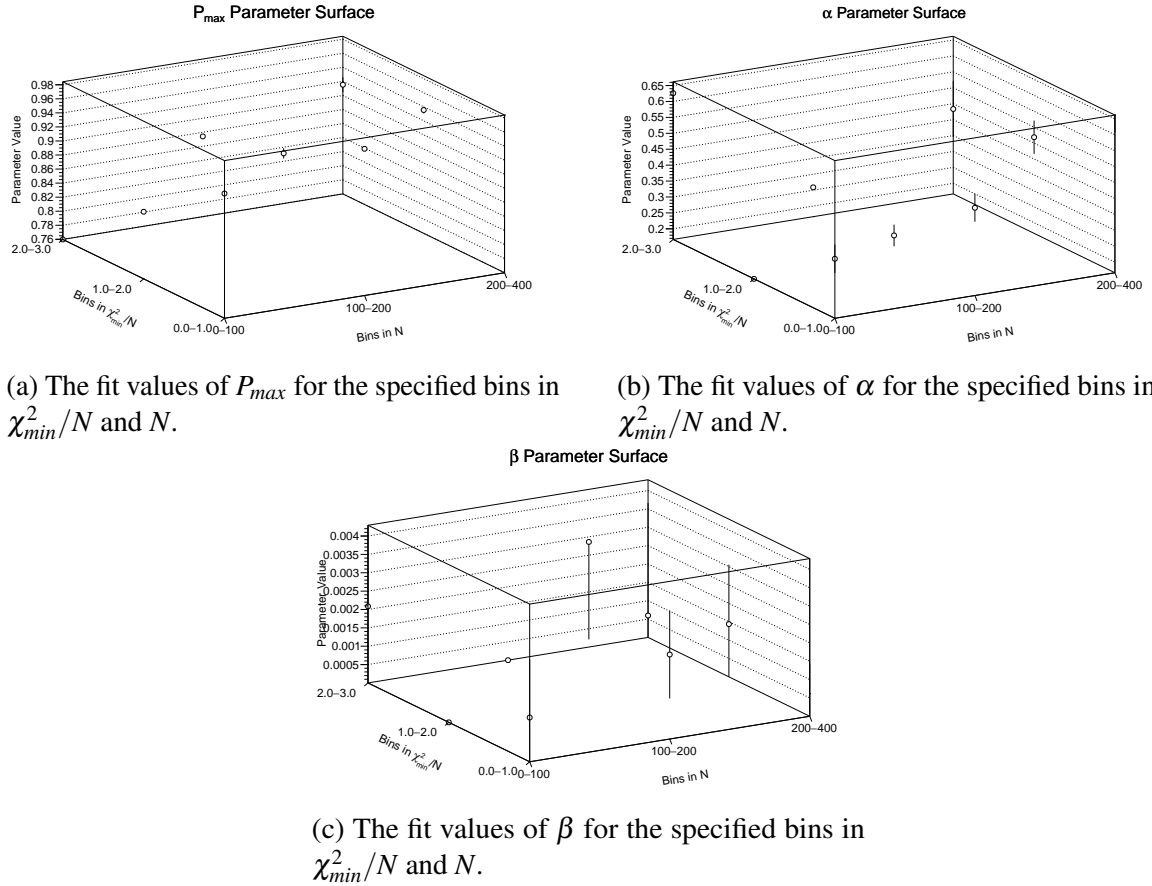



Fig. 6.11 The fit values for the parameters P_{max} , α and β for $75 \leq N \leq 225$ in bins of 75 and $0.25 \leq \chi_{min}^2/N \leq 2.25$ in bins of 0.5.

A simple surface of the form given in equation 6.8 can be fitted to the parameter surfaces of Figure 6.11. In equation 6.8, p_n describes the two parameters P_{max} and α ($n = 0, 1$), and a_n , b_n and c_n are three surface-specific fit parameters. This results in the surface fits depicted in Figure 6.12.

$$p_n = a_n + b_n N + c_n \chi_{min}^2/N \quad n = 0, 1 \quad (6.8)$$

Having obtained these surface fits, it is possible to compare the direct fits of equation 6.3 to the nine underlying parameter curves in each of the figures in Figure 6.12 to the probability

curves constructed from the parametrisation of the parameters P_{max} , α and β . This, in essence, is a visualisation of the impact that the deviation between the parameter surface points and the surface fits has on the resulting probability curves. A number of selected underlying probability curves are given in Figure 6.13.




(a) Parameter surface fit to the fit values of P_{max} .

(b) Parameter surface fit to the fit values of α .


Fig. 6.12 The fit values for the parameters P_{max} and α as shown in Figure 6.11, but with a fitted 2D surface of the form $p_n = a_n + b_n N + c_n \chi_{min}^2 / N$.

6.3 Directional Cosmic Ray Tagging


The directional variables $\Delta\chi_{DU}^2 N$ and P_c can be used to enhance the performance of cosmic ray tagging (CRT) procedures. In particular, this section demonstrates that it is possible to formulate a directional cosmic ray tagging (DCRT) procedure that aids in the removal of cosmic-ray tracks while minimally affecting the neutrino interactions in the sample. In general, the purpose of these variables is to enhance the cosmic ray tagging logic that already exists within the Pandora framework. The Pandora cosmic ray tagging logic uses the PandoraCosmic pass to reconstruct all particles as tracks or δ -rays, and creates associations between particles by determining whether the point of closest approach of the trajectories of potential daughter particles are near to the endpoints of a parent particle. All associated particles are assigned to their own detector slice, and each slice is flagged as being either an obvious cosmic ray particle, in which case it is removed, or a likely neutrino interaction, in which case it is dissolved into its constituent hits and reprocessed by both the PandoraCosmic and PandoraNu passes. To determine which slices potentially contain a neutrino interaction, margins in the x , y and z directions are defined to be 5, 20 and 10cm from the detector edges, respectively, and for each primary particle in the slice, a check is performed whether both endpoints are within certain ‘in-time’ x thresholds, are contained within y and z margins, and




(a) Underlying P_c distribution for $0 \leq N \leq 100$ and $0 \leq \chi^2_{min}/N \leq 1$.



(b) Underlying P_c distribution for $100 \leq N \leq 200$ and $0 \leq \chi^2_{min}/N \leq 1$.



(c) Underlying P_c distribution for $200 \leq N \leq 400$ and $0 \leq \chi^2_{min}/N \leq 1$.



(d) Underlying P_c distribution for $200 \leq N \leq 400$ and $2 \leq \chi^2_{min}/N \leq 4$.

Fig. 6.13 Underlying P_c distributions taken from the range $0 \leq N \leq 400$ in bins of 100 and 200 for the last bin, to increase the underlying number of events, and $0 \leq \chi^2_{min}/N \leq 4$ in bins of 1 and 2 for the last bin.

whether the polar angle $\cos(\theta) \leq 0.6$. If all these conditions are met, the slice is retained as a potential neutrino-containing slice. If a primary particle in a slice is not in a probable neutrino-containing slice, and is either not in-time, does not meet the polar angle cut, or if the particle extends beyond both y margins, the slice is tagged as probably containing a cosmic ray particle. Any slices that have not been assigned to either category after performing these checks are deemed to be ambiguous and are retained and processed as potential neutrino slices, to minimise the probability of discarding a neutrino interaction.

The CRT logic only flags a primary particle as an obvious cosmic ray if both its endpoints fall outside the fiducial region in y , such that stopping cosmic ray particles or neutrino-induced tracks that leave the detector near its top face are always identified as ambiguous. A DCRT procedure can help remove stopping cosmic ray tracks at the cosmic ray tagging stage, and in order to develop such a procedure a study is performed specifically on tracks that enter the detector near the top face and have not been flagged by the existing CRT logic as obvious cosmic ray particles. In order to define a region near the top of the detector into which a track must extend to be considered of interest, Figure 6.14 shows the distributions of y_{max} for both neutrino-induced and cosmic ray muons, which is the y coordinate of the highest- y endpoint of a track. In this figure and all plots hereafter, a cut has been applied to the number of hits so that $L \geq 25$ cm, designed to remove cosmic ray protons. Furthermore, this discussion focuses on cosmic-ray tracks, since direction fits cannot be constructed for shower-like particles. The cosmic ray sample therefore consists of muons, and is compared to true neutrino-induced muons. Space charge corrections have also been applied in all cases. Figure 6.14 shows that




Fig. 6.14 Distributions of y_{max} for neutrino-induced and cosmic-ray muons.

y_{max} peaks at $y = 116.5$ cm, which is the detector boundary. However, the distribution also extends beyond this point, since the PandoraCosmic pass utilises all constructed 3D hits, and

no quality cuts have been applied to these hits. As a result, 3D hits are constructed beyond the detector boundary, which are included in the PandoraCosmic track reconstruction. The distributions correspond to expectation: cosmic rays enter through the top face of the detector and so have large y_{max} values. Cosmic rays may also enter the detector near the top but through one of its side faces, accounting for values of $y_{max} < 116.5$ cm. For neutrino-induced muons, only those muons leaving the detector through its top face have y_{max} values that peak near 116.5cm. To define an appropriate ‘slice’ in the top region of the detector into which tracks must extend to be considered as potential cosmic rays, a cut is applied at $y_{max} = 100$ cm, so that most cosmic rays are retained. In other words, this analysis is limited to untagged track-like particles for which $y_{max} \geq 100$ cm and $L \geq 25$ cm.

6.3.1 Formulating Cuts

The target sample of untagged tracks with $y_{max} \geq 100$ cm and $L \geq 25$ cm still contains a large number of neutrino-induced particles: approximately 23.3% of untagged tracks are neutrino-induced. A number of variables can be constructed that can be used to create a purer sample of mainly cosmic ray particles to which DCRT can be applied, such that the likelihood of retaining as many neutrino-induced particles is maximised. The following variables capture the facts that cosmic ray particles point downwards in the detector and have higher average energies than neutrino-induced particles:

- Polar angle θ : the opening angle between the target track and the y axis is a useful variable, since neutrino-induced particles are expected to be forwards-pointing, i.e. they have a significant extent in z , and have larger θ values on average than cosmic ray particles;
- PFO charge Q_{PFO} : the summed input energy (in units of integrated ADC counts) of all the 3D hits in the PFO likely has separating power, since cosmic ray tracks are expected to have higher average energies than neutrino-induced particles, since neutrino-induced particle energies are constrained by the BNB energy spectrum. Calorimetric information from all views is used to make this quantity more robust to gaps in the W view, and the units are not scaled to MeV because this multiplication by a scalar has no impact on the distribution of the variable;
- Approximate neutrino momentum z -component $p_{v,z} \equiv \hat{p}_v \cdot \hat{z} / |\hat{p}_v \cdot \hat{z}|$: a useful quantity to construct is the momentum vector under the assumption that the primary PFO is a muon, and any daughter PFOs within 5cm of the lowest- y endpoint are protons. Cosmic

ray muons are expected to have at most a single Michel electron daughter near the contained endpoint, and neutrino-induced particles are expected to be forwards-pointing with predominantly protons near the interaction vertex. Summing the momenta under these assumptions should lead to a vector p_v that has a large z component for neutrino-induced PFOs.

Figure 6.15 shows the distributions of θ , Q_{PFO} and $p_{v,z}$ on the sample of untagged PFOs with $L \geq 25$ cm and $y_{max} \geq 100$ cm. The peak in 6.15a is due to the previously-described cut that is applied in the CRT logic at $\cos(\theta) \leq 0.6$, which corresponds to $\theta \approx 0.93$ rad, which modifies the phase space of the untagged sample. Furthermore, as expected, the θ distribution for neutrino-induced tracks peaks at higher values. The distribution of Q_{PFO} captures the fact that cosmic rays particles have higher average energies than neutrino-induced particles. The distribution of $p_{v,z}$ shows that neutrino-induced particles point predominantly along the z axis, unlike cosmic rays.

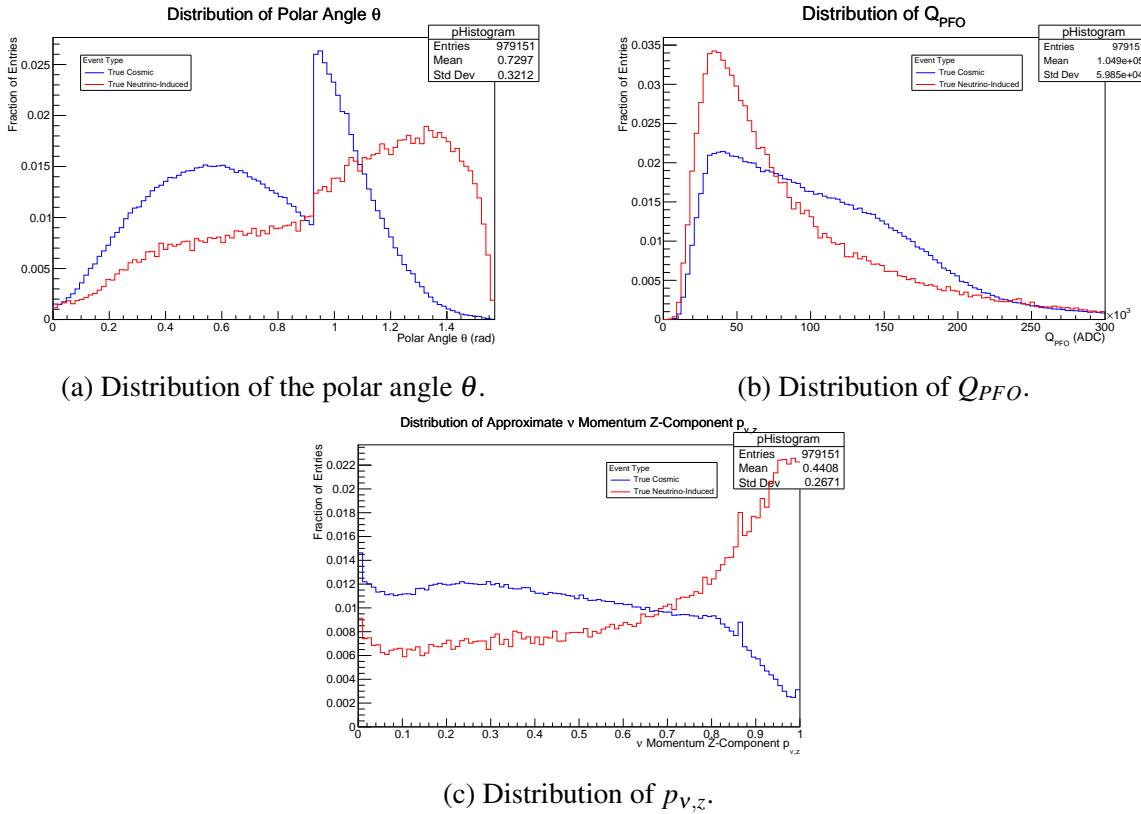


Fig. 6.15 Distributions of θ , Q_{PFO} and $p_{v,z}$ for neutrino-induced and cosmic-ray muons in the $y_{max} \geq 100$ cm, $L \geq 25$ cm untagged sample.

The distributions in Figure 6.15 can be used to create a very pure sample of cosmic ray particles. A bin-by-bin scan can be performed to find the first bin for which the integrated

event count contains $\geq 90\%$ of cosmic ray entries. This fraction has been chosen to approximate the cuts that one would make by eye, and to make these cuts systematic and consistent. The order in which these cuts should be applied is determined by the efficacy of each cut, where cut efficacy in this context refers to the amount of background rejection under fixed signal efficiency, i.e. the fraction of neutrino-induced particles removed while retaining 90% of cosmic-ray tracks. Table 6.2 shows the cut values obtained when determining the cut values sequentially, such that each cut affects the variable distribution through which the subsequent cut is found.

<i>Cut Values</i>			
Variable	Cut Value	Fraction CR Retained	Fraction ν Removed
Polar Angle θ	1.14	90.17%	56.50%
Momentum $p_{V,z}$	0.825	90.09%	66.56%
PFO Charge Q_{PFO}	31500	90.83%	80.86%

Table 6.2 Cut values retaining as close to 90% of the cosmic ray entries as possible. These cut values are obtained by applying each cut sequentially to the untagged $y_{max} \geq 100$ cm $L \geq 25$ cm sample, such that each cut affects the subsequent variable distribution.

Furthermore, Table 6.3 outlines the effect that each cut in Table 6.2 has on the sample of untagged muon tracks. In this table, a distinction has been made between neutrino-induced muon tracks that are downwards going, denoted ν_{\downarrow} , and neutrino-induced tracks that are upwards-going, denoted ν_{\uparrow} . The importance of this distinction is discussed in the next section. The initial sample of untagged muon tracks with $L \geq 25$ cm consists of approximately 17.80% neutrino-induced particles, which is reduced down to a fraction of 3.67% after applying the four additional cuts sequentially. This procedure also reduces the number of cosmic ray particles to 50.94% of the initial untagged sample with $L \geq 25$ cm, and 78.92% of the untagged sample with $L \geq 25$ and $y_{max} \geq 100$ cm. This simple cut-based procedure is therefore very effective at creating a pure sample of cosmic rays.

6.3.2 Applying Directionality

Table 6.3 shows that applying the three preselection cuts results in a sample of untagged muon tracks that predominantly consists of cosmic-ray tracks. However, it is possible to improve the procedure by incorporating directional information to remove the majority of the neutrino-induced tracks that still remain. A distinction was made in the previous section between ν_{\downarrow} and ν_{\uparrow} , where the downwards-going particles result from neutrino interactions with a true


<i>Number of Particles After Each Cut</i>			
Cut (Sequential)	Number CR	Number v_{\uparrow}	Number v_{\downarrow}
All	10003746	466997	1585667
$L \geq 25$ cm	9192450	138625	227309
Untagged	1517134	134988	193651
$y_{max} \geq 100$ cm	979151	30514	41378
Polar angle $\theta \leq 1.14$	889998	14131	27154
v Momentum $p_{v,z} \leq 0.825$	832325	11765	23694
PFO charge $Q_{PFO} \geq 31500$	772788	10394	19075

Table 6.3 The number of cosmic ray, v_{\uparrow} and v_{\downarrow} primary particles left after each cut is applied.

vertex y position $y_{vertex} \geq 100$ cm. The set of v_{\downarrow} tracks are directionally indistinguishable from cosmic-ray tracks, and are topologically very similar to such tracks, having passed the three preselection cuts. However, v_{\downarrow} have a non-fiducial vertex position by construction, since the fiducial volume threshold in y is placed at $y = 81.5$ cm, and as such these tracks are not of interest to physics analyses that consider fiducial events, including the analysis in Chapter 7. Therefore, in what follows, the aim is to retain v_{\uparrow} tracks through the use of directional information, while retaining as many cosmic-ray tracks as possible.

Figure 6.16 shows the normalised distributions of $\Delta\chi_{DU}^2/N$ and $\Delta\chi_{DU}^2$ after applying the various cuts, and when using the ‘Bragg peak finder’ (BPF) outlined in Section 6.2.1. Figure 6.16 shows that the distributions of $\Delta\chi_{DU}^2/N$ and $\Delta\chi_{DU}^2$ conform to expectation: downwards-going particles such as cosmic rays and v_{\downarrow} have predominantly negative values, whereas the v_{\uparrow} entries have mostly positive values. These distributions also demonstrate why directionality is such a useful quantity to aid cosmic ray tagging: it attains good separation between signal and background. It is also clear from this figure that the $\Delta\chi_{DU}^2/N$ distribution has a clearer structure, as it corrects for the fact that low- N tracks have small $\Delta\chi_{DU}^2$ values by dividing by N .

As was discussed in Section 6.2, there are multiple direction measures available, and it remains to be determined which measure is most effective in this context. The different direction measures that have been discussed are P_c , $\Delta\chi_{DU}^2/N$ and $\Delta\chi_{DU}^2$. Cuts can be applied on these direction measures, each of which removes a certain fraction of the remaining cosmic rays, and retain a fraction of v_{\downarrow} and v_{\uparrow} still present in the sample after applying the various cuts of the previous section. This procedure is visualised in Figure 6.17, which plots these



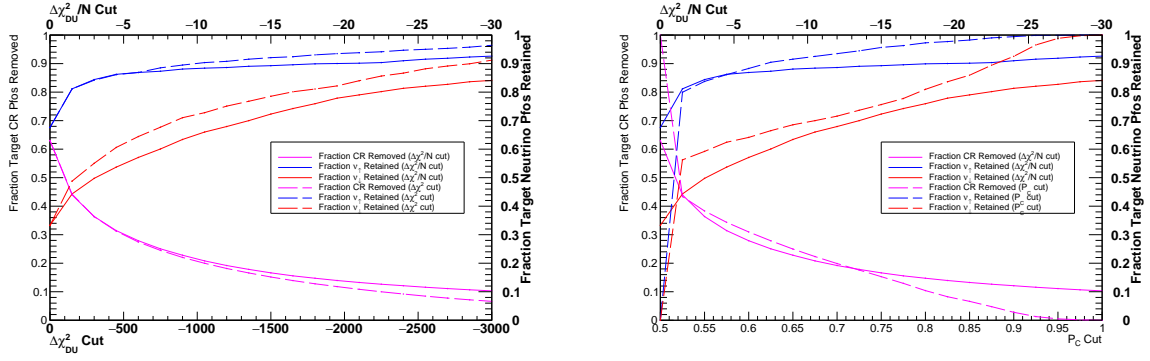
(a) Normalised distributions of $\Delta\chi^2_{DU}/N$ for cosmic rays, ν_{\downarrow} and ν_{\uparrow} .

(b) Normalised distributions of $\Delta\chi^2_{DU}$ for cosmic rays, ν_{\downarrow} and ν_{\uparrow} .

Fig. 6.16 A comparison of the normalised distributions of $\Delta\chi^2_{DU}$ and $\Delta\chi^2_{DU}/N$ for cosmic rays, ν_{\downarrow} and ν_{\uparrow} . The event sample is the untagged sample with $y_{max} \geq 100$ cm and all cuts in 6.2 applied in order. A Bragg peak finder has been applied, as described in the text.

three fractions as a function of cuts on P_c , $\Delta\chi^2_{DU}/N$ and $\Delta\chi^2_{DU}$. As can be seen, the curves have large gradients near the peaks of every distribution, and cutting on $\Delta\chi^2_{DU}/N$ and $\Delta\chi^2_{DU}$ results in very similar curves. In contrast, cutting on P_c forces the fraction of cosmic rays removed and the fraction of neutrino-induced particles to go to 0 and 1, respectively, since all $\Delta\chi^2_{DU}/N$ have been mapped to the range $[0.5, 1.0]$. In contrast, $\Delta\chi^2_{DU}/N$ and $\Delta\chi^2_{DU}$ have very long tails, such that almost all cuts remove a small fraction of cosmic rays. Regardless, if the cuts on these two variables were made over a very large range, their curves would converge to the curves due to P_c . Since these direction measures behave similarly, and because using a probabilistic quantity has no added benefit in this context, $\Delta\chi^2_{DU}/N$ is used as the directional variable in what follows.

The $\Delta\chi^2_{DU}/N$ cuts in Figure 6.17 were made on the untagged sample with all cuts applied. However, it is also possible to find an optimal $\Delta\chi^2_{DU}/N$ after applying each cut, in order to quantify how each cut improves the DCRT procedure. For instance, since each preselection cut creates a purer sample of cosmic rays, the possibility exists that after a certain point an optimised cut on $\Delta\chi^2_{DU}/N$ removes fewer cosmic rays than if the preselection cut had not been applied, while retaining the same number of ν_{\uparrow} . To illustrate that this is not the case, Table 6.4 keeps the number of ν_{\uparrow} fixed at 98% of the amount of ν_{\uparrow} in the untagged sample with $N \geq 25$ cm and $y_{max} \geq 100$ cm. This corresponds to maximally removing 2700 ν_{\uparrow} , and so after applying each cut, the value of $\Delta\chi^2_{DU}/N$ is found that retains as close to this number of ν_{\uparrow} as possible, by scanning the range $-30 \leq \Delta\chi^2_{DU}/N \leq 0$ with a resolution of 0.01. Table 6.4 shows that each subsequent cut allows for a larger additional number of cosmic-ray tracks



(a) Cut performance curves for $\Delta\chi^2_{DU}$ and $\Delta\chi^2_{DU}/N$.

(b) Cut performance curves for P_c and $\Delta\chi^2_{DU}/N$.

Fig. 6.17 The curves in this figure quantify what fraction of cosmic muon tracks are removed and what fraction of ν_\uparrow and ν_\downarrow are removed, as a function of cut values applied to directional variables. Figure 6.17a compares cuts on $\Delta\chi^2_{DU}$ and $\Delta\chi^2_{DU}/N$, and Figure 6.17b compares cuts on P_c and $\Delta\chi^2_{DU}/N$.

to be tagged, while removing the same fraction of ν_\uparrow . In addition, as each cut results in a purer sample of cosmics, the optimal cut value moves closer to $\Delta\chi^2_{DU}/N = 0$.

Number of Particles After Each Cut					
Number Cuts Applied	Variable	$\Delta\chi^2_{DU}/N$ Cut	CR Removed	ν_\uparrow Removed	ν_\downarrow Removed
1	y_{max}	-19.32	135947	2698	8173
2	θ	-4.5	286150	2696	12839
3	$p_{v,z}$	-2.16	345167	2696	12476
4	Q_{PFO}	-0.96	388859	2667	11993

Table 6.4 $\Delta\chi^2_{DU}/N$ cut values that retain as close to 98% of ν_\uparrow as possible, which corresponds to removing at most 2700 ν_\uparrow . The number of cuts applied is indicated, as well as the number of additional cosmic ray particles, actual ν_\uparrow and ν_\downarrow removed. The $\Delta\chi^2_{DU}/N$ is that retains closest to the target ν_\uparrow is found by scanning with a resolution of 0.01. The cut $L \geq 25$ cm is applied everywhere in this table.

6.3.3 Impact on the Reconstruction

In the previous section, the choice to maximally remove 2% of ν_\uparrow was somewhat arbitrary. To determine what number of ν_\uparrow is an ‘acceptable loss’, this section outlines how the fraction of ν_\uparrow removed number affects the reconstruction as a whole. The trade-off is that removing

additional cosmic rays makes it easier to identify the true neutrino interaction in an event, but that by applying stricter cuts more ν_\uparrow are lost. This, in turn, results in events in which the neutrino interaction is lost, and which can never be correctly reconstructed. A number of $\Delta\chi_{DU}^2/N$ cut values are selected from the range $-30 \leq \Delta\chi_{DU}^2/N \leq 0$, which are increasingly biased towards removing cosmic ray particles over retaining neutrino-induced particles. Table 6.5 lists these cuts, and also includes a cut at $\Delta\chi_{DU}^2/N \leq 0$, which removes any particles that appear to be downwards-going, and a cut at $\Delta\chi_{DU}^2/N = \infty$, which removes all cosmic ray particles that remain after applying the various non-directional cuts. Since the cut at $\Delta\chi_{DU}^2/N \leq 0$ retains roughly 97% of ν_\uparrow , cuts are also included at 98% and 99% in order to explore the entire range. Beyond 99%, the cut becomes so strict that very few cosmic rays are removed, and as such the impact on reconstruction performance is expected to be minimal. In this table, the column ‘Target ν_\uparrow Removed’ represents 2% of the total number of ν_\uparrow per row, and indicates the number of removed ν_\uparrow each cut may not exceed.

$\Delta\chi_{DU}^2/N$ Cut Values				
$\Delta\chi_{DU}^2/N$ Cut	Target Fraction ν_\uparrow Retained	Target ν_\uparrow Removed	ν_\uparrow Removed	Cosmic Rays Removed
13.98	99.00%	1350	1348	137659
-0.96	98.00%	2700	2667	388859
0	96.96%	4101	4101	510054
∞	77.40%	30514	30514	979151

Table 6.5 Various $\Delta\chi_{DU}^2/N$ corresponding to certain fractions of ν_\uparrow to be retained. These fractions refer to the fractions of ν_\uparrow remaining after applying the N and y_{max} cuts, so out of 30514 total ν_\uparrow . The impacts of these cuts on the reconstruction performance are compared in Table 6.6.

Table 6.6 shows that the best performance increase is due to the two cuts near $\Delta\chi_{DU}^2/N \approx 0$, since the other two cuts remove too many and too few particles. The cut at $\Delta\chi_{DU}^2/N \leq 0$ makes conceptual sense, since it amounts to removing any top-face intersecting particle that appears to be downwards-going. As such, this is the cut that is used for the DCRT procedure in the analysis of the next chapter, where the improved cosmic ray removal performance is used to formulate a more effective event selection.

<i>Fractions of Correct Events</i>						
Interaction Type	No DCRT	$\Delta\chi_{DU}^2/N \leq -13.98$	$\Delta\chi_{DU}^2/N \leq -0.96$	$\Delta\chi_{DU}^2/N \leq 0$	$\Delta\chi_{DU}^2/N \leq \infty$	Total
CCQEL μ^-	17315 (66.87%)	17514 (67.64%)	17766 (68.62%)	17844 (68.92%)	17876 (69.04%)	25892
CCQEL $\mu^- + p$	19250 (64.77%)	19241 (64.74%)	19272 (64.85%)	19140 (64.40%)	18345 (61.73%)	29720
CCRES μ^-	704 (67.18%)	705 (67.27%)	719 (68.61%)	719 (68.61%)	722 (68.89%)	1048
CCRES $\mu^- + p$	2059 (65.12%)	2053 (64.93%)	2079 (65.75%)	2074 (65.59%)	2028 (64.14%)	3162
CCMEC μ^-	3854 (67.65%)	3881 (68.12%)	3944 (69.23%)	3957 (69.46%)	3991 (70.04%)	5697
CCMEC $\mu^- + p$	9155 (68.99%)	9136 (68.85%)	9183 (69.20%)	9201 (69.34%)	8909 (67.14%)	13270

Table 6.6 The impact of the different $\Delta\chi_{DU}^2/N$ requirements on the Pandora reconstruction performance. The columns are arranged by the strictness of the cut (left to right, left being no cut). The indicated fractions are the fraction of correct events, using the Pandora correctness criterion.

Chapter 7

Selecting Charged-Current Muon Neutrino Final States

This chapter describes the use of the Pandora consolidated reconstruction output, including the improvements outlined in Chapter 6, to select events of the form $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ final states, where X denotes the final-state nucleus. This chapter will start, in Section 7.1, by discussing the notation used to denote neutrino interaction types and will give definitions of what constitutes contained and signal events. In Section 7.2, interaction type tables will outline the underlying distribution of interaction types for events with the a reconstructed primary particle multiplicity of one and two. In addition, several aspects of the reconstruction are ‘cheated’ (using the Monte Carlo truth information) to quantify how different aspects of the reconstruction affect the selection purity and efficiency, and the cosmic-ray background is analysed in detail. Section 7.3 defines a number of variables, which are divided into three categories: topological variables, directional variables and calorimetric/particle identification (PID) variables. These variables are designed to separate signal and background events, and are used in a multivariate analysis (MVA) to select signal events in Section 7.4. In this section, the impact of the directional variables on the performance of the MVA is quantified. The event selection is applied to the simulated sample in Section 7.5, and selection purities and efficiencies are quantified. Finally, the resulting variable distributions are compared to the distributions that result from applying the event selection to real data.

7.1 Definitions and Pre-Selection Cuts

This section gives definitions for concepts that are frequently used throughout this chapter. In particular, the notation used for interaction types is outlined, as well as how the Pandora reconstructability criterion is included in the interaction type definitions. In this chapter, the final reconstructed neutrino candidate is denoted as ν_{reco} , and the true neutrino particle in the simulation is indicated by ν_{true} . Furthermore, different interaction types are indicated using shorthand consisting of three parts:

- CC/NC: indicates whether an interaction is charged current (CC) or neutral current (NC);
- QEL/RES/MEC/COH/DIS: indicates whether an interaction is quasi-elastic scattering (QEL), resonant production (RES), meson exchange current (MEC), coherent production (COH), or deep inelastic scatter (DIS);
- $\mu^-/p/\pi^+, \dots$: primary Monte Carlo final state particles are included in the interaction type definition if they are reconstructable, according to the reconstructability definition criterion defined in Section 4.5. This criterion demands that a Monte Carlo particle deposits the majority of hit charge for ≥ 15 reconstructed hits, and for ≥ 5 reconstructed hits in at least two views.

For instance, CCQEL μ^- indicates a charged current quasi-elastic ν_μ interaction with one reconstructable muon: $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$.

To denote events in which the neutrino interaction produced no reconstructable Monte Carlo primaries, an additional interaction type is defined:

- ν NO RECONSTRUCTABLE represents events in which the neutrino does interact, but in which none of the Monte Carlo primary final state particles meet the reconstructability requirement (i.e. not enough hits are produced);

In addition to neutrino-induced particles, the cosmic-ray background may contaminate ν_{reco} . The reconstructed neutrino may consist entirely of cosmic-ray particles, or a cosmic-induced particle may occupy a region of detector space near the neutrino interaction, and may be folded into the reconstructed neutrino particle hierarchy. To represent such reconstruction outcomes, two additional interaction types are defined that denote reconstruction final states affected by cosmic ray particles:

- COSMIC RAY: in this case ν_{true} does induce reconstructable Monte Carlo particles, but none of the associated PFOs are selected as ν_{reco} . This interaction type typically occurs when the cosmic ray tagging does not remove a true cosmic ray, and it is selected by the ν ID logic as the most probable neutrino candidate;
- In the case where ν_{reco} contains both neutrino-induced and cosmic ray particles, the interaction type are denoted by the ν_{true} interaction type suffixed by $_K_CR$, where $_K_$ is the number of cosmic rays. This happens when a cosmic-ray track crosses the neutrino interaction or appears very close to it in 3D, so that the slicing procedure is unable to separate them.

The COSMIC RAY interaction type differs from other interaction types, since it is defined by the reconstruction output, and not the underlying Monte Carlo information.

This chapter describes two event selection procedures, which are designed to select events of the type $\nu_\mu \rightarrow \mu^- + X$ and $\nu_\mu \rightarrow \mu^- + p + X$. Table 7.1 outlines the interaction types in the MC BNB COSMIC sample that can result in these two final states, and their total events counts in this sample (without any containment cuts applied). The events in this table are categorised by their true charged particle multiplicity T . The reconstructed charged particle multiplicity is denoted by N in the rest of this chapter. Events of the types $\nu_\mu \rightarrow \mu^- + X$ and $\nu_\mu \rightarrow \mu^- + p + X$ are charged-current quasi-elastic-like events, but also include events due to resonance decays, as well as other channels. The performance of both event selections can be quantified by a selection efficiency η and a selection purity \mathcal{P} . For an event sample consisting of a set of signal events S and background events B from which a subset of events C is chosen by the event selection, the selection efficiency η and selection purity \mathcal{P} are defined as

$$\begin{aligned}\eta &= \frac{|S \cap C|}{|C|}, \\ \mathcal{P} &= \frac{|S \cap C|}{|S|}.\end{aligned}\tag{7.1}$$

An additional important signal criterion is whether the reconstructed neutrino particle hierarchy ν_{reco} is contained within the fiducial volume. The fiducial volume is defined to be the region of the detector that excludes both the detector boundaries and significant inactive detector regions. The fiducial volume is schematically represented in Figure 7.1 [96]. Since only reconstruction output is available to the event selection procedure, a pre-selection cut is applied that retains only those events for which ν_{reco} is contained. However, a number of

$T=1$ Target Channels		$T=2$ Target Channels	
Interaction Type	Event Count	Interaction Type	Event Count
CCQEL μ^-	64503	CCQEL $\mu^- + p$	39126
CCMEC μ^-	19921	CCMEC $\mu^- + p$	17883
CCRES μ^-	13860	CCRES $\mu^- + p$	5116
CCDIS μ^-	4324	CCDIS $\mu^- + p$	1147
CCCOH μ	142	-	-
TOTAL	102156	TOTAL	63162

Table 7.1 All target channels: these are the interaction types of the form $\nu_\mu \rightarrow \mu^- + X$ and $\nu_\mu \rightarrow \mu^- + p + X$, which result in the target $1\mu^-$ and $1\mu^- 1p$ final states.

possible definitions of containment are possible for neutrino particles. Possible criteria to use to define containment include:

- *Contained interaction vertex*: a contained event can be defined as an event for which the reconstructed neutrino interaction vertex is within the fiducial volume. This definition is conceptually simple, and is simple to implement for both track-like and shower-like particles. However, particularly for long tracks, a fiducial vertex does not guarantee that all neutrino-induced particles lie within the fiducial volume. Since a fiducial reconstructed interaction vertex is not a strict requirement, this definition is likely to retain the largest number of signal events;
- *Contained endpoints*: another approach, which is quite common, is to define a contained event as an event for which every neutrino-induced particle has contained endpoints. However, this definition is ill-defined for shower-like particles, although it is possible to demand that the shower envelope is fiducial. This stricter requirement will exclude a number of additional events;
- *Contained hit fraction*: a third approach is to consider what fraction of the hits induced by the neutrino fall within the fiducial volume. A containment fraction F_c can be defined as the fraction of contained hits, and an event can be considered fiducial when F_c is above a pre-defined threshold. While this approach is equivalent for track-like and shower-like particles, it is dependent on the impact of the cosmic ray removal procedure, since different removal efficiencies will result in different sets of reconstructed 3D hits.




Fig. 7.1 The fiducial volume, which avoids regions with a significant number of unresponsive channels. Taken from [96].

This analysis will use the containment criterion that the interaction vertex must lie within the fiducial volume, since it is conceptually simple and consistent between tracks and showers. An additional motivation is that, when comparing the effect of each definition has on the simulated analysis sample, the containment of the interaction vertex is the least stringent of the three requirements and as such retains the largest number of signal events.

The directional cosmic ray tagging (DCRT) cut, developed in Chapter 6, is used as a pre-selection cut. The effect this pre-selection cut has on the distribution of the number of events per interaction type in the simulated event sample is discussed in Section 7.2.

To summarise, an event is considered to be a true signal event if the following criteria are met:

- The ν_{true} interaction type is one of the interaction types listed in Table 7.1.
- The final state particles used to define each interaction type meet the Pandora reconstructability criterion, which retains Monte Carlo particles if they deposit the majority of hit charge for ≥ 15 reconstructed hits and ≥ 5 reconstructed hits in at least two views;
- The event is contained: the containment definition used here is that the true neutrino vertex must be within the fiducial volume, as defined in Figure 7.1.

7.2 Reconstruction Studies

A useful starting point for the event selection is to consider the primary particle multiplicity T of an event. The target interaction types $\nu_{\mu} + \text{Ar} \rightarrow \mu^{-} + X$ and $\nu_{\mu} + \text{Ar} \rightarrow \mu^{-} + p +$

X have associated true primary particle multiplicities of $T = 1$ and $T = 2$, respectively. However, the reconstructed primary particle multiplicity N of the reconstructed neutrino ν_{reco} may differ from the true multiplicity, such that $N \neq T$. This can be due to reconstruction effects, such as when the reconstructed muon track in an event of type CCQEL μ^- is erroneously reconstructed as two separate tracks, resulting in $N = T + 1$. In addition, the cosmic background may contaminate ν_{reco} , such that $N > T$. Furthermore, a cosmic ray particle may be identified as the most likely ν_{reco} candidate, and the distribution of N for this subset of events will peak sharply at 1, regardless of T . Lastly, in $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ final states, the proton may be merged with the muon, resulting in a reconstructed multiplicity of $N = T - 1$. In this discussion, the event sample is divided into two subsets for which $N = 1$ and $N = 2$, since these subsets contain the majority of the signal events. A very small fraction of signal events is contained in the subsets with $N = 0$ and $N > 2$, and the exclusion of such events can be considered to be an additional pre-selection cut.

In addition to the containment and reconstructability cuts, the directional cosmic ray tagging (DCRT) cut developed in Section 6.3 is applied as a pre-selection cut. This cut increases the purity of the $N = 1$ sample subset from 13.61% to 15.69% and the purity of the $N = 2$ sample subset from 32.68% to 35.48%. In addition, the total number of signal events increases by 1.03% for $N = 1$ and by 0.48% for $N = 2$. The DCRT pre-selection cut reduces the number of COSMIC RAY events in the $N = 1$ sample by 14.74% (from 84650 to 72173) and in the $N = 2$ sample by 19.47% (from 26308 to 21186). This significant reduction in cosmic background in each sample facilitates the identification of the true neutrino interaction by the reconstruction, increasing both the total number of signal events that are correctly reconstructed with $T = N$ and the sample signal purity. This increased number of signal events at the pre-selection stage improves the purities and efficiencies the $N = 1$ and $N = 2$ selection procedures eventually achieve. Section 7.4.1 demonstrates that, in addition to the substantial effectiveness of applying directional information at the pre-selection stage, the event selection is able to further utilise the directional information in the remaining events to improve the performance of both selections.

Table 7.2 outlines all the interaction types that result in reconstructed primary particle multiplicities of $N=1$ and $N=2$ after applying the containment, reconstructability and DCRT pre-selection cuts. In Table 7.2, the *Selected Events* column indicates the number of events with the indicated reconstructed primary particle multiplicity, categorised by interaction type. These events are labelled ‘selected’, since applying a cut on N is the basis of the event selection procedures, which is expanded upon in later sections. Indentation in the rows of the table indicate subsets: for instance, NO RECONSTRUCTABLE is a subset of the COSMIC RAY category, which in turn is a subset of the BACKGROUND category. The *Total*

Events column indicates the total number of fiducial events per interaction type, regardless of their multiplicity. In the *Selected Events* and *Total Events* columns, the row labelled TOTAL indicates the total number of selected events, and the total number of fiducial events in the sample, respectively. The percentages in each column indicate what fraction of the selected and total event counts each interaction type constitutes. Furthermore, interaction types that contribute $\leq 1\%$ of the total number of selected events individually are combined into the SMALL event class. The OTHER INTERACTION event class contains events that have not been explicitly enumerated, such as events with ≥ 6 protons, or deep inelastic scatters with ≥ 3 photons.

<i>N=1 Events</i>					<i>N=2 Events</i>				
Interaction Type	Selected Events		Total Events		Interaction Type	Selected Events		Total Events	
SIGNAL	15015	15.69%	19867	19.36%	SIGNAL	20116	35.48%	24774	11.67%
CCQEL μ	11684	12.21%	12322	5.80%	CCQEL $\mu + p$	12632	22.28%	14605	6.88%
CCMEC μ	2595	2.71%	2788	1.31%	CCRES $\mu + p$	1267	2.23%	1537	0.72%
-	-	-	-	-	CCMEC $\mu + p$	5928	10.45%	6799	3.20%
SIGNAL SMALL ($< 1\%$)	736	0.77%	-	-	SIGNAL SMALL ($< 1\%$)	289	0.51%	-	-
BACKGROUND	80693	84.31%	192491	90.64%	BACKGROUND	36588	64.52%	187584	88.33%
COSMIC RAY	72173	75.41%	126293	59.47%	COSMIC RAY	21186	37.36%	126293	59.47%
NO RECONSTRUCTABLE	48114	50.27%	83923	39.52%	NO RECONSTRUCTABLE	13867	24.46%	83923	39.52%
CCQEL $\mu + p$	1396	1.46%	14605	6.88%	CCQEL μ	628	1.11%	12322	5.80%
NCQEL p	1257	1.31%	1415	0.67%	CCRES $\mu + \pi^+$	3088	5.45%	3794	1.79%
-	-	-	-	-	CCRES $\mu + p + \pi^+$	687	1.21%	4828	2.27%
-	-	-	-	-	NCRES π^0	636	1.12%	1232	0.58%
-	-	-	-	-	CCDIS $\mu + \pi^+$	1177	2.08%	1485	0.70%
-	-	-	-	-	OTHER INTERACTION	1419	2.50%	5819	2.74%
BACKGROUND SMALL ($< 1\%$)	5869	6.13%	-	-	BACKGROUND OTHER ($< 1\%$)	6619	11.67%	-	-
TOTAL	95708	100.00%	212358	100.00%	TOTAL	56704	100.00%	212358	100.00%
PURITY	15015/95708	15.69%			PURITY	20116/56704	35.48%		
EFFICIENCY	15015/19867	75.58%			EFFICIENCY	20116/24774	81.20%		

Table 7.2 A breakdown of interaction types that result in reconstructed primary particle multiplicities of $N=1$ and $N=2$ in a simulated sample with neutrino interaction generated by GENIE and cosmic rays simulated by CORSIKA. The vertex containment cut, described in Section 7.1, has been applied in these tables. Furthermore, a directional cosmic ray tagging (DCRT) cut has been applied, details of which are outlined in the text. Blank entries have been introduced in the $N=1$ table for row consistency between tables. The event sample contains 520448 events in total.

The predominant background for both the $N=1$ and $N=2$ interaction type tables is the COSMIC RAY category. The fraction of COSMIC RAY events is especially large for events with $N=1$, which is due to two reasons:

- Virtually every NO RECONSTRUCTABLE event across the event sample, regardless of N , will end up in the $N = 1$ category. Well-reconstructed cosmic-ray tracks (i.e. not split due to δ -rays, for instance) are reconstructed as a single particle. If such a particle

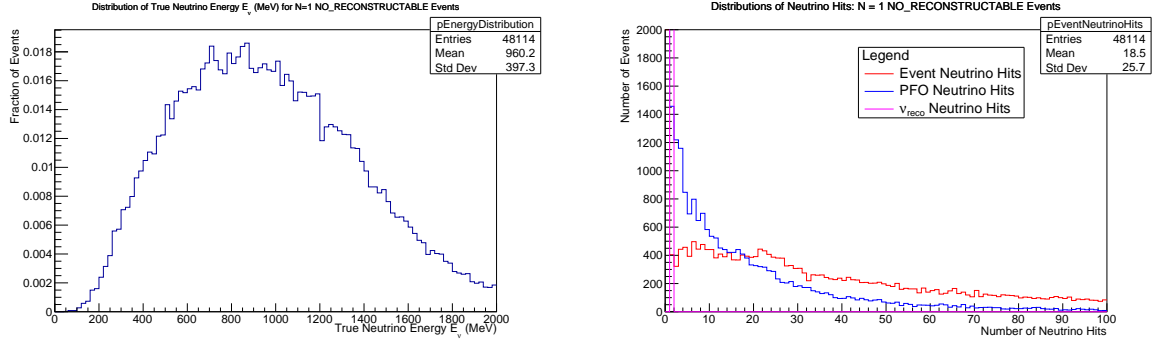
is selected as the reconstructed neutrino (v_{reco}), which is highly probable in the case of NO RECONSTRUCTABLE, it will contribute to the $N = 1$ category;

- The neutrino interactions in the $N = 1$ category predominantly produce single tracks, which are topologically very difficult to distinguish from cosmic ray muons.

In the case where $N=2$, v_{reco} can be associated to two cosmic ray tracks if the reconstruction splits a single cosmic ray track, although this is less likely than when $N = 1$. To account for split cosmic ray tracks, the opening angle ζ between both particles in an $N = 2$ event is included as a variable in the event selection described in Section 7.3. The distribution of ζ peaks at π radians for the COSMIC RAY and NO RECONSTRUCTABLE interaction types, indicating a back-to-back split. In addition to these mechanisms, two cosmic rays can have significant overlap in the detector, so that their topology resembles a neutrino interaction, or an interaction like CCQEL $\mu^- + p$ can be mistaken for a cosmic ray muon and a Michel electron.

The NO RECONSTRUCTABLE interaction type consists of events that contain no particles that meet the Pandora reconstructability criterion. Note that this does not imply that these events are induced by neutrinos with low energies. Rather, these events contain neutrino-induced hits that are very diffuse, and cannot be resolved to individual particles. Figure 7.2a displays the distribution of the true neutrino energy E_ν for $N = 1$ NO RECONSTRUCTABLE events. Figure 7.2b shows the distribution of the number of hits for which the majority of the hit charge was deposited by a neutrino-induced particle. Three distributions are shown: the true number of neutrino hits in i) the entire event, ii) in all event PFOs and iii) in the reconstructed neutrino. These distributions demonstrate that a non-negligible amount of reconstructed hits result from neutrino-induced Monte Carlo primary final states, a small fraction of which ends up in PFOs. Almost no true neutrino hits are contained in the reconstructed neutrino. This is due to events in which energetic neutrons are induced in the neutrino interaction, travel far in the detector without inducing tracks, resulting in a sparse ‘cloud’ of hits. An example event is given in Figure 7.3.

The predominant background in the $N = 1$ channel is the cosmic-ray background, but it is also apparent that CCQEL $\mu^- + p$ events that are reconstructed as a single particle contribute to the background for this multiplicity. In the $N=2$ interaction type table, one of the backgrounds is the CCRES π^0 channel, $\nu_\mu + \text{Ar} \rightarrow \mu^- + \pi^0 + X$, where the π^0 decays into two photons, resulting in two electromagnetic showers. The smaller of the two showers can be merged with the larger shower, or not reconstructed at all if it is very small, leading to a reconstructed multiplicity of $N = 2$. This observation illustrates that the reconstructed primary particle multiplicity N is a useful pre-selection variable, but that more sophisticated



(a) The distribution of E_ν for $N = 1$ NO RECONSTRUCTABLE events.

(b) Distribution of true neutrino-induced hits in the event, event PFOs and v_{reco} .

Fig. 7.2 Plots that characterise the NO RECONSTRUCTABLE interaction type. Figure 7.2a indicates that events of this type are not just low energy, and Figure 7.2b demonstrates that such events contain few true neutrino-induced hits, of which only a subset are incorporated into PFOs. This is due to events containing neutrons that induce sparse clouds of hits, as discussed in the text.

variables are needed to improve the event selection. To largely remove events containing showers, P_{track} is used in the multivariate analysis, which is the probability that a particle is a track, as calculated by the Pandora track/shower identification functionality. A second prominent source of background events for $N=2$ is due to interaction channels that produce μ and a π^+ or a π^- , since the pions typically produce short tracks that topologically resemble protons. However, the particle mass of pions and protons is very different, and this fact is used in Section 7.3.4 to improve the event selection.

7.2.1 Characterising Reconstruction Effects

In order to gain an understanding of how the reconstruction itself affects the distribution of interaction types in Table 7.2, this section will investigate the impact of cheating three aspects of the reconstruction sequentially: the neutrino ID (ν ID), the cosmic ray tagging (CRT) and the slicing. Cheating is the process of using Monte Carlo truth information to inform reconstruction decisions, thereby mimicking a perfect reconstruction. Here, the cheating of different reconstruction aspects is sequential, so that the three reconstruction configurations that are investigated are as follows:

- *Cheating the ν ID*: The cheating ν ID tool will select the reconstructed particle hierarchy containing the most neutrino-induced hits. As such, cheating this step will reveal the impact of the ν ID on the interaction type tables;




Fig. 7.3 An example event in which a neutron induces a diffuse cloud of hits, leading to the event being classified as a ν NO RECONSTRUCTABLE event. The cosmic ray muon indicated in blue is selected as ν_{reco} .

- *Cheating the CRT and the ν ID*: the cheated cosmic ray tagging functionality will remove any particle if $\geq 25\%$ of its hits are predominantly induced by cosmic rays. Cheating the cosmic ray tagging as well as the ν ID will remove all cosmic ray particles, and will reveal the distribution of interaction types in the absence of cosmic rays;
- *Cheating the slicing, the CRT and the ν ID*: the cheated event slicing functionality creates a new slice for each Monte Carlo particle hierarchy and assigns each 2D hit to the Monte Carlo particle that deposited the majority of its charge. Cheating the slicing will have an impact on the interaction type tables when imperfect slicing separates hits deposited by the same particle into separate slices, or when hits from separate particles end up in the same slice, effects which affect the reconstructed particle multiplicity of V_{reco} .

Tables 7.3 and 7.4 show the effect of these different reconstruction configurations on the $N = 1$ and $N = 2$ interaction type tables. The Pandora cheating ν ID tool removes a large fraction of the NO RECONSTRUCTABLE and COSMIC RAY categories. The total number of selected events is smaller for this configuration, because when none of the slices in the event contain any neutrino-induced hits, all slices are identified as cosmic ray hypothesis slices, and no V_{reco} is produced. The event counts for each interaction type change when cheating the reconstruction, since event classes such as COSMIC RAY are dissolved into other channels, and previously incorrect events with $N \neq T$ can, through cheating, become correct with $N = T$.

When cheating the ν ID, some events in both the NO RECONSTRUCTABLE and COSMIC RAY categories still remain because some neutrino-induced hits may end up in cosmic ray PFOs, even when the cosmic ray tagging procedure has erroneously removed PFOs containing the majority of neutrino hits. In this case, the cheating ν ID tool will select the cosmic ray PFO with the largest number of neutrino hits as the neutrino hypothesis slice, which will then become V_{reco} . This occurs for COSMIC RAY events, including NO RECONSTRUCTABLE events, in which the neutrino can induce many hits, as shown in Figure 7.2b. In particular, in events with sparse true neutrino hits, such as the example event in Figure 7.3, a few hits are likely to be associated to a cosmic ray particle. This is why NO RECONSTRUCTABLE is the dominant contribution to the COSMIC RAY category for this configuration.

Cheating the cosmic ray tagging, in addition to the ν ID, removes any remaining NO RECONSTRUCTABLE and COSMIC RAY events, since cosmic ray contamination is eliminated, such that the remaining neutrino-induced hits can be reconstructed with no risk of being

merged with or removed in favour of a cosmic ray. However, neutrino-induced hits may still be assigned to different detector slices, such as when photons induce electromagnetic showers far away from the interaction vertex. When the slicing is also cheated, the number of correctly-reconstructed events increases further. In general, the aspects of the reconstruction that deal with cosmic ray reconstruction have a dramatic impact on the interaction type tables, and also on the purity and efficiency of the resulting sample. Reducing the cosmic-ray background is therefore paramount in developing both event selections.

Table 7.3 also reveals the dominant non-cosmic background channels. For events with $N = 1$, the two largest neutrino-induced backgrounds are NCQEL p and CCQEL $\mu^- + p$. Events that contain a single proton are addressed in the event selection by introducing PID variables, as well as topological variables such as track length that are correlated with particle stopping power. Events in which CCQEL $\mu^- + p$ events are reconstructed with $N = 1$ are taken into account by introducing the particle splitting $\delta\chi_{BA}^2$ introduced in Section 5.5 as a variable in the multivariate analysis, which will allow the event selection to learn that events with large values of $\delta\chi_{BA}^2$ are likely to have an incorrect reconstructed multiplicity $N \neq T$.

<i>N=1 Impact of Cheating Reconstruction Functionality</i>								
Interaction Type	Actual Reconstruction		Cheated ID		Cheated ID, CRT		Cheated ID, CRT, Slicing	
SIGNAL	15015	15.69%	21891	52.54%	26048	61.38%	26308	62.57%
CCQEL μ	11684	12.21%	11684	12.21%	9915	41.88%	10096	45.73%
CCRES μ	556	0.58%	1529	3.67%	1923	4.53%	1923	4.57%
CCDIS μ	149	0.16%	471	1.13%	591	1.39%	606	1.44%
CCMEC μ	2595	2.71%	3891	9.34%	4640	10.93%	4698	11.17%
SIGNAL SMALL ($\leq 1\%$)	29	0.03%	86	0.21%	100	0.24%	97	0.23%
BACKGROUND	80693	84.31%	19777	47.46%	16388	38.62%	15740	37.43%
COSMIC RAY	72173	75.41%	7362	17.67%	-	-	-	-
ν NO RECONSTRUCTABLE	48114	50.27%	5027	12.06%	-	-	-	-
CCQEL $\mu + p$	1396	1.46%	2011	4.83%	2358	5.56%	2438	5.80%
NCQEL p	1257	1.31%	1641	3.94%	2475	5.83%	2553	6.07%
CCRES $\mu + \pi^+$	351	0.37%	508	1.22%	614	1.45%	544	1.29%
NCRES p	341	0.36%	452	1.08%	676	1.59%	696	1.66%
NCRES π^+	272	0.28%	348	0.84%	462	1.09%	450	1.07%
NCRES π^0	394	0.41%	552	1.32%	626	1.48%	543	1.29%
OTHER INTERACTION	749	0.78%	1562	3.75%	2210	5.21%	2075	4.93%
CCMEC $\mu + p$	645	0.67%	911	2.19%	1100	2.59%	1031	2.45%
CCMEC p	536	0.56%	745	1.79%	1102	2.60%	1142	2.72%
BACKGROUND SMALL	2581	2.70%	3690	8.86%	4666	11.00%	4250	10.11%
TOTAL	95708	100.00%	41668	100.00%	42436	100.00%	42048	100.00%
PURITY	15015/95708	15.69%	21891/41668	52.54%	26048/42436	61.38%	26308/42048	62.57%
EFFICIENCY	15015/19867	75.58%	21891/26940	81.26%	26048/29510	88.27%	26308/29657	88.71%

Table 7.3 Comparison of the $N=1$ interaction type tables when cheating several aspects of the reconstruction. From left to right: cheating the ν ID, cheating the ν ID and the cosmic ray tagging, and cheating the ν ID, the cosmic ray tagging and the slicing.

<i>N=2 Impact of Cheating Reconstruction Functionality</i>								
Interaction Type	Actual Reconstruction		Cheated ID		Cheated ID, CRT		Cheated ID, CRT, Slicing	
SIGNAL	20116	35.48%	20414	49.65%	23582	58.75%	23975	59.49%
CCQEL $\mu + p$	12632	22.28%	12780	31.08%	14908	37.14%	15071	37.40%
CCRES $\mu + p$	1267	2.23%	1331	3.24%	1516	3.78%	1560	3.87%
CCMEC $\mu + p$	5928	10.45%	5999	14.59%	6803	16.95%	6980	17.32%
SIGNAL SMALL ($\leq 1\%$)	289	0.51%	304	0.74%	355	0.88%	364	0.90%
BACKGROUND	36588	64.52%	20705	50.35%	16559	41.25%	16324	40.51%
COSMIC RAY	21186	37.36%	2068	5.03%	-	-	-	-
ν NO RECONSTRUCTABLE	13867	24.46%	1426	3.47%	-	-	-	-
CCQEL μ	628	1.11%	822	2.00%	948	2.36%	951	2.36%
CCRES $\mu + \pi^+$	3088	5.45%	3208	7.80%	3575	8.91%	3686	9.15%
CCRES $\mu + p + \pi^+$	687	1.21%	695	1.69%	791	1.97%	757	1.88%
NCRES $p + \pi^-$	395	0.70%	402	0.98%	474	1.18%	489	1.21%
NCRES π^0	636	1.12%	717	1.74%	873	2.17%	971	2.41%
NCRES $p + \pi^0 \pi$	373	0.66%	428	1.04%	459	1.14%	445	1.10%
CCDIS $\mu + \pi^+$	1177	2.08%	1258	3.06%	1441	3.59%	1459	3.62%
OTHER INTERACTION	1419	2.50%	1769	4.30%	2176	5.42%	2151	5.34%
CCMEC $\mu + p + p$	387	0.68%	382	0.93%	447	1.11%	451	1.12%
BACKGROUND SMALL	6615	11.67%	8960	21.79%	5362	13.36%	4954	12.29%
TOTAL	56704	100.00%	41119	100.00%	40141	100.00%	40299	100.00%
PURITY	20116/56704	35.48%	20414/41119	49.65%	23582/40141	58.75%	23975/40299	59.49%
EFFICIENCY	20116/24774	81.20%	20414/26053	78.36%	23582/29320	80.43%	23975/29645	80.87%

Table 7.4 Comparison of the $N=2$ interaction type tables when cheating several aspects of the reconstruction. From left to right: cheating the ν ID, cheating the ν ID and the cosmic ray tagging, and cheating the ν ID, the cosmic ray tagging and the slicing.

7.3 Multivariate Analysis for Event Selection

This section will outline a number of variables designed to separate signal events of the form $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ from background events. These event selection variables are divided into three categories: topological, calorimetric and directional. Furthermore, in the case where $N = 2$, a distinction is made between the variable distributions of the longest and the shortest PFO in the event, which are given the subscripts L and S , respectively. This distinction is designed to ameliorate signal/background separation by identifying the muon and proton in $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ interactions.

7.3.1 Topological Variables

The first set of variables is based on the event topology, the detector geometry and on quantities that can be queried directly from Pandora objects. For both $N=1$ and $N=2$, the following topological variables are correlated with the different underlying energy distributions and kinematics of signal and background events:

- *Total charge Q_{PFO}* : the charge of a PFO is defined as

$$Q_{PFO} = \sum_{k=1}^3 \sum_{i=1}^{N_k} Q_{k,i}, \quad (7.2)$$

where the subscript k identifies the three detector views U, V and W, and N_k indicates the total number of hits in view k . The subscript PFO makes the difference from hit charge Q explicit, but is omitted where it is implied through context. Here, Q is the input energy of the hit object in integrated ADC counts scaled to units of MeV, defined in equation 5.7. Calorimetric information from all three views is used to make this variable more robust to the inactive detector regions in each view;

- *Track length L* : the 3D track length is retrieved by summing the distances between neighbouring 3D hits in a hit list sorted in z . Note that this quantity has little meaning for a shower-like object, but other variables are introduced to remove shower-like objects;
- *Polar and azimuthal angles θ and ϕ* : these angles contain information about event kinematics, and are calculated by creating a 3D sliding fit (defined in Section 4.4.2) to the PFO object and calculating the opening angle between the longitudinal sliding fit axis and the y and x axes, respectively;

- *Begin- and endpoint x , y and z coordinates $\{x, y, z\}_{min/max}$* : these coordinates represent the smallest and largest x , y and z coordinates of both track endpoints. These coordinate variables capture the fact that cosmic rays enter the detector near its top face, and that neutrino-induced particles can have significant extent in the z direction.

Figures 7.4, 7.5 and 7.6 show stacked distributions of a subset of topological variables for $N = 1$ and $N = 2$. In both cases, the distributions that are shown correspond to the dominant signal interaction types listed in Table 7.1, as well as the cosmic-ray background channels NO RECONSTRUCTABLE and COSMIC RAY. All other interaction types are accumulated in the OTHER distribution in each stack. Throughout this chapter, cosmic-ray background channels are plotted as the bottom distribution in each stacked distribution.

Figures 7.4 and 7.5 show the stacked plots of the previously described variables for $N=1$, including a subset of the endpoint coordinate variables. It is clear from Figure 7.4 that because neutrino final states and cosmic ray particles have different underlying energy distributions, their charge and length distributions also differ. Furthermore, Figure 7.5 shows that cosmic ray interactions have very high y_{max} values, since they originate in the atmosphere. Conversely, many neutrino-induced PFOs have high z_{max} coordinates, as the long forwards-pointing muons leave the detector through the downstream detector face. Generally, the distributions for these simple topological variables already have discrimination power. The same stacked plots as in Figure 7.4 for the longest PFO in $N=2$ are very similar, and are omitted here. The variables based on the shortest PFO with the most discriminating power for $N=2$ are that of Q_{PFO} and L_{PFO} , which are shown in 7.6. In the case of $N=2$, in addition to the above PFO-based variables, a number of other variables can be constructed that describe the relationship between the two particles in v_{reco} :

- *Opening angle ζ* : the opening angle between the two particles in the event will also carry information related to event kinematics. For a fixed polar angle, the opening angle between the two particles is dictated by momentum conservation, in the absence of additional scatters on argon nuclei.
- *Reconstructed momentum components $p_{v,x}$, $p_{v,y}$ and $p_{v,z}$* : momentum conservation makes it possible to formulate the components of the neutrino momentum p_v under the hypothesis that the longest particle in v_{reco} is a muon and the shortest particle is a proton. In $N=2$ background channels such as CCRES $\mu^- + \pi^+$, the mass of the shorter particle will dictate the reconstructed neutrino momentum p_v . This variable is effective at distinguishing signal events from cosmic-ray backgrounds since the neutrino momentum is known to be along the beam line, whereas cosmic ray momenta

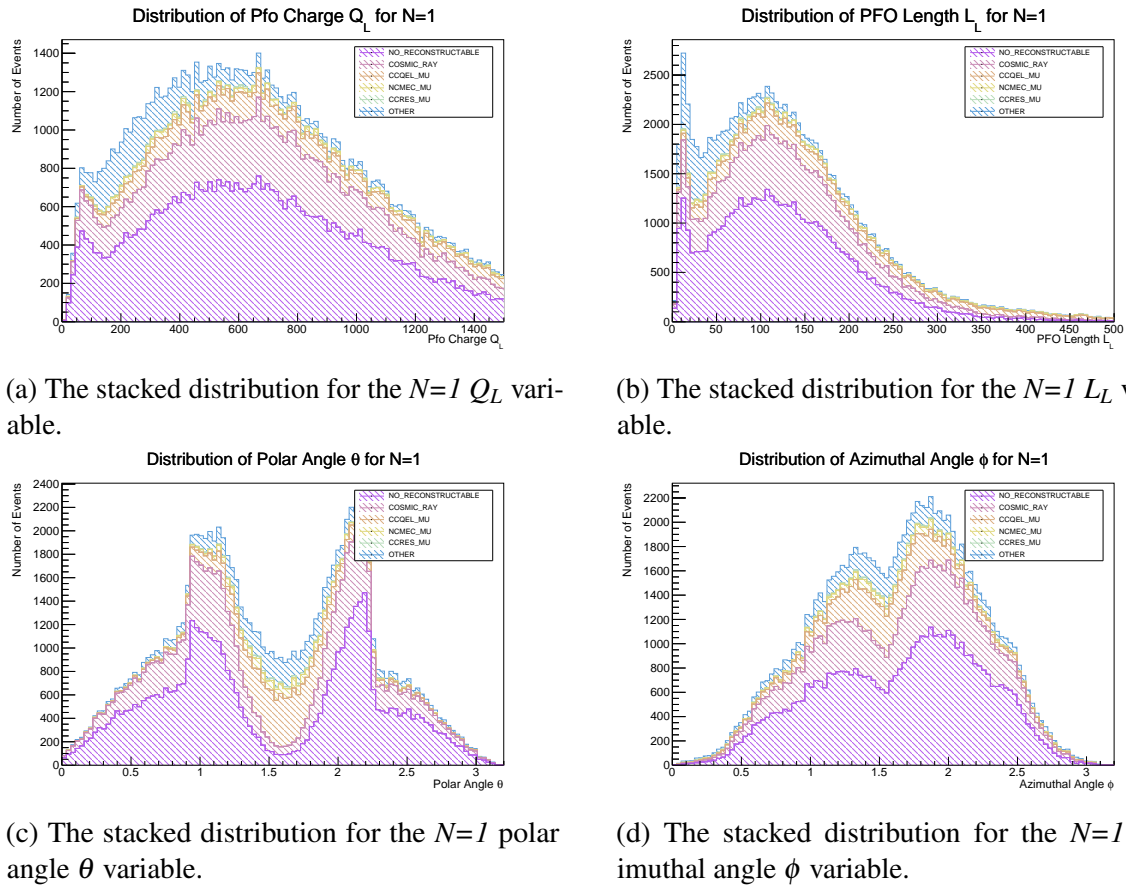


Fig. 7.4 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$.

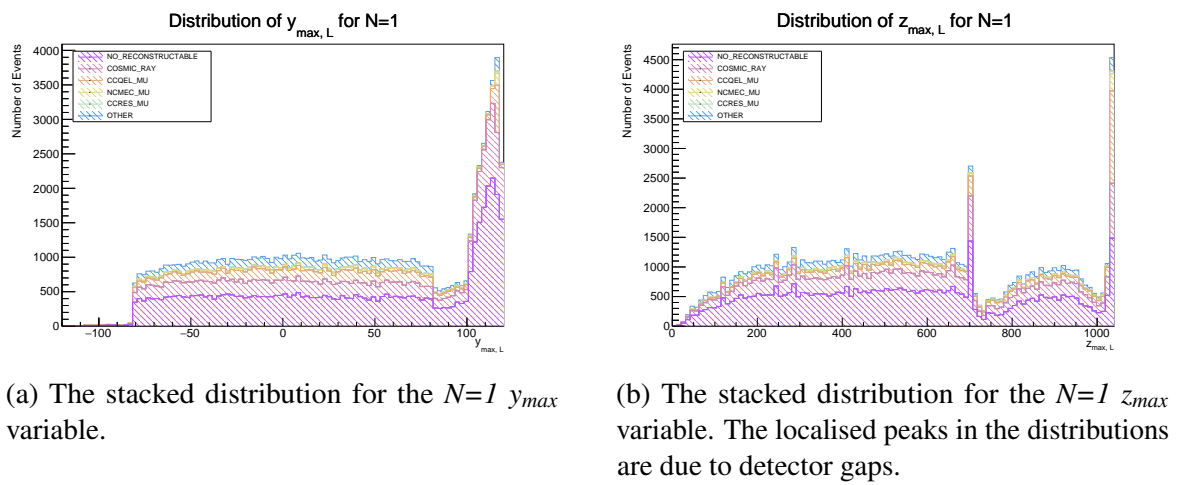



Fig. 7.5 Stacked plots of the distributions of the PFO y_{max} and z_{max} variables.



(a) The stacked distribution for the $N=2$ Q_S variable.

(b) The stacked distribution for the $N=2$ L_S variable.

Fig. 7.6 Stacked plots of the distributions of the $N=2$ shortest PFO summed charge and length variables.

are directed along the y axis. In addition, p_v is effective at separating out other background channels, since incorrectly setting the mass of the shortest particle to the proton mass will lead to spuriously large x - and y -components in p_v . To calculate p_v in practice, one can leverage the fact that total particle energy is the sum of kinetic energy and rest mass $E = E_k + m_0$, and that the kinetic energy of the particle is converted to calorimetric deposits in the detector, such that $E_k \approx Q_{PFO}$, where Q_{PFO} is defined in equation 7.2. Combined with the energy momentum relationship, this yields the expression

$$p_{PFO} = \sqrt{Q_{PFO}^2 + 2m_0Q_{PFO}}, \quad (7.3)$$

where m_0 is the rest mass of the particle, and p_{PFO} is the momentum. For convenience, this momentum quantity is normalised, such that $p_{v,z} \equiv |p_v \cdot \hat{z}|/|p_v|$, for instance.

The stacked distributions of the opening angle ζ , as well as the x -, y - and z -components of p_v are shown in Figure 7.7. It is clear from the distributions of ζ is useful for identifying split cosmic ray tracks with $\zeta \approx 0$ or $\zeta \approx \pi$ radians, and that the components of p_v are particularly well-suited to picking out neutrino interactions from a background of cosmic ray interactions. As expected, the z -component of p_v peaks around 1 for neutrino interactions, and p_v has a large y -component for cosmic ray interactions.

7.3.2 Directional Information and Cosmic Ray Background

The previous section has outlined how a number of topological variables are correlated with the probability that a given event is a background event. The direction finding tool that

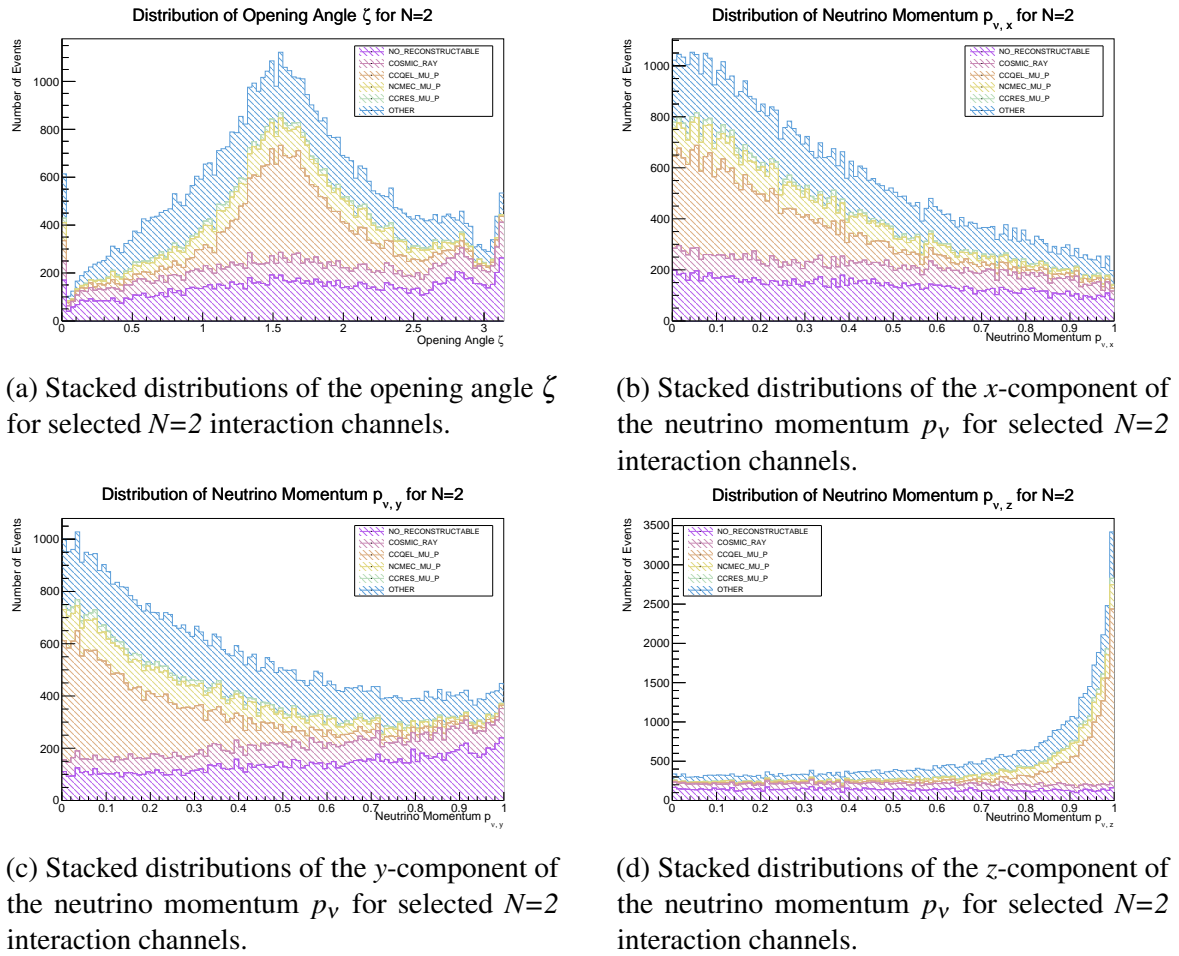



Fig. 7.7 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$.

was described in Chapter 5 can be used to capture directional information, which is highly useful in separating neutrino-induced particles from cosmic ray particles, since the latter usually travel downwards. In addition, the mass-dependence of $\langle -dE/dx \rangle$ can be used to help formulate PID variables, as is discussed in Section 7.3.3. The per-particle variables that are provided directly by the tool are:

- $\Delta\chi_{FB}^2/N$: the difference in χ^2/N between the forwards and backwards fits. This variable contains information about whether the particle is likely to be forwards-going or backwards-going. This is a useful quantity because neutrino-induced particles are known to be forwards-going in the absence of backwards scatters, whereas cosmic ray particles have an equal probability to be forwards- or backwards-going;
- $\Delta\chi_{DU}^2/N$: the difference in χ^2/N between the downwards and upwards fits. Analogously to the previous variable, this variable will contain information about whether the particle is likely to be downwards- or upwards-going. This is useful because cosmic ray particles are known to be downwards-going, whereas neutrino-induced particles have an equal probability of being upwards- or downwards-going;
- P_C : as described in Chapter 6, the cosmic probability P_C is a more sophisticated variable that is based on $|\Delta\chi_{DU}^2/N|$, N and χ_{min}^2/N and is designed specifically to separate neutrino-induced particles that leave the detector through its top face, as was described in Section 6.2.1;
- χ_{min}^2/N : this variable is a measure of goodness of fit of the previous forwards and backwards (and hence upwards and downwards) fits;
- $\delta\chi_{min,BA}^2/N$: the difference in χ_{min}^2 resulting from finding the best split position, if it is present, as described in section 5.5. This variable is constructed for the longest particle in \mathbf{v}_{reco} , and will help identify long clusters that consist of hits induced by separate particles, leading to the reconstructed primary particle multiplicity being smaller than the true primary multiplicity: $N < T$.

Figure 7.8 shows the stacked distributions of $\Delta\chi_{FB}^2/N$ and $\Delta\chi_{DU}^2/N$, with the Bragg peak filter of Section 6.2.1 applied to clarify the underlying distributions. The distributions are shown for $N=1$ events, where the cosmic-ray background is most prominent. As expected, the distribution of $\Delta\chi_{FB}^2/N$ is predominantly negative for neutrino-induced events, since the resulting interactions are boosted along the beam direction. Similarly, the distribution of $\Delta\chi_{DU}^2/N$ is predominantly negative for cosmic rays, indicating that they are downwards-going.

(a) The $\Delta\chi^2_{FB}/N$ stacked distribution for $N=1$.(b) The $\Delta\chi^2_{DU}/N$ stacked distribution for $N=1$.Fig. 7.8 Stacked distributions for the $N=1$ $\Delta\chi^2_{FB}/N$ and $\Delta\chi^2_{DU}/N$ variables.

7.3.3 Calorimetric Information and Particle ID

In addition to cosmic-ray background events, Table 7.2 contains neutrino-induced background channels in which the neutrino interaction topologically resembles the target channel, but for which a single particle is different from the target final state. For instance, an NCQEL p event can be mistaken for a CCQEL μ^- event with a low-energy muon, or a CCRES $\mu^- + \pi^+$ event can be mistaken for a CCQEL $\mu^- + p$ event. This type of background is particularly prominent in the case of $N=2$, where interaction types with a muon and a pion are common. To help resolve such events, a number of simple calorimetric variables can be formulated that are correlated with the underlying particle mass. Furthermore, it is demonstrated that in many cases, the collection of fit charges \tilde{Q}_{fit} produced by the direction fitting tool provide a less noisy set of hits on which such variables can be based. Lastly, the direction fitting procedure described in Chapter 5 relies on the Bethe equation, and the mass-dependence of $\langle -dE/dx \rangle$ can be utilised to formulate PID variables. To this end, the direction fitting tool is modified to take 3D clusters as input and a proton χ^2 variable can be constructed by fixing the fit mass to be the proton mass.

The example direction fit in Figure 7.9 indicates that the observed charges \tilde{Q} and fitted charges \tilde{Q}_{fit} can be used to formulate simple PID variables, where the quantities \tilde{Q} and \tilde{Q}_{fit} were defined in equations 5.8 and 5.11. A study on proton identification in the MicroBooNE detector [97] has demonstrated that the sums of the first and last three hit charges in a track

are useful variables:

$$\begin{aligned}\tilde{Q}_A &= \sum_{i=1}^3 \tilde{Q}_i/3 \\ \tilde{Q}_B &= \sum_{i=N-2}^N \tilde{Q}_i/3\end{aligned}\tag{7.4}$$

where $i = 1 \cdots N$ indexes the hits in each particle, ordered by z , and N is the number of hits in the particle. In particular, the quantities $\tilde{Q}_{min} = \min(\tilde{Q}_A, \tilde{Q}_B)$ and $\tilde{Q}_{max} = \max(\tilde{Q}_A, \tilde{Q}_B)$ can be used as a basis for a number of additional variables. Since the fit charges \tilde{Q}_{fit} are much more robust to effects such as charge-sharing between hits near the interaction vertex, they are more suitable to formulate these quantities than reconstructed observed hits. The set of simple PID variables based on, and including, \tilde{Q}_{min} and \tilde{Q}_{max} are:

- *Mean \tilde{Q} for the first and last 3 hits in z , \tilde{Q}_{min} and \tilde{Q}_{max} :* larger masses lead to larger values of $\langle -dE/dx \rangle$, and hence also to larger values of \tilde{Q}_{min} and \tilde{Q}_{max} . The rationale behind retaining both variables, despite the fact that they are correlated, is that either track endpoint may be obscured;
- *Ratio of \tilde{Q}_{max} and \tilde{Q}_{min} , $\tilde{Q}_R = \tilde{Q}_{max}/\tilde{Q}_{min}$:* while this variable introduces a minor redundancy, this variable is highly-correlated with the presence of a Bragg peak in the detector. The absence of a Bragg peak will lead to a flat direction fit, leading to $\tilde{Q}_R \approx 1$;
- *Extent in \tilde{Q}_{fit} normalised by 2D cluster length, $\tilde{Q}_{R/L} \equiv \tilde{Q}_{range}/L_{2D}$:* here, \tilde{Q}_{range} is the difference between the smallest and largest values of \tilde{Q} in the W view cluster, such that $\tilde{Q}_{range} = \max(\tilde{Q}) - \min(\tilde{Q})$. As before, L_{2D} is the 2D track length in the W view, and dividing by L_{2D} turns \tilde{Q}_{range}/L_{2D} into a measure of $\langle -dE/dx \rangle$;
- *The mean value of \tilde{Q}_{fit} in the particle, $\langle \tilde{Q}_{fit} \rangle$:* for higher values of $\langle -dE/dx \rangle$, the \tilde{Q}_{fit} values resulting from the direction fit assume higher values.

Figure 7.10 shows the stacked distributions for the most effective simple PID variables for the shortest particle in $N=2$ events. In general, when the shortest particle is a proton, $\tilde{Q}_{min,S}$ and $\tilde{Q}_{R/L,S}$ assume higher values.






Fig. 7.9 An example CCQEL μ^- event, in which the muon observed charge deposits \tilde{Q}_{obs} and the magenta fitted charges \tilde{Q}_{fit} are shown. The empty region for $15 \lesssim L_{2D} \lesssim 50$ cm is due to an inactive detector region in the collection plane. The spike in \tilde{Q}_{obs} near $L_{2D} = 0$ cm is due to overlapping hits near the interaction vertex.



(a) Stacked distribution of $\tilde{Q}_{min,S}$ for $N=2$.



(b) Stacked distribution of $\tilde{Q}_{R/L,S}$ for $N=2$.

Fig. 7.10 Stacked plots the two most important simple PID variables, both for the shortest PFO in $N = 2$ events.

7.3.4 Proton PID

The direction fitting method of Chapter 5 can be modified to work in three dimensions, which enables the formulation of more sophisticated PID variables. When performing direction fits in three dimensions, the 3D track length is known. In Chapter 5, an additional fit parameter had to be introduced to resolve L_{3D} . In the 3D case, the value of dx for a measured value of $-dE/dx$ can be resolved explicitly, by using the wire spacing and the angle the track segment makes with the wire pitch direction. The fundamental quantity of interest is $d\tilde{Q}/dx$, the amount of charge collected on a wire normalised by the 3D track segment it corresponds to. The assumption is made here that the track direction is constant on the length scale corresponding to individual wires. In addition to having access to 3D track information, the known positions of 3D hits within the detector makes it possible to calibrate the $d\tilde{Q}/dx$ values to account for various effects:

- *Misconfigured or cross-connected TPC channels* [98]: effects such as ASIC misconfiguration can change the electronic gain of electronic channels, and cross-connected or touching wires, can distort the electric field between wire planes. Such effects affect roughly 10% of channels and introduce a location dependence on measured charge values;
- *Space charge effects (SCE)* [82, 99]: the large amount of cosmic rays causes an accumulation of positive argon ions, which distort the uniformity of the electric field by affecting its magnitude and direction. Changes in the electric field affect both the recombination effect and the reconstructed trajectories of particles, effectively contracting near the cathode leading to smaller reconstructed dx values, which increases the measured value of dQ/dx ;
- *Electron attenuation* [100]: electronegative contaminants such as H_2O and O_2 can capture drift electrons, reducing the measured value of dQ/dx ;
- *Diffusion* [91]: clouds of ionisation electrons diffuse in both the longitudinal and transverse directions. Transverse diffusion can cause drift electrons to end up on neighbouring wires, and longitudinal diffusion can change the distribution of collected charge with time, especially for signals produced at longer drift distances;
- *Recombination* [91]: a fraction of ionisation electrons recombine with nearby positive argon ions, or do not fully liberate. This effect, if not taken into account, leads to an underestimation of particle energy.


Much of this calibration work is outlined in [97] and has recently been superseded by the study in [91]. In this paper, the $d\tilde{Q}/dx$ calibration is separated into three parts: the corrections in the x direction, in the yz plane and temporal variations in the data. The latter effect is small, and are neglected in this study. Furthermore, the assumption is made that the corrections in the x coordinate and the yz plane are separate, which is a simplification, but one that works well in practice.

Figure 7.11 shows the corrections calculated in [97] for the x coordinate (Figure 7.11a) and the yz plane (Figure 7.11c). Furthermore, the x direction correction has been parametrised using a fit of the form $y = a \cdot x + b \cdot \sin(x)$, where a and b are fit parameters, and is shown in Figure 7.11b. This parametrisation of the correction is computationally efficient, and is used to correct in the drift time coordinate in what follows. In addition, Figure 7.11c shows how the correction in the yz plane is applied: the detector region outlined by the green dashed lines are set to the mean correction value in this region, and 1 elsewhere. This correction is necessary due to a cluster of wires on the U plane for which the response signals are observed to be too small in magnitude.


In addition to calibrating the $d\tilde{Q}/dx$ values, $d\tilde{Q}/dx$ must be mapped to dE/dx . So far, only global detector effects have been taken into account, such as space charge effects, and faulty TPC channels, but not recombination, which is inversely proportional to the local electric field $|\vec{E}|$. Because the recombination effect depends on dE/dx , it cannot be accounted for by a global correction. To convert $d\tilde{Q}/dx$ values into dE/dx the ‘modified box model’ is used [101], which calculates dE/dx as

$$\left(\frac{dE}{dx}\right)_{calibrated} = \frac{\exp\left(\left(\frac{d\tilde{Q}}{dx}\right) \frac{\beta' W_{ion}}{C \rho \mathcal{E}}\right) - \alpha}{\frac{\beta'}{\rho \mathcal{E}}}. \quad (7.5)$$


In equation 7.5, $C = 0.00507689$ is a calibration constant used to convert ADC values to number of electrons, $W_{ion} = 23.6 \times 10^{-6}$ eV is the work function of argon, $\mathcal{E} = 0.273$ kV/cm is the MicroBooNE drift electric field, $\rho = 1.38$ g/cm² is the density of liquid argon at a pressure of approximately 124.11 kPa, $\beta' = 0.212$ (kV/cm)(g/cm²)/MeV and $\alpha = 0.93$ is a model constant. Most of these parameters are defined by physics, but some, like α and β were determined empirically by the ArgoNeuT collaboration, the experiment where the modified box model was developed originally. The modified box model is an empirical modification of the ‘box model’ (through the parameters α and β), which in turn is based on the Jaffe diffusion equations. The electron diffusion and ion mobility terms in the Jaffe diffusion equations can be neglected in liquid argon, leading to the original box model



(a) The charge calibration values for the drift time coordinate.



(b) A parametrization of the calibration of the x coordinate.



(c) The charge calibration values in the yz plane.

Fig. 7.11 The charge calibration values for the x coordinate (7.11a) and the yz plane (7.11c). A fit to the distribution of Figure 7.11a of the form $y = a \cdot x + b \cdot \sin(x)$ is shown in 7.11b. Images 7.11a and 7.11c from [97].

[102]. Using the detector calibration and the modified box model approach, it is possible to calculate dE/dx values for individual hits, and a 3D direction fit can be constructed for these values. The fitting tool performs a direction fit on the 10 cm of the track that is most likely to contain the Bragg peak, since stopping power values in this region of residual track length are significantly affected by particle mass.

The fitting procedure itself is performed analogously to the 2D case described in Chapter 5. Because there are so many similarities, only a brief overview of the fitting procedure is given here. First, the hit filtering procedure described in Section 5.3 is applied to the target hit collection. The next step is to select the track segment that is most likely to contain the Bragg peak, by selecting the 10 cm of track length nearest the endpoint with the highest summed dE/dx value. Fitted dE/dx values are calculated for this observed subset of dE/dx values, analogously to the \tilde{Q}_{fit} objects in the 2D fit. The particle mass in this fit is fixed to the proton mass, and the fit is given one degree of freedom that allows the particle energy at $L = L_{max}$ to be different from zero. This will give the fit the freedom to prevent the last dE/dx value from exhibiting asymptotic behaviour and will enable it to fit truncated tracks. Since the fitted values of dE/dx are determined by the particle mass, the χ_p^2 value of the fit are small for real protons, and very large for other particles. For muons, for example, the observed values of dE/dx are much smaller than the fitted values, leading to a large value of χ_p^2 . This idea is analogous to the ideas presented in [103]. This variable is a useful particle ID variable, as is clear from Figure 7.12, which shows the χ_p^2 stacked distribution for $N = 2$ cosmic and signal channels.

7.4 Event Selection

The variables formulated in Section 7.3 are well-suited for use in a multivariate analysis (MVA) that is designed to select signal events of the types $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\bar{\nu}_\mu + \text{Ar} \rightarrow \mu^+ + X$ from the various background channels discussed in Section 7.2. The multivariate analysis is performed in the TMVA framework, which is a ROOT-integrated environment for the processing, parallel evaluation and application of multivariate classification and regression techniques [104]. This package contains implementations for many common machine learning algorithms and data preprocessing techniques makes it possible to easily create training and validation sets and has built-in functionality to compare the performance of different methods.

Many machine learning techniques have been successfully applied to signal selection problems in particle physics, and for the purposes of this study a subset of the TMVA




Fig. 7.12 The stacked distribution of χ_p^2 for $N = 2$. The OTHER channel collection has been added to the stack to demonstrate the separating power of this variable when distinguishing between signal and background.

algorithms has been selected that range in complexity. The selected methods are maximum likelihood estimation (MLE), k-nearest neighbours (KNN), support vector machines (SVM), boosted decision trees (BDT) and multi-layer perceptrons (MLP). Traditional statistical techniques such as MLE and KNN can, for certain problem formulations, separate signal from background events with good signal efficiency and background rejection. The advantage of such techniques is that they are interpretable (they are not ‘black boxes’) and are robust. More sophisticated machine learning techniques such as SVMs and BDTs can often attain better signal/background separation, but also run the risk of overfitting on their training samples. In addition, artificial neural networks and deep learning techniques have been successfully applied in the context of LArTPCs in recent years [105]. A comparison can be made between the performance of each technique by comparing their receiver operating characteristic (ROC) curves, which plot background rejection as a function of signal efficiency for binary classifiers. The area under the ROC curve, indicated by the shorthand \int_{ROC} , is a good measure of the efficacy of each technique. Note that a perfect classifier would have $\int_{ROC} = 1$.

Figure 7.13 shows the ROC curves for the five outlined classification techniques trained to select signal events of the form $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ (Figure 7.13a) and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ (Figure 7.13b). This comparison has been made without any model hyperparameter optimisation: the default values of the TMVA implementations have been used. Each classifier has been checked for overtraining by applying the Kolmogorov-Smirnov test to the signal

and background response distributions for the training and validation sets, ensuring that the distance measure between each pair of training and validation distributions is small. It is clear from Figure 7.13 that default implementation of the BDT is the most effective classifier. For this reason, and because BDTs are quick to train and to apply, the BDT has been selected as the classification technique for this event selection.




Fig. 7.13 ROC curves for several methods for the $N=1$ and $N=2$ contained event sample.

7.4.1 Impact of Correlated Variables and Directional Information on BDT Performance

The ROC curves in Figure 7.13 were attained using every variable outlined in Section 7.3. However, some of the information contained in these variables is highly-correlated. To quantify these correlations, the matrix containing the linear correlation coefficients between all variables is investigated in the case that $N=2$, since the $N=1$ variables are a subset of the $N=2$ variables. The TMVA package supplies a colour-coded correlation matrix containing the the linear correlation coefficient for each variable pair, where the linear correlation coefficient $\rho_{X,Y}$ between two random variables X and Y is defined as

$$\rho_{X,Y} = \frac{\text{Cov}(X,Y)}{\sigma_X \sigma_Y}, \quad (7.6)$$

where $\text{Cov}(X,Y)$ is the covariance of two random variables X and Y . The resulting correlation matrix for $N=2$ is shown in Figure 7.14a. There are a number of highly-correlated submatrices with dimensions 2×1 , 2×2 and a very large submatrix with dimensions 5×5 . These correlations are all very intuitive and are enumerated below. Furthermore, for each correlated

- The correlated 2×2 submatrices are due to the fact that the x , y and z coordinates of the longest and shortest PFOs are strongly correlated, e.g. $z_{min,s}$, $z_{max,s}$, $z_{min,l}$ and $z_{max,l}$. This is unsurprising: for both the long and short PFOs, one endpoint is approximately the position of the reconstructed neutrino vertex. In addition, most short PFOs are very short, and so their maximal coordinates are very similar to both their minimal coordinates, and one of the long PFO endpoints. Removing the short PFO coordinate variables does not adversely affect the BDT performance;
- $\Delta\chi^2_{FB}/N$ and $\Delta\chi^2_{DU}/N$ are highly-correlated because they are formulated using the same direction fits. As was explained in Chapter 6, $\Delta\chi^2_{FB}/N = \Delta\chi^2_{DU}/N$ when the high- z particle endpoint is the low- y endpoint. Both variables are kept, as they contain different directional information;
- $\Delta\chi^2_{DU}/N$ and the cosmic probability P_C are highly-correlated, since P_C is based on $|\Delta\chi^2_{DU}/N|$. P_C is designed to remove cosmic rays in an uncontained context, and can safely be removed when a containment cut is applied.
- The large correlated submatrix is caused by two factors:
 - First, the $\Delta\chi^2$ variables are highly-correlated with the simple calorimetric variables, since the latter are also based on the direction fit \tilde{Q}_{fit} points. For instance, the ratio between \tilde{Q}_{min} and \tilde{Q}_{max} is approximately 1 if the particle is not contained, which will also drive $\Delta\chi^2$ to 0;
 - Second, the individual simple calorimetric variables are correlated. For instance, when a particle is not contained, \tilde{Q}_{min} and \tilde{Q}_{max} will have approximately the same value. In this case, \tilde{Q}_{range}/L will also be approximately zero. In fact, most of the simple calorimetric variables are redundant: only \tilde{Q}_{min} really has an impact. It is simple to see why: \tilde{Q}_{max} is in the Bragg peak, where the individual charge values, while large, are subject to hit reconstruction effects and the particle filter introduced in 5. Furthermore, the ratio between \tilde{Q}_{min} and \tilde{Q}_{max} contains no additional information. However, \tilde{Q}_{range}/L does empirically appear to contain some uncorrelated information.

The TMVA framework provides variable importance rankings, which, in the case of a BDT, are obtained by recording how frequently a variable is used to split a decision tree node, and by weighting each split occurrence by the squared separation gain it has achieved and by the number of events in the node. Querying this information reveals that a number of variables do not contribute to the performance of the BDT:

- Whereas y_{max} is very useful for cosmic removal, y_{min} has virtually no separating power and can be removed. Similarly, z_{min} and z_{max} provide little discrimination;
- The azimuthal angles are very correlated with $p_{v,x}$ in the correlation matrix in Figure 7.14b and can be removed. While the polar angles are similarly correlated with $p_{v,y}$, removing the polar angle variable removes information about event kinematics and negatively impacts the performance of the BDT;
- $\tilde{Q}_{10cm,s}$ provides little discrimination: in most cases this is simply \tilde{Q}_{PFO} , since the shortest PFO is so short;
- The variable Δy_s provides little discrimination: whereas Δy_l contains some information relating to whether the particle is likely to be a cosmic ray, no additional information is contained in Δy_s

Based on the variable correlation coefficients ρ and the variable rankings, the variables $\langle \tilde{Q} \rangle_{L/S}$, $z_{min/max,L}$, $y_{min,L}$, $P_{C,L/S}$, $\tilde{Q}_{R,L/S}$ and $\tilde{Q}_{max,L/S}$ as well as all coordinate variables associated to the short PFO, $x/y/z_{min/max,S}$, are removed. This reduces the total number of variables from 47 to 30 for $N = 2$ and from 27 to 20 for $N = 1$. This new set of variables is referred to as the ‘reduced variable set’. The $N=2$ signal correlation matrix for the reduced variable set is shown in Figure 7.14b.

The impact of the reduction of the variable set and of the exclusion of directional information can be quantified by calculating the values of \int_{ROC} for the BDT for $N = 1$ and $N = 2$. Table 7.5 compares the BDT performance for $N = 1$ and $N = 2$ in terms of \int_{ROC} for three configurations: i) when using all variables, ii) when using the reduced variable set and iii) when using the reduced variable set without directional information (labelled ND). It is clear from Table 7.5 that including directional information improves the performance of the BDT, despite the fact that $\Delta\chi^2_{DU}/N$ was already used in the application of the DCRT pre-selection cut. In addition, the removal of correlated variables has a negligible impact on the performance of the BDT.

7.4.2 Impact of the Event Selection on Simulated Data

Tables 7.6 and 7.7 show the effect of applying the BDT selection procedure. As a result of applying the event selection, the purity increases from 15.69% to 57.91%, and the efficiency is reduced from 75.58% to 63.09% for $N = 1$, and the purity increases from 35.48% to 78.30% and the efficiency is reduced from 81.20% to 62.65% for $N = 2$, compared to the

BDT Performance Comparisons			
	All Variables	Reduced Variable Set	Reduced Variable Set (ND)
$N = 1$	0.952	0.948	0.938
$N = 2$	0.973	0.971	0.965

Table 7.5 A comparison of the values of \int_{ROC} when using all variables, the reduced variable set and the reduced variable set without directional information (labelled ‘ND’ for ‘no direction’). The reduced variable set consists of the variables that remain after removing highly-correlated variables and variables with little discriminative power.

event samples before applying the BDT outlined in Table 7.2. The purity increase is larger for $N = 2$, since the cosmic-ray background is topologically more similar to $N = 1$ events, as was discussed in Section 7.2. Furthermore, the addition of directional information leads to a purity increase of 4.97% while also increasing the efficiency by 1.26% for $N = 1$, and a purity increase of 0.98% and an efficiency increase of 0.33% for $N = 2$. This increased selection performance due to directional information is in addition to the purity increase due to the DCRT cut, as was also detailed in Section 7.2.

<i>N=1 Impact of Event Selection</i>								
Interaction Type	Total		Before Selection		After Selection		With Direction	
SIGNAL	19867	9.36%	15015	15.69%	12534	52.94%	12785	57.91%
CCQEL μ	12322	5.80%	11684	12.21%	9915	41.88%	10096	45.73%
CCRES μ	661	0.31%	556	0.58%	412	1.74%	417	1.89%
CCMEC μ	2788	1.31%	2595	2.71%	2081	8.79%	2143	9.71%
SIGNAL SMALL ($\leq 1\%$)	-	-	178	0.19%	126	0.53%	129	0.58%
BACKGROUND	192491	90.64%	80693	84.31%	11143	47.06%	9294	42.09%
COSMIC RAY	126293	59.47%	72173	75.41%	8526	36.01%	6769	30.66%
ν NO RECONSTRUCTABLE	83923	39.52%	48114	50.27%	5072	21.42%	4031	18.26%
CCQEL $\mu + p$	14605	6.88%	1396	1.46%	747	3.15%	706	3.20%
CCMEC $\mu + p$	6799	3.20%	645	0.67%	382	1.61%	349	1.58%
BACKGROUND SMALL	-	-	6481	6.77%	1489	6.29%	1471	6.66%
TOTAL	212358	100.00%	95708	100.00%	23677	100.00%	22079	53.19%
PURITY	-	-	15015/95708	15.69%	12534/23677	52.94%	12785/22079	57.91%
EFFICIENCY	-	-	15015/19867	75.58%	12534/19867	63.09%	12785/19867	64.35%

Table 7.6 The impact of applying the event selection on the $N = 1$ interaction type tables.

<i>N=2 Impact of Event Selection</i>								
Interaction Type	Total		Before Selection		After Selection		With Direction	
SIGNAL	24774	11.67%	20116	35.48%	15440	77.32%	15522	78.30%
CCQEL $\mu + p$	14605	6.88%	12632	22.28%	10519	52.67%	10544	53.19%
CCRES $\mu + p$	1537	0.72%	1267	2.24%	724	3.63%	743	3.75%
CCMEC $\mu + p$	6799	3.20%	5928	10.45%	4030	20.18%	4069	20.52%
SIGNAL SMALL ($\leq 1\%$)	-	-	289	0.51%	167	0.84%	166	0.84%
BACKGROUND	187584	88.33%	36588	64.52%	4530	22.68%	4303	21.70%
COSMIC RAY	126293	59.47%	21186	37.36%	624	3.12%	523	2.64%
ν NO RECONSTRUCTABLE	83923	39.52%	13867	24.46%	420	2.10%	340	1.72%
CCRES $\mu + \pi^+$	3794	1.79%	3088	5.45%	994	4.98%	969	4.89%
CCRES $\mu + p + \pi^+$	4828	2.27%	687	1.21%	309	1.55%	303	1.53%
NCRES $p + \pi^-$	505	0.24%	395	0.70%	186	0.93%	202	1.02%
CCDIS $\mu + \pi^+$	1485	0.70%	1177	2.08%	359	1.80%	340	1.72%
CCMEC $\mu + p + p$	3407	1.60%	387	0.68%	237	1.19%	239	1.21%
OTHER INTERACTION	5819	2.74%	1419	2.50%	349	1.75%	329	1.66%
BACKGROUND SMALL	-	-	8252	14.55%	1473	7.38%	1400	7.06%
TOTAL	212358	100.00%	56704	100.00%	19970	100.00%	19825	53.19%
PURITY	-	-	20116/56704	35.48%	15440/19970	77.32%	15522/19825	78.30%
EFFICIENCY	-	-	20116/24774	81.20%	15440/24774	62.32%	15522/24774	62.65%

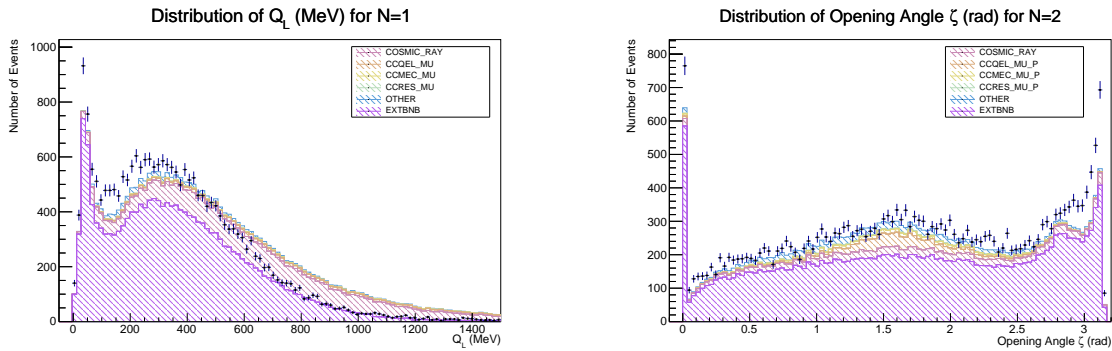
Table 7.7 The impact of applying the event selection on the $N = 2$ interaction type tables.

7.5 Simulated and Real Data Comparison

The Monte Carlo distributions of the previous section can be compared to the distributions of the variables observed in real data. Each recorded event in MicroBooNE corresponds to a BNB neutrino spill, although only about 1 in 600 events produce a neutrino interaction. To reduce the amount of data stored on disk, events are filtered by the software trigger, which is described in Chapter 3. The software trigger mainly looks at PMT activity that coincides with the $1.6 \mu\text{s}$ beam spill window. Despite the software trigger, the majority of recorded events still do not contain a neutrino interaction. However, the simulated neutrino event sample, henceforth labelled MC BNB + COSMIC, contains a neutrino interaction in every event, and so to compare the simulated event sample to the real data, the distributions must be normalised. In order to preserve the information contained in the BNB DATA sample, this sample is never rescaled. Instead, a Run 1 sample is used that contains data from when the BNB was turned off, resulting in a collection of cosmic-only events. This event sample is referred to as the EXTBNB sample. The variable distributions from the EXTBNB sample are added to those of the MC BNB COSMIC to make the data and simulation distributions comparable. This procedure will result in an event sample containing one event with a simulated neutrino interaction and simulated cosmic rays for every 600 EXTBNB events,

reproducing the expected frequency of neutrino interactions in the BNB sample. The EXTBNB and MC BNB COSMIC samples are then scaled in order to correspond to the exposure (protons-on-target) of the BNB data sample. This normalised, combined sample is referred to as EXTBNB + MC.

Quantitatively, the simulated MC BNB COSMIC variable distributions are rescaled by a factor $\alpha_{DATA} = \text{POT}_{BNB}/\text{POT}_{MC}$, such that they correspond to the same number of protons on target as the BNB sample distributions. The normalisation of off-beam to on-beam data is done by normalising to the same amount of spills/pulses received before the software trigger was applied. The EXTBNB sample is scaled by a factor $\beta_{DATA} = \text{TRIG}_{EXTBNB}/\text{TRIG}_{BNB}$, the ratio of the number of EXTBNB triggers and the number of BNB beam triggers. The rescaled EXTBNB samples are then added to the simulated distributions, to ensure the fraction of neutrino interactions in the sample is consistent. In addition to this, each simulated event is weighted by an additional factor $\gamma_{DATA} = 1.028$, since it has also been shown that the TPC starts at $z = 463.6$ m, while the flux is generated considering a start position $z = 470$ m, leading to a correction factor of $(470/463.6)^2 = 1.028$ [106]. Furthermore, the neutrino flux in the simulation is known to be slightly deficient, leading to a small energy-dependent per-event correction factor to the ν_μ flux [106]. Lastly, the signal response on the detector sense wires is lower in this version of the simulation compared to data; a study performed in [97] demonstrates that the ADC to electron count calibration factor F_{ADC} is $(5.076 \pm 0.001) \times 10^{-3}$ in simulation and $(4.113 \pm 0.011) \times 10^{-3}$ in data, such that the electron count (and hence charge) is underestimated in simulation. To account for this, a correction factor of $4.113/5.076 \approx 0.810$ is applied to charge values in data.



(a) Distributions of $N = 1$ Q_L for simulation and data with a data-driven cut at $Q_L < 100$ MeV applied.

(b) Distributions of $N = 2$ opening angle ζ for simulation and data with a data-driven cut applied that retains $0.1 \leq \zeta \leq 3.04$ rad.

Fig. 7.15 The data-driven pre-selection cuts, excluding $Q_L < 100$ MeV (7.15a) for $N = 1$ and retaining $0.1 \leq \zeta \leq 3.04$ rad for $N = 2$.

However, there are known differences between the distributions due to the version of the simulation in this study and the data distributions. These differences are most prominent in the distributions of the longest PFO charge Q_L and the opening angle ζ . Figure 7.15a shows the distributions of Q_L for $N = 1$ and the opening angle ζ for $N = 2$, where the differences between the simulated and data distributions are most pronounced. The version of the simulation used in this analysis uses CORSIKA as the generator for cosmic-ray interactions [77]. However, it is known now that this cosmic ray simulation differs from the observed cosmic ray (EXTBNB) data. In particular, CORSIKA underestimates the number of cosmic ray particles in the low-energy $E < 100$ MeV regime [107]. This effect will affect the simulated cosmic charge distributions, and in part explains the reduced peak in the simulated Q_L distributions in Figure 7.15a where $Q_L < 100$ MeV. The difference between the cosmic ray simulation and data has been a significant source of systematic uncertainty for other MicroBooNE studies [81]. Furthermore, it has been observed that the background-subtracted reconstructed charged particle multiplicities differ for simulated events generated by the MicroBooNE default GENIE configuration, used in this study, and data events [108]. For this reason, MicroBooNE is in the process of replacing the CORSIKA simulation with cosmic ray data overlays, using EXTBNB cosmic-ray backgrounds that are added to the BNB simulated neutrino interaction in the next production of the simulation, currently ongoing. In addition, the modelling of dynamic-induced charge (DIC) effects in the simulation is absent. For tracks that are highly inclined (nearly orthogonal) to the wire planes, signals are induced on neighbouring wires when ionisation electrons travel past the ‘primary’ wire on the induction planes. This can cause destructive interference when super-imposing induction waveforms, causing hits to effectively be lost and reducing the total measured charge. The next generation of the Monte Carlo simulation takes this effect into account through an updated deconvolution process [109, 110].

These differences between simulation and data would cause a difference in the total number of selected events for simulation and data, since events near $\zeta \sim 0$ rad and $\zeta \sim \pi$ rad, for instance, resemble cosmic-ray tracks and are likely to be removed. To exclude the regions of the phase space in which there are understood differences between simulation and data, two new data-driven pre-selection cuts are applied, based on Figures 7.15a and 7.15b. For $N = 1$, a cut on $Q_L > 100$ MeV is applied, and for $N = 2$ a cut retaining $0.1 \leq \zeta \leq 3.04$ rad is applied. The cuts on Q and ζ are empirically formulated to exclude the regions of largest tension between simulated and data variable distributions, and remove 13.11% of simulated events and 17.42% of data events for $N = 1$ and 9.64% of simulated events and 10.72% of data events for $N = 2$. The simulated and data distributions with the data-driven pre-selection cuts applied are shown in Figures 7.16 and 7.17, and reasonable agreement is generally seen.

The BDT is retrained on the simulated training set with these data-driven pre-selection cuts applied, and the results of applying the retrained event selection to the normalised event samples are shown in figures 7.19, 7.20 and 7.21. There is general agreement between the simulation and data post-selection distributions. Where differences remain, they are in part due to differences between the data and simulation variable distributions that the pre-selection cuts did not correct.

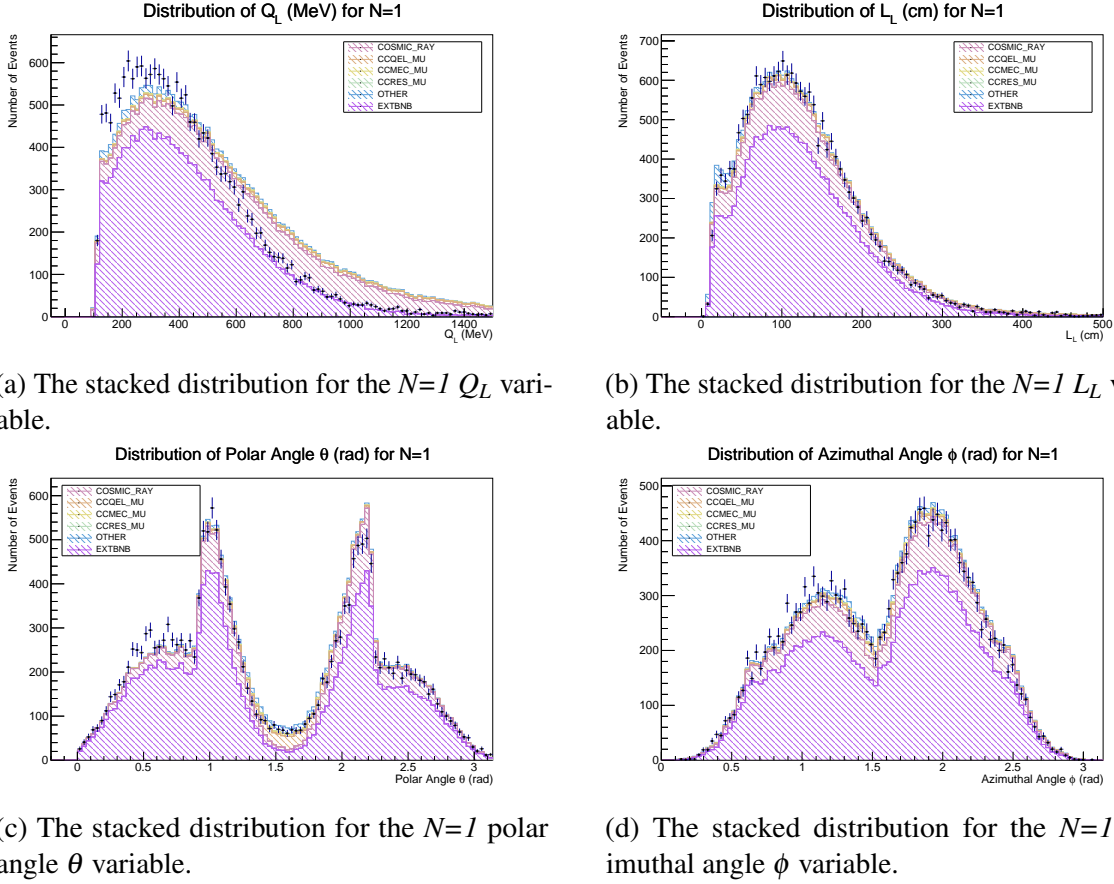


Fig. 7.16 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$ with data points overlaid before applying the event selection. A data-driven pre-selection cut excluding $Q_L < 100$ MeV is applied in these distributions.

Table 7.8 shows the number of selected events for $N = 1$ and $N = 2$ for the normalised EXTBNB + MC and BNB data samples, before and after applying the data-driven preselection cuts. This table shows that without excluding poorly-modelled regions of the simulation phase space, the selection retains more simulated events than data events for $N = 1$. These differences can be quantified by considering the total efficiency $\eta_{tot} \equiv N_B/N_A$, the total fraction of events retained by the selection, where N_B is the number of events before the selection and N_A is the number of events after applying the selection. For $N = 1$, the

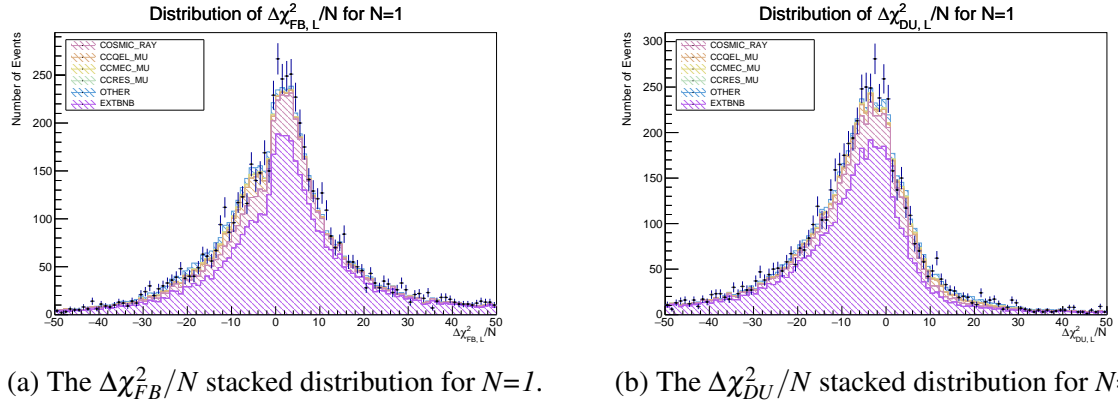


Fig. 7.17 Stacked distributions for the $N=1$ $\Delta\chi^2_{FB}/N$ and $\Delta\chi^2_{DU}/N$ variables with data points overlaid before applying the event selection. A data-driven pre-selection cut excluding $Q_L < 100$ MeV is applied in these distributions.

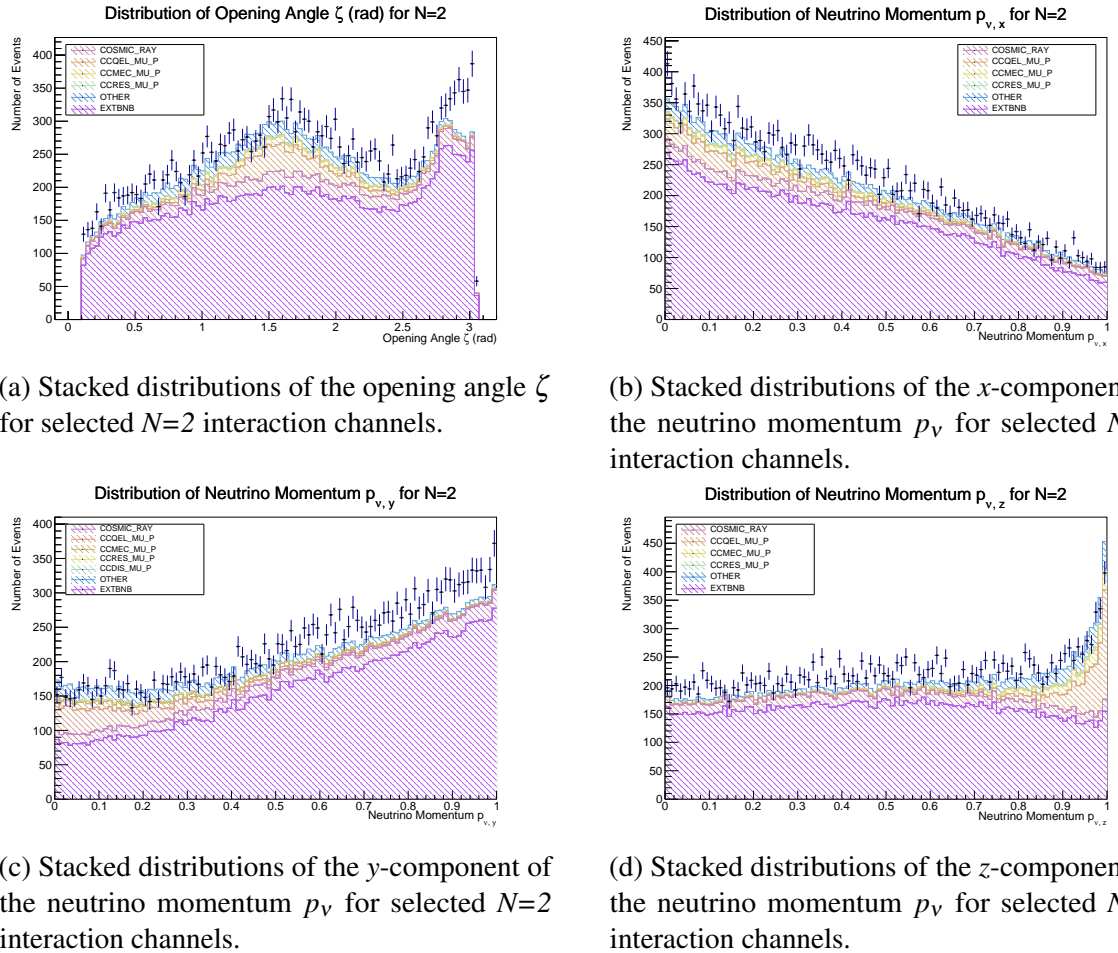





Fig. 7.18 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=2$ with data points overlaid before applying the event selection. A data-driven pre-selection cut retaining $0.1 \leq \zeta \leq 3.04$ rad is applied in these distributions.




(a) The stacked distribution for the $N=1$ Q_L variable.



(b) The stacked distribution for the $N=1$ L_L variable.




(c) The stacked distribution for the $N=1$ polar angle θ variable.




(d) The stacked distribution for the $N=1$ azimuthal angle ϕ variable.

Fig. 7.19 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=1$ after applying the event selection, with data points overlaid.




(a) The $\Delta\chi^2_{FB}/N$ stacked distribution for $N=1$.




(b) The $\Delta\chi^2_{DU}/N$ stacked distribution for $N=1$.


Fig. 7.20 Stacked distributions for the $N=1$ $\Delta\chi^2_{FB}/N$ and $\Delta\chi^2_{DU}/N$ variables with data points overlaid, after applying the event selection.




(a) Stacked distributions of the opening angle ζ for selected $N=2$ interaction channels.



(b) Stacked distributions of the x -component of the neutrino momentum p_v for selected $N=2$.



(c) Stacked distributions of the y -component of the neutrino momentum p_v for selected $N=2$ interaction channels.



(d) Stacked distributions of the z -component of the neutrino momentum p_v for selected $N=2$ interaction channels.

Fig. 7.21 Stacked plots of the distributions of the PFO summed charge, length and polar and azimuthal angles for $N=2$ after applying the event selection, with data points overlaid.

application of the data-driven pre-selection cuts changes η_{tot} for the BNB sample from 10.33% to 13.43% and for the EXTBNB + MC sample from 14.23% to 13.59%. Equivalently, for $N = 2$, the application of the pre-selection cuts changes η_{tot} for the BNB sample from 7.33% to 7.65% and for the EXTBNB + MC sample from 7.54% to 7.96%. The difference in selected events in Table 7.8 for $N = 1$ can be seen to be due to a much increased η_{tot} for the EXTBNB + MC sample, which is corrected by the pre-selection cuts. For $N = 2$, the value of η_{tot} is comparable before and after the cuts for both samples, indicating that the event selection is less affected by the differences between simulation and data for this multiplicity.

Dataset	$N=1$ Post-Selection Event Counts				$N=2$ Post-Selection Event Counts			
	Before Cuts		After Cuts		Before Cuts		After Cuts	
BNB	2662	10.33%	2561	13.43%	1837	7.33%	1713	7.65%
EXTBNB + MC	3286	14.23%	2782	13.59%	1919	7.54%	1855	7.96%
SIGNAL	799	3.46%	552	2.85%	1115	4.98%	1111	5.44%
OTHER ν	186	0.81%	131	0.68%	264	1.18%	261	1.28%
OTHER BACKGROUND	2475	10.72%	2100	10.84%	540	2.41%	483	2.36%

Table 7.8 Event counts after applying the event selection for $N = 1$ and $N = 2$. Event counts are normalised by protons-on-target and trigger counts, as outlined in the text. The ‘Before Cuts’ column indicates the number of events that pass the selection before applying the data-driven pre-selection cuts on \tilde{Q}_L and ζ . The ‘After Cuts’ column indicates event counts with these cuts applied. Within the EXTBNB + MC category, a distinction is made between i) signal events (SIGNAL) ii) neutrino-induced background events (OTHER ν) and iii) all other background (OTHER BACKGROUND). Indented rows in this table indicate subsets, such that e.g. SIGNAL is the subset of EXTBNB + MC that contains signal events. Percentages indicate the fractions of events remaining in the normalised sample after applying the event selection.

7.6 Examples of Selected Events

Figures 7.22 and 7.23 give examples of $N = 1$ events selected by the event selection when applied to the BNB data sample. These figures show the 3D view from the Pandora event display, as well as the projections in the U, V and W views. These figures demonstrate that the event selection is capable of selecting both short (7.22) and long (7.23) muon-like tracks. These examples also demonstrate the difficulty of selecting $N = 1$ μ^- final states, as the cosmic-ray background is extremely topologically similar. Regardless, the $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ event selection is able to achieve a purity of 64.35% on the simulated BNB MC COSMIC sample, and this section has shown that the event selection obtains reasonable agreement between simulated and data event samples when accounting for known simulation

deficiencies. Similarly, Figure 7.24 shows an $N = 2$ interaction identified by the event selection. This interaction is a clear neutrino interaction, and illustrates the type of interaction selected in the highly-pure, post-selection $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ sample. As was outlined in Table 7.7, the main background that remains after this event selection is neutrino-induced $\mu^- + \pi$ interactions, or neutrino interactions with true multiplicity $T = 3$ in which one (small) proton is not resolved by the reconstruction.


7.7 Discussion

This chapter has presented a fully-automated event selection for MicroBooNE $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ interactions, capable of achieving purities of 64.35% and 78.30% and efficiencies of 57.91% and 62.65%, respectively. It was found that the selection procedure is sensitive to differences between the simulation and the Run 1 dataset, due to known deficiencies in the cosmic ray and charge response simulation. Because of these significant discrepancies it was not possible to produce a quantitative cross section measurement. When accounting for such effects with additional empirical pre-selection cuts, the post-selection distributions for simulation and data are in reasonable agreement. Fully-automated event selections of $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ final states for use in LArTPC surface detectors with large amounts of cosmic-ray background have never before been developed. Single muon final state interaction channels represent some of the most difficult neutrino interactions to identify in LArTPC surface detectors, due to the ubiquitous cosmic-ray background that topologically, directionally and calorimetrically resembles these target final states. Despite this fact, the event selections outlined in this chapter have achieved very high purities and good efficiencies.


The ability to formulate these event selections is founded on the utility and accuracy supplied by the Pandora multi-algorithm approach to reconstructing neutrino events in liquid argon. Pandora has seen significant development in recent years, and represents the first fully-automated reconstruction framework for LArTPCs. This event selection therefore also further demonstrates the fact that high-quality output based on a fully-automated reconstruction within LArTPC surface detectors is a feasibility and a reality. The automated reconstruction procedure developed at MicroBooNE represents a foundation upon which other LArTPC detectors, such as DUNE and ProtoDUNE-SP are built.

The novel $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ event selections described in this chapter complement the other event selections developed by the MicroBooNE collaboration, such as those in [111], [112] and [113]. In addition, the identification of μ^-


and $\mu^- + p$ neutrino final states is crucial for several cross section measurements that are carried out by the MicroBooNE collaboration. The accuracy of such measurements inherently relies on the accurate separation of neutrino final states from the cosmic-ray background, of which the μ^- and $\mu^- + p$ final states are two of the most challenging. Potential future cross section measurements, such as measurements of CC ν_μ quasi-elastic-like interactions will furthermore rely on the accurate identification of $\mu^- + p$ final states and distinguishing them from $\mu^- + e^-$ Michel electron decay final states, as was done in this study.




(a) 3D view (with zoomed in view).



(b) U view.




(c) V view.




(d) W view.


Fig. 7.22 A short $N = 1$ track chosen by the event selection (7.22a). Cosmic ray hierarchies are indicated in red, and ν_{reco} is indicated in blue. Figures 7.22b, 7.22c and 7.22d show the projections of ν_{reco} in the U, V and W views.




(a) 3D view (with zoomed in view).



(b) U view.




(c) V view.




(d) W view.


Fig. 7.23 A long $N = 1$ track chosen by the event selection (7.23a). Cosmic ray hierarchies are indicated in red, and v_{reco} is indicated in blue. Figures 7.23b, 7.23c and 7.23d show the projections of v_{reco} in the U, V and W views.




(a) 3D view (with zoomed in view).



(b) U view.



(c) V view.



(d) W view.

Fig. 7.24 An $N=2$ interaction chosen by the event selection (7.24a). Cosmic ray hierarchies are indicated in red, and v_{reco} is indicated in blue. Figures 7.24b, 7.24c and 7.24d show the projections of v_{reco} in the U, V and W views.

Chapter 8

Conclusion

The MicroBooNE experiment was designed to resolve the low-energy excess of electron-like events at LSND and MiniBooNE. To do so, the MicroBooNE experiment uses a LArTPC detector, and is one of a series of present and future experiments to employ this technology. The construction and operation of the MicroBooNE experiment has been a major achievement both in terms of research and development (R&D) of LArTPC detectors, and furthering our understanding of neutrino physics. MicroBooNE has seen a very stable and successful data-taking period, exceeding the projected number of delivered protons-on-target with consistently high liquid argon purity levels. In addition, several unforeseen effects have been identified, studied and successfully incorporated into the Monte Carlo simulation, such as the characterisation and filtering of noise in the LArTPC [98] and space charge effects [99]. Future LArTPC experiments will benefit greatly from the studies performed in the context of the MicroBooNE experiment, and the experience gained through the operation of the MicroBooNE detector.

The MicroBooNE detector is part of the Short-Baseline Neutrino (SBN) programme at Fermilab. The purpose of SBN is to significantly develop the physics community's understanding of the possibility of the existence of sterile neutrinos, by utilising the well-characterised FNAL Booster Neutrino Beam (BNB). The SBN programme consists of three detectors: MicroBooNE, SBND and ICARUS-T600. The MicroBooNE experiment constitutes the first phase of the SBN programme, and aims to understand the nature of the MiniBooNE low-energy excess. The second phase aims to search for short baseline oscillations both in appearance and disappearance channels. To do this, the second phase uses a near detector, the Short-Baseline Near Detector (SBND) and a far detector, the ICARUS T600 detector. SBND is located just 110 meters from the Booster Neutrino Beam target, and

has 112 tons of liquid argon within the active volume of its detection systems. The ICARUS T600 detector, which comprises two cryostats holding liquid argon time projection chamber modules and photodetectors, will serve as the Short-Baseline Program Far Detector. It is the farthest from the BNB target, at a distance of 600 meters, and it is the largest of the detectors with 500 tons of liquid argon in the active volumes. Both detectors will contribute further to LArTPC R&D, and will measure neutrino-Ar cross sections at energies relevant to the DUNE experiment. A large mass far detector and a near detector of the same technology reduces both statistical and systematic uncertainties, and the SBN detectors enable 5σ coverage of the 99% C.L. allowed region of the LSND signal and global best fit values. The ICARUS T600 detector is expected to start taking data in 2019, and SBND is to be commissioned next year.

In the context of these future experiments, the development of a fully-automated reconstruction procedure for particle interaction in LArTPCs is indispensable. The Pandora software development kit (SDK) aims to provide a general framework in which pattern recognition algorithms can be developed. Pandora functionality includes over a hundred algorithms that have been developed for use in LArTPC detectors, including MicroBooNE and ProtoDUNE. The development of algorithms within Pandora is motivated by a set of well-defined performance metrics and correctness criteria. This thesis has outlined improvements that have been made to the Pandora vertex reconstruction by introducing two new types of vertex candidates. It was shown that the introduction of these new vertex candidates has a significant beneficial impact on the vertex offset (ΔR) distributions, and increases the fraction of correct events in charged-current quasielastic, resonant, deep inelastic and muon exchange current channels, which are the main signal channels in the development of the event selection procedure.

A distinct advantage of LArTPC detectors is the superior calorimetric resolution they provide, which allows analysers to exploit calorimetric information in reconstruction logic and analyses. To this end, a tool has been developed that constructs direction fits to cluster objects, and formulates the probability that the particle is propagating along an axis. This direction finding tool is based on the predicted dE/dx values from Bethe theory. Individual hits were observed to occasionally have spuriously high charge values, which was determined to be predominantly due to secondary ionisation from δ rays, and charge sharing between hits near the interaction vertex. A hit filtering procedure was developed that can remove spuriously-high hit charges with a large degree of accuracy, and the performance of the filter has been quantified. In addition, the direction fitting tool was applied to the splitting of merged tracks, introducing logic to identify potential split positions and to compare the

change in fit χ_{min}^2/N due to each split. The frequency of errors was quantified as a function of the number of hits in the track, and a sophisticated error reduction procedure was developed. Lastly, a measure of the forwards probability P_F of tracks in a contained charged-current quasi-elastic sample was developed, and it was shown that both the hit filtering and particle splitting logic have a beneficial effect on the probability curves for such events.

The MicroBooNE detector, as a surface detector, is subject to a large number of cosmic background events, at a rate of approximately 5 kHz. This results in an average of approximately 24 cosmic rays per DAQ readout window, and as such accurate cosmic ray rejection procedures are of the utmost importance for MicroBooNE and any other LArTPC surface detector. The MicroBooNE cosmic ray tagger (CRT) system was installed mid-2017, providing roughly 85% coverage of the LArTPC volume. In addition, flash information from the PMT system may be matched with readout windows to reject cosmic-ray tracks. Nonetheless, the sheer quantity of cosmic rays makes their removal a difficult problem. In order to provide analysers with information that can aid in cosmic ray removal that is distinct from the CRT and PMT systems, the direction fitting tool was used to construct a chi-squared measure $\Delta\chi_{DU}^2/N$ and a cosmic probability P_C . The probabilistic quantity was developed using a sample of simulated cosmic ray muons, and the correlations of P_C with a number of reconstruction quantities was taken into account. In addition, a fit to the P_C probability curves was developed and a parametrisation of this fit was developed which can predict the fit with a high degree of accuracy, at reduced computational cost. In addition, it was shown that the quantity $\Delta\chi_{DU}^2/N$ can be leveraged to improve the performance of the Pandora cosmic ray tagging algorithms. A study was performed on the cosmic ray muons that remained untagged by the existing Pandora cosmic ray tagging logic, and a number of preselection cuts were developed to create a very pure sample of cosmics. It was shown that the introduction of directional information allows for the selective removal of cosmic ray tracks that enter the detector near its top face, which has a significant beneficial impact on the reconstruction performance. The impact of the directional cosmic ray tagging (DCRT) procedure on the reconstruction performance was quantified by studying the impact on the Pandora performance metrics, and it was shown that cuts that incorporated directional information were most effective.

An event selection has been created to select charged-current ν_μ -induced interactions containing either a single muon or a single muon and a single proton in the final state. This event selection has been developed in a simulated dataset that corresponds to 5.37×10^{20} protons on target. The event selection was validated using Run 1 beam-on and beam-off data, gathered by the MicroBooNE detector between February and October 2016. The real data

samples have had data quality cuts applied to them, and the total on-beam analysis sample after data quality cuts corresponds to 4.52×10^{19} protons on target and 1.00×10^8 beam triggers. The off-beam analysis sample also originates from Run 1 data, when the beam was turned off, and corresponds to 7.19×10^7 EXT triggers. For this event selection, a number of preselection cuts were formulated, including a containment cut, a Pandora reconstructability cut and a cut on reconstructed primary particle multiplicity. The impact of the directional cosmic ray tagging cuts on the simulated analysis sample was quantified, showing an increase in sample purity. Furthermore, the contribution of each interaction channel to the number of signal and background events in the simulated sample was quantified in interaction type tables, keyed on reconstructed primary particle multiplicity. In addition, the impact of cheating the different stages of the Pandora reconstruction was quantified using interaction type tables. Both studies showed the large impact of the cosmic ray background on the sample purity and efficiency.

The event selection procedure itself was based on the formulation of a number of topological, calorimetric and directional quantities. It was shown how the fit points resulting from the direction fitting tool can be used to formulate more robust simple calorimetric quantities. Furthermore, the direction fitting tool was modified to work in 3D and space charge and calorimetric corrections were applied, such that the particle mass could be fixed to the proton mass. This modification resulted in a χ_{proton}^2 quantity, which is highly correlated with particle mass. A Boosted Decision Tree (BDT) was trained on the variables in the simulated sample, and it was shown that the trained model is more effective at resolving signal from background events when directional information is included in the training set. The resulting $\nu_\mu + \text{Ar} \rightarrow \mu^- + X$ and $\nu_\mu + \text{Ar} \rightarrow \mu^- + p + X$ event selections achieve purities of 64.35% and 78.30% and efficiencies of 57.91% and 62.65%, respectively. Lastly, the trained model was applied to a Run 1 BNB dataset, and the resulting variable distributions are compared to the equivalent stacked distributions resulting from the simulated sample. Reasonable agreement was seen in the resulting distributions, demonstrating that these event selections work on experimental data as well as in simulation.

References

- [1] LSND, D. A. Whitehouse, *LSND: A search for neutrino oscillations at LAMPF*, in *Proceedings, The Vancouver Meeting, Particles & Fields '91: Vancouver, Canada, August 18-22, 1991*, pp. 949–951, 1991.
- [2] MiniBooNE, A. A. Aguilar-Arevalo *et al.*, *Unexplained Excess of Electron-Like Events From a 1-GeV Neutrino Beam*, Phys. Rev. Lett. **102**, 101802 (2009), arXiv:0812.2243.
- [3] MiniBooNE Collaboration, Aguilar-Arevalo *et al.*, *Significant Excess of Electron-like Events in the MiniBooNE Short-Baseline Neutrino Experiment*, Phys. Rev. Lett. **121**, 221801 (2018).
- [4] M. Sorel, J. Conrad, and M. Shaevitz, *Combined analysis of short-baseline neutrino experiments in the $(3+1)$ and $(3+2)$ sterile neutrino oscillation hypotheses*, Physical review D: Particles and fields **70** (2003).
- [5] E. Wang, L. Alvarez-Ruso, and J. Nieves, *Single photon events from neutral current interactions at MiniBooNE*, Phys. Lett. **B740**, 16 (2015), arXiv:1407.6060.
- [6] J. L. Rosner, *Low-energy photon production in neutrino neutral-current interactions*, Physical Review D **91** (2015).
- [7] E. D. Church, *LArSoft: A Software Package for Liquid Argon Time Projection Drift Chambers*, (2013), arXiv:1311.6774.
- [8] E. L. Snider and G. Petrillo, *LArSoft: Toolkit for Simulation, Reconstruction and Analysis of Liquid Argon TPC Neutrino Detectors*, J. Phys. Conf. Ser. **898**, 042057 (2017).
- [9] K. C. Wang, *A Suggestion on the Detection of the Neutrino*, Phys. Rev. **61**, 97 (1942).
- [10] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, *Detection of the free neutrino: A Confirmation*, Science **124**, 103 (1956).
- [11] M. Perl and R. F., *The Nobel Prize in Physics 1995*, 2009, "The Nobel Prize in Physics 1995 was divided equally between Martin L. Perl for the discovery of the tau lepton and Frederick Reines for the detection of the neutrino.
- [12] M. Goldhaber, L. Grodzins, and A. W. Sunyar, *Helicity of Neutrinos*, Phys. Rev. **109**, 1015 (1958).

- [13] R. Davis, Jr. and D. S. Harmer, *Attempt to observe the $Cl^{37}(\bar{\nu}e^-)Ar^{37}$ reaction induced by reactor antineutrinos*, Bull. Am. Phys. Soc. **4**, 217 (1959).
- [14] G. Danby *et al.*, *Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos*, Phys. Rev. Lett. **9**, 36 (1962).
- [15] M. L. Perl *et al.*, *Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation*, Phys. Rev. Lett. **35**, 1489 (1975).
- [16] DONUT, K. Kodama *et al.*, *Observation of tau neutrino interactions*, Phys. Lett. **B504**, 218 (2001), arXiv:hep-ex/0012035.
- [17] LEP, ALEPH, DELPHI, L3, OPAL, LEP Electroweak Working Group, SLD Heavy Flavor Group, SLD Electroweak Group, D. Buskulic *et al.*, *A Combination of preliminary electroweak measurements and constraints on the standard model*, (2001), arXiv:hep-ex/0103048.
- [18] B. T. Cleveland *et al.*, *Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector*, The Astrophysical Journal **496**, 505 (1998).
- [19] B. T. Cleveland *et al.*, *Measurement of the solar electron neutrino flux with the Homestake chlorine detector*, Astrophys. J. **496**, 505 (1998).
- [20] SAGE, J. N. Abdurashitov *et al.*, *Results from SAGE*, Phys. Lett. **B328**, 234 (1994).
- [21] GALLEX, P. Anselmann *et al.*, *GALLEX solar neutrino observations: The Results from GALLEX-I and early results from GALLEX-II*, Phys. Lett. **B314**, 445 (1993).
- [22] Super-Kamiokande, Y. Fukuda *et al.*, *The Super-Kamiokande detector*, Nucl. Instrum. Meth. **A501**, 418 (2003).
- [23] Super-Kamiokande, T. Kajita, E. Kearns, and M. Shiozawa, *Establishing atmospheric neutrino oscillations with Super-Kamiokande*, Nucl. Phys. **B908**, 14 (2016).
- [24] SNO, J. Boger *et al.*, *The Sudbury neutrino observatory*, Nucl. Instrum. Meth. **A449**, 172 (2000), arXiv:nucl-ex/9910016.
- [25] SNO, J. R. Klein, *Solar neutrino results from the Sudbury Neutrino Observatory*, Int. J. Mod. Phys. **A17**, 3378 (2002), arXiv:hep-ex/0111040, [,470(2001)].
- [26] Z. Maki, M. Nakagawa, and S. Sakata, *Remarks on the unified model of elementary particles*, Prog. Theor. Phys. **28**, 870 (1962), [,34(1962)].
- [27] B. Pontecorvo, *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*, Sov. Phys. JETP **26**, 984 (1968), [Zh. Eksp. Teor. Fiz.53,1717(1967)].
- [28] V. N. Gribov and B. Pontecorvo, *Neutrino astronomy and lepton charge*, Phys. Lett. **28B**, 493 (1969).
- [29] B. Kayser, *On the Quantum Mechanics of Neutrino Oscillation*, Phys. Rev. **D24**, 110 (1981).

- [30] M. Thomson, *Modern Particle Physics*, Cambridge University Press (2013), ISBN: 9781107034266.
- [31] Daya Bay, Z. Wang, *Daya Bay Neutrino Experiment: Goal, Progress and Schedule*, (2011), arXiv:1109.3253.
- [32] Daya Bay, F. P. An *et al.*, *Observation of electron-antineutrino disappearance at Daya Bay*, Phys. Rev. Lett. **108**, 171803 (2012), arXiv:1203.1669.
- [33] RENO, J. K. Ahn *et al.*, *RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang*, (2010), arXiv:1003.1391.
- [34] RENO, J. K. Ahn *et al.*, *Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment*, Phys. Rev. Lett. **108**, 191802 (2012), arXiv:1204.0626.
- [35] Double Chooz, Y. Abe *et al.*, *Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment*, Phys. Rev. Lett. **108**, 131801 (2012), arXiv:1112.6353.
- [36] KamLAND, D. M. Markoff, *An overview of the KamLAND experiment*, J. Phys. **G29**, 1481 (2003).
- [37] KamLAND, S. Abe *et al.*, *Precision Measurement of Neutrino Oscillation Parameters with KamLAND*, Phys. Rev. Lett. **100**, 221803 (2008), arXiv:0801.4589.
- [38] MINOS, I. Ambats *et al.*, *The MINOS Detectors Technical Design Report*, (1998).
- [39] MINOS, P. Adamson *et al.*, *Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam*, Phys. Rev. Lett. **101**, 131802 (2008), arXiv:0806.2237.
- [40] KATRIN, M. Beck, *The KATRIN Experiment*, J. Phys. Conf. Ser. **203**, 012097 (2010), arXiv:0910.4862.
- [41] J. Schechter and J. W. F. Valle, *Neutrino Masses in $SU(2) \times U(1)$ Theories*, Phys. Rev. **D22**, 2227 (1980).
- [42] E. Ma, *Verifiable radiative seesaw mechanism of neutrino mass and dark matter*, Phys. Rev. **D73**, 077301 (2006), arXiv:hep-ph/0601225.
- [43] F. T. Avignone, III, S. R. Elliott, and J. Engel, *Double Beta Decay, Majorana Neutrinos, and Neutrino Mass*, Rev. Mod. Phys. **80**, 481 (2008), arXiv:0708.1033.
- [44] J. Schechter and J. W. F. Valle, *Neutrinoless Double beta Decay in $SU(2) \times U(1)$ Theories*, Phys. Rev. **D25**, 2951 (1982), [,289(1981)].
- [45] CUORICINO, S. Capelli, *CUORICINO latest results and background analysis*, Nucl. Phys. Proc. Suppl. **221**, 334 (2011).
- [46] D. Casper, *The nuance neutrino physics simulation, and the future*, Nuclear Physics B - Proceedings Supplements **112**, 161 (2002).
- [47] J. A. Formaggio and G. P. Zeller, *From eV to EeV: Neutrino cross sections across energy scales*, Rev. Mod. Phys. **84**, 1307 (2012).

- [48] C. H. Llewellyn Smith, *Neutrino Reactions at Accelerator Energies*, Phys. Rept. **3**, 261 (1972).
- [49] O. Benhar, N. Farina, H. Nakamura, M. Sakuda, and R. Seki, *Electron- and neutrino-nucleus scattering in the impulse approximation regime*, Phys. Rev. **D72**, 053005 (2005), arXiv:hep-ph/0506116.
- [50] R. A. Smith and E. J. Moniz, *Neutrino Reactions on Nuclear Targets*, Nucl. Phys. **B43**, 605 (1972), [Erratum: Nucl. Phys. **B101**, 547 (1975)].
- [51] R. Bradford, A. Bodek, H. Budd, and J. Arrington, *A New Parameterization of the Nucleon Elastic Form Factors*, Nuclear Physics B - Proceedings Supplements **159**, 127 (2006), Proceedings of the 4th International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region.
- [52] S. Moraghe, J. Amaro, C. Garcia-Recio, and A. Lallena, *Meson-exchange current effects in inelastic electron scattering from polarized nuclei*, Nuclear Physics A **576**, 553 (1994).
- [53] A. Bodek, H. S. Budd, and M. E. Christy, *Neutrino Quasielastic Scattering on Nuclear Targets: Parametrizing Transverse Enhancement (Meson Exchange Currents)*, Eur. Phys. J. **C71**, 1726 (2011), arXiv:1106.0340.
- [54] T. Katori, *Meson Exchange Current (MEC) Models in Neutrino Interaction Generators*, AIP Conf. Proc. **1663**, 030001 (2015), arXiv:1304.6014.
- [55] M. Martini, M. Ericson, and G. Chanfray, *Neutrino quasielastic interaction and nuclear dynamics*, Phys. Rev. C **84**, 055502 (2011).
- [56] D. Rein and L. M. Sehgal, *Coherent π^0 production in neutrino reactions*, Nuclear Physics B **223**, 29 (1983).
- [57] T. Golan, *Modeling nuclear effects in NuWro Monte Carlo neutrino event generator*, PhD Thesis (2014).
- [58] LSND, A. Aguilar-Arevalo *et al.*, *Evidence for neutrino oscillations from the observation of anti-neutrino(electron) appearance in a anti-neutrino(muon) beam*, Phys. Rev. **D64**, 112007 (2001), arXiv:hep-ex/0104049.
- [59] MiniBooNE, A. A. Aguilar-Arevalo, *The MiniBooNE Experiment*, eConf **C0406271**, MONP01 (2004), arXiv:hep-ex/0408074.
- [60] KARMEN, B. Zeitnitz, *KARMEN: Neutrino physics at ISIS*, Prog. Part. Nucl. Phys. **32**, 351 (1994).
- [61] OPERA, G. Giacomelli and M. Giorgini, *The OPERA experiment*, in *Proceedings, Vulcano Workshop 2006 : Frontier Objects in Astrophysics and Particle Physics: Vulcano, Italy, May 22-27, 2006*, 2006, physics/0609045.
- [62] MicroBooNE, R. Acciarri *et al.*, *Design and Construction of the MicroBooNE Detector*, JINST **12**, P02017 (2017), arXiv:1612.05824.

- [63] ICARUS, F. Varanini, *The ICARUS experiment*, PoS **EPS-HEP2017**, 142 (2017).
- [64] P. Adamson *et al.*, *The NuMI Neutrino Beam*, Nucl. Instrum. Meth. **A806**, 279 (2016), arXiv:1507.06690.
- [65] The ArgoNeuT Collaboration, R. Acciarri *et al.*, *Measurements of inclusive muon neutrino and antineutrino charged current differential cross sections on argon in the NuMI antineutrino beam*, Phys. Rev. D **89**, 112003 (2014).
- [66] T. D. Collaboration, *The DUNE Far Detector Interim Design Report Volume 1: Physics, Technology and Strategies*, (2018), arXiv:1807.10334.
- [67] T. D. Collaboration, *The DUNE Far Detector Interim Design Report Volume 2: Single-Phase Module*, (2018), arXiv:1807.10327.
- [68] T. D. Collaboration, *The DUNE Far Detector Interim Design Report Volume 3: Dual-Phase Module*, (2018), arXiv:1807.10340.
- [69] The DUNE Collaboration, J. Dawson, *The ProtoDUNE large demonstrator of the Liquid Argon double phase TPC program at CERN*, Journal of Physics: Conference Series **888**, 012158 (2017).
- [70] I.-W. MicroBooNE, LAr1-ND, *A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam*, (2015), arXiv:1503.01520.
- [71] V. Shiltsev, *Fermilab Proton Accelerator Complex Status and Improvement Plans*, Mod. Phys. Lett. **A32**, 1730012 (2017), arXiv:1705.03075.
- [72] MiniBooNE, A. A. Aguilar-Arevalo *et al.*, *The Neutrino Flux prediction at Mini-BooNE*, Phys. Rev. **D79**, 072002 (2009), arXiv:0806.1449.
- [73] MicroBooNE, C. Adams *et al.*, *Booster Neutrino Flux Prediction at MicroBooNE*, MicroBooNE Public Note MICROBOONE-NOTE-1031-PUB (2018).
- [74] MicroBooNE, V. Meddage, *Electron Attenuation Measurement using Cosmic Ray Muons at the MicroBooNE LArTPC*, in *Proceedings, Meeting of the APS Division of Particles and Fields (DPF 2017): Fermilab, Batavia, Illinois, USA, July 31 - August 4, 2017*, 2017, 1710.00396.
- [75] B. Baller, *Liquid argon TPC signal formation, signal processing and reconstruction techniques*, JINST **12**, P07010 (2017), arXiv:1703.04024.
- [76] F. Cavanna, A. Ereditato, and B. T. Fleming, *Advances in liquid argon detectors*, Nucl. Instrum. Meth. **A907**, 1 (2018).
- [77] C. Andreopoulos *et al.*, *The GENIE Neutrino Monte Carlo generator*, Nucl. Instrum. Methods Phys. Res., Sect. A **614**, 87 (2010).
- [78] R. Brun *et al.*, *GEANT Detector Description and Simulation Tool*, (1994).
- [79] MicroBooNE, C. Adams *et al.*, *Design and construction of the MicroBooNE Cosmic Ray Tagger system*, JINST **14**, P04004 (2019), arXiv:1901.02862.

- [80] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw, *CORSIKA: A Monte Carlo code to simulate extensive air showers*, (1998).
- [81] M. Del Tutto, *First Measurements of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon at 0.8 GeV Average Neutrino Energy with the MicroBooNE Detector*, PhD thesis, Oxford U., 2019.
- [82] MicroBooNE, A. A. Aguilar-Arevalo *et al.*, *Study of Space Charge Effects in MicroBooNE*, MicroBooNE Public Note MICROBOONE-NOTE-1018-PUB (2016).
- [83] L. Linssen, A. Miyamoto, M. Stanitzki, and H. Weerts, *Physics and Detectors at CLIC: CLIC Conceptual Design Report*, (2012), arXiv:1202.5940.
- [84] J. S. Marshall and M. A. Thomson, *The Pandora software development kit for pattern recognition*, The European Physical Journal C **75**, 439 (2015).
- [85] MicroBooNE, C. Adams *et al.*, *The Pandora Multi-Algorithm Approach to Automated Pattern Recognition of Cosmic Ray Muon and Neutrino Events in the MicroBooNE Detector*, The European Physical Journal C **78**, 82 (2018).
- [86] H. Bethe and J. Ashkin, *Passage of Radiations through Matter*, Experimental Nuclear Physics , 253 (1953).
- [87] N. Mokhov *et al.*, *The Density Effect for the Ionization Loss in Various Materials*, Atom.Data Nucl.Data Tabl. **78**, 183 (2001).
- [88] Particle Data Group, C. Patrignani *et al.*, *Passage of particles through matter*, Chin. Phys. C **40** (2016 and 2017 update).
- [89] R. M. Sternheimer, *The Density Effect for the Ionization Loss in Various Materials*, Phys. Rev. **88**, 851 (1952).
- [90] MicroBooNE, C. Adams *et al.*, *Ionization electron signal processing in single phase LArTPCs. Part II. Data/simulation comparison and performance in MicroBooNE*, JINST **13**, P07007 (2018), arXiv:1804.02583.
- [91] MicroBooNE, C. Adams *et al.*, *Calibration of the charge and energy response of the MicroBooNE liquid argon time projection chamber using muons and protons*, (2019), arXiv:1907.11736.
- [92] M. R. F. James, *Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations*, Comput.Phys.Commun. **10**, 343 (1975).
- [93] Particle Data Group, M. Tanabashi *et al.*, *Review of Particle Physics*, Phys. Rev. D **98**, 030001 (2018).
- [94] R. Acciarri *et al.*, *Measurement of cosmic-ray reconstruction efficiencies in the MicroBooNE LArTPC using a small external cosmic-ray counter*, Journal of Instrumentation **12**, P12030 (2017).
- [95] MicroBooNE, C. Adams *et al.*, *Rejecting cosmic background for exclusive neutrino interaction studies with Liquid Argon TPCs; a case study with the MicroBooNE detector*, (2018), arXiv:1812.05679.

- [96] MicroBooNE, C. Adams *et al.*, *Proton Track Identification in MicroBooNE Simulation for Neutral Current Elastic Events*, MicroBooNE Public Note MICROBOONE-NOTE-1025-PUB (2017).
- [97] MicroBooNE, C. Adams *et al.*, *Detector calibration using through going and stopping muons in the MicroBooNE LArTPC*, MicroBooNE Technical Note (2018).
- [98] MicroBooNE, R. Acciarri *et al.*, *Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC*, Journal of Instrumentation **12**, P08003 (2017).
- [99] M. Mooney, *The MicroBooNE Experiment and the Impact of Space Charge Effects*, (2015), arXiv:1511.01563.
- [100] MicroBooNE, C. Adams *et al.*, *Measurement of the Electronegative Contaminants and Drift Electron Lifetime in the MicroBooNE Experiment*, MicroBooNE Public Note MICROBOONE-NOTE-1003-PUB (2016).
- [101] ArgoNeuT, R. Acciarri *et al.*, *A Study of Electron Recombination Using Highly Ionizing Particles in the ArgoNeuT Liquid Argon TPC*, JINST **8**, P08005 (2013), arXiv:1306.1712.
- [102] J. Thomas and D. A. Imel, *Recombination of electron-ion pairs in liquid argon and liquid xenon*, Phys. Rev. **A36**, 614 (1987).
- [103] MicroBooNE, C. Adams *et al.*, *Selection of ν_μ charged-current induced interactions with $N > 0$ protons and performance of events with $N = 2$ protons in the final state in the MicroBooNE detector from the BNB*, (2018).
- [104] A. Hocker *et al.*, *TMVA - Toolkit for Multivariate Data Analysis*, (2007), arXiv:physics/0703039.
- [105] R. Acciarri *et al.*, *Convolutional neural networks applied to neutrino events in a liquid argon time projection chamber*, Journal of Instrumentation **12**, P03011 (2017).
- [106] MicroBooNE, C. Adams *et al.*, *First Muon-Neutrino Charged-Current Inclusive Differential Cross Section Measurement for MicroBooNE Run 1 Data*, (2018), arXiv:1811.02700.
- [107] MicroBooNE, C. Adams *et al.*, *Technical Note: MCC9 Overlay GENIE simulated BNB and Cosmic data Version 1.0*, MicroBooNE Internal Technical Note MicroBooNE-doc-20737-v2 (2019).
- [108] MicroBooNE, C. Adams *et al.*, *Comparison of ν_μ -Ar multiplicity distributions observed by MicroBooNE to GENIE model predictions*, Eur. Phys. J. **C79**, 248 (2019), arXiv:1805.06887.
- [109] C. Adams *et al.*, *Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation*, Journal of Instrumentation **13**, P07006 (2018).
- [110] C. Adams *et al.*, *Ionization electron signal processing in single phase LArTPCs. Part II. Data/simulation comparison and performance in MicroBooNE*, Journal of Instrumentation **13**, P07007 (2018).

- [111] MicroBooNE, C. Adams *et al.*, *Selection of ν_μ Events for the MicroBooNE Deep Learning Low Energy Excess Analysis*, MicroBooNE Public Note MICROBOONE-NOTE-1051-PUB (2018).
- [112] MicroBooNE, C. Adams *et al.*, *Electron-neutrino selection and reconstruction in the MicroBooNE LArTPC using the Pandora multi-algorithm pattern recognition*, MicroBooNE Public Note MICROBOONE-NOTE-1038-PUB (2018).
- [113] MicroBooNE, C. Adams *et al.*, *Selection of ν_μ charged-current induced interactions with $N > 0$ protons and performance of events with $N = 2$ protons in the final state in the MicroBooNE detector from the BNB*, MicroBooNE Public Note MICROBOONE-NOTE-1056-PUB (2018).