

**Study of Charged Current Neutral Pion
Production in the T2K Near Detector**

Matthew David Lawe

**Submitted for the Degree of Doctor of
Philosophy in Physics**

Department of Physics and Astronomy

September 2014

Supervisor: Prof. Lee Thompson

The University of Sheffield

ABSTRACT

T2K is a long-baseline neutrino oscillation experiment based at the J-PARC research facility in eastern Japan. The main systematic uncertainties dominating the neutrino oscillation analyses for T2K are related to the beam flux, cross-section and final state interaction uncertainties. Within this thesis, a first attempt is made to develop an analysis to reconstruct charged-current single neutral-pion interactions within the tracking detector region of the T2K near detector. The analysis presented is divided into six topologies and produces a selection efficiency, over all the event topologies of interest, of 9.6% with 12.1% purity. When considering the most effective selection topology, labelled FGD2DsDs, an efficiency of 9.4% with purity of 22.6% is achieved. This leads to a FGD2DsDs Monte-Carlo prediction of 58.5 ± 3.5 (stat.) ± 15.4 (syst.) events being selected across the T2K Run 2, Run 3 and Run 4 data taking periods, with a total of 58 events being selected within the real data.

With further development in the future, it is expected that the presented analysis will contribute towards the data samples used to constrain the T2K experiment oscillation analysis and lead to a $CC1\pi^0$ final state cross-section measurement.

Acknowledgements

Firstly, I'd like to gratefully acknowledge the support and guidance provided over the last four years by my supervisor Prof. Lee Thompson. His assistance has been crucial in directing my work, and the detailed feedback provided during the preparation of this thesis has been invaluable.

I'd also like to express thanks to Dr. Jon Perkin for his ceaseless aid when dealing with the idiosyncrasies of both the ND280 and ROOT software packages, along with the countless helpful discussions to which he has contributed during the development of this analysis. Likewise, the advice and insight offered by Dr. Susan Cartwright and the students here at Sheffield, Callum Wilkinson, Leon Pickard and Mike Wallbank, has been enormously beneficial; I could have not asked for a better group of people to work with during my time at Sheffield University.

Within the wider T2K collaboration I must give special thanks to Prof. Clark McGrew, Dr. Helen O'Keeffe, Dr. Sam Short and the rest of the ND280 π^0 group past and present for all the assistance they have provided. Thanks must also go to Dr. Anselmo Cervera, Dr. Ben Smith and all others involved in the development of the highland framework for their tireless effort.

As part of the PhD study it was my great pleasure to live and work in Tōkai-mura, Japan. I would like to thank the residents of Tōkai for making me feel so welcome. Particular thanks goes to the Japanese language teachers for their invaluable instruction, along with the staff at J-PARC, Masago International Lodging, Avúe Climbing Gym and Pono Pono, all of whom made living in Japan so much easier and enjoyable. Life in Japan would not have been complete without the kind friendships of Dr. Sarah Foley, Dr. Megan Friend, Dr. Ben Smith, Dr. Jevon Keane-Brennan, Andrew Furmanski, Linda Cremonesi and great many others who it was an honour to share that time with.

Finally, I must thank my family, friends, and most importantly my partner Sam for their unwavering support over the years.

Contents

1	Introduction to Neutrino Physics	1
1.1	Historic Background	1
1.2	Neutrino Oscillations	5
1.2.1	Oscillation Theory	8
1.3	Current Status	10
1.4	Neutrino Interaction Cross-Sections	13
1.4.1	Interaction Processes	13
1.4.2	CC1 π^0 Cross-Section Measurements	16
2	The T2K Experiment	20
2.1	The T2K Neutrino Beam	20
2.1.1	J-PARC	20
2.1.2	Accelerator Complex	21
2.1.3	Neutrino Beamline	22
2.1.4	MUMON	26
2.1.5	Off-Axis Design	27
2.1.6	Neutrino Flux Prediction	29
2.2	Near Detector Complex	30
2.2.1	INGRID	30

2.2.2	ND280	34
2.3	Far Detector	46
2.3.1	Čerenkov Radiation for Particle Identification	47
3	The ND280 Software Suite and Calibration Techniques	50
3.1	Data and Monte Carlo Processing Chain	50
3.1.1	ND280 Data Processing	50
3.1.2	ND280 Monte Carlo Processing	52
3.2	Calibration	53
3.2.1	ECal Energy Calibrations	53
3.2.2	ECal Timing Calibrations	56
3.3	Event Reconstruction	58
3.3.1	Tracker Reconstruction	58
3.3.2	Tracker ECal Reconstruction	60
3.3.3	Global Reconstruction	70
3.4	High Level Analysis Package	72
3.4.1	Core Packages	72
3.4.2	Standard Analyses	76
4	ECal Timing Calibration	77
4.1	Calibration Methodology	77
4.1.1	Inter-TFB Timing Calibration	78
4.1.2	Inter-RMM Timing Calibration	79
4.1.3	Timing Calibration Methodology Discussion	79
4.2	Initial Timing Calibration Result	81
4.2.1	TFB Calibration Results	82
4.2.2	RMM Calibration Results	82

4.2.3	Track Resolution	85
4.3	Addressing RMM Timing Instability	85
4.3.1	Cause Of RMM Timing Instability	85
4.3.2	Considerations For RMM Time Slip Detection	87
4.3.3	Detection of RMM Time Slips	87
4.3.4	Validation of RMM Time Slips	96
4.4	Application Of RMM Time Slip Correction	98
5	Selecting Charged-Current Single π^0 Interactions in the FGDs	101
5.1	Analysis Motivation	101
5.2	Analysis Samples	102
5.2.1	Data Samples	102
5.2.2	Monte Carlo Samples	105
5.3	Signal Definition	105
5.4	Muon Selection	110
5.4.1	Muon Selection Criteria	110
5.4.2	Muon Selection Parameter Space	112
5.4.3	Muon Selection Results	118
5.5	π^0 Selection	126
5.5.1	π^0 Selection Criteria	126
5.5.2	Multivariate Analysis	129
5.5.3	Boosted Decision Tree Analysis Implementation	137
5.5.4	Boosted Decision Tree Training Performance	147
5.5.5	Data and Monte-Carlo Comparisons for the Boosted Decision Tree	155
5.5.6	π^0 Selection Performance	159

5.5.7	π^0 Selection Results	159
5.6	Systematic Uncertainty	178
5.6.1	Multivariate Analysis Systematic Error	178
5.6.2	Neutrino Flux Prediction Systematic Error	186
5.6.3	Cross-Section Systematic Error	192
5.6.4	Detector Systematic Error	197
5.7	Summary	199
6	Discussion and Conclusions	202
	Bibliography	207
A	Discriminating Variable Distributions and Correlations from BDT Training Samples	219
B	Data and Monte-Carlo Comparisons for BDT Discriminating Variables	240
C	Data and Monte-Carlo BDT Discriminating Variable Comparisons for Each Event Topology	256
D	π^0 Invariant Mass Calculation	272
E	Correlation and Covariance Matrices for Flux Systematic Error Calculation	275

List of Tables

1.1	Current best fit neutrino oscillation parameters under the assumption of the normal and (inverted) mass hierarchy [38].	11
5.1	Experimental operation periods being used for this analysis showing PØD water status and the good POT collected. Good POT requirements are described in Section 5.2.1.	103
5.2	MC POT used to compare with each data taking period, along with the multiple of good data POT that it is equates to, the PØD water status and the simulated beam power.	105
5.3	Codes used to describe the six event topologies of interested in this analysis based upon the interaction FGD and where the π^0 decay photons shower.	106
5.4	Final state breakdown for FGD fiducial volume charged-current neutrino interactions.	109
5.5	Further breakdown of CC1 π^0 final states by interaction vertex and π^0 decay topology.	109
5.6	MC prediction for the number of events selected as a function of cut, and associated CC inclusive selection efficiency and purity for all running periods to be analysed.	115
5.7	Breakdown of true particle type for the track selected as the muon in the CC inclusive-like selection.	118

5.8	The number of selected events in data and (MC) for each data taking period as a function of analysis cut.	123
5.9	The total number of selected events in data and (MC) for each data taking period as a function of analysis cut. The data over MC ratio is along with its statistical error.	126
5.10	Breakdown in the MC samples used to select signal and background events for the BDT training.	140
5.11	The number of signal and background pairwise ECal combinations that were used to train the BDT discriminators for each signal topology.	141
5.12	Variable importance ranking from the BDT training for each of the event topologies with an FGD1 vertex. The number in brackets next to each variable represents its MVA method independent ranking.	149
5.13	Variable importance ranking from the BDT training for each of the event topologies with an FGD2 vertex. The number in brackets next to each variable represents its MVA method independent ranking.	150
5.14	Optimal BDT cut values as suggested by the MVA training and the newly optimised cut vales that are applied to the final analysis.	155
5.15	MC prediction for the number of events selected as a function of cut, $CC1\pi^0$ selection efficiency and purity, and π^0 decay photon purity across all running periods to be analysed.	160
5.16	Final state contributions to final event selection.	161
5.17	The number of selected events in data and (MC) for each data taking period as a function of analysis cut. The total across all running periods is also shown along with the data over MC ratio with its statistical error.	171

5.18	Data over MC ratio, with statistical error, for the total number of higher and lower energy isolated ECal clusters reconstructed in each ECal module after the Reconstructed Mass cut has been applied, broken down by event vertex location.	173
5.19	Data over MC ratio, with statistical error, for the total number of events that could contribute to each topology after the Reconstructed Mass cut has been applied.	174
5.20	Nominal MC event selection and associated multivariate analysis implementation systematic error with and (without) the variable importance reweighting.	183
5.21	Nominal MC event selection and associated flux systematic error.	192
5.22	NIWG cross-section parametrisation and associated errors with applicable neutrino energy ranges for the NEUT MC generator. . . .	193
5.23	MC prediction for the number of events selected for each topology of interest, CC1 π^0 selection efficiency and purity for the specific event topology of interest, the selected purity from all CC1 π^0 final states, and π^0 decay photon purity, across all running periods. . .	200
5.24	p -values calculated from the mass peak distributions for each event topology using the statistical uncertainties on each histogram bin as error input.	201

List of Figures

1.1	Variation in hadron production cross-section around Z^0 the mass as a function of the number of active neutrino species [18].	4
1.2	Solar neutrino energy spectrum with 1σ uncertainty values, as predicted by the Standard Solar Model. Solid line spectra represent those neutrino emissions from the <i>proton-proton</i> chain of fusion reactions, dashed line spectra are as a result of the <i>CNO cycle</i> chain of fusion reactions [24].	6
1.3	Data points show the ratio of fully contained data to Monte Carlo predicted (in the absence of neutrino oscillations) events as a function of distance over energy for the Super-Kamiokande atmospheric neutrino sample. Dashed lines show the predicted shape with the inclusion of $\nu_\mu \leftrightarrow \nu_\tau$ oscillations with $\Delta m^2 = 2.2 \times 10^{-3}$ eV and $\sin^2 2\theta = 1$ [29].	8
1.4	Mass splitting, Δm^2 , as a function of oscillation angle, $\tan^2 \theta$, for four types of neutrino oscillation. Lines show exclusion limits and coloured regions show allowed parameter space for a range of named experiments [39].	12
1.5	The two possible neutrino mass hierarchies, normal on the left and inverted on the right, given that m_2 is known to be greater than m_1 from observations of solar neutrinos. The colours indicate the weak eigenstate contributions to each of the mass eigenstates. The absolute neutrino mass scale is also unknown but constrained by cosmology and astrophysics, and ${}^3\text{H}$ β -decay experiments [42]. . .	13

1.6	Illustrative Feynman diagrams for charged-current quasi-elastic (CCQE), resonant (CCRES) and deep inelastic (CCDIS) neutrino-nucleon interaction processes.	15
1.7	Charged-current cross-sections per nucleon for neutrino, 1.7a, and antineutrino, 1.7b, interactions as a function of neutrino energy. Cross-section measurements are shown for quasi-elastic interactions along with the total across all interactions for a broad selection of experiments. The NUANCE [44] Monte-Carlo generator cross-section prediction is shown for CCQE, CCRES and CCDIS interactions, along with their combined total [43].	17
1.8	Historic $CC1\pi^0$ cross-section measurements compared to the NUANCE [44] Monte-Carlo generator cross-section prediction [43]. . .	18
1.9	MiniBooNE $CC1\pi^0$, flux averaged, differential cross-section measurement on a mineral oil target compared to the NUANCE [44] Monte-Carlo generator cross-section prediction [45].	19
2.1	Schematic view of the T2K experiment. Neutrinos are produced at J-PARC on the east coast of Japan and recorded at the Near Detector and Super-Kamiokande far detector [51].	21
2.2	J-PARC facility with the three accelerators labelled. The RCS is labelled as 3 GeV PS and the MR is labelled as 50 GeV PS. The near detector complex is located a point FD [52].	22
2.3	Schematic diagram of the neutrino beamline indicating the sub components of the primary and secondary beamlines [51].	23
2.4	Schematic of the primary beamline indicating the location of the beamline monitors [51].	24
2.5	Schematic of the secondary beamline indicating its sub components [51].	25

2.6	Neutrino flux spectrum as a function of energy at different off-axis angles, along with the muon-neutrino survival probability at 295 km. At a 2.5° off-axis angle the neutrino energy spectrum peaks at 0.6 GeV, corresponding to the first oscillation maximum [55].	28
2.7	FLUKA2008.3b neutrino flux prediction for the T2K ND280 detector for ν_μ , 2.7a, and ν_e , 2.7b, both re-weighted to the NA61 thin target data. The contribution from pion, kaon and muon decays to the total neutrino flux are each shown [57].	30
2.8	ND280 detector complex. The INGRID detector is formed of the vertical stack of modules rising from the lowest floor in the foreground, spanning the lower and central floors, along with the horizontal span of modules on the central floor behind the vertical stack. The ND280 detector is located on the top floor, shown in the magnet open configuration. The magnetic coils and flux returns are opened to each side of the pit revealing the central basket portion of the detector [51].	31
2.9	Position of the INGRID detector modules. Seven modules are used in both the vertical and horizontal branches of the detector cross, these monitor the beam position. A further two modules are placed symmetrically off the cross axis to measure the axial symmetry of the beam [51].	32
2.10	Exploded view of an INGRID module. 2.10a shows the iron planes in blue, separated by the grey scintillator planes. 2.10b shows a full module, with the surrounding veto planes shown in black [51].	33
2.11	Exploded view of the ND280 detector [51].	34
2.12	Image of an SMRD paddle prior to installation with the characteristic snaking WLS fibre running through it [65].	36
2.13	Cross sectional view of the PØD showing the position of the upstream and downstream calorimeters and intervening water target region along with the position of the triangular shaped bars used through the sub-detector [66].	38

2.14	Simplified cut-away diagram of a TPC showing the detection volume, cathode components, read out components, neutrino beam direction, and electric and magnetic field direction [68].	40
2.15	Overview of electronics components and their connections for the ND280 [51].	44
2.16	Photograph of an MPPC and its active pixel face. The sensitive face is 1.3 by 1.3 mm ² , and the 26 by 26 grid of individual pixels and electrode attachments can be clearly seen [72].	45
2.17	Sketch of the SK detector and its position under the Ikenoyama mountain [74].	46
2.18	Example SK event displays from a muon-neutrino interaction, 2.18a, and electron-neutrino interaction, 2.18b, showing the characteristic sharp and “fuzzy” Čerenkov radiation rings of these types of interactions [51].	48
3.1	The ND280 offline software suite structure the most important or representative packages are shown [51].	51
3.2	ADC response from the high and low gain read out channels as a function of injected charge [69].	54
3.3	Hit charge distributions for cosmic ray muons for the six barrel ECal modules at different stages in the calibration chain. 3.3a shows the hit charge distribution after the application of the pedestal subtraction, linearity correction and MPPC gain calibration. 3.3b also includes the application of the MPPC response calibration [69].	55
3.4	Lines showing the MC predictions for $\frac{dE}{dx}$ as a function of momentum for muons, electrons, protons and pions. Those are superimposed upon the data for negatively charged, 3.4a, and positively charged, 3.4b, particles from the first T2K experimental data taking period [68].	61

-
- 3.5 Pull distributions obtained for a large MC sample of true muons. 3.5a shows the pull distributions under the muon, electron, proton, pion and kaon particle hypotheses. 3.5b shows the difference in muon and electron pull as a function of measured particle momentum in the TPCs. 62
- 3.6 True energy spectra for higher (blue) and lower (red) energy decay photons from π^0 's, with the absolute (solid line) and reconstructed (dashed lines) spectra indicated. The reconstruction efficiency as a function of true photon energy is also shown. These and the energy spectra are both presented as a function of module orientation. . . 66
- 3.7 True angular distribution for higher (blue) and lower (red) energy decay photons from π^0 's, with the absolute (solid line) and reconstructed (dashed lines) distributions indicated. The reconstruction efficiency as a function of true photon incidence angle is also shown, and these and the angular distributions are both broken down by module orientation. 68
- 3.8 Fractional energy difference calculated between the reconstructed and true energy of π^0 decay photons as a function of true energy for three different ECal module orientations. The right hand plots are the same as the left hand plots, but with a logarithmic z axis. 69
- 3.9 Fractional angular difference calculated between the reconstructed thrust axis and true incidence angle of π^0 decay photons into the ECal modules as a function of true angle for three different ECal module orientations. The right hand plots are the same as the left hand plots, but with a logarithmic z axis. 71

3.10	Beam spill event number 54642 from ND280 run 6462, subrun 11, shown looking at the ND280 from the side; the beam comes from the left. 3.10a shows the calibrated hits within the event; two tracks spanning multiple sub-detectors are prominently seen along with some noise hits. 3.10b displays the globally reconstructed tracks overlaying the calibrated hits. The track starting towards the upper centre of the PØD is probably a muon produced during a CCQE interaction. The track starting from the upstream face of the PØD is probably a muon produced by a neutrino interacting in the earth upstream of the ND280 detector pit. Both tracks have been globally reconstructed passing through the PØD, TPCs, FGDs and DsECal.	73
4.1	4.1a shows weighted combinations between multiple TFBs; 4.1b shows the minimal spanning tree that would be extracted as a result [87].	80
4.2	TFB calibration constants as a function of ND280 run number for each RMM in the tracker region.	83
4.3	TFB residuals as a function of ND280 run number for each RMM in the tracker region.	84
4.4	RMM calibration constants as a function of ND280 run number.	85
4.5	Mean hit time for cosmic ray muons across each TFB, prior to (blue) and after (red) the timing calibration has been applied. The data shown is for ND280 run 8115.	86
4.6	Time relative to trigger distribution for ECal RMM 5 during part of the third T2K experimental running period. By fitting a Gaussian function to the data sample it is easy to reject the noise, most easily seen in 4.6b, from the time slip analysis samples. One ND280 Clock Tick equates to 2.5 ns.	88

4.7	Mean time relative to trigger for five minute data taking periods as a function of Unix time. Data for ECal RMM 5 during the third T2K experimental running period.	89
4.8	Progressive smoothing of mean time relative to trigger for five minute data taking periods as a function of Unix time. Data for ECal RMM 11 during the third T2K experimental running period.	92
4.9	Standard deviation spectrum obtained from Figure 4.8d shown in black, Figure 4.8d is shown superimposed above in red.	93
4.10	Selected time slips superimposed upon Figure 4.8d. The red bands indicate the lower and upper edges of the time slip regions as located by the standard deviation peak finding algorithm. The blue lines with arrow indicate the selected time slip position at the centre of the red bands, with the arrow indicating the direction of each time slip.	95
4.11	Resultant mean time relative to trigger timing distribution once the RMM time slip corrections, as identified in Figure 4.10 have applied to the data.	97
4.12	The final mean time relative to trigger data spectra, once all relevant time slip corrections have been applied, for four different RMMs across three different Trip-T detectors. All the data spectra are now flat with time, except for the presence of the CTM to MCM time slips which occur simultaneously for all RMMs in all Trip-T sub-detectors.	99
4.13	RMM calibration constants, with pre-applied RMM time slip corrections, as a function of ND280 run number for the third T2K experimental running period.	100
5.1	Accumulated total POT, good spill data quality POT, and good spill and ND280 data quality POT as a function of time.	104

5.2	Example event displays for simulated interactions where both π^0 decay photons shower within the DsECal and/or BrECal modules for FGD1. Detector hits are shown with true particle trajectories overlaid in green for the muon and yellow for the decay photons. .	107
5.3	Example event displays for simulated interactions where both π^0 decay photons shower within the DsECal and/or BrECal modules for FGD2. Detector hits are shown with true particle trajectories overlaid in green for the muon and yellow for the decay photons. .	108
5.4	Number of TPC vertical clusters (nodes) crossed by tracks passing the Negative FGD Fiducial Volume cut. The red arrow indicates the tracks which pass the Track Quality cut criteria.	111
5.5	TPC pull hypothesis distributions for tracks passing the Forward Going track cut. Tracks must have a pull within the selection criteria, as indicated by the red arrows, for at least one of the TPC PID hypotheses in each TPC crossed.	113
5.6	TPC muon pull hypothesis, 5.6a, and electron pull hypothesis, 5.6b, distributions for tracks passing the Track PID Quality cut. Tracks must have a pull within the selection criteria, as indicated by the red arrows, for both the muon and electron TPC PID hypotheses simultaneously in at least one TPC that is crossed. . . .	114
5.7	The efficiency and purity for selecting CC inclusive-like events, along with the efficiency for selecting $CC1\pi^0$, events as a function of cut.	115
5.8	The true momentum and direction of FGD FV muons prior to any selection criteria have been applied. 5.8c and 5.8d show the same information but with linear and logarithmic z axis scales, in both cases the z axis unit is the percentage of entries in each histogram bin.	116

5.9	The true momentum and direction of FGD FV muons selected by the criteria as described in Section 5.4. 5.9c and 5.9d show the same information but with linear and logarithmic z axis scales, in both case the z axis unit is the percentage of entries in each histogram bin.	117
5.10	The efficiency for selecting FGD FV muons, using the criteria as described in Section 5.4, as a function of their true momentum and direction. 5.10c and 5.10d show the same information but with linear and logarithmic z axis scales, in both case the z axis unit is the selection efficiency in each histogram bin.	119
5.11	The measured momentum and direction for the track selected as the muon candidate for running periods Run 2 Air to Run 3b. The integral values at the top of each histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. . . .	121
5.11	The measured momentum and direction for the track selected as the muon candidate for running periods Run 3c to Run 4 Air. The integral values at the top of each histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. . . .	122
5.12	5.12a shows the measured momentum of the track selected as the muon candidate for all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.12b shows the data over MC ratio as a function of track momentum from 5.12a.	124
5.13	5.13a shows the measured direction of the track selected as the muon candidate for all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.13b shows the data over MC ratio as a function of track momentum from 5.13a.	125

-
- 5.14 Number of isolated ECal clusters, as broken down by final state, after all muon selection criteria have been previously applied. The red arrow indicates the bunches which pass the ECal Multiplicity criteria. 127
- 5.15 Receiver Operating Characteristic (ROC) curve for a range of MVA techniques. The better the technique performance, the closer to the top right hand corner (*signal efficiency = background rejection = 1*) of the plot the ROC curve will extend. The above curves are produced from training against the discriminating variables of the FGD2DsBr event topology training sample, described in detail in Section 5.5.3. The discriminating performance of the different MVA techniques for the FGD2DsBr event topology is typical of all event topologies. 130
- 5.16 From the discriminating variables available at the root node, branching cuts are applied one at a time until samples are sufficiently segregated into signal-like and background-like leaf nodes [93]. . . 131
- 5.17 Separation, 5.17a, and Significance, 5.17b, achieved from training multiple BDTs with varied numbers of trees in the forest and beta learning rate factors for the FGD2DsBr topology. 135
- 5.18 BDT response for FGD2DsBr training and testing samples with $\beta = 0.5$ and 1200 trees in the training forest. The results of a KS test to compare the training and testing sample show the two signal distributions are in near perfect agreement, but some minor disagreement is found for the background samples. 136
- 5.19 Strength of the matching between the training and testing sample response distributions for signal, 5.19a, and background, 5.19b, as obtained with a KS test from training multiple BDTs with varied numbers of trees in the forest and beta learning rate factors for the FGD2DsBr topology. 137

5.20	BDT cut tuning for FGD2DsBr training and testing samples with $\beta = 0.5$ and 1200 trees in the training forest. The optimal performance is achieved when applying a cut of -0.05, yielding a significance, $S/\sqrt{S+B}$, of 27.48. Note, this is a different significance definition to that described by Equation 5.9.	138
5.21	Energy and Direction of π^0 's from the $CC1\pi^0$ MC sample, 5.21a, and <i>particle gun</i> simulation, 5.21b.	141
5.22	Discriminating variables used in the training of the FGD2DsBr topology BDT.	144
5.23	Correlation between the signal, 5.23a, and background, 5.23b, discriminating variables used in the training of the FGD2DsBr topology BDT.	148
5.24	BDT response and cut optimisation for the FGD1 vertex analysis topologies.	151
5.24	BDT response and cut optimisation for the FGD2 vertex analysis topologies.	152
5.25	Efficiency and Purity of final selection sample for each of the six signal topologies of interest, as a function of BDT response cuts. The dashed black lines and arrows, indicate the position of the optimised cut values.	154
5.26	Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection.	156
5.27	The efficiency and purity for selecting $CC1\pi^0$ signal topology events, along with the purity for selecting π^0 decay photon events as a function cut.	161

5.28	5.28a shows the reconstructed π^0 invariant mass distribution for the FGD1DsDs event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.28b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.28a.	163
5.29	5.29a shows the reconstructed π^0 invariant mass distribution for the FGD1DsBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.29b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.29a.	164
5.30	5.30a shows the reconstructed π^0 invariant mass distribution for the FGD1BrBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.30b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.30a.	165
5.31	5.31a shows the reconstructed π^0 invariant mass distribution for the FGD2DsDs event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.31b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.31a.	166

5.32	5.32a shows the reconstructed π^0 invariant mass distribution for the FGD2DsBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.32b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.32a.	167
5.33	5.33a shows the reconstructed π^0 invariant mass distribution for the FGD2BrBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.33b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.33a.	168
5.34	Reconstructed π^0 invariant mass distribution for each event topology over all running periods, broken down by final state.	169
5.35	Reconstructed π^0 invariant mass distribution for each event topology over all running periods, broken down by π^0 decay photon reconstruction.	170
5.36	Data-MC comparison plots for events that could contribute to the final FGD1BrBr event topology, after the Reconstructed Mass cut has been applied.	175
5.37	Example of binned data over MC ratios used in calculating the Multivariate Analysis systematic error. The error bar shown on each data point is the statistical error.	179
5.38	Example of reconstructed π^0 mass as a function of BDT inputs for all events contributing to the final selection within those topologies.	180
5.39	Data over MC ratio multiplying distributions from Figure 5.37 and the result of their application to Figure 5.38.	181

5.40	Reconstructed π^0 mass peak after multiplying through by the variable importance reweighted data over MC ratio for each of the six event topologies of interest.	184
5.41	Reconstructed π^0 mass peak after multiplying through by the data over MC ratio for each of the six event topologies of interest. No variable importance reweighting is applied for these distributions.	185
5.42	Fractional covariance matrix for the neutrino flux uncertainty as a function of neutrino flavour and detector.	187
5.43	Cholesky decomposition of the fractional covariance matrix for the neutrino flux uncertainty as a function of neutrino flavour and detector.	189
5.44	Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD1DsDs event topology.	191
A.1	Discriminating variables used in the training of the FGD1DsDs topology BDT.	220
A.2	Correlation between the signal, A.2a, and background, A.2b, discriminating variables used in the training of the FGD1DsDs topology BDT.	223
A.3	Discriminating variables used in the training of the FGD1DsBr topology BDT.	224
A.4	Correlation between the signal, A.4a, and background, A.4b, discriminating variables used in the training of the FGD1DsBr topology BDT.	227
A.5	Discriminating variables used in the training of the FGD1BrBr topology BDT.	228
A.6	Correlation between the signal, A.6a, and background, A.6b, discriminating variables used in the training of the FGD1BrBr topology BDT.	231

A.7	Discriminating variables used in the training of the FGD2DsDs topology BDT.	232
A.8	Correlation between the signal, A.8a, and background, A.8b, discriminating variables used in the training of the FGD2DsDs topology BDT.	235
A.9	Discriminating variables used in the training of the FGD2BrBr topology BDT.	236
A.10	Correlation between the signal, A.10a, and background, A.10b, discriminating variables used in the training of the FGD2BrBr topology BDT.	239
B.1	Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the left side BrECal. . . .	241
B.2	Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the right side BrECal. . . .	244
B.3	Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the top BrECal.	247
B.4	Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the bottom BrECal. . . .	250
B.5	Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the DsECal.	253
C.1	Data-MC comparison plots for events that could contribute to the final FGD1DsDs event topology, after the Reconstructed Mass cut has been applied.	257
C.2	Data-MC comparison plots for events that could contribute to the final FGD1DsBr event topology, after the Reconstructed Mass cut has been applied.	260
C.3	Data-MC comparison plots for events that could contribute to the final FGD2DsDs event topology, after the Reconstructed Mass cut has been applied.	263

C.4	Data-MC comparison plots for events that could contribute to the final FGD2DsBr event topology, after the Reconstructed Mass cut has been applied.	266
C.5	Data-MC comparison plots for events that could contribute to the final FGD2BrBr event topology, after the Reconstructed Mass cut has been applied.	269
D.1	The schematic decay of a π^0 in flight into a pair of photons. The π^0 is initially travelling in the positive z direction, and the two decay photons move away from the decay vertex at angles α and β from the z direction in the xz plane.	272
E.1	Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD1DsBr event topology. . .	276
E.2	Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD1BrBr event topology. . .	277
E.3	Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD2DsDs event topology. . .	278
E.4	Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD2DsBr event topology. . .	279
E.5	Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD2BrBr event topology. . .	280

Chapter 1

Introduction to Neutrino Physics

1.1 Historic Background

The history of neutrino physics initially takes us back to 1930 and the proposal of their existence by Wolfgang Pauli. In a letter [1] to participants of the Tübingen conference on radioactivity, Pauli proposed the existence of an unseen neutrally charged particle as a means to conserve energy (along with momentum and spin) during the process of β -decay. Initially named the neutron, this particle would be imparted with the missing energy from β -decays, as inferred from the observed continuous energy spectrum of emitted electrons [2], therefore preventing the violation of energy conservation that would have occurred were electrons to be the only particle emitted during β -decay. James Chadwick's discovery in 1932 of the neutron [3] lead to Enrico Fermi renaming Pauli's neutron the neutrino in his theory of β -decay [4].

It took over two decades from its existence being first postulated for the antineutrino to be experimentally observed. This was first achieved by Clyde Cowan and Frederick Reines in 1953 [5], and was confirmed in 1956 [6], via the search for the products of inverse β -decay caused by reactor antineutrino interactions on protons in water, $\bar{\nu}_e + p^+ \rightarrow e^+ + n$. In their 1956 experimental design, Cowan and Reines used a 7.5 cm wide water target doped with cadmium chloride sandwiched between liquid scintillator to detect the antineutrino interactions. The signal for

an interaction was the coincident detection of two 0.51 MeV gamma-rays from the annihilation of the emitted positron within the scintillator followed by the capture of the moderated neutron on cadmium. The neutron capture resulted in the emission of several gamma-rays, with energies totalling 9 MeV, delayed by 5 μ s relative to the initial coincident positron annihilation signal. The results of the experiment showed an event rate which varied with reactor power and was consistent with the predicted neutrino interaction cross-section.

Whilst Cowan and Reines were conducting their research Raymond Davis was also working to detect reactor antineutrinos via inverse β -decay on chlorine, $\bar{\nu}_e + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}$. The method, first proposed by Bruno Pontecorvo in 1946 [7] and described as a “breathtakingly beautiful proposal” by John Bahcall [8], employed carbon tetrachloride as the target material and detected events by filtering out any argon that was produced and looking for its subsequent radioactive decay [9]. The decay of ${}^{37}\text{Ar}$ occurs through electron capture with a half-life of 35 days; this process produces an X-ray photon which the Davis experiment detected. Davis did not detect any excess of ${}^{37}\text{Ar}$ above cosmic-ray induced backgrounds and so this was the first indication that the neutrino and antineutrino were not identical. The work also placed the first upper limit on the neutrino flux from the sun, a topic that would dominate Davis’ work for the rest of his career.

Nearly a decade after the discovery of the (electron-anti)neutrino, the muon-neutrino was discovered by a group working at Brookhaven National Laboratory. The work, lead by Leon Lederman, Melvin Schwartz and Jack Steinberger [10], utilised the Brookhaven Alternating Gradient Synchrotron (AGS) to accelerate protons to 15 GeV before colliding them upon a beryllium target. The collision produced a flux of particles dominated by charged pions. The charged pions then decayed in flight to produce muons and neutrinos via $\pi^\pm \rightarrow \mu^\pm + \overset{(-)}{\nu}$. This was the first accelerator produced neutrino beam experiment, an approach which has been utilised by many subsequent experiments including T2K. The muons produced by the pion decay were then stopped by a 13.5 m thick iron shield placed prior to a 10 ton aluminium spark chamber. The neutrinos were then detected by their charged-current interactions within the spark chamber. If the neutrinos

produced coupled equally to electrons and muons the resultant interactions would be equally distributed between muon track-like and electron shower-like events. This was not observed, instead the events produced were consistent with the production of muons leading to the conclusions that neutrinos had two distinct types.

Evidence for the likely existence of a third generation of neutrino was inferred from the discovery of the tau-lepton during the latter half of the 1970s. It was initially observed by Martin Perl and his colleagues using the Mark I detector at the SPEAR e^+e^- colliding ring [11], and later confirmed by the PLUTO experiment at DESY [12]. The established link between the electron- and muon-neutrinos and their counterpart leptons implied the existence of a counterpart tau-neutrino to the tau-lepton. Some of the observed missing energy and momentum from the decays of the tau-lepton was then expected to be being carried away by this unobserved tau-neutrino.

Further evidence for the existence of a third generation of neutrino came in 1989 from the invisible decay width (inverse lifetime) of the Z^0 boson decay as measured by the four experiments sited around the Large Electron-Positron Collider (LEP) at CERN [13][14][15][16] and the SLD detector at the Stanford Linear Collider [17]. The experiments measured the width of the Z^0 production peak and compared that with the branching fraction of the observable states and the Standard Model prediction for the unobservable neutrinos. Under the assumption that all the invisible width of the Z^0 production peak is due entirely to decays to active neutrino states with $m_\nu \leq \frac{m_Z}{2}$, the total number of neutrino families was found to be 2.9840 ± 0.0082 [18]. Figure 1.1 shows the results from the LEP experiments and demonstrates how the hadron production cross-section would vary as a function of the number of active neutrino species.

The discovery of the 3rd neutrino generation, the tau-neutrino itself, would have to wait until it was observed by the DONUT experiment in 2000 [19]. DONUT employed the Tevatron accelerator at Fermilab to create a neutrino beam by impinging 800 GeV protons onto the accelerator beam dump. The DONUT detector was placed 36 m downstream of the beam dump and most particles from the proton interactions were removed from the beam through magnetic sweeping

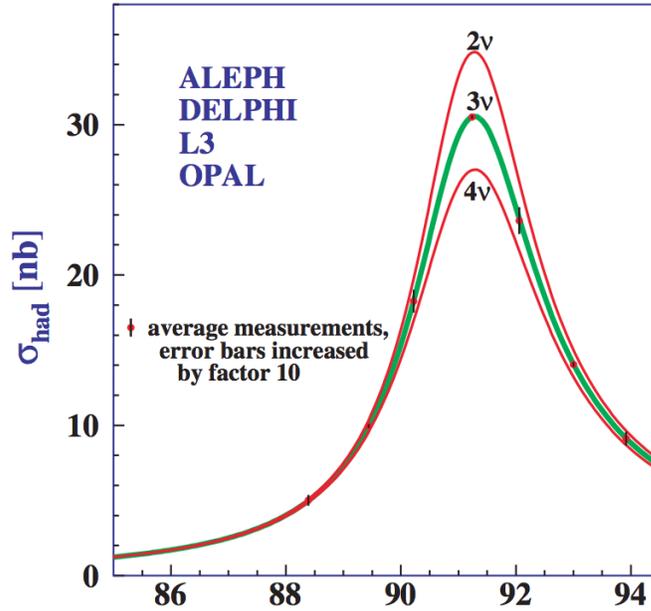


Figure 1.1: Variation in hadron production cross-section around Z^0 the mass as a function of the number of active neutrino species [18].

or arrested by the concrete, iron and lead shielding that preceded the experiment. After shielding, a mixed flavour beam of neutrinos entered the DONUT detector with the tau-neutrino component coming primarily from the decay of charmed mesons, particularly the leptonic decay of the D_S which accounted for 85% of tau-neutrinos produced. The detection of tau-neutrinos was then attained by searching for their charged-current interactions in nuclear emulsion targets, interleaved with scintillators and drift chamber detectors. The scintillators and drift chambers allowed the electronic tagging of events, identifying the relevant nuclear emulsions for photographic development. A tau-neutrino event was then characterised by the observation of charged tracks emerging from an interaction vertex, in which one track would show a kink in its trajectory due to the rapid decay of a tau lepton into a single charged daughter and unobserved neutral particles (86% branching fraction). DONUT observed four tau-neutrino candidate events on a predicted background of 0.34 ± 0.05 events, thus concluding that they had made the first observation of the tau-neutrino.

1.2 Neutrino Oscillations

Returning to Davis, his experiments continued with neutrino capture on chlorine, but expanded in size and moved to the Homestake gold mine in South Dakota. Davis was then able to measure the solar neutrino flux in a low cosmic ray background environment in an attempt to confirm the predicted neutrino flux from the Standard Solar Model (SSM) as calculated by Bahcall. The Homestake experiment and the SSM were continually updated and their results were continuously at odds, the experiments always measuring a neutrino flux below that expected by the SSM. The discrepancy was dubbed the ‘Solar Neutrino Problem’ and the final results of the Homestake experiment measured a solar neutrino flux of about one-third [20] of that predicted by the SSM [21].

One of the potential causes of the discrepancy between the SSM and Homestake experiment may have been caused by the sensitivity of the measurement to the uncertainties in the flux prediction for the higher energy electron-neutrinos that chlorine is sensitive to, see Figure 1.2. Two later experiments, GALLEX and SAGE were able to address this possibility using neutrino capture on gallium. Gallium, through its lower neutrino capture energy threshold ($E_\nu > 0.233$ MeV compared to $E_\nu > 0.814$ MeV for chlorine [8]), was able to probe the large, better modelled flux of electron-neutrinos produced by the primary *proton-proton* fusion reaction. The final results from both GALLEX [22] and SAGE [23] also measured a deficit in the expected solar neutrino flux of about 50 - 60%, leaving the Solar Neutrino Problem intact.

Further evidence for a problem within the neutrino sector came from two proton-decay search experiments, IMB and Kamiokande. Both experiments employed a large volume of ultra-pure water surrounded by an array of Photo-Multiplier Tubes (PMTs) to detect the light produced during particle interactions. Such experiments can detect neutrinos through the Čerenkov radiation produced by their daughter leptons from charged-current interactions within the water. Using ring-imaging techniques allowed the two experiments to distinguish between electron-like and muon-like neutrino interactions. After analysis of their fully contained neutrino interactions, both IMB [25] and Kamiokande [26][27] observed a deficit

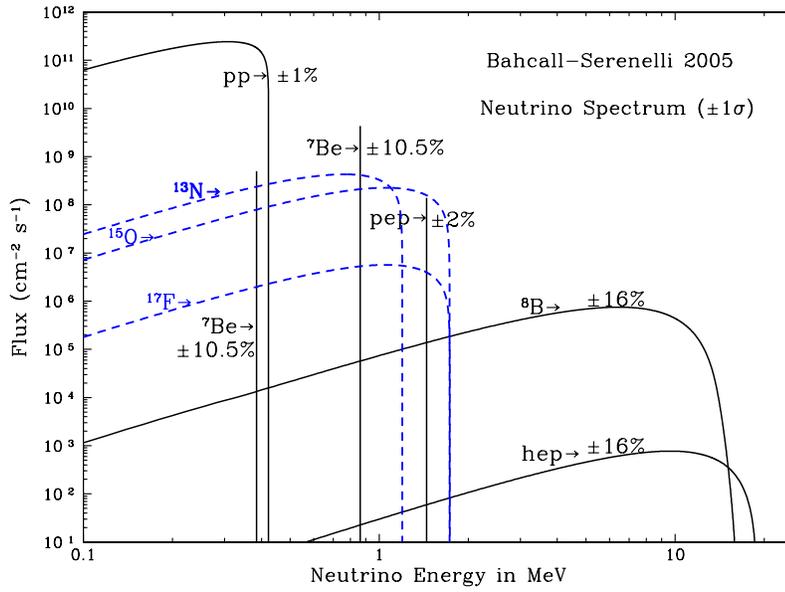


Figure 1.2: Solar neutrino energy spectrum with 1σ uncertainty values, as predicted by the Standard Solar Model. Solid line spectra represent those neutrino emissions from the *proton-proton* chain of fusion reactions, dashed line spectra are as a result of the *CNO cycle* chain of fusion reactions [24].

in the expected number of cosmic ray induced muon-neutrino events. However, there was no deficit in the rate of electron-neutrino events from neutrinos produced through cosmic ray interactions in the Earth's atmosphere. This difference was known as the 'Atmospheric Neutrino Problem'.

Kamiokande showed in their 1994 paper [27] that there was a slight angular, and therefore distance, dependence to the atmospheric muon-neutrino deficit. This indicated that the cause of the problem might be due to neutrinos oscillating from one type to another in flight. This angular dependence was confirmed in 1998 by the Super-Kamiokande experiment [28], a significantly larger successor to the Kamiokande experiment, which will be discussed in detail in Section 2.3. A clear deficit was observed in the number of upwards-going muon-like neutrino events, with the ratio of upwards to downwards events recorded as $0.52^{+0.07}_{-0.06}$ (stat.) ± 0.01 (sys.) against an expectation of 0.98 ± 0.03 (stat.) ± 0.02 (sys.). Super-Kamiokande later binned the result as a function of distance and neutrino energy, see Figure 1.3, showing clearly how this could be interpreted as muon-neutrinos oscillating to tau-neutrinos [29].

Given the evidence that the Atmospheric Neutrino Problem might be explained by neutrino oscillations, it was clear that the same effect might be reducing the flux of electron-neutrinos from the sun. This was confirmed in 2002 when the SNO experiment measured the total flux of neutrinos from the sun through the measurement of charged-current interactions of electron-neutrinos ($\nu_e + D \rightarrow e^- + 2p$), along with neutral-current interactions ($\nu_l + D \rightarrow \nu_l + n + p^+$) and elastic scatterings ($\nu_l + e^- \rightarrow \nu_l + e^-$) of all active neutrino species. SNO was a heavy-water Čerenkov detector capable of probing neutrinos produced by the ^8B fusion process, see Figure 1.2. Charged-current interactions and elastic scatterings were detected by the Čerenkov radiation produced by the electrons released during the interactions, whereas neutral-current interactions were detected by the roughly 6 MeV gamma-ray emitted during the capture of the free neutron upon deuterium. Both charged-current and elastic scattering events give some directional information for the incident neutrino, allowing their trajectories to be matched back towards the position of the sun, adding further certainty to the origin of such events. SNO showed that the charged-current interaction rate was well below that

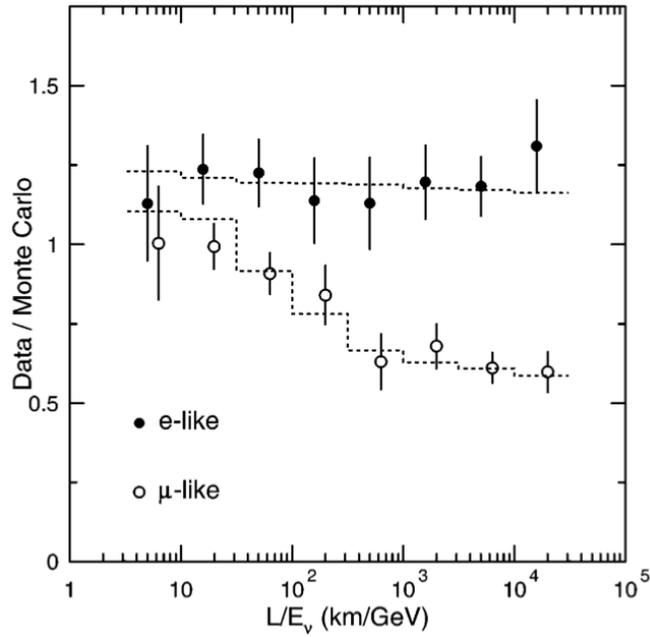


Figure 1.3: Data points show the ratio of fully contained data to Monte Carlo predicted (in the absence of neutrino oscillations) events as a function of distance over energy for the Super-Kamiokande atmospheric neutrino sample. Dashed lines show the predicted shape with the inclusion of $\nu_\mu \leftrightarrow \nu_\tau$ oscillations with $\Delta m^2 = 2.2 \times 10^{-3}$ eV and $\sin^2 2\theta = 1$ [29].

predicted by the SSM, however the flux of all active neutrino species was found to be in line with the SSM predicted value, thus providing very strong evidence that solar electron-neutrinos were undergoing flavour transformations [30].

1.2.1 Oscillation Theory

Neutrino oscillations were first proposed by Pontecorvo in 1957 [31], although his proposal considered oscillations between neutrino and antineutrino states. The modern view of three (or more) flavour oscillations was developed by Maki, Nakagawa and Sakata [32], and describes how the neutrino flavour eigenstates ($\nu_{l=e,\mu,\tau}$), those which couple via the weak interaction to the charged leptons, are a superposition of the mass eigenstates ($\nu_{i=1,2,3}$) that propagate through space

$$|\nu_l\rangle = \sum_i U_{li}^* |\nu_i\rangle, \quad (1.1)$$

Conversely the mass eigenstates can be represented as a superposition of the neutrino flavour eigenstates

$$|\nu_i\rangle = \sum_l U_{li} |\nu_l\rangle. \quad (1.2)$$

U_{li} is the the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix, as decomposed into its three axial rotations [33]

$$U_{li} \equiv \overbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}^{\text{'Atmospheric' term}} \overbrace{\begin{pmatrix} c_{13} & 0 & e^{-i\delta} s_{13} \\ 0 & 1 & 0 \\ -e^{i\delta} s_{13} & 0 & c_{13} \end{pmatrix}}^{\text{'Reactor/Accelerator' term}} \overbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}^{\text{'Solar' term}} \quad (1.3a)$$

$$\equiv \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & e^{-i\delta} s_{13} \\ -s_{12}c_{23} - e^{i\delta} c_{12}s_{23}s_{13} & c_{12}c_{23} - e^{i\delta} s_{12}s_{23}s_{13} & s_{23}c_{13} \\ s_{12}s_{23} - e^{i\delta} c_{12}c_{23}s_{13} & -c_{12}s_{23} - e^{i\delta} s_{12}c_{23}s_{13} & c_{23}c_{13} \end{pmatrix}, \quad (1.3b)$$

where $c_{ij} \equiv \cos \theta_{ij}$, $s_{ij} \equiv \sin \theta_{ij}$ and δ is a Charge-Parity violating phase.

The mixing angles, θ_{12} , θ_{13} and θ_{23} are traditionally explored via measurements using solar, reactor/accelerator and atmospheric neutrinos respectively. As such it is common to describe each axial component via these terms as shown in Equation 1.3a.

If we evolve the flavour eigenstate from Equation 1.1 to any position and time we obtain the wavefunction

$$|\varphi(x, t)\rangle = \sum_i U_{li}^* |\nu_i\rangle e^{-i\mathbf{p}_\nu \cdot \mathbf{x}_\nu}, \quad (1.4)$$

where \mathbf{p}_ν and \mathbf{x}_ν are the neutrino four momentum $\mathbf{p}_\nu = (E_\nu, \vec{p}_\nu)$, and four position $\mathbf{x}_\nu = (t_\nu, \vec{x}_\nu)$ respectively.

Under the assumption that the neutrino mass is small compared to its momentum then $E_i \simeq \mathbf{p}_\nu + \frac{m_i^2}{2\mathbf{p}_\nu}$. Using natural units, $c = \hbar = 1$, then Equation 1.4 becomes

$$|\varphi(x, t)\rangle = \sum_i U_{li}^* |\nu_i\rangle e^{-i\frac{m_i^2}{2\mathbf{p}_\nu} x}. \quad (1.5)$$

Substituting in with Equation 1.2 to express mass eigenstates in terms of a flavour eigenstate α , we find

$$|\varphi(x, t)\rangle = \sum_i \sum_\alpha U_{li}^* e^{-i\frac{m_i^2}{2p_\nu}x} U_{\alpha i} |\nu_\alpha\rangle. \quad (1.6)$$

We can then calculate the probability for a neutrino of flavour state l , to be later observed as a flavour state α having propagated some distance x (or time t under the assumption that $v_\nu = c = 1$)

$$P(l \rightarrow \alpha) = |\langle \nu_l | \nu_\alpha(x) \rangle|^2 \quad (1.7a)$$

$$= \left[\sum_j U_{lj} e^{-i\frac{m_j^2}{2p_\nu}x} U_{\alpha j}^* \right] \left[\sum_i U_{li}^* e^{-i\frac{m_i^2}{2p_\nu}x} U_{\alpha i} \right] \quad (1.7b)$$

$$= \sum_{i,j} U_{li}^* U_{lj} U_{\alpha i} U_{\alpha j} e^{i\frac{m_i^2 - m_j^2}{2p_\nu}x}. \quad (1.7c)$$

From Equation 1.7c we can therefore see that the probability for one neutrino flavour eigenstate to oscillate to another is dependent on the squared difference in the mass eigenstates, the neutrino momentum and distance travelled. Under the ultra relativistic assumption of $p_\nu \approx E_\nu$, it is possible to produce a three-flavour approximation of muon-neutrino oscillation probability, as is relevant for the T2K experiment

$$P(\nu_\mu \rightarrow \nu_\mu) \simeq 1 - \cos^4 \theta_{13} \sin^2 2\theta_{23} \sin^2 \left[1.27 \frac{\Delta m_{23}^2 L(km)}{E(GeV)} \right] \quad (1.8a)$$

$$P(\nu_\mu \rightarrow \nu_e) \simeq \sin^2 2\theta_{13} \sin^2 \theta_{23} \sin^2 \left[1.27 \frac{\Delta m_{13}^2 L(km)}{E(GeV)} \right]. \quad (1.8b)$$

Neutrino-matter interaction effects are considered negligible for the T2K baseline and so are ignored in the above three-flavour approximations.

1.3 Current Status

Neutrino physics is now in a position where a plethora of experiments have successfully measured all the mixing angles and mass squared differences to varying degrees of accuracy as summarised in Table 1.1 and Figure 1.4. The most recent

discovery to be made was that of θ_{13} ; this was first found likely to be non-zero in 2011 by both T2K, with 2.5σ significance [34], and Double Chooz, at the 94.6% confidence level [35]. These results were confirmed by RENO [36] and Daya Bay [37] the following year with 4.9σ and 5.2σ significance respectively.

Parameter	Best Fit ($\pm 1\sigma$)
Δm_{21}^2 [$10^{-5} eV^2$]	$7.54^{+0.26}_{-0.22}$
$ \Delta m_{31}^2 $ [$10^{-3} eV^2$]	2.43 ± 0.06 (2.38 ± 0.06)
$\sin^2 \theta_{12}$	0.308 ± 0.017
$\sin^2 \theta_{23}$	$0.437^{+0.033}_{-0.023}$ ($0.455^{+0.039}_{-0.031}$)
$\sin^2 \theta_{13}$	$0.0234^{+0.0020}_{-0.0019}$ ($0.0240^{+0.0019}_{-0.0022}$)

Table 1.1: Current best fit neutrino oscillation parameters under the assumption of the normal and (inverted) mass hierarchy [38].

There continues to be an effort to measure the mixing parameters with ever increasing accuracy, whilst also addressing many other outstanding questions within the field of neutrino physics. The most obvious is the value of δ , the Charge-Parity violating phase that is now accessible given that θ_{13} has been found to be non-zero, as can be seen in Equation 1.3. This is a question that will be addressed by T2K in conjunction with the NO ν A experiment [40], and other future high intensity accelerator experiments. A non-zero value of δ equates to CP violation in the lepton sector, potentially leading to an explanation for the scale of the matter-antimatter asymmetry of the universe [41].

Other outstanding questions in the field include discovering the nature of neutrinos, be they Dirac or Majorana particles; the sign of the mass squared splitting Δm_{31}^2 , better known as the hierarchy problem, this is the question as to whether m_3 is greater than or less than m_1 and m_2 , as illustrated in Figure 1.5; the absolute neutrino masses; the octant of θ_{23} ; and the possible existence of sterile neutrinos. These questions are being addressed by a range of experimental techniques which continue to add to our overall understanding of neutrino physics.

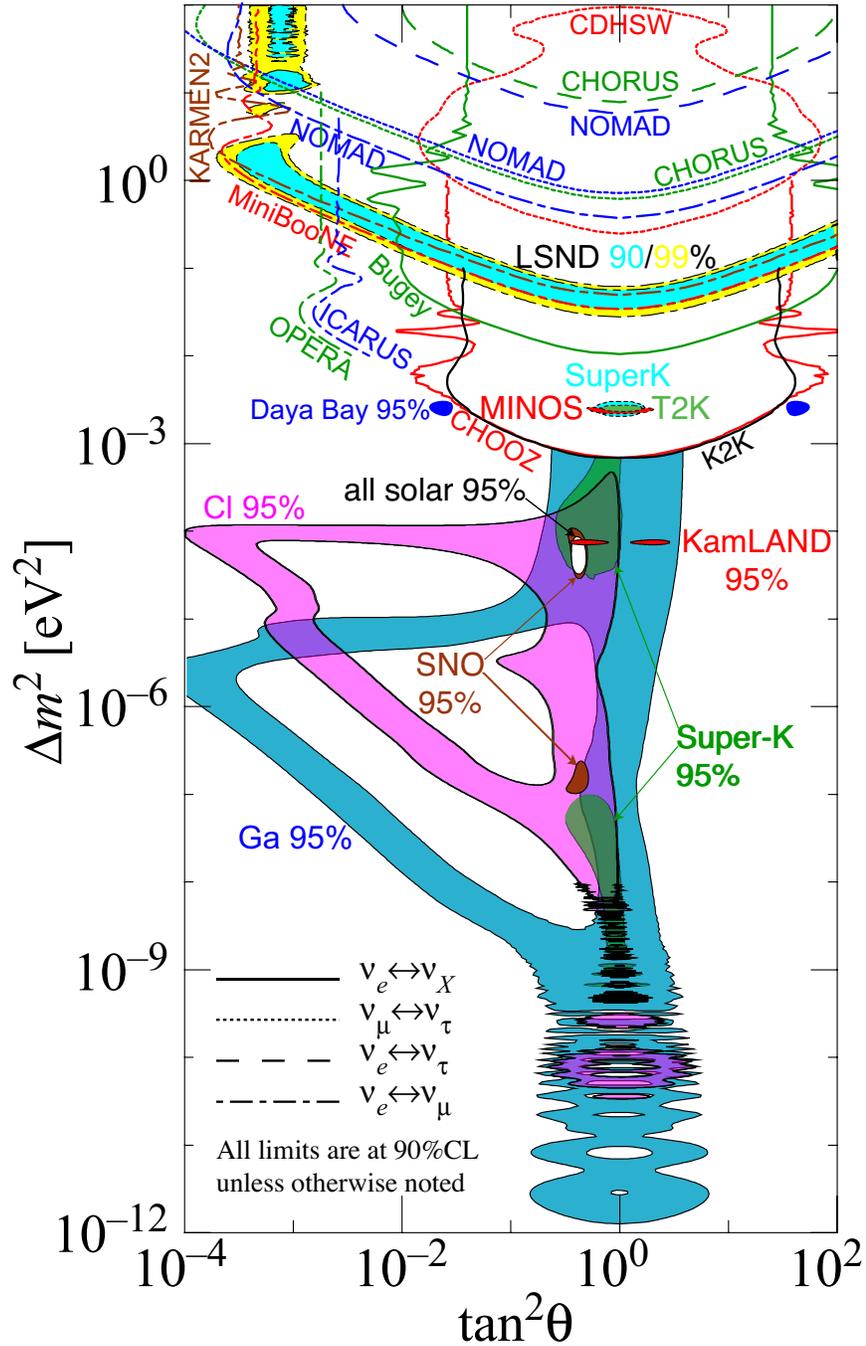


Figure 1.4: Mass splitting, Δm^2 , as a function of oscillation angle, $\tan^2 \theta$, for four types of neutrino oscillation. Lines show exclusion limits and coloured regions show allowed parameter space for a range of named experiments [39].

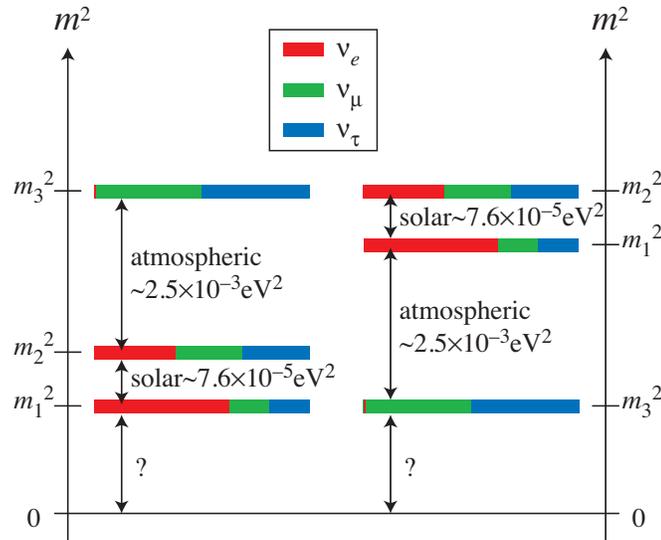


Figure 1.5: The two possible neutrino mass hierarchies, normal on the left and inverted on the right, given that m_2 is known to be greater than m_1 from observations of solar neutrinos. The colours indicate the weak eigenstate contributions to each of the mass eigenstates. The absolute neutrino mass scale is also unknown but constrained by cosmology and astrophysics, and ${}^3\text{H}$ β -decay experiments [42].

1.4 Neutrino Interaction Cross-Sections

The study of neutrino oscillations is now entering the era of precision measurements, as the mixing parameters become better constrained, neutrino beams become increasingly intense, and experimental design improves. Experiments are no longer limited by the statistics which they can acquire, given by the product of the neutrino flux and neutrino interaction cross-sections. Increasingly, the greatest constraint on experimental measurements is arising from the uncertainty on the neutrino-nucleus cross-sections for the range of target materials, neutrino energy, flavour, and interaction final states to which the experiments are sensitive.

1.4.1 Interaction Processes

The neutrino, being a neutral lepton, only interacts via the weak nuclear force. Direct observations of neutrinos is therefore impossible, instead experiments mea-

sure the particles outgoing from the exchange of a Z or W boson between a neutrino and nuclear target. The exchange of a Z boson produces a “neutral-current” (NC) interaction and preserves the neutrino within the final state. The exchange of a W boson leads to a “charge-current” (CC) interaction, leading to the presence of a charged lepton in the final state with the same flavour as that of the interacting neutrino. Around the GeV energy scale, both neutral- and charged-current neutrino interactions with nuclear targets exhibit a range of interaction processes with energy dependent cross-sections. The dominant interaction processes are known as (quasi-)elastic scattering, resonant interactions, and deep inelastic scatterings.

Below neutrino energies of approximately 2 GeV, neutrino-hadron interactions are dominated by scattering processes. These are called elastic scatters (ES) for neutral-current interactions and quasi-elastic scatters (QES) for charged-current interactions. Both processes are characterised by the neutrino elastically scattering off an entire nucleon, which may then be ejected from the nucleus. QES interactions also feature a charged lepton in the final state due to charge exchange with the target nucleon; Figure 1.6a shows a Feynman diagram to illustrate this process. Charged-current quasi-elastic (CCQE) interactions are the most useful interactions to experimentally measure as the kinematics of the outgoing lepton reveals information about the energy of the incoming neutrino, and because the relatively simple nature of the interaction process makes it a good test for different interaction models.

Within the 1 to 5 GeV neutrino energy region, resonant interactions (RES) become increasingly important. Such processes are characterised by the neutrino again interacting with a nucleon as a whole, but this time promoting the nucleon to an excited energy state such as a Δ^+ . The nucleon then quickly decays to the ground state, usually emitting a neutral or charged meson in the process. Such processes are common to both neutral- and charged-current interactions, with an example of the later shown in Figure 1.6b.

Finally, above neutrino energies of approximately 5 GeV, neutrino-nucleon interactions become dominated by the deep inelastic scattering (DIS) process. At these energies, the neutrino no longer interacts with a nucleon as a whole, but

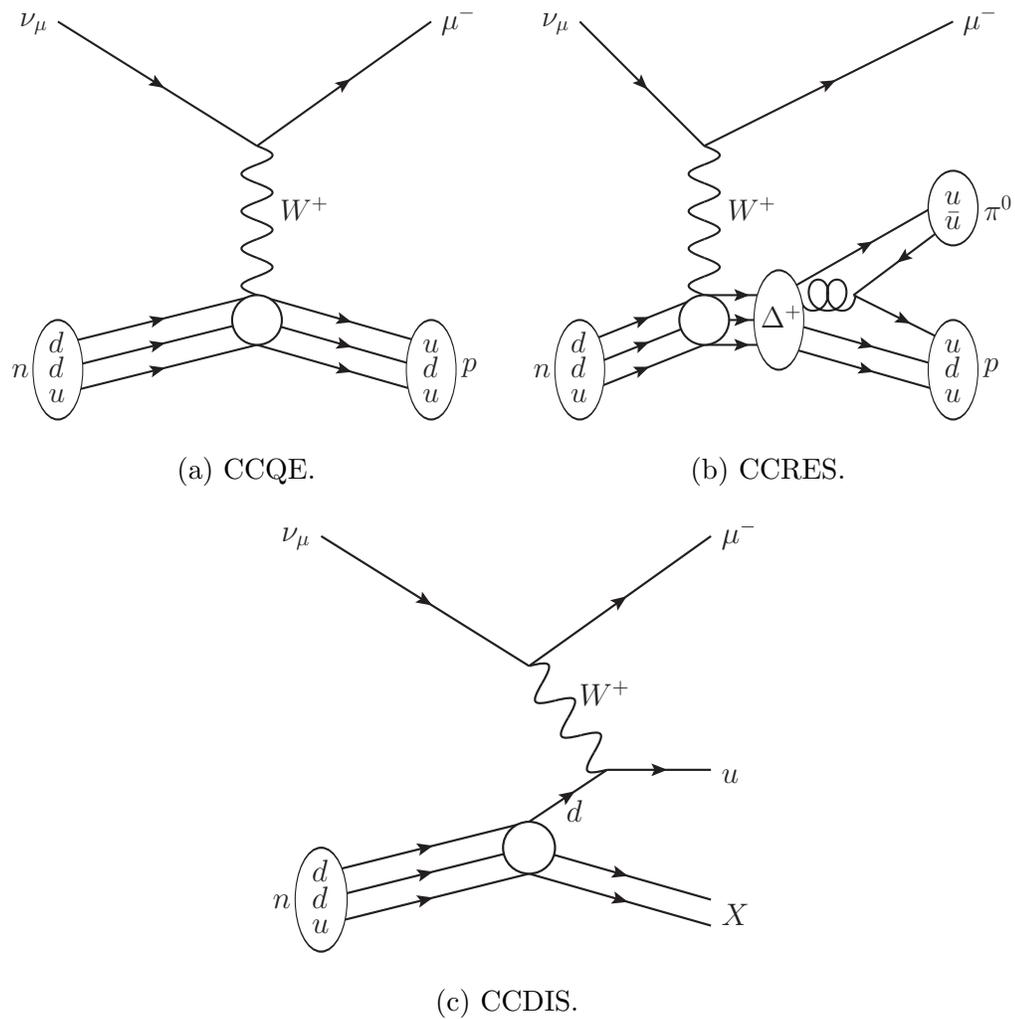


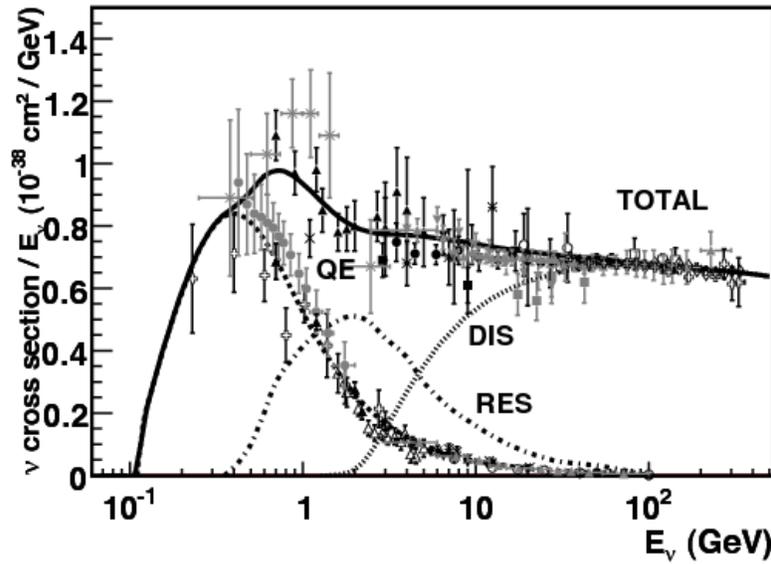
Figure 1.6: Illustrative Feynman diagrams for charged-current quasi-elastic (CCQE), resonant (CCRES) and deep inelastic (CCDIS) neutrino-nucleon interaction processes.

instead directly interacts with an individual quark. The release of the quark from a nucleon produces a final state featuring hadronic showers; see Figure 1.6c for a charged-current example.

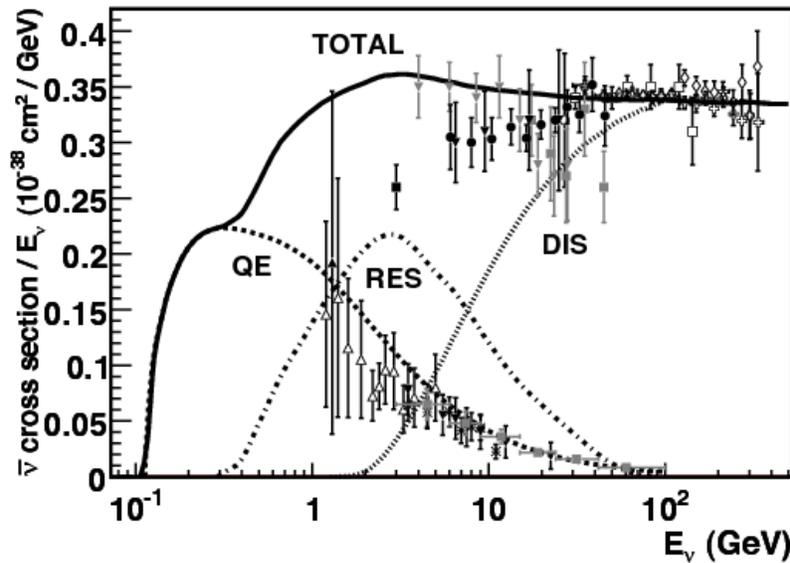
There are a range of models available to describe the different neutrino interaction processes, and a range of datasets available to compare them with. However, as can be seen in Figure 1.7, the data available tends to be for higher energy interactions, and for those data sets which are available around the 1 GeV region, the associated uncertainties are rather large, particularly for the antineutrino measurements. Additionally, although models are available for each of the separate interaction modes, there is no coherent model for the transition between the separate interaction regions. Furthermore, models that are available tend to make predictions at the nucleon level, but these are bound within nuclei of the neutrino target, for example oxygen, carbon and argon. Therefore it is increasingly important that consideration is given to accounting for processes that occur within the full nuclear environment. Interactions within the nuclear environment can affect both the kinematics and composition of the final state particles observed and so both theorists and experimentalists must account for this when comparing data with model predictions. For a comprehensive overview of the subject see Reference [43].

1.4.2 $CC1\pi^0$ Cross-Section Measurements

Within Chapter 5 an analysis will be described which aims to select $CC1\pi^0$ final states, with the hope that future work will lead to an exclusive $CC1\pi^0$ muon-neutrino cross-section measurement like that illustrated by the Feynman diagram in Figure 1.6b. The majority of previous measurements for this interaction process come from experiments undertaken in the 1980's with hydrogen and deuterium targets, see Figure 1.8. As can be seen the data covers a fairly comprehensive range of energies, but shows significant tension between the data and NUANCE prediction, particularly at higher energies. There is also a lack of significant data at the lowest energies applicable to modern accelerator based oscillation experiments, and the use of simple target nuclei make the application of such data to complex target nuclei difficult as subsequent interactions within



(a) Neutrino.



(b) Antineutrino.

Figure 1.7: Charged-current cross-sections per nucleon for neutrino, 1.7a, and antineutrino, 1.7b, interactions as a function of neutrino energy. Cross-section measurements are shown for quasi-elastic interactions along with the total across all interactions for a broad selection of experiments. The NUANCE [44] Monte-Carlo generator cross-section prediction is shown for CCQE, CCRES and CCDIS interactions, along with their combined total [43].

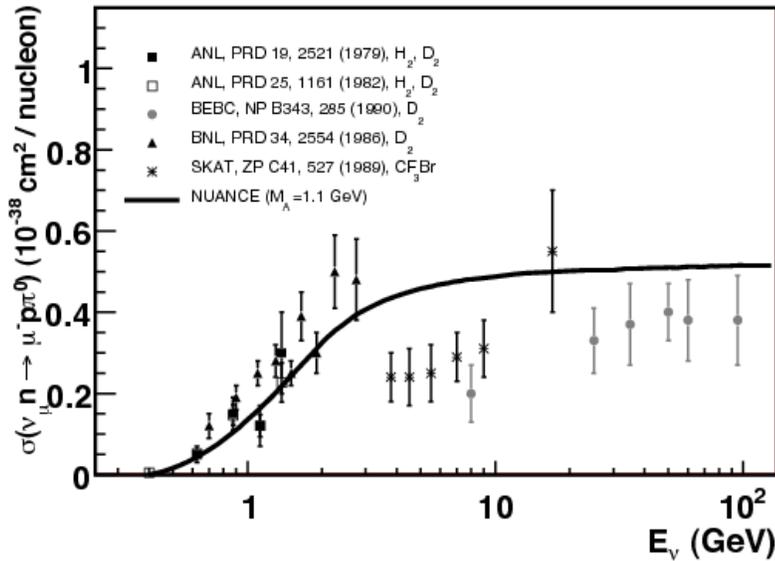


Figure 1.8: Historic $CC1\pi^0$ cross-section measurements compared to the NUANCE [44] Monte-Carlo generator cross-section prediction [43].

the target nucleus become important.

In 2011, the previous concerns were addressed by a measurement conducted by the MiniBooNE experiment, which released a flux averaged $CC1\pi^0$ differential cross-section on mineral oil (CH_2) [45], as seen in Figure 1.9. This measurement is the current state of the art and shows a systematically larger cross-section measurement when compared to the NUANCE prediction. This data is currently used by the T2K experiment to constrain the model input parameters used to evaluate the cross-section uncertainty for its oscillation results.

There is also limited additional information available from the $CC\pi^0$ inclusive cross-section measurements made by the SciBooNE [46] and K2K [47] experiments. These are useful results but do not cover the neutrino energy region around 1 GeV as comprehensively, or with as small an uncertainty, as the MiniBooNE result.

Additionally, significant numbers of related cross-section measurements are available for $NC\pi^0$ interactions. See for example K2K [48], SciBooNE [49] and MiniBooNE [50], where again the MiniBooNE result is considered the state of the art.

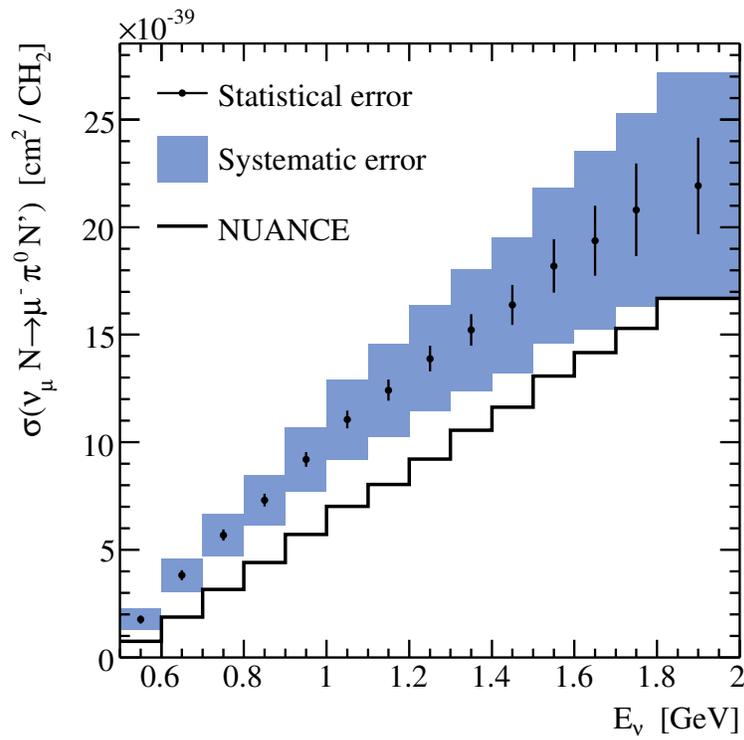


Figure 1.9: MiniBooNE CC1 π^0 , flux averaged, differential cross-section measurement on a mineral oil target compared to the NUANCE [44] Monte-Carlo generator cross-section prediction [45].

Chapter 2

The T2K Experiment

T2K (Tōkai to Kamioka) is a long-baseline neutrino oscillation experiment in Japan. The experiment is spread across two sites, the accelerator and near detector experimental complex at the Japan Proton Accelerator Research Complex (J-PARC) in the village of Tōkai-mura, Ibaraki Prefecture, and the far detector at the Kamioka Observatory within the Mozumi Mine, Gifu Prefecture, see Figure 2.1.

The primary aim of the T2K experiment is the precision measurements of the oscillation parameters, θ_{13} , θ_{23} and Δm_{23}^2 relevant to ν_e appearance and ν_μ disappearance from a ν_μ beam.

The following chapter will describe the production of the ν_μ beam and its characteristics, along with the design and purpose of the suite of near and far detectors.

2.1 The T2K Neutrino Beam

2.1.1 J-PARC

The J-PARC facility is situated on Japan's east coast and is a joint venture run by the High Energy Accelerator Research Organization (KEK) and the Japan Atomic Energy Agency (JAEA). The facility is home to accelerators and experimental halls for research into material and life sciences, nuclear physics and

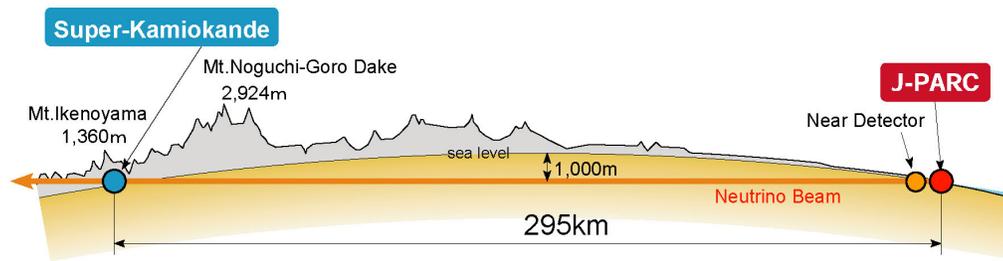


Figure 2.1: Schematic view of the T2K experiment. Neutrinos are produced at J-PARC on the east coast of Japan and recorded at the Near Detector and Super-Kamiokande far detector [51].

technologies, and particle physics.

J-PARC is home to the accelerators and target which produce the neutrino beam used by the T2K experiment, along with the ND280 (Near Detector at 280 m) complex.

2.1.2 Accelerator Complex

J-PARC has three accelerators, a 400 MeV Linear Accelerator (LINAC), a 3 GeV Rapid Cycling Synchrotron (RCS) and a 30 GeV Proton Synchrotron (PS), better known as the Main Ring (MR), see Figure 2.2. Acceleration of protons to the full 30 GeV is a three step process. First H^- anions are accelerated through the LINAC to 180 MeV before being stripped to H^+ cations by charge-stripping foils as they enter the RCS. Next the RCS further accelerates the protons to 3 GeV whilst also separating them into bunches. The RCS can contain two bunches at a time with a cycle rate of 25 Hz. Finally, about 5% of bunches from the RCS are then injected into the MR and accelerated to the full 30 GeV, with the remaining bunches being supplied to the Materials and Life Science Facility. Future upgrades to the MR will increase the maximum beam power to 50 GeV. The MR can contain up to nine proton bunches, although operates with only eight during neutrino running, and these can be extracted from the MR at two locations. At the slow extraction point bunches are diverted for use in the hadron

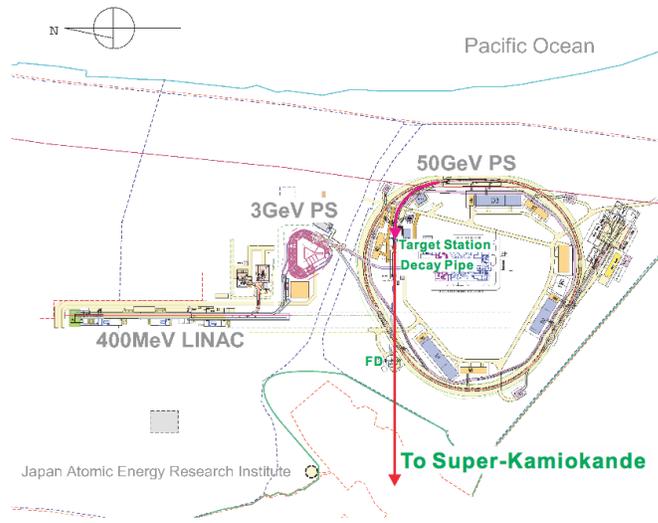


Figure 2.2: J-PARC facility with the three accelerators labelled. The RCS is labelled as 3 GeV PS and the MR is labelled as 50 GeV PS. The near detector complex is located at a point FD [52].

beamline experimental facility, and at the fast extraction point they are diverted into the neutrino beamline.

Fast extraction diverts all eight proton bunches into the neutrino beamline within a single turn, this constitutes one beam spill. The bunches are diverted by five kicker magnets which turn on during the passing of the unoccupied, ninth bunch position. Each bunch is constrained within 58 ns and the bunches are separated from one another by 581 ns [53]. The bunch timing structure allows effective rejection of cosmic ray induced background events at the detector complexes.

2.1.3 Neutrino Beamline

Once extracted from the MR by the kicker magnets the proton bunches travel through the two sections of the neutrino beamline. The primary beamline consists of a Preparation Section, Arc Section and Final Focusing Section which act to direct the proton beam towards Kamioka. The secondary beamline consists of the Target Station, Decay Volume and Beam Dump, this is where the neutrino beam itself is produced, see Figure 2.3.

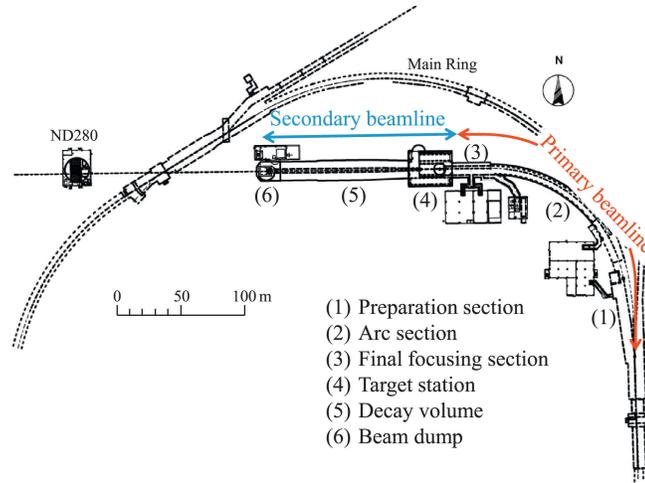


Figure 2.3: Schematic diagram of the neutrino beamline indicating the sub components of the primary and secondary beamlines [51].

Primary Beamline

The initial Preparation Section of the primary beamline is 57 m long and constitutes 11 normal conduction magnets which ensures the bunches are ready for acceptance into the Arc Section. The Arc Section bends the beam by 80.7° over 147 m (104 m radius of curvature) so that it is horizontally aligned towards Kamioka. The bending is done by 14 doublets of superconducting combined function magnets, whilst a further three pairs of superconducting vertical and horizontal steering magnets correct the beam orbit. The Final Focusing Section is 37 m in length and contains 10 normal conducting magnets which apply the final guiding and focusing of the beam onto the target, along with a final angular correction which diverts the beam angle down by 3.637° with respect to the horizontal to align the beam vertically with Kamioka.

Once all angular corrections are applied the resultant beam configuration does not point directly to Kamioka, but is angled 2.5° away. This produces a neutrino beam, which when observed by the far detector has a narrow energy spectrum, peaked to optimise the oscillation probability. This off-axis design is discussed in detail in Section 2.1.5.

It is essential for stable, low loss, beam operation that the proton beam is monitored throughout the primary beamline. This is accomplished by a suite of passive

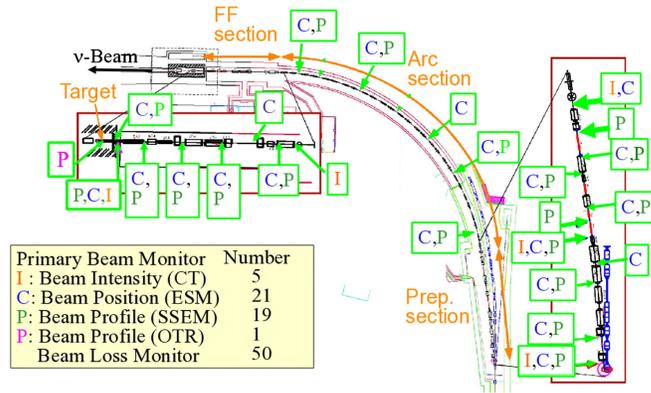


Figure 2.4: Schematic of the primary beamline indicating the location of the beamline monitors [51].

and active detector systems in and around the beamline as shown by Figure 2.4. Beam intensity is measured by 5 Current Transformers (CTs), beam position by 21 Electrostatic Monitors (ESMs), beam profile by 19 Segmented Secondary Emission Monitors (SSEMs) and losses by 50 Beam Loss Monitors (BLMs).

Secondary Beamline

As the proton beam enters the secondary beamline it reaches the Target Station as seen on the left hand side of Figure 2.5. The beam initially passes through two 0.3 mm thick titanium-alloy windows which separate the vacuum of the primary beam line from the Target Station. The beam then enters the baffle, a 1.7 m long graphite block with a 30 mm diameter hole running down its length. Beyond this lies the Optical Transition Radiation (OTR) Monitor and then the target itself.

The OTR is designed to precisely measure the position and direction of the proton beam just prior to it striking the target. A thin titanium-alloy film is held in the beam line at 45° to the beam direction, as the proton beam passes through this film light is emitted due to the differing dielectric constants of the materials it crosses. The emitted light is then collected and directed by a series of parabolic mirrors through the radiation shielding surrounding the target station to a camera that images the beam. The beam imaging from the OTR, along with measurements from the SSEMs and ESMs allows the position and direction of the proton beam to be measured to an accuracy of less than 1 mm and 0.5 mrad

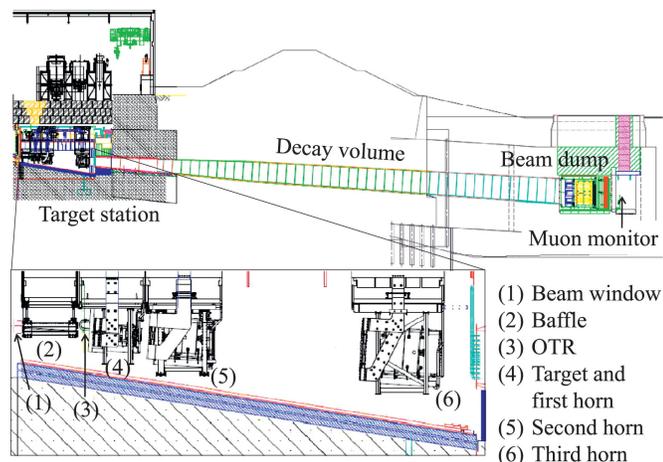


Figure 2.5: Schematic of the secondary beamline indicating its sub components [51].

respectively, as it strikes the target [54]

The target core is a graphite rod, 91.4 cm long, 2.6 cm in diameter and with a density of 1.8 g cm^{-3} ; this equates to 1.9 interaction lengths. The core is surrounded by a further 2 mm thick tube of graphite, and this is itself held within a 0.3 mm titanium case. Helium gas flows between the target layers at a rate of 250 ms^{-1} cooling the system to prevent the target from breaking down.

The proton beam, when striking the target, interacts with the carbon nuclei producing a flux of charged mesons, predominately pions but also some kaons, along with some other particles. The target sits within the upstream end of the inner conductor of the first of three magnetic horns. Each magnetic horn is made of an aluminium alloy and consists of a pair of coaxial conductors. These produce a toroidal field within each horn volume with a field strength which varies as r^{-1} , where r is the radial distance from the field centre. The first magnetic horn is designed to collect the charged mesons, directing them into the second and third horns. These subsequent horns then focus the mesons into a beam. Depending upon the polarity of the horns, positive or negatively charged mesons can be selected whilst the oppositely signed mesons are deflected out of the beam. By selecting positively (negatively) charged mesons a predominately neutrino (antineutrino) beam is produced.

The charged meson beam then enters an approximately 96 m long helium filled decay volume. The decay volume has a 1.4 m wide by 1.7 m high (3.0 m by 5.0 m) profile at its upstream (downstream) end, and is constructed of 16 mm thick steel walls surround by concrete shielding to a thickness of 6 m. Water cooling runs along the beam dump walls keeping the steel and concrete below 100 °C.

Most of the charged meson beam decays within this volume, predominately via the reaction

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

A small contribution arises from the decay of kaons, producing higher energy ν_μ via

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

$$K^+ \rightarrow \pi^0 + \mu^+ + \nu_\mu, \quad (2.2b)$$

along with the small background contamination produced mostly from

$$K^+ \rightarrow \pi^0 + e^+ + \nu_e, \quad (2.3a)$$

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu. \quad (2.3b)$$

At the end of the decay volume lies the beam dump. This is a water cooled block of graphite, 1.94 m wide, 4.69 m high and 3.17 m deep, weighing 75 tons, along with 15 iron plates with a combined depth of 2.40 m. Only muons with an energy upwards of approximately 5 GeV are capable of penetrating the beam dump and reaching the downstream muon monitor (MUMON).

2.1.4 MUMON

The T2K neutrino beam direction needs to be accurate to less than 1 mrad. To ensure this, MUMON measures the position profile of muons which are produced during the meson beam decay, and are sufficiently energetic to penetrate the beam dump. The beam direction is measured by taking the vector between the target and the measured muon profile centre.

MUMON consist of two independent detector planes, an upstream array of silicon photodiodes and a downstream array of ionisation chambers, separated by 1.2 m. The detection plane in both the silicon photodiodes and the ionisation chambers is composed of 7 by 7 sensor arrays at intervals of 25 cm. The muon profile is reconstructed in 2D by the distribution of reconstructed charge across each detector array.

MUMON is capable of a monitoring the beam direction to an accuracy of 0.25 mrad and beam intensity to a precision of 2.9% on a bunch by bunch basis [53].

2.1.5 Off-Axis Design

T2K utilises an off-axis neutrino beam configuration. This means that the angle between the beam focusing axis and far detector are offset from one another, in this case by 2.5°. This off-axis design ensures that the neutrino beam, as measured at the far detector, has a narrower spread in energy with a reduced peak energy of 0.6 GeV, than would be detected by an on-axis detector. This approach minimises the muon-neutrino survival probability at the far-detector, see Figure 2.6.

The technique was first proposed in 1995 by Beavis *et al.* as part of a proposed oscillation experiment which would use a neutrino beam produced by the Brookhaven National Laboratory in the United States [56]. The angular energy dependence on the neutrino beam comes about through conservation of momentum and energy in the two body decay $\pi^\pm \rightarrow \mu^\pm + \bar{\nu}$. For a pion decaying along the beam axis, the resultant neutrino energy is given by (assuming the neutrino to be massless)

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - |\mathbf{p}_\pi| \cos \theta)}, \quad (2.4)$$

where E_ν and E_π are the energy of the neutrino and pion respectively, m_π and m_μ are the mass of the pion and muon respectively, \mathbf{p}_π is the pion three-momentum and θ is the angle between the neutrino and the pion's direction of travel.

As θ increases the maximum possible neutrino energy decreases

$$E_\nu^{max} = \frac{m_\pi^2 - m_\mu^2}{2E_\pi \sin^2 \theta}. \quad (2.5)$$

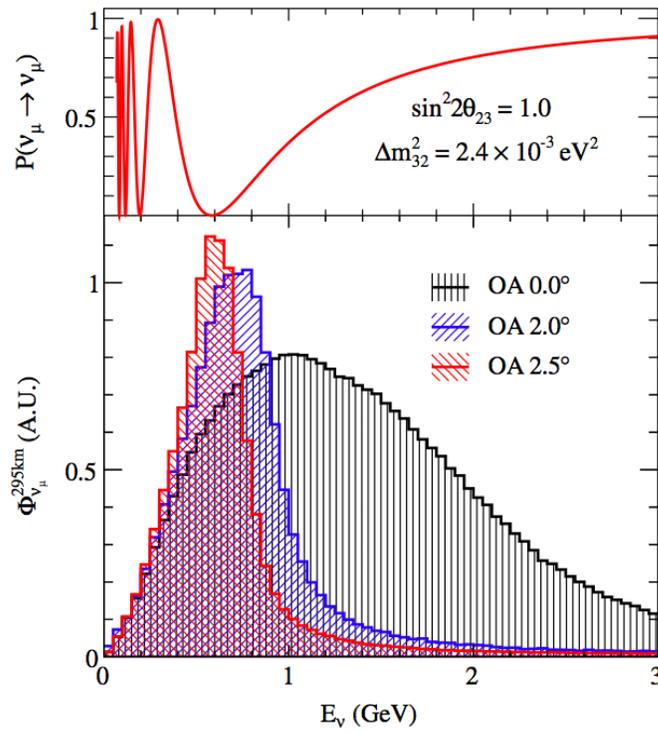


Figure 2.6: Neutrino flux spectrum as a function of energy at different off-axis angles, along with the muon-neutrino survival probability at 295 km. At a 2.5° off-axis angle the neutrino energy spectrum peaks at 0.6 GeV, corresponding to the first oscillation maximum [55].

This places a constraint on the neutrino flux and energy spectrum so that, for a given off-axis angle, a pion with $E_\pi \neq E_\pi^m$ will produce a decay neutrino with $E_\nu < E_\nu^{max}$.

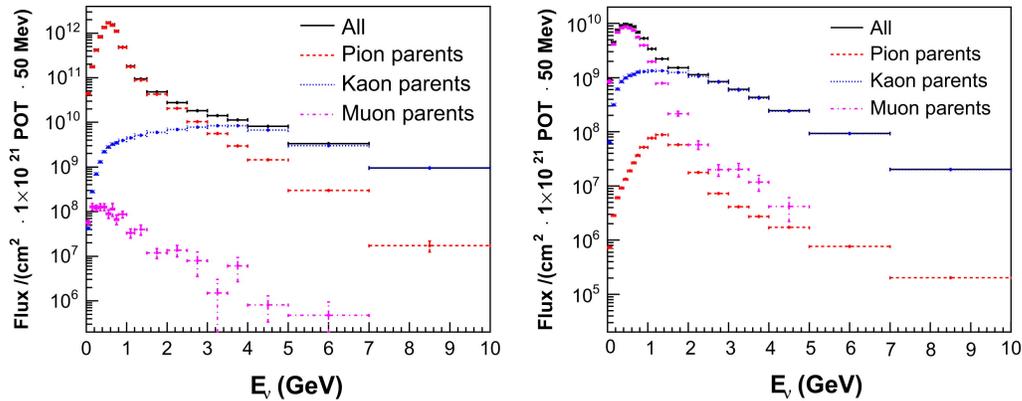
2.1.6 Neutrino Flux Prediction

The absolute neutrino flux produced by the T2K experiment is unknown, what is measured is the number of protons on target (POT). Total integrated POT is then used as a proxy measure for the total integrated neutrino flux. The flux itself is predicted through data driven Monte Carlo (MC) simulations, along with measurements of the proton beam profile and the magnetic horn fields.

The primary source of uncertainty in the neutrino flux predictions comes from the uncertainty in modelling the hadron production due to proton-nuclear interactions within the target. This is being addressed by the dedicated experiment NA61/SHINE [57] at CERN, which measures hadron production with a 30 GeV proton beam impinging upon a graphite target for the entire kinematic acceptance of T2K.

T2K has produced JNUBEAM, a MC simulation package which covers the whole secondary beamline. JNUBEAM is based upon GEANT3 [58], whilst drawing upon FLUKA [59] to model interactions within the target and GEANT3 with GCALOR [60] for interactions outside the target. Of particular importance is the in-target simulations, which although simulated with FLUKA, are weighted to the NA61/SHINE results to replicate the true neutrino flux as accurately as possible. The resultant T2K neutrino flux prediction can be seen in Figure 2.7.

Through the careful monitoring of the proton beam, magnetic horn fields, and hadron production measurements, the absolute T2K neutrino flux is predicted with an uncertainty of 15% at the peak neutrino energy. The uncertainty in the predicted far to near detector flux ratios is down to 2% at the peak neutrino energy, and 6% across all energies [55].



(a) Muon-neutrino flux prediction. (b) Electron-neutrino flux prediction.

Figure 2.7: FLUKA2008.3b neutrino flux prediction for the T2K ND280 detector for ν_μ , 2.7a, and ν_e , 2.7b, both re-weighted to the NA61 thin target data. The contribution from pion, kaon and muon decays to the total neutrino flux are each shown [57].

2.2 Near Detector Complex

Approximately 280 m from the beam target lies the T2K near detector complex which houses two detectors, INGRID and ND280. The Interactive Neutrino GRID (INGRID) is an on-axis neutrino beam and direction profiling detector. The Near Detector at 280 m (ND280) is a magnetised, 2.5° off-axis detector capable of measuring the neutrino flux, energy spectrum, beam contamination (from ν_e) and specific neutrino final states for cross-section measurements on a range of target materials.

Both of the near detectors are located on the J-PARC site, see Figure 2.2, in a specially excavated cylindrical pit, 37 m deep with a 17.5 m diameter which comprises three floor levels, see Figure 2.8.

2.2.1 INGRID

The INGRID detector [61] is composed of 16 identical modules, 14 of which form the vertical and horizontal axes of a cross, seven per axis, with the central module on each axis aligned with the beam axis. The remaining two modules are placed

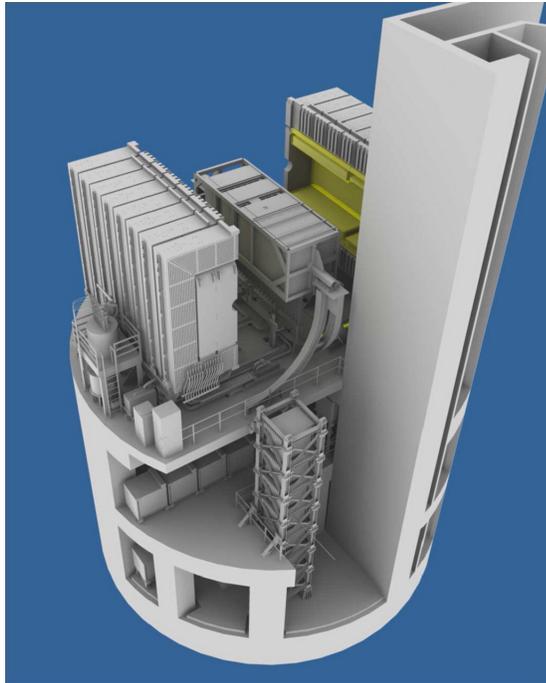


Figure 2.8: ND280 detector complex. The INGRID detector is formed of the vertical stack of modules rising from the lowest floor in the foreground, spanning the lower and central floors, along with the horizontal span of modules on the central floor behind the vertical stack. The ND280 detector is located on the top floor, shown in the magnet open configuration. The magnetic coils and flux returns are opened to each side of the pit revealing the central basket portion of the detector [51].

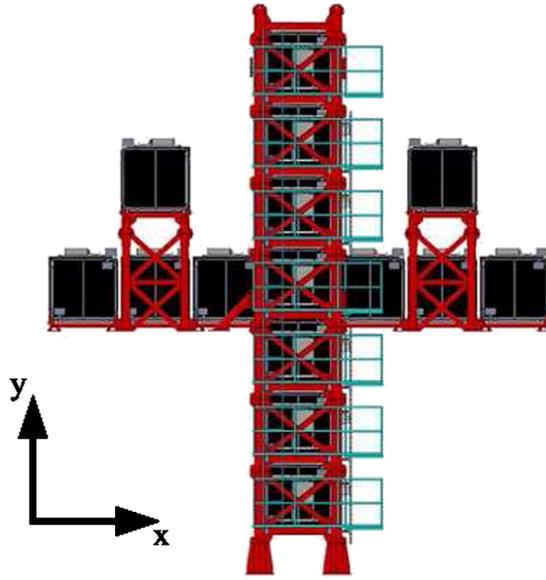


Figure 2.9: Position of the INGRID detector modules. Seven modules are used in both the vertical and horizontal branches of the detector cross, these monitor the beam position. A further two modules are placed symmetrically off the cross axis to measure the axial symmetry of the beam [51].

symmetrically off-axis from the main cross to measure the beam's axial symmetry, see Figure 2.9.

Each INGRID module is composed of 9 iron and 11 scintillator planes sandwiched between one another, with no iron plate being placed between the 10th and 11th scintillator planes. Each iron plane is square with a length of 124 cm on each side and a thickness of 6.5 cm, this gives a total target mass of 7.1 t of iron per module. Each of the tracking scintillation planes is composed of two layers of plastic scintillator bars, with one plane orientated horizontally, the other vertically. 24 bars are used in each layer, and an individual scintillator bar has a cross sectional area of 1 cm by 5 cm. Each module is surrounded by a veto plane on each side to reject incoming particle tracks. The veto planes are constructed of 22 scintillator bars aligned along the beam axis, see Figure 2.10.

The scintillation bars were produced at Fermilab [62] and are made of extruded polystyrene doped with 1% of 2,5-diphenyloxazole (PPO) and with 0.03% of 1,4-di-(5-phenyl-2-oxazolyl)-benzene (POPOP). Each bar is surrounded by a co-

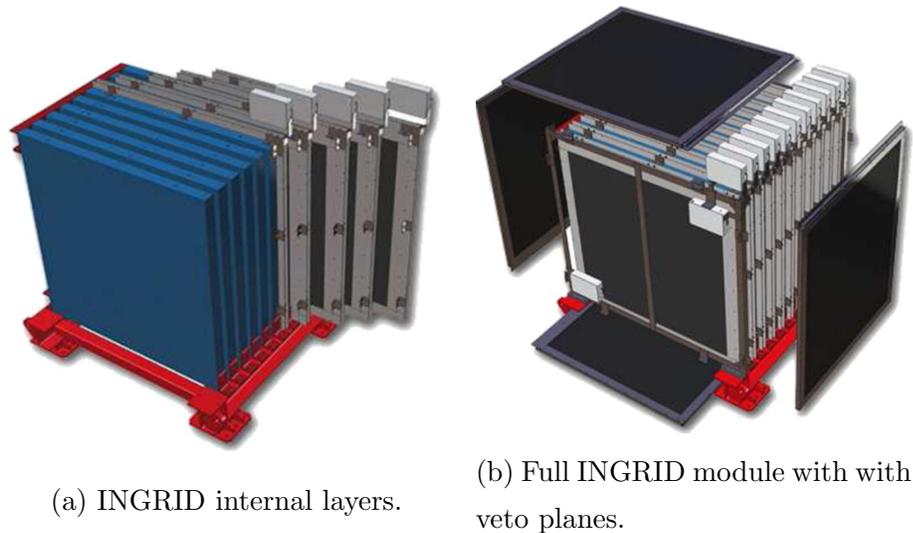


Figure 2.10: Exploded view of an INGRID module. 2.10a shows the iron planes in blue, separated by the grey scintillator planes. 2.10b shows a full module, with the surrounding veto planes shown in black [51].

extruded white reflective jacket made of polystyrene infused with TiO_2 . A 3 mm diameter hole runs the length of the scintillator bars for the insertion of a wavelength shifting (WLS) fibre which extracts the scintillation light produced in the bars. The WLS fibres are Kuraray Y11(200)M and have a diameter of 1.0 mm which couples to Multi-Pixel Photon Counters (MPPCs) for electronic read out. MPPCs are described in detail in Section 2.2.2. Only one end of each WLS fibre is readout by an MPPC, with the uninstrumented end of each fibre and bar being painted with a reflective coating of ELJEN[®] EJ-510 to increase total light yield.

For the beam power currently achieved, INGRID has measured the beam centre position on a month by month basis. This was found to be consistently stable within 28 cm, corresponding to beam direction being known to within 1 mrad as required for precision oscillation measurements. The neutrino event rate is also measured, this time on a day by day basis, and this is stable within statistical uncertainty which is typically 1.7% [61].

In addition to the standard INGRID modules, is an iron-free, higher granularity scintillator module. This is known as the proton module and is placed between the two central modules on each axis of the cross. The proton module is designed

to measure both muons and protons from charged current interactions through its higher granularity tracking capabilities. The desire is to isolate quasi-elastic interactions in this module to compare with Monte Carlo simulations.

2.2.2 ND280

The ND280 lies slightly downstream of INGRID within the detector pit and is offset from the beam axis by 2.5° . The detector is housed in the refurbished magnetic coils and flux return yokes previously used by the UA1 [63] and NOMAD [64] experiments (kindly donated by CERN). It is composed of two regions, the upstream π^0 -detector (P0D) region and downstream tracker region, both of which consist of several individual sub-detectors as seen in Figure 2.11.

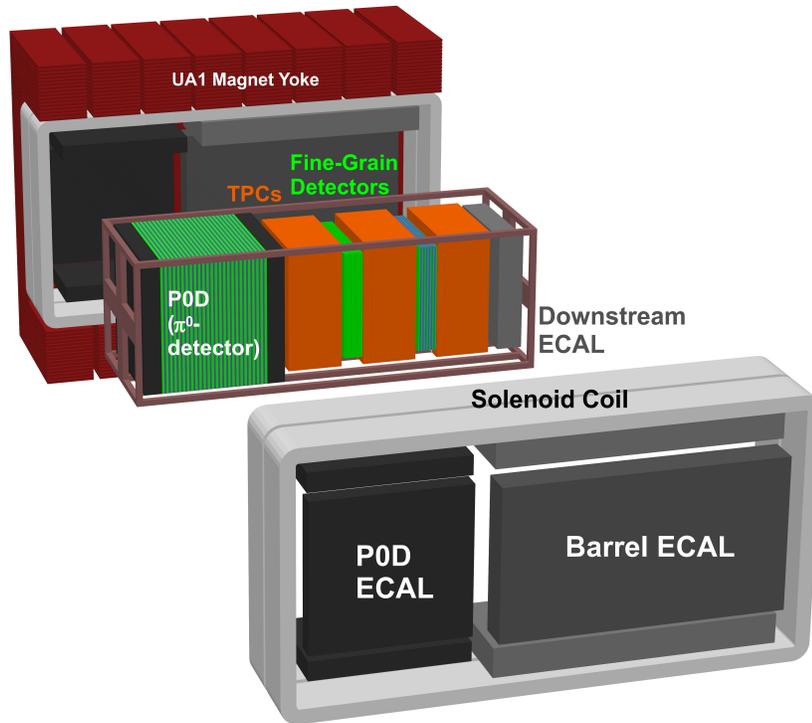


Figure 2.11: Exploded view of the ND280 detector [51].

The ND280 has external dimensions of 7.6 m (l) by 5.6 m (w) by 6.1 m (h), as constrained by the size of the magnetic yoke. The internal basket is 6.5 m by

2.6 m by 2.5 m and is supported from each end by two beams which pass through the former beam pipe holes at each end of the magnetic coil. In this way the two magnet clams can be opened up allowing access to the internal detectors for installation and maintenance work, and then securely closed around the basket during experimental running.

The magnetic coils of the ND280 are held within the external magnetic yoke, with the Barrel Electromagnetic Calorimeters (BrECals) and π^0 -detector Electromagnetic Calorimeters (PØDECals) being attached to the internal surface of the magnetic coils. The π^0 -detector (PØD) resides within the upstream end of the basket region and is surrounded by the PØDECals when the magnet is closed. Behind the PØD are the tracking sub-detectors; these are the three Time Projection Chambers (TPCs) which are separated by the two Fine Grained Detectors (FGDs). The Downstream Electromagnetic Calorimeter (DsECal) makes up the downstream face of the basket and is the final component of the tracking region, with the BrECals then surrounding the entire tracking region when the basket is closed.

Magnet

The ND280 magnet produces a 0.2 T dipole magnetic field; this allows high resolution measurements of the charge and momenta of charged particles within the detector's tracker region. The field is produced by four water-cooled aluminium coils which sit, two per side, within the the return yoke. Each side of the return yoke is known as a clam and these are mirror-symmetric. Each yoke is segmented into eight C-shaped elements, each of which is made of low-carbon steel plates. 16 steel plate layers make up each yoke element, these are 48 mm thick and separated by 17 mm air gaps [65]. The total mass of the magnet and yoke is 850 t, dominating the total mass of the ND280 and therefore the site of the vast majority of neutrino interactions within the detector.

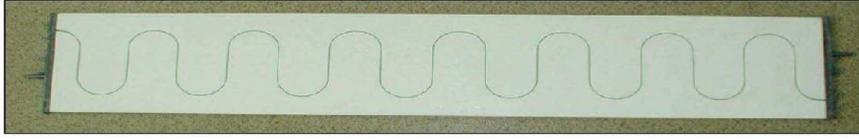


Figure 2.12: Image of an SMRD paddle prior to installation with the characteristic snaking WLS fibre running through it [65].

SMRD

The Side Muon Range Detector is a series of 440 scintillation paddles that are placed within some of the air gaps of the magnetic yoke. The SMRD has three primary functions; firstly it measures the momentum of muon tracks exiting the ND280 at high angles, secondly it acts as a veto for particles entering the detector from the outside and neutrino interactions within the magnetic yoke and surrounding rock, and finally it works as a trigger for incoming cosmic rays which are used in detector calibration and validation studies.

For the eight yoke elements that make up each side of the clam, the top and bottom portion of each yoke element have the three most internal layers (closest to the magnetic coil) instrumented with the scintillator paddles. Moving from upstream to downstream, the first five yoke elements have the three most internal layers of their sides instrumented, the sixth element has the four most internal layers instrumented, and the seventh and eighth layers have the six most internal layers instrumented. This bias in read out towards the downstream end of the detector gives better reconstruction for tracks in the forward going direction.

Each scintillator paddle is composed of the same combination of scintillating material, wavelength shifting fibre and read out as previously described for the INGRID detector, and is used by all the other scintillator sub-detectors within the ND280. The SMRD is unique though in the size of the scintillator paddles that it uses, with each being 7 mm deep, 875 mm long and with a width of 167 mm or 175 mm. This large width means that the WLS fibre snakes through the paddles as can be seen in Figure 2.12 and is read out from only one end.

PØD

The π^0 detector is the most upstream sub-detector within the ND280 and is primarily tasked with detecting neutral-current neutrino interactions with associated π^0 production ($\text{NC}\pi^0$). An accurate measurement of the $\text{NC}\pi^0$ cross-section on water is important, as this would constrain the uncertainty in the ν_e appearance studies at the T2K far detector, for which $\text{NC}\pi^0$ are a major background.

The PØD has an active region measuring 2103 mm wide by 2239 mm high by 2400 mm deep, and is built up of layers of scintillator bars, brass or lead sheets, and high density polyethylene (HDPE) water bags [66]. The target water bags can be filled and emptied as required, changing the detector mass between 15,800 kg and 12,900 kg, allowing a water-in/water-out $\text{NC}\pi^0$ rate difference to be calculated, leading to an on-water cross-section measurement.

The upstream and downstream ends of the PØD are calorimeter regions, these are composed of alternating direction (x and y plane) scintillator bars and lead sheets. The centre of the PØD is then the water target region, this has a similar design, but substitutes the lead for brass and an additional water bag layer, see Figure 2.13. Throughout the PØD triangular scintillator bars, each 33 mm wide by 17 mm high, are used for increased tracking precision. Each bar is read out from a WLS fibre at one end, with the opposing end being mirrored; the scintillator material, WLS fibre and read out is the same as the other ND280 scintillator sub-detectors.

Separating the detector out into the water target and calorimetry regions allows good containment of electromagnetic showers emanating from the water target region, whilst vetoing interactions occurring in other regions of the ND280. A similar function is also played by the PØD Electromagnetic Calorimeter (PØDECal) which surrounds the PØD; the PØDECal will be discussed in detail shortly.

FGD

The tracker region of the ND280 uses two Fine Grained Detectors as the target mass which are sandwiched between three Time Projection Chambers. The FGDs are each 1.1 t and have external dimensions of 230 cm wide by 240 cm high

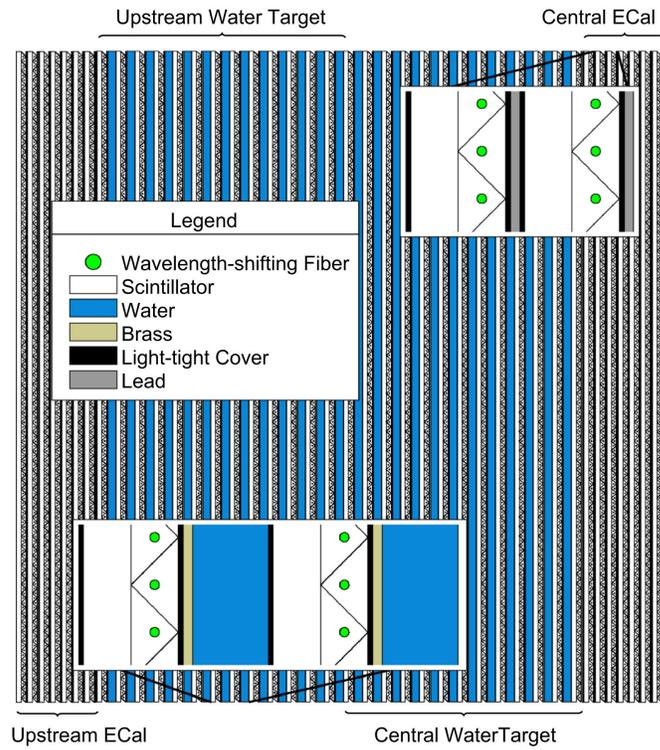


Figure 2.13: Cross sectional view of the PØD showing the position of the upstream and downstream calorimeters and intervening water target region along with the position of the triangular shaped bars used through the sub-detector [66].

by 36.5 cm depth [67]. The upstream of the two FGDs is called FGD1 and is composed purely of scintillator, the downstream FGD, FGD2, is composed of scintillator sandwiched around water layers. This gives each FGD approximately $0.85X_0$ of material depth.

Like the PØD, the FGD uses single ended read out from the scintillator bars with the opposing end being mirrored for increased light yield. The bars have a square end, 9.61 mm by 9.61 mm, and are all 1864.3 mm in length; the scintillator material, WLS fibre and read out is the same as the other ND280 scintillator sub-detectors. The bars form layers in the x and y direction, with each layer made up of 192 bars. A pair of layers, one in x and one in y , is called a module and there are 15 modules in FGD1 and 7 in FGD2.

The presence of water in FGD2 is important for making on-water cross-section measurements for a range of neutrino interaction final states. These can then be used in limiting the uncertainty on the predicted event rate at the far detector. The water is held within target layers that have a thickness of 2.5 cm and are made of rigid, corrugated Sunlite[®] polycarbonate panels. The water in the target layers are held at below atmospheric pressure to prevent any water leaks that may occur from seeping into the detector.

The FGD characteristics of continuous read out, the highest precision timing of any sub-detector, completely active detector material, and a high enough granularity to allow highly accurate vertexing and track reconstruction allows very precise final states to be selected. This will allow a vast array of electron- and muon-(anti)neutrino cross-sections to be measured by the ND280 on water and scintillator target materials.

TPC

Surrounding the FGDs are the three Time Projection Chambers of the ND280 tracker region. The TPCs are all identical in construction and are named TPC1, TPC2 and TPC3 moving from the upstream to the downstream ends of the tracker. The TPCs are designed to precisely track charged particles in three dimensions that emerge from neutrino interactions in the FGDs. The 0.2 T mag-

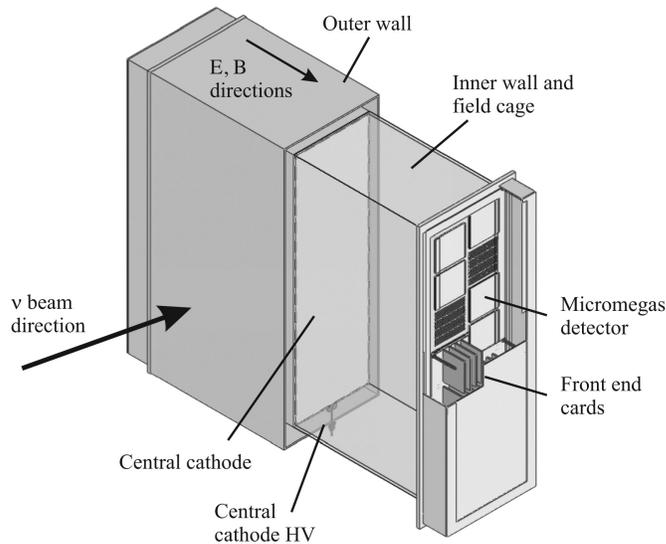


Figure 2.14: Simplified cut-away diagram of a TPC showing the detection volume, cathode components, read out components, neutrino beam direction, and electric and magnetic field direction [68].

netic field curves the path of charged particles as they pass through the TPC chambers, this allows the momenta of particles, along with the energy loss as a function of distance travelled (dE/dx), to be measured. Combining this information with the known particle energy loss rates as a function of momentum allows the TPC to produce particle identification hypotheses as described in Section 3.3.1.

Each TPC is built of an inner and outer volume, with the outer volume being 2302 mm wide by 2400 mm high and 974 mm deep [68]. The inner volume contains the Argon-based drift gas ($\text{Ar}:\text{CF}_4:\text{iC}_4\text{H}_{10}$ in the ratio 95:3:2), whilst the outer volume contains CO_2 as an insulator. The inner volume is split down the centre (in the zy plane) by the central cathode of the detector, and combined with a strip pattern machined into the copper walls of the inner-volume produces a uniform electric field across the TPC volumes which is precisely aligned with the detector's magnetic field. See Figure 2.14 for a simplified diagram of a TPC module.

As a charged particle passes through a TPC it ionises the gas, releasing electrons which drift away from the central cathode to the read out planes. The electrons

are incident upon micromegas (MM) modules which act to amplify and subsequently measure the ionisation. There are twelve MM modules for each read out plane and these are organised into two, slightly offset, vertical columns. Each MM module is 342 mm by 359 mm and is segmented into 1728 anode pads which are 7.0 mm (vertical) by 9.8 mm (horizontal) each.

The horizontal and vertical positions read out from the MM modules, combined with the hit times and known ionisation drift velocity in the modules allows the precise tracking of multiple simultaneous tracks through the TPCs, and allows matching of tracks to objects reconstructed in the other ND280 sub-detectors.

ECal

The Electromagnetic Calorimeters of the ND280 surround the detector on all faces except the upstream end of the detector, and are broken down into three primary sections. These are the Downstream ECal (DsECal), the Barrel ECal (BrECal), which together make up the tracker ECal, and the PØD ECal. The PØDECAL resides between the PØD and the magnetic coils, similarly the BrECAL resides between the tracker region of the detector (FGDs and TPCs) and the magnetic coils, and the DsECal makes up the downstream face of the basket behind TPC3. All the ECal modules are scintillator and lead sampling calorimeters which provide measurement of nearly all particles exiting the basket detectors.

The BrECAL and PØDECAL each have six constituent modules, whereas the DsECAL is just composed of a single module. All the modules use the same scintillator material, WLS fibre and read out as the other ND280 scintillator sub-detectors, but employ rectangular bars with a cross-section 10 mm by 40 mm which vary in length depending upon module and bar orientation [69]. All the ECal module materials are housed between carbon fibre sheets with aluminium support structures. The edges are then walled with aluminium sheets to which the electronics, power distribution bars, dry-air circulation and water cooling systems are attached.

The PØDECAL serves to tag particles exiting the PØD and distinguish between e^\pm/γ and μ , but is not required to do full track/shower reconstruction as this

is covered by the PØD itself. As such the PØDECal modules have the lowest depth, with only six lead-scintillator layers in each module, and all of the bars are aligned with the beam direction and read out from one end. To compensate for the low modules depth, the lead used in the detector is the thickest of the three ECal sections at 4 mm per sheet. This gives each module $4.3 X_0$ of material, this quantity was decided upon through Monte Carlo simulations which aimed to optimise photon detection efficiency, shower containment and particle discrimination [70]. The modules are arranged so that there are two above the PØD (one in each side of the clam), two below, and one larger module on each side, all of which use bars that are 2454 mm in length.

The BrECal and DsECal are designed to complement the tracking capabilities of the TPCs and FGDs, and as such are large tracking ECal modules capable of full three-dimensional reconstruction of tracks and particle showers. This means neutral particles, including π^0 decay photons, can be searched for and their energy measured, along with the charged particles from interactions in the tracker region. The BrECal modules have 31 lead-scintillator layers, corresponding to $9.7 X_0$, and the DsECal has 34 layers, corresponding to $10.6 X_0$. The lead in the tracker modules is 1.75 mm thick and this was chosen following studies of π^0 decay photon detection efficiency. The BrECal modules are laid out around the tracker region in the same manner as the PØDECal modules, but the bars in the BrECal and DsECal modules are orientated in alternating directions, rotated by 90° to one another. The bars in the DsECal are aligned alternately in the x and y directions and are all 2000 mm long and read out from both ends. For the BrECal there are long bars that align with the beam direction (z direction), which are 3840 mm long and read out from both ends, these are then alternated with shorter, single end read out bars in the perpendicular direction. In the top and bottom modules the perpendicular bars align with the x direction and are 1520 mm long, for the side modules these bars align with the y direction and are 2280 mm long.

The DsECal module was the first to be constructed and in 2009 it was shipped to CERN to be placed in the T9 testbeam to check the operation of the integrated module systems and collect a reference data sample. During the testing the module was subjected to a mixed beam of pions, electrons and protons with

momenta in the range 0.3 to 15 GeV/c (data acquired up to 5 GeV/c). This data, in conjunction with more recent control samples, has been used to tune the particle identification algorithms that are currently used by the experiment. The DsECal module was then transported and installed within the ND280 for use during T2K data taking Run 1. The remaining BrECal and PØDECAL modules were installed the following year for use during T2K data taking Run 2.

Electronics

INGRID and the ECals, SMRD and PØD sub-detectors of the ND280 all use the Trip-T electronics read out system [71] in conjunction with MPPCs (as discussed in the next section) to record data. The electronics hierarchy of the Trip-T detectors is such that individual MPPCs are readout by a Trip-T chip, these chips are capable of reading data from up to 16 individual MPPCs simultaneously. Each Trip-T chip is then read out by a Trip-T Frontend Board (TFB), a single TFB is capable of reading out up to 4 Trip-T chips. Each TFB then has its data collated by a Readout Merger Module (RMM), and a single RMM can merge data from up to 48 TFBs. The RMMs pass the data from each event to Frontend Processing Nodes (FPNs), two RMMs per FPN, which work to combine and compress the data from all sub-detectors for each event taken.

The RMMs also work to control when the data is read out from the MPPCs by issuing read out triggers down the electronics hierarchy. The Master Clock Module (MCM) registers triggers issued by various sources and fans them out through optical connections to the Slave Clock Modules (SCMs), of which there is one for each sub-detector, that controls the triggering of the RMMs, again through optical connections. It is possible to run any of the sub-detectors individually by making a sub-detector's SCM act like the MCM to issue triggers directly to that sub-detector. This is used when running specialised calibration runs for a specific sub-detector, tracking down detector problems or testing new settings.

Triggers are issued when a signal is received from the beamline GPS clock system which indicates that a beam spill has been sent; these triggers have priority over all other triggers. Outside of beam spills, the MCM cycles through a range of other triggers such as pedestal read out, light injections and cosmic ray triggers. A

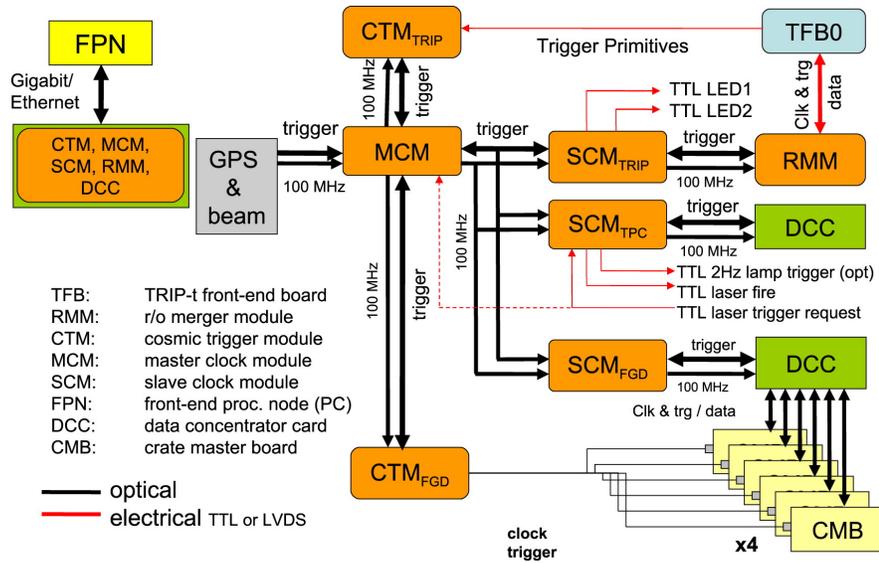


Figure 2.15: Overview of electronics components and their connections for the ND280 [51].

separate board, the Cosmic Trigger Module is directly linked between the Trip-T detector RMMs and the MCM to issue look back triggers to read out events that it decides are cosmic ray events based on coincidence selection criteria between two cosmic trigger RMMs. See Figure 2.15 for a schematic view of the electronics hierarchy.

When a trigger is received by an MPPC the number of photoelectrons registered within an integration period, after a threshold of 2.5 to 3.5 photoelectrons is reached, is recorded by two electronics read out channels. The high gain and low gain read out channels saturate at 50 and 500 photoelectrons respectively. The read out integration period is programmable and set to 480 ns during nominal running, with a subsequent 100 ns read out dead time. The high gain channel is used to register the timing of an event when the minimum photoelectron threshold is exceeded, this is done by a 2.5 ns clock on the TFB boards and sets the timing resolution of the individual channels.

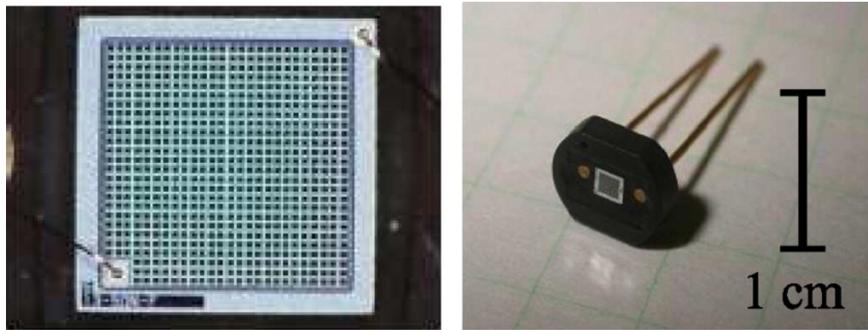


Figure 2.16: Photograph of an MPPC and its active pixel face. The sensitive face is 1.3 by 1.3 mm^2 , and the 26 by 26 grid of individual pixels and electrode attachments can be clearly seen [72].

MPPCs

T2K is the first experiment to deploy the large-scale use of the Hamamatsu Multi-Pixel Photon Counter [72]. Each MPPC has an array of 667 Silicon avalanche photodiodes operating in Geiger mode. The photodiode pixels are arranged into a 26 by 26 square, with 9 pixels lost to a electrode on the sensitive face, see Figure 2.16. Their novel design is ideal for use in the ND280 and INGRID detectors as they are compact, insensitive to the magnetic field of the ND280 and well matched to the emission spectra of the WLS fibres (420 nm absorption leading to 476 nm emission to MPPCs [51][73]).

Each MPPC pixel is a reverse-biased diode, held between 0.8 and 1.5 V above the breakdown voltage. When a photoelectron hits a pixel an ionisation avalanche is produced inducing a large reverse current with a gain of about $(0.5 - 1.5) \times 10^6$. Each pixel acts as a binary read out device, so the total signal is derived from a sum over the number of fired pixels on the MPPC. Once fired a pixel is quenched and the voltage drops to the breakdown voltage before rising back to the bias voltage. It takes about 100 ps for the voltage to drop, and a further 10 ns to return back to the bias voltage.

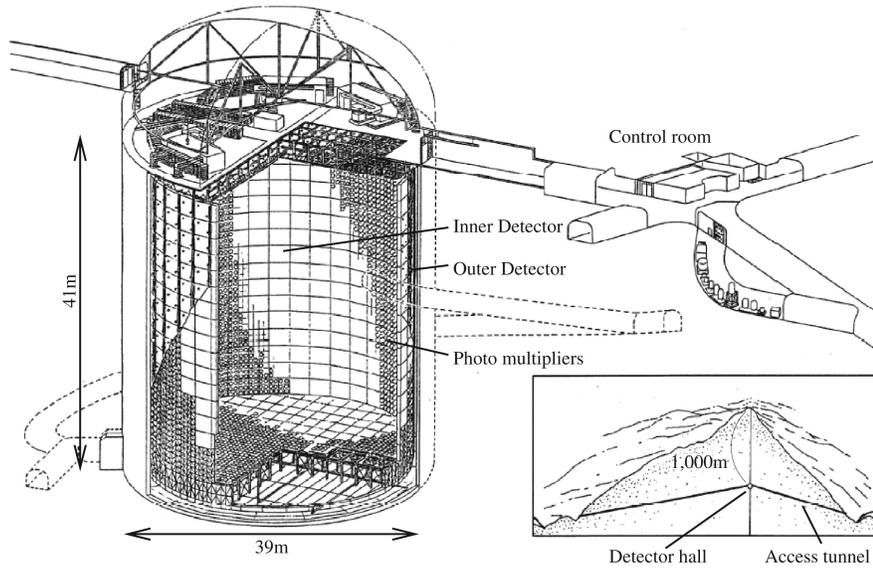


Figure 2.17: Sketch of the SK detector and its position under the Ikenoyama mountain [74].

2.3 Far Detector

The T2K experiment employs the Super-Kamiokande (Super-K, SK) detector [74] as the far detector for the experiment. Located 295 km west of Tōkai-mura, SK lies at a depth of 1 km below Ikenoyama mountain, a water equivalent depth of 2.7 km. The detector itself is a cylindrical, 50 kT (22.5 kT fiducial volume) ultra-pure water Čerenkov detector, readout by just over 13,000 Photo Multiplier Tubes (PMTs), see Figure 2.17. The detector started operation in 1996, was used as the far detector for the K2K experiment, and is famous for its contribution to neutrino oscillation results for observations of solar and atmospheric neutrinos [75][76]. Due to its long operation, SK is a well understood detector with an energy scale accuracy, and calibrated data to Monte Carlo agreement, both down to the percent level.

The SK cylinder is separated into two sections, an inner detector (ID) and outer detector (OD) volumes. The ID is 33.8 m in diameter and 36.2 m high, the OD extends a further 2 m beyond the walls of the ID, and the two volumes are separated from one another by a 50 cm thick steel scaffold wall with black panelling. The outer walls of the ID are instrumented with 11,129, 50 cm diameter, R3600

hemispherical Hamamatsu PMTs. These provide a total coverage of 40% and have a photon detection efficiency of 20%. The inter-PMT walling is covered in black plastic sheeting to prevent light reflecting within the ID and the transmission of light from the ID to the OD. The ID fiducial volume requires that any event vertex be reconstructed a minimum of 2 m from any internal wall.

The OD of SK is instrumented on the inner wall with 1,885 outward facing Hamamatsu PMTs, each with a 20 cm diameter. These provide 7% detector coverage, so to maximise light yield the outer wall of the OD is coated with a highly reflective coating. The OD is there to act as a veto for particles entering the ID from outside interactions and cosmic rays. Even with its sparse coverage, the OD veto is almost 100% effective in rejecting these external backgrounds.

SK uses cosmic ray particles and a laser light injection system to calibrate the detector. T2K neutrino beam event read out is triggered by a GPS timing signal from the neutrino beamline. All events within 1 ms either side of this beamline signal are read out, with further data reduction then applied during the analysis chain eventually leading to the measurement of signal neutrino events.

2.3.1 Čerenkov Radiation for Particle Identification

As a particle passes through a medium, if it exceeds the speed of light within that medium, it produces a cone of Čerenkov radiation. During a charged-current neutrino interaction within SK, a proton and electron or muon are produced, the charged lepton produced being dependant upon the interacting neutrino type. Other final state interaction effects are ignored in this example, that is to say we assume a charge-current quasi-elastic (CCQE) event type. By looking at the Čerenkov light emitted by the charged lepton it is possible to record the particle's energy and produce a particle identification hypothesis.

As a muon produces Čerenkov radiation, its relatively high mass prevents its trajectory from being significantly deviated as it traverses the SK detector. As such the Čerenkov radiation ring it produces, as detected by the ID PMTs, has a sharp, well defined edge to it. On the other hand, as an electron traverses the SK detector, its relatively low mass means the electron will multiply scatter,

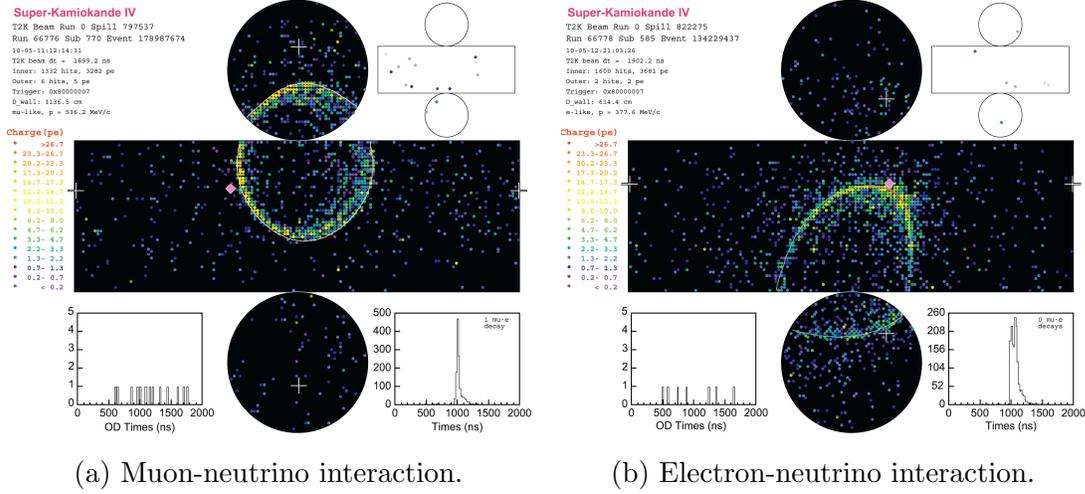


Figure 2.18: Example SK event displays from a muon-neutrino interaction, 2.18a, and electron-neutrino interaction, 2.18b, showing the characteristic sharp and “fuzzy” Čerenkov radiation rings of these types of interactions [51].

and at the energy of the T2K beam, induce an electromagnetic shower in the water medium. These effects cause an electron to produce multiple, overlapping rings of Čerenkov radiation, giving a “fuzzy” edge to the final Čerenkov radiation ring, as detected by the ID PMTs. Over the lifetime of SK, algorithms have been produced which can distinguish between these two Čerenkov radiation ring formations with very high efficiency, examples of each type can be seen in the event displays of Figure 2.18.

By integrating the photoelectrons collected by the PMTs, it is also possible to get a measure of the incident neutrino energy as the Čerenkov photon yield is related to the total lepton energy, as described by the Frank-Tamm formula [38]

$$\frac{d^2N}{dx d\lambda} = \frac{2\pi\alpha z^2}{\lambda^2} \left(1 - \frac{c^2}{v^2 n^2(\lambda)} \right), \quad (2.6)$$

where $\frac{d^2N}{dx d\lambda}$ is the number of photons produced per unit length per unit wavelength, z and v are the charge and velocity of the radiating particle, n is the refractive index of the medium, and α is the fine-structure constant.

This allows binned neutrino energy spectra, for both electron- and muon-neutrino interactions to be produced.

It is sometimes possible to detect the electron Čerenkov radiation ring produced

by the Michel electrons produced during the decay of muons which stop within the detector volume. π^0 particles can also sometimes be reconstructed by SK, as the two photons produced by its decay induce separate electromagnetic showers, which produce electron-like Čerenkov radiation ring formations. These rings can then be combined to reconstruct the original π^0 particle. However, if one of the decay photons is not sufficiently energetic to produce an electromagnetic shower above the Čerenkov threshold, or the two electron-like Čerenkov radiation rings indistinguishably overlap, then the event may be accepted as a electron-neutrino induced candidate event. Therefore π^0 interactions are one of the most significant backgrounds to electron-neutrino appearance searches by the T2K experiment.

Chapter 3

The ND280 Software Suite and Calibration Techniques

The ND280 offline software packages form a coherent suite of tools that takes raw event data, or produces Monte Carlo (MC) simulation data, applies associated calibrations, reconstructs particle interactions and saves the final data in a user friendly format; Figure 3.1 shows the package structure. The software packages are written in C++ and built around the ROOT [77] and GEANT4 [78] software libraries, and employ predominantly Python scripts to control the work flow.

3.1 Data and Monte Carlo Processing Chain

3.1.1 ND280 Data Processing

The ND280 and INGRID detectors write data to disk in the MIDAS [79] format. This data is then saved on tape at KEK in Tsukuba, Japan and distributed across the GRID to the Tier-1 computing centres at the Rutherford Appleton Laboratory (RAL) in Didcot, UK and TRIUMF in Vancouver, Canada. Data files are then converted into the “oaEvent” format by the oaUnpack ND280 software package. oaEvent format files are used throughout the ND280 software processing chain and store all data in ROOT files as various types of TObject that have been

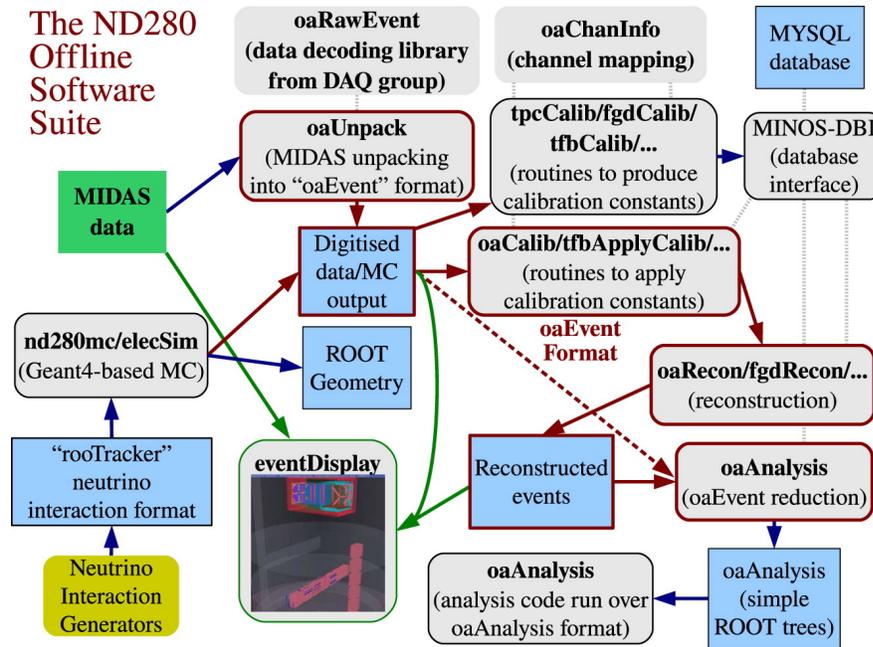


Figure 3.1: The ND280 offline software suite structure the most important or representative packages are shown [51].

specifically designed for the ND280 software suite.

The unpacked events are then passed through the ND280 calibration chain, controlled by the `oaCalib` package. `oaCalib` calls upon several sub-detector specific calibration methods, along with globally utilised calibration methods, to apply corrections to the data. Calibration corrections can be time-independent, such as channel mapping corrections, which are hard coded into the software; or time-dependent, such as timing offsets or charge response, which are stored in a MySQL database and vary with experimental running period.

Post-calibration the data is then fed through to the ND280 event reconstruction software, as controlled by the `oaRecon` package. `oaRecon` calls upon the independent sub-detector specific reconstruction methods, which aim to reconstruct tracks, showers and vertices within each event. Subsequently the objects reconstructed in each sub-detector are combined and refitted with the `RecPack` [80] software package, developed at CERN, to produce globally reconstructed objects which cross multiple sub-detectors. Both sub-detector level and globally reconstructed objects are retained by the event reconstruction, and it is possible to

switch between both levels of reconstruction for any reconstructed objects at the analysis stage.

Finally the calibrated and reconstructed data is passed through the `oaAnalysis` package, which produces the final user friendly output files. `oaAnalysis` removes all excess information from the `oaEvent` files and greatly reduces the size of resultant files, leaving them containing only ROOT-based physics objects. The resultant `oaAnalysis` files can then be analysed by the end user, independently of the rest of the ND280 software suite.

3.1.2 ND280 Monte Carlo Processing

As described in Section 2.1.6, a neutrino flux prediction is produced using the JNUBEAM simulation. This neutrino flux is extrapolated to the ND280 where neutrino interactions are simulated on the various nuclei within the detector using both the GENIE [81] and NEUT [82] MC generators. The final state particles emerging from the nucleus at the interaction vertex are then propagated through the detector using the `nd280mc` software package, via the GEANT4 interaction libraries. The trajectories, energy deposition and interactions of all resultant particles above an energy threshold of 25 MeV, and that interact, or produce daughter particles which interact within the active regions of the ND280 are stored in an `oaEvent` format file.

After the MC event propagation has been executed the resultant files are passed through the `elecSim` software package. This bespoke software simulates the electronics response of the detector to the energy deposition within the sensitive components of the detector, as produced by `nd280mc`. It also includes the simulation of detector noise, read out saturation, light attenuation in the bars and fibres of the scintillator detectors, and drift of resultant ionisation within the TPCs. The final output is then a digitised output for each sub-detector, in line with that read out by the true electronics systems. The result of this processing is to produce event files in the same format as unpacked real data. The MC simulated data then goes through the same processing as outlined above for real data, but with MC specific calibration constants applied instead.

3.2 Calibration

The calibration of the ECal sub-detector of the ND280 is broken down into two areas, the energy calibration and the timing calibration, with each having several subcomponents that must be individually accounted for as will be described in following section. The calibration of the other scintillator based sub-detectors follow very similar procedures as described here for the ECal, see [68] for some discussion on the calibration of the TPC.

3.2.1 ECal Energy Calibrations

To obtain a uniform energy response across all the ECal modules a large number of separate calibrations need to be applied. These are required to normalise the response across the entire ECal, and to ensure agreement between the data and MC; each step is described below.

Pedestal Subtraction

The first step in the energy calibration for the ECal is to remove the pedestal response from the ADC read out on each channel. The pedestal response is simply the charge that accumulates on each channel due to noise within the electronics read out, in combination with some additional bias introduced to avoid the lowest channels where the response is less linear. The pedestal values are continuously monitored by the DAQ system and vary in time due to the different Trip-T chip capacitors, and for each channel. Once the pedestal ADC response has been subtracted from each channel we are left with the true zero response on each channel.

Linearity Correction

The ADC response of the channels needs to be converted into a read out charge, this requires the signal to be linearised across the low and high gain read out from each channel. When the beam is off, known levels of charge are injected into the

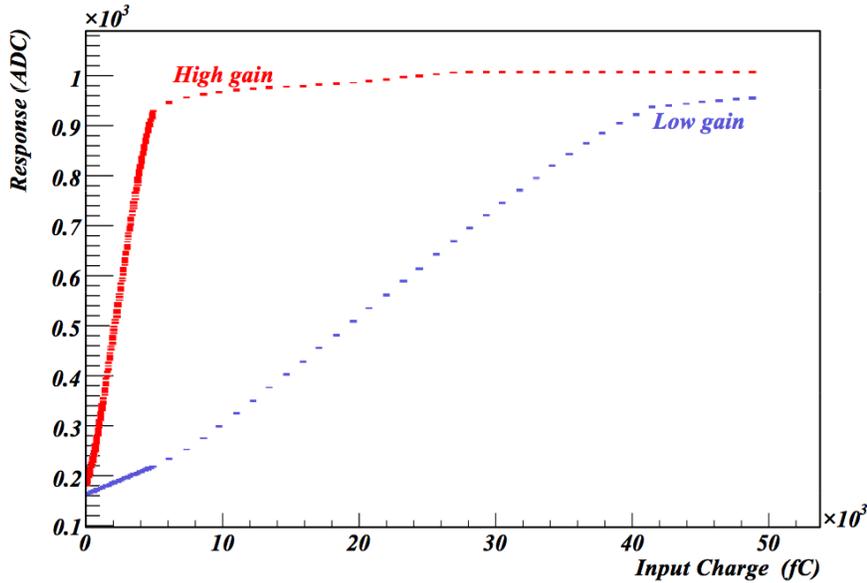


Figure 3.2: ADC response from the high and low gain read out channels as a function of injected charge [69].

capacitors of each channel by the TFB boards. Through incremental increases in the injected charge, a linearity curve can be built up, which converts read out ADC response to charge on each gain channel as shown in Figure 3.2. This also dictates which gain channel to use, depending on the read out charge from each channel.

MPPC Gain Calibration

After the channel corrections have been applied, the corrections to the response of each MPPC are made. To get an accurate measurement of the number of photoelectrons (p.e.) produced in the MPPC it is crucial to measure the gain of each device. This is calculated by measuring the difference in ADC counts between the pedestal peak and the subsequent 1 p.e. peak, produced when only one pixel on a MPPC has fired, as measured by the high gain channel. The gain of the MPPC is highly temperature dependent, this is due to the breakdown voltage on the MPPC varying linearly with temperature as approximately $50 \text{ mV } ^\circ\text{C}^{-1}$. As such, diurnal temperature variations can have a large effect on the channel gain, particularly for the DsECal, and so the gain is continuously measured along

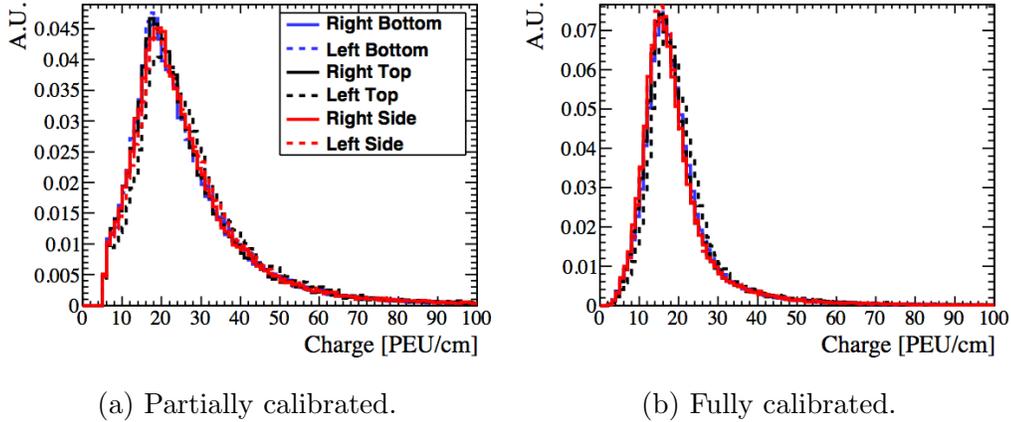


Figure 3.3: Hit charge distributions for cosmic ray muons for the six barrel ECal modules at different stages in the calibration chain. 3.3a shows the hit charge distribution after the application of the pedestal subtraction, linearity correction and MPPC gain calibration. 3.3b also includes the application of the MPPC response calibration [69].

with the pedestal level so that it can be corrected for on a three hourly time scale. Figure 3.3a shows the cosmic ray hit charge distribution after all of the above calibration steps have been applied.

MPPC Response Calibration

After the gain calibration there are a number of remaining subtle correlated noise effects that still need to be taken into account. These include the crosstalk between the individual pixels of an MPPC and after-pulsing which, like gain, varies with temperature. In addition there is the possibility of multi-photon incidence upon an individual MPPC pixel prior to its voltage being reset, along with the characteristics of the WLS fibre transmission and coupling to each MPPC, that can lead to variation in the photon distribution across the MPPC face. These effects were calculated by test bench measurements in the laboratory and produced parametrisations of the expected calibrated MPPC response to a given p.e. count, at idealised operating over-voltage for the MPPCs. Figure 3.3b shows the cosmic ray hit charge distribution after the MPPC response has also been applied.

Bar to Bar Correction

Once all channels have been corrected it is then necessary that the response across all ECal bars is uniform. This is achieved by fitting a Landau-Gaussian convolution to the Minimum Ionising Particle (MIP) energy spectrum from cosmic ray data as seen in Figure 3.3b. A correction is then applied on a bar-by-bar basis to align the fitted peaks to the average value on a per module, per bar orientation basis.

Fibre Scaling Correction

Finally, to align the response of all the different bar orientations and modules, a fibre length scaling factor is applied so that the peak MIP charges align to that measured in the DsECal. This fibre scaling factor is simply the fitted MIP peak value for the DsECal divided by the fitted MIP peak value for each combination of module and bar length individually.

3.2.2 ECal Timing Calibrations

The other stage in calibrating the ECal is correcting the timing of hits across all channels. This is important for the clustering of hits, determining the direction of tracks and showers, and rejecting noise hits. Differences in the timing of hits arise from three main sources; how they arise and are corrected for is described below.

Electronic Timing Offset Calibration

Each electronics element in the read out hierarchy is connected by electrical or optical cabling. These cables are of varying length and so lead to read out timing offsets between the different electronics elements. These offsets can be calculated by measuring the difference in expected and observed hit times on different electronics elements for cosmic ray muons. A more detailed description of this measurement is provided in Chapter 4.

Electronics Time-walk Correction

The Trip-T chips house capacitors that facilitate the creation of the time stamp for an event once a charge threshold of 2.5 p.e. equivalent has been exceeded. Every time an MPPC pixel fires the voltage across a capacitor increases leading it to charge. The more pixels that simultaneously fire, the greater the potential difference across the capacitor, and the sooner the charge threshold will be breached. This leads to the production of delayed time stamps for low charge events and is known as the electronics time-walk. By injecting incrementally increasing known charge levels into the capacitors, it is possible to produce a correlation between charge and the measured time stamp. This can then be used to apply a correction to the time stamps as a function of read out charge.

Fibre Time-walk Correction

The larger the energy deposit within the scintillator bars, the greater the number of photons that will be produced. The scintillation photons are absorbed by the fluors of the Y-11 fibres, the resultant production rate of wavelength shifted photons follows an exponential decay function. The scintillation photons will have some timing distribution, say normally distributed, around some mean value. For a low energy deposition, few photons will be produced, and so the earliest recorded photons from that interaction will have times close to the mean photon time. However for a high energy deposition, more photons are produced, and so the earliest recorded photons from that interaction are likely to come from the wings of the normal distribution, and therefore at a significantly earlier time than the mean photon time. This means that the timing of high energy interactions gets skewed to earlier interaction times and is known as the fibre time-walk. The process is purely statistical, so can vary significantly between identical energy depositions. The fibre time-walk can be analytically modelled, producing a correction which flattens the timing dependence as a function of charge.

3.3 Event Reconstruction

There are two stages to the reconstruction of events within the ND280. Initially each sub-detector group is tasked with producing and maintaining algorithms that reconstruct the individual interactions throughout their sub-detector. Subsequently, the RecPack software package is employed to, where possible, globally match reconstructed objects between the sub-detectors, such that a continuous track is produced spanning multiple sub-detectors.

A summary of the reconstruction within the tracking detectors and surrounding ECals, and their performance follows, along with some further detail on the global reconstruction implementation, as these are relevant to the analysis discussed in Chapter 5.

3.3.1 Tracker Reconstruction

Reconstruction within the TPCs and FGDs starts with track identification within each of three TPCs individually. Hits upon the micromegas modules are clustered together to form track segments when they appear close together in both space and time. Subsequent pattern recognition techniques work to match together these track segments into full tracks. A likelihood fit is then applied to the TPC tracks to calculate the track start and end position, initial direction and curvature.

The tracks reconstructed in the TPCs are then extrapolated into the FGDs using a Kalman Filter [83]. This incrementally adds hits in both the xz and yz planes to the original TPC track to produce a track which spans one or more FGD and TPC. Hits not added to the TPC tracks are then dealt with by the independent FGD reconstruction. This starts by matching hits into quasi-straight lines in the xz and yz planes independently to produce single view tracks. These single view tracks are then projected onto one another in the hope of matching and combining the two views to produce full 3D track objects.

The absolute timing of TPC tracks comes from a range of sources depending on the specifics of the track topology. If the TPC track has extrapolated into

the FGD then the t_0 of the track is obtained from the FGD. For tracks which cross the TPC central cathode, but are not FGD matched, the maximum drift time is used to calculate the t_0 , as the track is split and drift occurs in opposing directions from the cathode. Finally, tracks that are neither FGD matched nor cathode crossers obtain their timing through later matching to the other sub-detectors, or if this is not possible, a default t_0 is assigned. Knowing the t_0 of the tracks is important as it is used to provide the track starting position in the drift plane.

TPC Particle Identification

Particle identification (PID) within the TPCs is done through a measure of the energy loss per unit length of each track, $\frac{dE}{dx}$, and compared to theoretical predictions, as a function of particle momentum, for five different types of charged particle.

First the 70% of clusters with the lowest charge values are used to calculate a truncated mean charge for each track. A correction is applied to take into account the number of clusters used and the track angle in the truncated mean calculation. The 70% cut off value was selected through optimisation of MC and test beam studies. The truncated mean is then a measure of the energy loss of the charged particle within the TPC gas mixture.

A particle PID hypothesis can be formed by calculating a “pull”,

$$pull_{\alpha} = \frac{\left(\frac{dE}{dx}\right)^{meas} - \left(\frac{dE}{dx}\right)_{\alpha}^{exp}}{\sigma_{\alpha}}, \quad (3.1)$$

which measures the deviation (in terms of number of sigma) between the measured truncated mean, $\left(\frac{dE}{dx}\right)^{meas}$, and the five theoretically calculated values, $\left(\frac{dE}{dx}\right)_{\alpha}^{exp}$, where α is either a muon, electron, proton, kaon or pion, and σ_{α} is the resolution for the hypothesis α .

In Figure 3.4 theoretically calculated $\left(\frac{dE}{dx}\right)_{\alpha}^{exp}$ curves as a function of particle momentum are superimposed upon data from the first T2K data taking period. Example pull distributions for true MC muons are shown in Figure 3.5. It can be seen in Figure 3.5a how the $pull_{\mu}$ hypothesis forms a tight Gaussian distribution

about zero as it should. It is also seen how the difference in $\left(\frac{dE}{dx}\right)$ curves for electrons lead to a significant pull_e offset from zero, thus allowing muons to be distinguished from electrons at most energies in the TPCs. This is clearly shown in Figure 3.5b, where at most energies the difference in pull_μ and pull_e is significant, except at around 150 MeV where the $\left(\frac{dE}{dx}\right)$ curves for electrons and muons cross, resulting in them appearing indistinguishable in the TPCs.

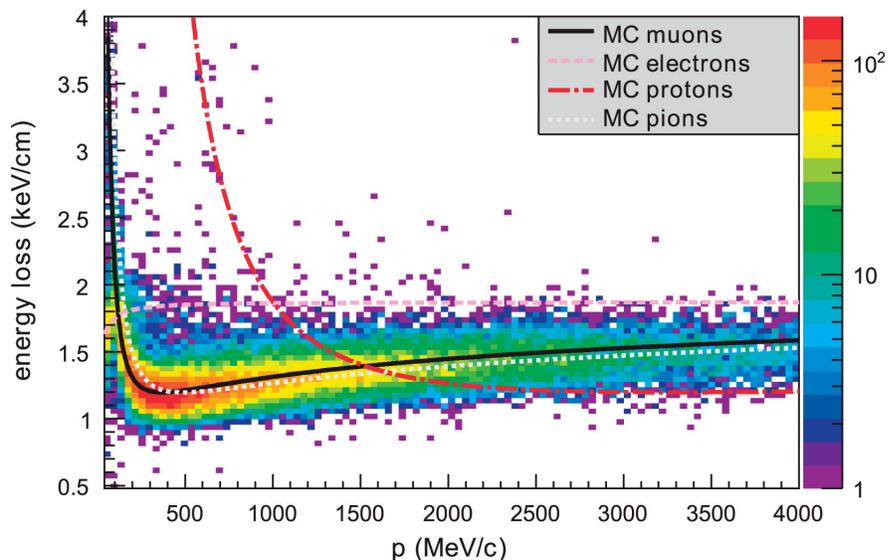
3.3.2 Tracker ECal Reconstruction

Good reconstruction of electromagnetic showers is key to searching for the decay photons from π^0 's as required by the analysis detailed in Chapter 5. The reconstruction of tracks and showers within the ECal is performed by the ecalRecon software package and is broken down into many steps as described below.

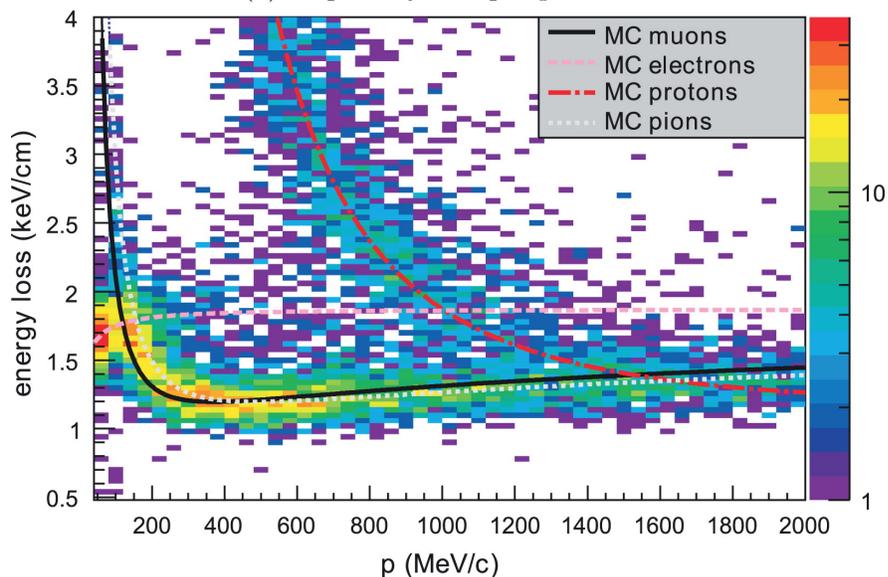
Hit Preparation

The initial stages of ecalRecon take the hits passed by oaCalib, on a spill by spill basis, and orders them in time. These time ordered hits are then searched for intervening gaps of greater than 50 ns and then split apart at these points to break the spill up into the individual beam bunches. The hits are then sorted by sensor to locate hits which occur on opposite ends of the double ended read out bars. If such hits are close in time they are combined into a single hit, and the time difference between the read out at each end is used to estimate the hit position along the length of the bar.

Subsequently a correction is applied to the hit amplitudes which removes the effect caused by light attenuation as the photons are propagated down the WLS fibres. The charge is then converted from Pixel Equivalent Units (PEU) to the MIP Equivalent Units (MEU), where 1 MEU is equivalent to the charge deposition from a MIP passing through an ECal bar at a perpendicular angle 1 m from the MPPC. The scaling factor to MEU is monitored through the collection of cosmic ray data.

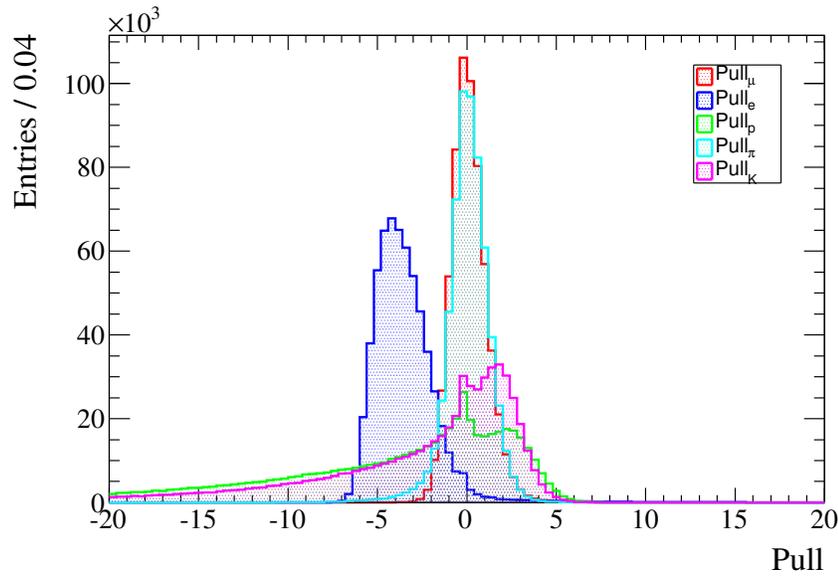


(a) Negatively charged particles.

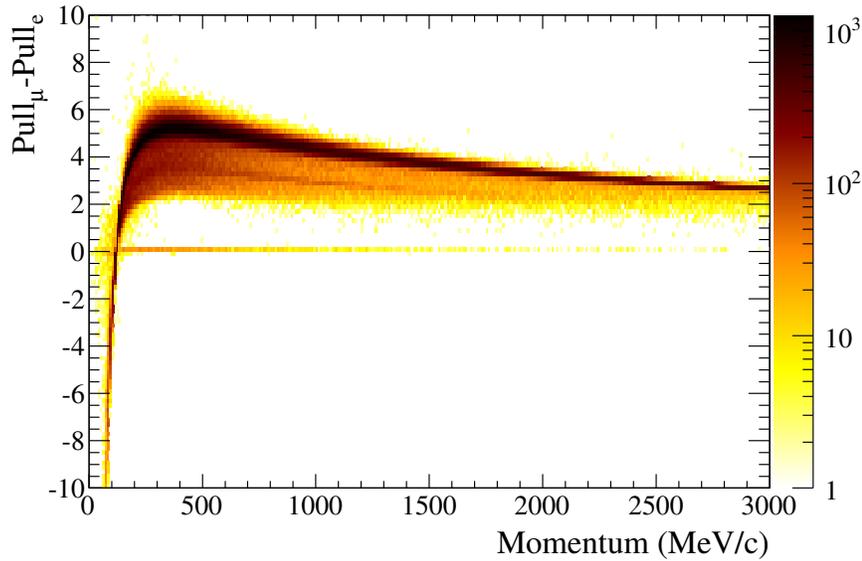


(b) Positively charged particles.

Figure 3.4: Lines showing the MC predictions for $\frac{dE}{dx}$ as a function of momentum for muons, electrons, protons and pions. Those are superimposed upon the data for negatively charged, 3.4a, and positively charged, 3.4b, particles from the first T2K experimental data taking period [68].



(a) Different particle pulls for true muons.



(b) Difference between muon and electron pulls for true muons.

Figure 3.5: Pull distributions obtained for a large MC sample of true muons. 3.5a shows the pull distributions under the muon, electron, proton, pion and kaon particle hypotheses. 3.5b shows the difference in muon and electron pull as a function of measured particle momentum in the TPCs.

Basic Clustering

Once all the hits have been prepared they are passed through a basic clustering algorithm. This selects the highest charge hit and opens up a 30 ns window around it, this is then the cluster seed. Hits within that time window are then incrementally added to the cluster if they lie within no more than one intervening bar and two intervening layers of any hit previously added to the cluster. Any cluster must contain a minimum of three hits, and this process occurs for the two separate read out planes independently, leading to the formation of an array of 2D clusters for each view in each module.

Combine Clusters

Next the basic clusters are passed to the combine clusters algorithm. This attempts to combine any showers or tracks that have not been fully formed by the basic clustering stage by combining these together to produce larger clusters. The first stage is to find the largest basic cluster, in terms of number of cluster hits, and calculate a principal direction based on a charge weighted principal component analysis (PCA) [84]. The smallest distance is then calculated between the PCA axes of the primary basic cluster and any other basic clusters. If this distance is less than 80 mm and the difference in average hit time between the two clusters is less than 40 ns the two clusters are combined into a single combined cluster. This process is repeated until no further combinations can be formed.

Expand Clusters

The final stage with the 2D clusters is to combine the unused hits left over from the basic clustering stage and add them into the combined clusters. The first stage here is to calculate the PCA axes of the combined clusters, along with the 1σ width of the cluster along both the primary and secondary axes. The distance of each unmatched hit is calculated, in units of σ , from each axis of the PCA and then combined in quadrature to give a matching weighting between each cluster and the unmatched hit. Each unmatched hit is then combined with the cluster to which it has the lowest matching weight, provided the matching weight is less

than 80 and the hit lies no more than 40 ns prior to the earliest or 40 ns after the latest hit within the combined cluster.

3D Clustering

Once all the possible hits have been combined in each view the 2D clusters need to be combined into 3D clusters, and this achieved through two approaches.

First, tracks from the tracker region are extrapolated onto the inner face of each ECal module, acting as a seed to which 2D clusters may be matched. Clusters, in each view, which contain hits within the first four layers of the ECal module are then combined to form a combination likelihood based on the relative difference in position and time between one another and the seed track.

Second, all combinations of 2D clusters from each view are combined to form a likelihood based upon the ratio of total charge within each cluster and the differences between both the innermost and outermost layers hit. If a cluster combination also has a likelihood from the seeded combination, then the two likelihoods are combined into a single likelihood.

Finally the cluster combination with the greatest likelihood is combined into a final 3D combination, this continues until all possible combinations have been considered.

3D Hit Positioning

Once a full 3D object has been formed, the unknown position of each hit can be calculated. To do this the charge weighted average position of the hits in each layer is calculated. The position of each individual hit is then found by first making a linear fit through the average hit positions from the four surrounding layers in the opposing view. The unknown coordinate for each hit is then taken as the point where the fit line crosses the layer to which the hit being considered belongs.

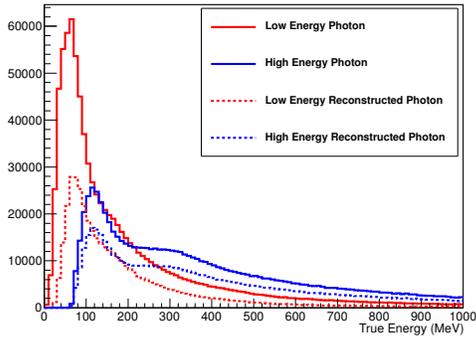
Low Energy Clustering

A final reconstruction step is called in an attempt to combine any unused 2D clusters and remaining isolated hits. This works by attempting to match single hits to clusters which lie within an adjacent layer to a hit and no more than 10 ns prior to the earliest or 10 ns after the latest hit within the combined cluster. This is particularly important for reconstructing the low energy clusters from π^0 decay photons as these frequently fail to produce the minimum of two clusters required in each view by the standard reconstruction algorithms.

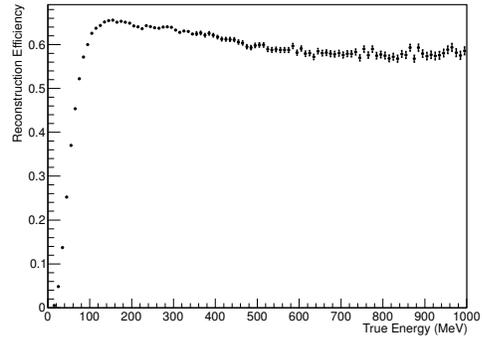
ECal Reconstruction Performance

The analysis discussed in Chapter 5 relies upon the ECal to reconstruct the photons produced during the decay of π^0 's from charged-current neutrino interactions. Figure 3.6 shows the true energy spectrum of the higher and lower energy photon produced during the decay of π^0 's produced within the FGDs. Also shown is the associated reconstruction efficiency as a function of true photon energy. Both the energy spectra and efficiency plots are shown on a per module orientation basis, using photons incident upon the Downstream, Barrel South side and Barrel Bottom ECal modules. The reconstruction within the Barrel South side is equally representative of the Barrel North Side, and likewise reconstruction within the Barrel Bottom is equally representative of the ECal Top modules. It is seen that the efficiency for reconstructing photons in the ECals is poor below approximately 150 MeV, but plateaus at about 70% in the DsECal and 80% in the Barrel ECal modules for higher energies. The low reconstruction efficiency below 150 MeV is to be expected due to the low number of hits produced in the ECal modules at such low energies, this is rather unfortunate though, as this is the energy regime in which the majority of π^0 decay photons occur.

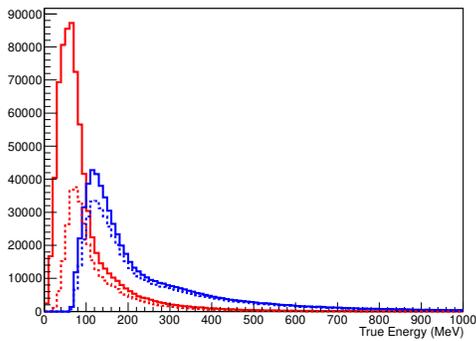
Figure 3.7 shows the true incidence angle distribution and reconstruction efficiency for the same data sample. In all cases the incidence angle is defined between the direction of the photon into the ECal module and the local z axis, where the local z axis is the one lying perpendicular to the internal face of each



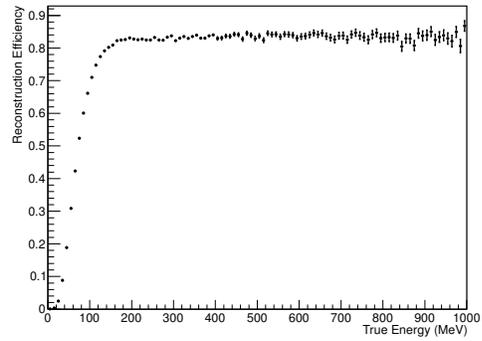
(a) DsECal true energy spectrum.



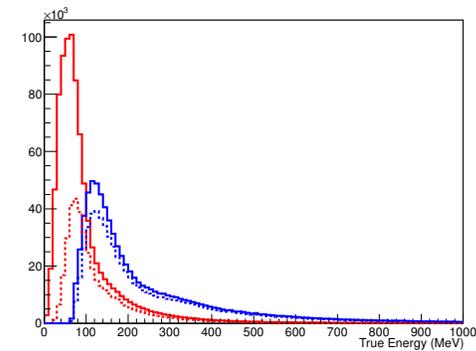
(b) DsECal reconstruction efficiency.



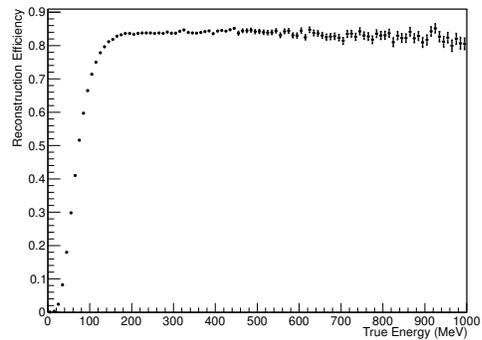
(c) South Side ECal true energy spectrum.



(d) South Side ECal reconstruction efficiency.



(e) Bottom ECal true energy spectrum.



(f) Bottom ECal reconstruction efficiency.

Figure 3.6: True energy spectra for higher (blue) and lower (red) energy decay photons from π^0 's, with the absolute (solid line) and reconstructed (dashed lines) spectra indicated. The reconstruction efficiency as a function of true photon energy is also shown. These and the energy spectra are both presented as a function of module orientation.

module (closest to the tracking detectors). The reconstruction efficiency as a function of incidence angle is fairly flat at around 60% across most of the angular distributions, falling away slowly towards the higher incidence angles.

Energy Calculation

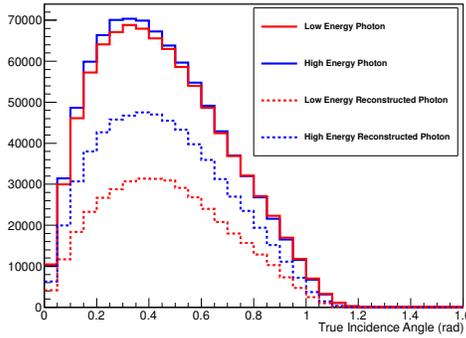
The energy of 3D clusters reconstructed in the ECals is calculated from the total cluster charge, the charge RMS and charge skewness. By minimising a likelihood function based on these input variables an estimate of the cluster energy is produced within the range 25 MeV to 20 GeV. The energy estimate to reconstructed input relation is obtained through splines which relate true energy to each of the likelihood variables. These splines are produced from an in-depth MC study which covers the entire energy range considered, but with extra focus on the most common true particle energies which are below 2 GeV.

Figure 3.8 shows the fractional difference between true energy and reconstructed energy, as a function of true energy, for the same π^0 decay photon sample as shown in Figures 3.6 and 3.7. It is seen that at lower energies, below about 200 MeV, the reconstructed energy of photons can vary significantly above and below the true value. This is unsurprising given the difficulty in reconstructing such low energy objects, where the likelihood inputs will show the largest variation between clusters and the effects of detector noise will be at their most significant.

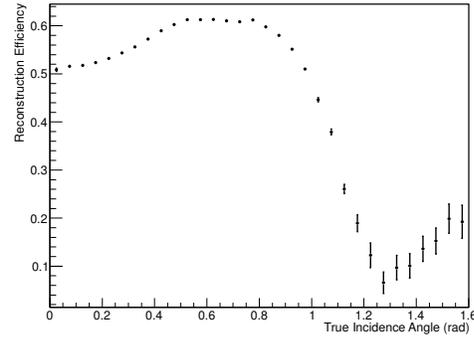
Angular Reconstruction

The angular reconstruction of ECal objects is done via two primary methods, one runs a principal component analysis over the hits within each reconstructed cluster, to obtain a direction from the major axis of the calculated ellipse. The second method, and that which is used in the analysis described in Chapter 5, relies on the calculation of the cluster “thrust”, as adapted from jet physics reconstruction in collider experiments. This was specifically adapted for use in searches for π^0 decay photons as described in [85].

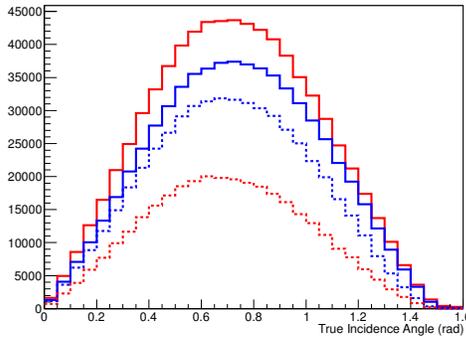
First a thrust origin is found through calculating the charge-weighted mean po-



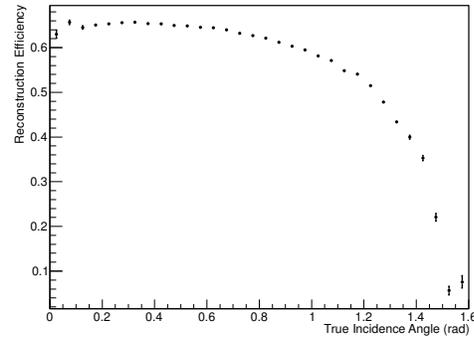
(a) DsECal true angular distribution.



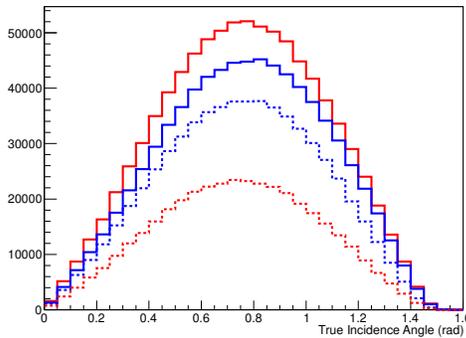
(b) DsECal reconstruction efficiency.



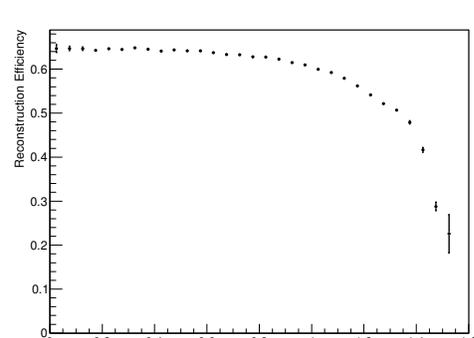
(c) South Side ECal true angular distribution.



(d) South Side ECal reconstruction efficiency.



(e) Bottom ECal true angular distribution.



(f) Bottom ECal reconstruction efficiency.

Figure 3.7: True angular distribution for higher (blue) and lower (red) energy decay photons from π^0 's, with the absolute (solid line) and reconstructed (dashed lines) distributions indicated. The reconstruction efficiency as a function of true photon incidence angle is also shown, and these and the angular distributions are both broken down by module orientation.

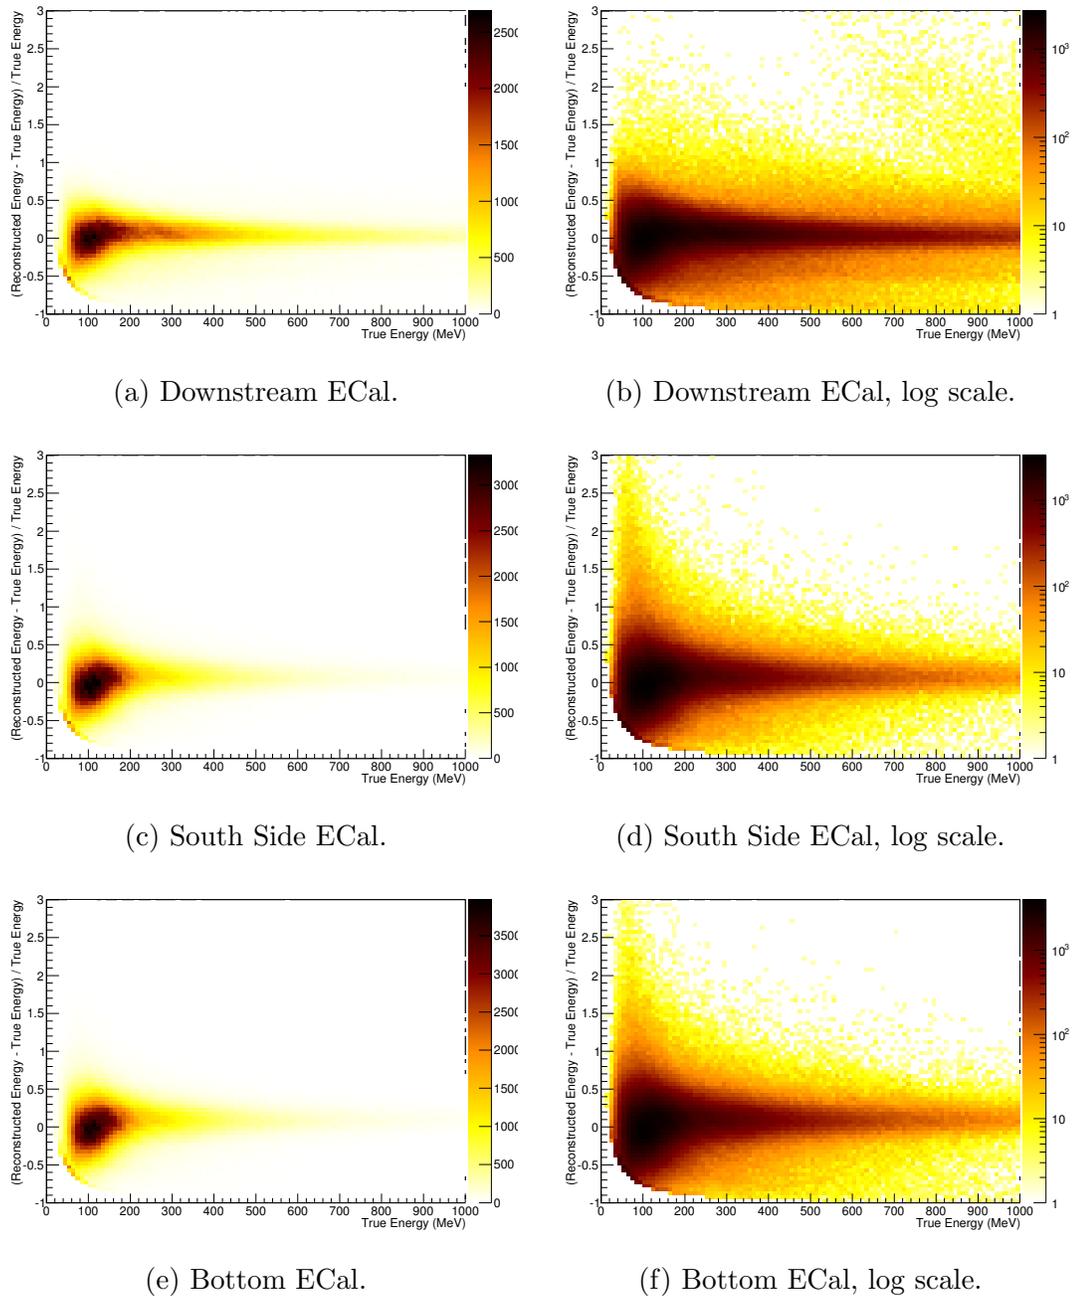


Figure 3.8: Fractional energy difference calculated between the reconstructed and true energy of π^0 decay photons as a function of true energy for three different ECal module orientations. The right hand plots are the same as the left hand plots, but with a logarithmic z axis.

sition of hits in the innermost ECal layer with cluster hits. The cluster thrust is then calculated by maximising $t(\theta, \varphi)$ in the calculation

$$t(\theta, \varphi) = \frac{\sum_i q_i |\vec{n}_{(\theta, \varphi)} \cdot (\vec{x}_i - \vec{o})|}{\sum_i q_i |\vec{x}_i - \vec{o}|}, \quad (3.2)$$

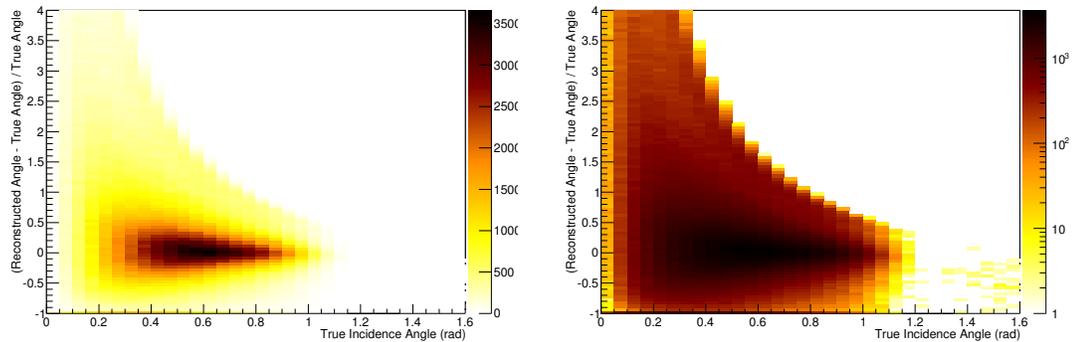
where q_i is the charge and \vec{x}_i is the position of the i th hit in the cluster, \vec{o} is the thrust origin and $\vec{n}_{(\theta, \varphi)}$ is a unit vector with polar angles θ and φ .

The unit vector at maximal $t(\theta, \varphi)$ then defines the thrust axis, the direction of the cluster. This does have the disadvantage of assuming that all cluster objects originate within the tracking detectors and are propagating outwards, but this is a sensible assumption for beam neutrino interactions.

Figure 3.9 shows the fractional angular difference between the true incidence angle and reconstructed thrust axis, as a function of true angle, for the same π^0 decay photon sample as shown in Figure 3.8. It is seen that the reconstructed angle in the Barrel modules appears to be better than that of the DsECal, which at lower angles shows a significant spread in the reconstructed angular bias compared to the higher angles in the Barrel modules.

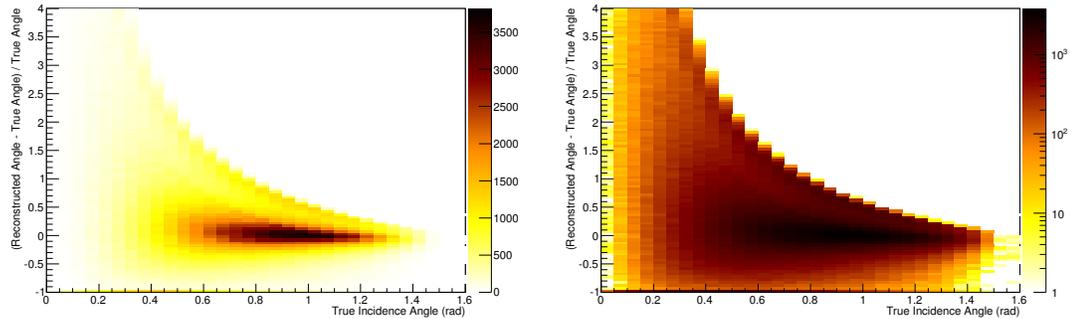
3.3.3 Global Reconstruction

The reconstruction of globally fitted tracks starts by initially extrapolating tracks from the tracking region of the detector onto the faces of the surrounding sub-detectors. Tracks are then searched for around the extrapolated entry position. If a candidate track is found for matching, and is within 300 ns of the extrapolated track, then a χ^2 fit is applied between the two based upon the track's direction and position. If the χ^2 fit result is less than 100, then the two tracks are combined into one. A Kalman filter is then applied to refit the new global track. This takes into account the expected energy losses and previous momentum measured from the sub-detector level reconstruction, to produce a position, direction and momentum measure for the new global track. The track matching and refitting is then repeated until all the possible global track objects are formed, including tracks that do not have a tracking detector component, such as those passing between the P \emptyset D and the P \emptyset DECal and/or SMRD.



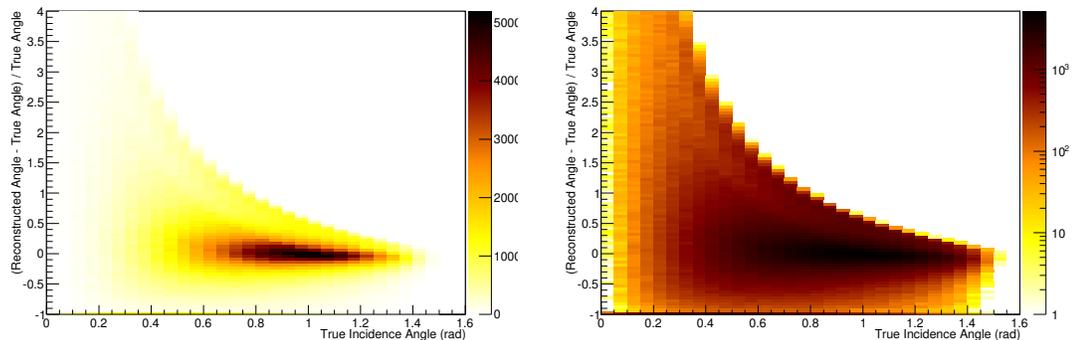
(a) Downstream ECal.

(b) Downstream ECal, log scale.



(c) South Side ECal.

(d) South Side ECal, log scale.



(e) Bottom ECal.

(f) Bottom ECal, log scale.

Figure 3.9: Fractional angular difference calculated between the reconstructed thrust axis and true incidence angle of π^0 decay photons into the ECal modules as a function of true angle for three different ECal module orientations. The right hand plots are the same as the left hand plots, but with a logarithmic z axis.

After the global matching is complete, a Kalman filter is also used to locate the global event vertex. If only one global track exists, this will be the upstream end of that track. Otherwise the Kalman filter will combine information from multiple tracks to find a best fit vertex. If multiple vertices are located, the one associated to the highest momentum track within the bunch is selected as the primary vertex.

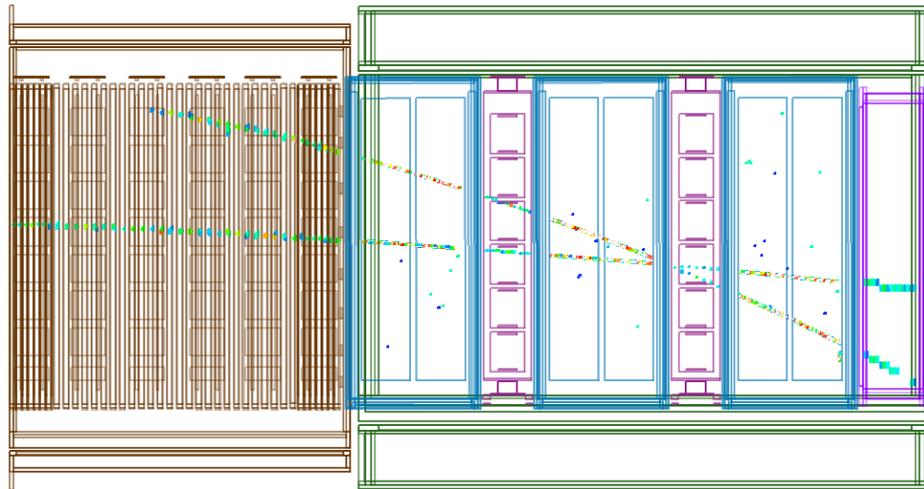
Figure 3.10 shows an example of a beam spill event with calibrated hits and global tracks fitted through them.

3.4 High Level Analysis Package

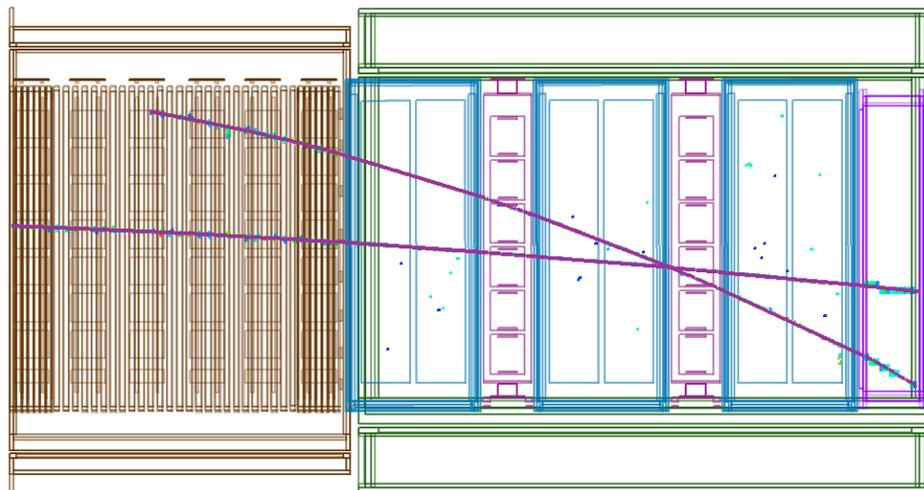
Once all data is calibrated and reconstructed it is down to the end user to develop an analysis to select specific interactions within the ND280. To streamline this process, and ensure consistency of approach between analyses, the **high level analysis** at the **near detector** (highland) framework has been developed. This analysis framework contains a C++-based analysis base class from which user analysis can inherit predefined cuts and actions, and to which user-specific ones can be added. It provides functionality for flattening the oaAnalysis data files into skimmed, lightweight flattened data tree files containing only the variables relevant to each user's analysis, and the ability to add corrections and evaluate the impact on systematic errors arising from them. Tools for quickly drawing analysis variables after each stage in the analysis cut chain, and the production of efficiency and purity figures are also included.

3.4.1 Core Packages

There are five core packages within the highland framework, one of which is optional to the end user. Each package will be described briefly below.



(a) Calibrated hits within the ND280.



(b) Globally reconstructed tracks overlaying the calibrated hits within the ND280.

Figure 3.10: Beam spill event number 54642 from ND280 run 6462, subrun 11, shown looking at the ND280 from the side; the beam comes from the left. 3.10a shows the calibrated hits within the event; two tracks spanning multiple sub-detectors are prominently seen along with some noise hits. 3.10b displays the globally reconstructed tracks overlaying the calibrated hits. The track starting towards the upper centre of the PØD is probably a muon produced during a CCQE interaction. The track starting from the upstream face of the PØD is probably a muon produced by a neutrino interacting in the earth upstream of the ND280 detector pit. Both tracks have been globally reconstructed passing through the PØD, TPCs, FGDs and DSECal.

highLevelAnalysis

The highLevelAnalysis package is the simple master package which specifies all other highland packages to be included in the highland build. If any other highland packages are not sufficiently maintained to be up to date with changes to the highland framework then they will be removed from the highLevelAnalysis build.

highlandTools

The highlandTools package defines several classes which contain the core tools all other highland packages build upon. These include:

- DataClasses, representation of reconstructed objects as converted from the oaAnalysis files into the reduced highland data types.
- AnalysisBase, the base analysis which deals with the event loop structure, data read and write, and filling of automatic calculated and stored data.
- StepBase, which handles the implementation of user defined cuts and actions.
- InputVariation, controls the adding of corrections and systematics to variables stored by the reconstruction.
- Parameters, allows access to user defined analysis parameter files.
- DrawingTools, contains a suite of tools for analysing and drawing results from the user analysis.
- MultiAnalysisTools, for interfacing multiple data sets to the drawing tools such that comparisons can be drawn between them.

highlandIO

The highlandIO package provides the interface between oaAnalysis files or highland flattened data tree files, and the DataClasses of the highlandTools package. It contains three primary classes which are:

- `oaAnalysisTreeConverter`, which converts `oaAnalysis` files into the `DataClasses` used by highland analyses.
- `FlatTreeConverter`, which similarly converts highland flattened data trees into the required `DataClasses` format.
- `CreateFlatTree`, which converts `oaAnalysis` files into the highland flattened data trees.

baseAnalysis

The `baseAnalysis` package draws together the tools and interfaces provided by the core packages and presents them to the user as a basic analysis framework that includes the functionality which all other analyses require. The functionality includes data quality checks, the application of corrections and systematics prior to the data being analysed, storage of intermediary analysis level quantities and the filling of standard variables that are useful to all analyses. `baseAnalysis` also provides the ability to add additional variables beyond the scope of the standard `DataClasses`, this includes sub-detector specific quantities such as the ECal PID variables that are not stored by the global reconstruction as standard. Finally there is vast array of utility functions that could be of use to many independent analyses. These include utilities to find the interaction type for an event in MC truth, select all tracks that match certain criteria, be it charge or track location, and the ability to calculate neutrino energy assuming a selected charged track is from a CCQE interaction.

highlandRecPack

The `highlandRecPack` package is an optional package which provides an interface between highland and the RecPack software. This allows access to additional, non-standard, operations such as propagating reconstruction changes due to alterations to the detector materials or magnetic field; calculating the momentum of charged particles based on their range; and associating broken tracks to one another.

3.4.2 Standard Analyses

Beyond the core packages is a suite of standard analysis packages which inherit from `baseAnalysis` or another standard analysis, and from which any new analysis can choose to inherit if it is advantageous. Some of the standard analyses which inherit from `baseAnalysis` include:

- `numuCCAnalysis`, the standard ν_μ inclusive analysis selection which uses FGD1 as the interaction target.
- `nueCCAnalysis`, the standard ν_e inclusive analysis selection which uses FGD1 as the interaction target.
- `antiNumuCCAnalysis`, the $\bar{\nu}_\mu$ inclusive analysis selection which uses FGD1 as the interaction target.
- `gammaAnalysis`, an e^-e^+ selection, trying to locate pairs from γ conversions.
- `numuCCOuterDetAnalysis`, a ν_μ inclusive analysis selection which uses the BrECal and SMRD as the interaction target.

And then there are additional analyses which have more complex inheritance, these include:

- `numuCCQEAnalysis`, an exclusive ν_μ CCQE selection which inherits from the `numuCCAnalysis`.
- `numuCCQExsecAnalysis`, a selection for measuring the ν_μ CCQE cross-section and inherits from the `numuCCQEAnalysis`.

Chapter 4

ECal Timing Calibration

It is necessary to perform a calibration to the timing of hits within the ECal as this is a key stage in optimising the reconstruction of interactions. This is because hit timing is used in the matching of hits together to form clusters, rejecting spurious noise hits and globally matching objects between sub-detectors. This chapter will cover the timing calibration of the Barrel and Downstream ECals. The calibration of the PØDECals, due to their differing design, is done separately and shall not be discussed here.

4.1 Calibration Methodology

The timing calibration described is used to remove the timing offsets introduced by differences in cable length between TFBs and their parent RMM, and each RMM and the ECal SCM. There are two stages to the calibration, first the offsets between each TFB on an RMM are calculated as described in Section 4.1.1, before the inter-RMM offsets are calculated as described in Section 4.1.2. The author was responsible for maintaining the code used for the ECal timing calibration and applying it to data from the third T2K experimental running period (Run 3). The author also developed the code to add an additional timing calibration step that occurs prior to the calibration as described in this section, which accounts for jumps in the RMM timing distribution as described in Section 4.3.

4.1.1 Inter-TFB Timing Calibration

The timing calibration for the ECal uses inter-spill cosmic ray muons to determine the timing offsets. These are triggered by the coincidence of hits between any two sub-detector components from the SMRD towers, the DsECal and the PØD. The sub-detectors are all separated into towers which delineate separate physical regions of each sub-detector for the purpose of issuing cosmic ray triggers.

Due to the differences in cable lengths connecting the triggering modules and variation in cosmic ray path length between separate triggering tower combinations, only the the difference between hit times on separate TFBs can be used in the timing calibration, not their absolute hit times. Therefore the TFB timing offsets are calculated by measuring the difference in hit times between TFBs on each RMM individually.

First code is used to fit a simple track through hits in an ECal module in each view. To prevent any reconstruction confusion, only events with one track per module are used, the separate views are then combined to produce a single 3D track (as was described in Section 3.3.2) per ECal module. This fitting provides the unknown coordinate for each view, allowing the hit position of the track to be located along each bar. The recorded time of each hit in the track is then adjusted to take into account the time of flight of the muon across the detector and the light propagation time down the length of the ECal bars. An assumption is made that the cosmic ray muons travel at the speed of light and originate at the top of the detector, propagating downwards. This process should then adjust the hit times such that, were there no electronics timing offsets, the hits would all be aligned in time.

Next, for each reconstructed track the combinations of TFBs which were hit on each RMM are found. The TFB combinations are then weighted based on the frequency at which they occur and Prim's algorithm [86] is used to produce a minimal spanning tree between all the TFB combinations on each RMM. Then each hit is iterated over and the mean hit time on each TFB is calculated per event. The difference in mean hit times between the TFB combinations specified by the minimal spanning tree are then recorded on an event by event basis. Figure

4.1 shows an example of the resultant minimal spanning tree selected from initial weighted TFB combinations.

This process builds up a distribution of hit time differences for TFBs on an RMM. To find the timing offset between any two selected TFBs, it is necessary to sum the mean difference between all intervening TFB combinations in the minimal spanning tree. A mean offset for each RMM is then calculated between each TFB and a reference TFB (usually TFB 0). The final TFB offsets are then calculated by taking the difference between each TFB and the reference TFB offset, and the mean offset for the RMM.

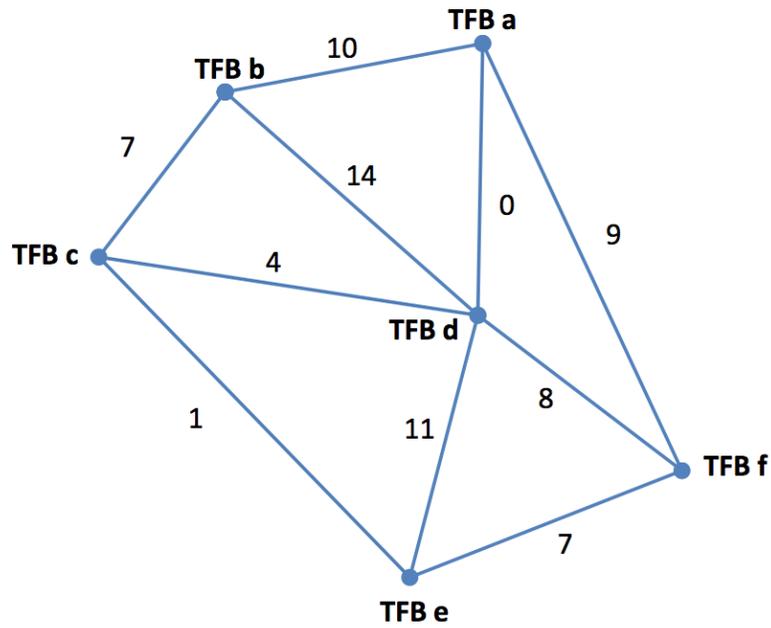
4.1.2 Inter-RMM Timing Calibration

Once the TFB offsets have been calculated they are applied back to the data used to calculate them. For events where a cosmic ray muon track is split across two ECal modules, the upper track is extrapolated down to the lower one in an attempt to match the sections together. If a match is made, then the hits in the lower module have their hit times adjusted to account for the time of flight between the ECal modules. Once these corrections are applied, were there no RMM offsets, the hits would then be aligned in time.

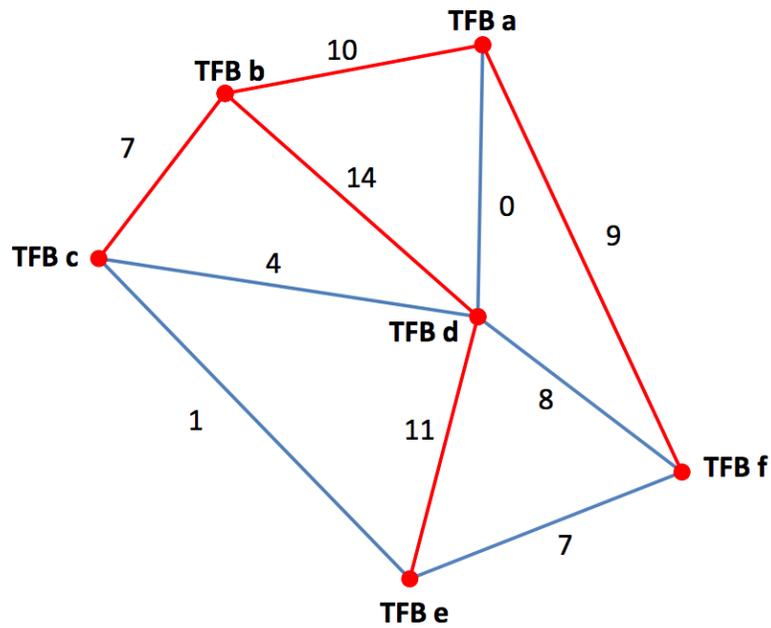
To calculate the offsets between RMMs, the mean hit time on RMMs 0 and 1 are calculated, and an offset is found for each which aligns them to the mean time between them. RMMs 0 and 1 are the two which read out the DsECal, and this mean hit time then dictates the mean hit time to which the remaining BrECal modules are aligned.

4.1.3 Timing Calibration Methodology Discussion

The methodology used is not perfect, and fails to account for some factors. For instance, the effect of the ND280 magnetic field is never taken into account, this being potentially important for the extrapolation of tracks between the ECal modules during the RMM offset calculations. It is also assumed that the TFB to Trip-T, and Trip-T to MPPC cable lengths are all identical and therefore do



(a) Weighted TFB combinations.



(b) Weighted TFB combinations and resultant minimal spanning tree in red.

Figure 4.1: 4.1a shows weighted combinations between multiple TFBs; 4.1b shows the minimal spanning tree that would be extracted as a result [87].

not effect the distribution of hit times across each TFB; this is unlikely to be the case. As such, there has been much successful work undertaken to develop a new timing calibration methodology based on that used by the MINOS experiment [88], however this work was not completed in time to calibrate the data used in the analysis discussed in Chapter 5.

To complete the calibration as described, enough data needs to be obtained to give sufficient weight to all the TFB combinations. To produce the necessary statistics for Prim's algorithm and the resultant offsets, approximately 20 ND280 subruns are required. Therefore the timing calibration is run over any ND280 run with this minimum number of subruns. Due to the computationally intensive nature of the processing, if an ND280 run has greater than the minimum required number of subruns, only a subset of them will be used for the offset calculation.

Once offsets have been calculated they will be applied to all data until a subsequent ND280 run with sufficient statistics is obtained and the offsets are recalculated. This would not be an issue if the timing distributions were stable with time, however it has been found that this is not the case for the RMM level hit times. This leads to offsets being applied to datasets for which they are not appropriate due to a shift in the RMM timing distributions. This is an issue the author attempted to address as discussed in Section 4.3.

For a more detailed discussion of the timing calibration methodology described here see Reference [87].

4.2 Initial Timing Calibration Result

Using the methodology outlined in Section 4.1, the TFB and RMM timing calibration constants were calculated for the third T2K experimental running period. This was done for all ND280 runs with sufficient statistics to accurately calculate the timing offsets, and for which the quality of the data taking had been deemed good.

4.2.1 TFB Calibration Results

Figure 4.2 shows the calculated TFB calibration constants for the third T2K experimental running period. Figure 4.3 shows the residuals for the same data set. The residuals are calculated as the difference between each calibration constant for a TFB, and the mean of all calibration constants on that TFB.

In theory, the TFB calibration constants should be constant with time as they use a steady electronic connection between the each TFB and RMM. From Figures 4.2 and 4.3 it is seen that these calibration constants are generally consistent to within ± 2.5 ns, although some, notably RMM 9, are seen to vary by as much as ± 5 ns. Given that one clock tick on the TFB boards is 2.5 ns and the variation in statistics and TFB hit occupancy between the different ND280 runs, such a spread is not unexpected. Quite why the residuals appear to be worse on some RMMs than others is still somewhat unclear.

4.2.2 RMM Calibration Results

Figure 4.4 shows the calculated RMM calibration constants for the third T2K experimental running period. The calibration constants for RMMs 0 and 1 mirror one another about the line of zero offset. The other RMMs then all have offsets that align back to this zero point, and so changes in the offset between RMMs 0 and 1 are reflected as additional contributions to the offsets of the other RMMs.

It is clear from Figure 4.4 that the stability of the RMM calibration constants is significantly worse than the TFB level calibration constants. The calibration constants can show jumps in their calculated value of order 10 ns, between separate runs, seemingly at random. As previously stated in Section 4.1.3, this variation in the RMM timing distribution can lead to RMM calibration constants being applied to data for which it is inappropriate and this can in turn lead to a degradation in the quality of event reconstruction.

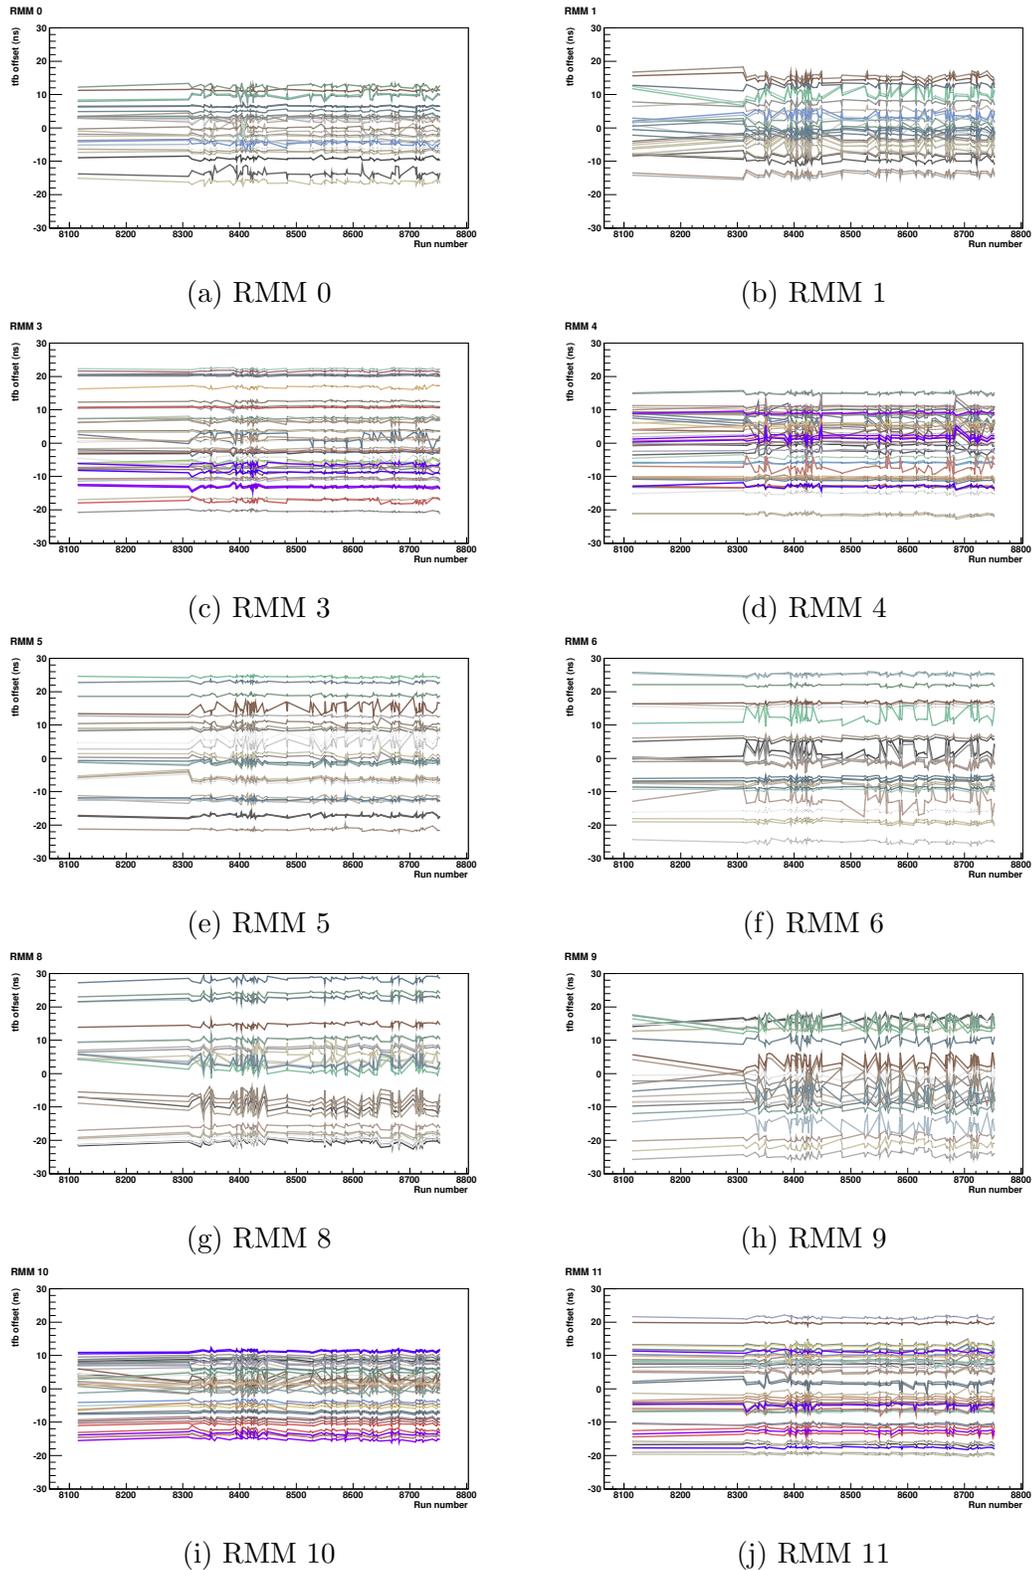
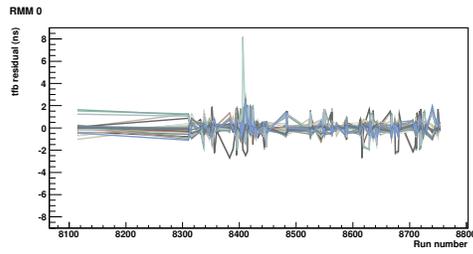
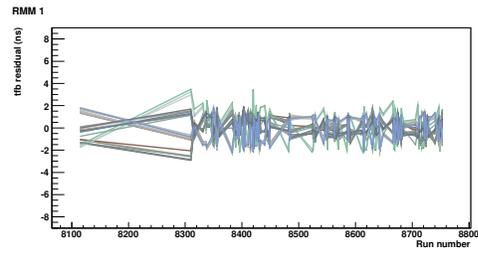


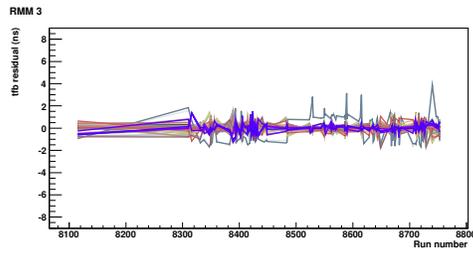
Figure 4.2: TFB calibration constants as a function of ND280 run number for each RMM in the tracker region.



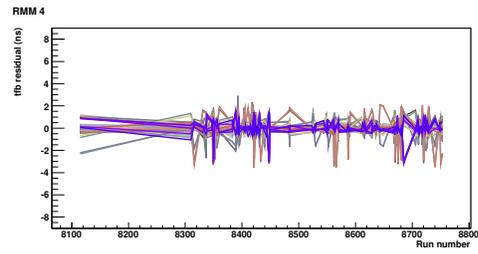
(a) RMM 0



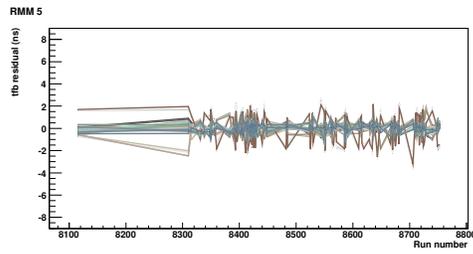
(b) RMM 1



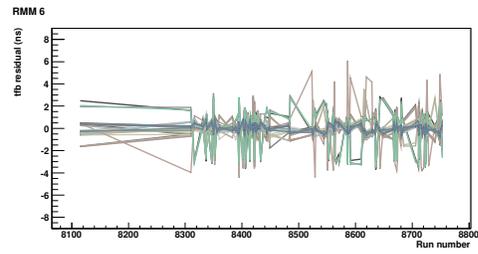
(c) RMM 3



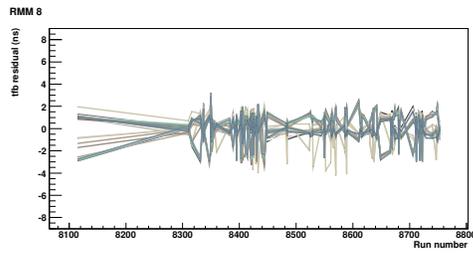
(d) RMM 4



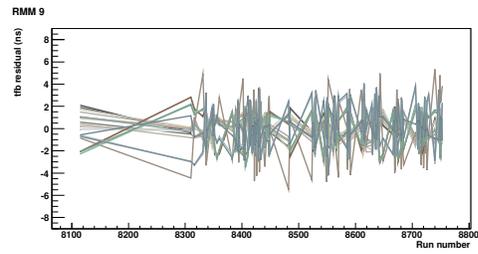
(e) RMM 5



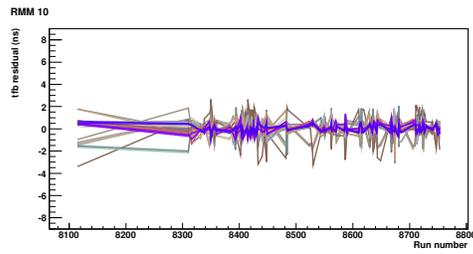
(f) RMM 6



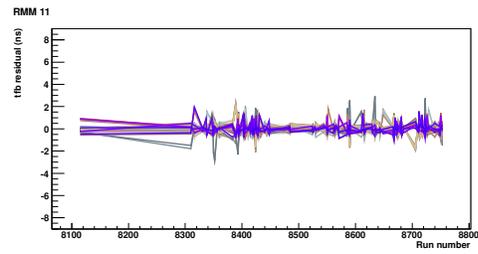
(g) RMM 8



(h) RMM 9



(i) RMM 10



(j) RMM 11

Figure 4.3: TFB residuals as a function of ND280 run number for each RMM in the tracker region.

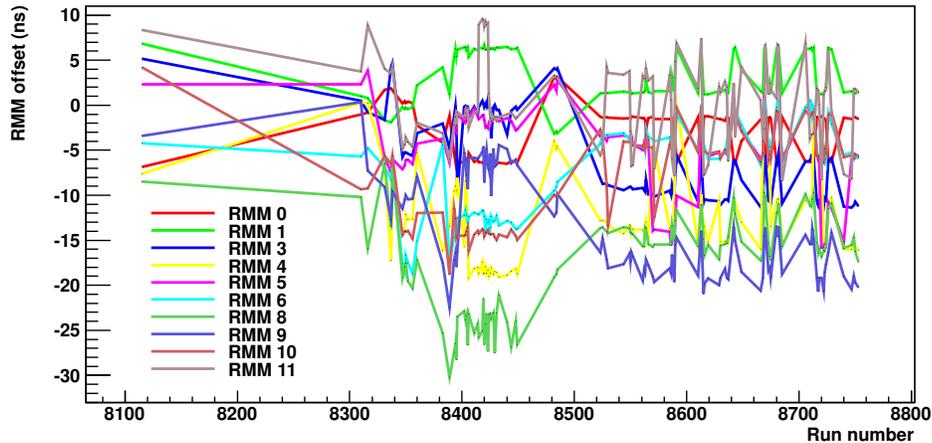


Figure 4.4: RMM calibration constants as a function of ND280 run number.

4.2.3 Track Resolution

When offsets are correctly applied to data, as can be seen in the mean cosmic ray hit time distribution shown in Figure 4.5, the range in mean hit time across all TFBs is seen to significantly reduce, with the RMS of the distribution halving. The central value of the distribution is negative due to the look-back nature of the cosmic trigger, with the introduced delay being due the time difference between a cosmic ray event occurring and a trigger being issued by the MCM. The idealised resolution of the MCM triggering is 4 ns RMS, but the true width of the distribution is dominated by the 10 ns granularity in the MCM clock phase-lock and fibre time walk [69].

4.3 Addressing RMM Timing Instability

4.3.1 Cause Of RMM Timing Instability

The source of the RMM timing instability is the nature of the connection between the ECal SCM and the RMMs. These read out boards communicate via a Xilinx RocketIO™ transceiver [89] along optical fibre links. The connection has a 100MHz phase-lock loop, one quarter the rate at which the clock on the

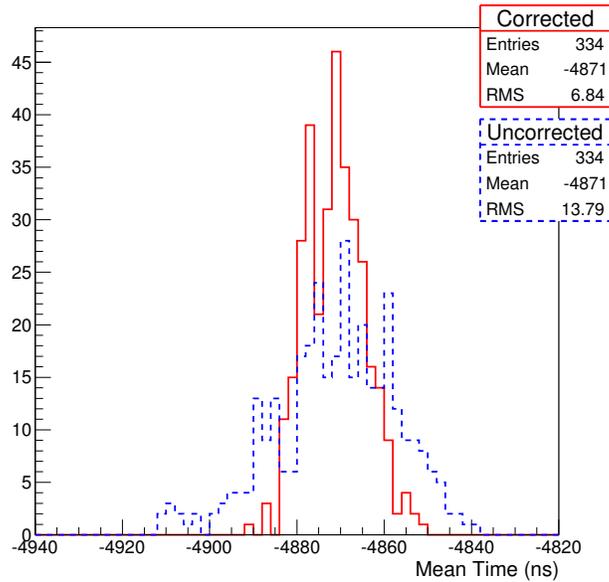


Figure 4.5: Mean hit time for cosmic ray muons across each TFB, prior to (blue) and after (red) the timing calibration has been applied. The data shown is for ND280 run 8115.

TFB boards operate. Each time the connection between an RMM and SCM is established, the phase-lock between boards must also be established. Each time this occurs it is possible for the phase-lock to shift the timing between boards by 10 ns, thus leading to 10 ns shifts in the recorded hit times. Phase-locking occurs when the electronics boards are power cycled, but can also occur randomly during detector operations due to link loss down the optical fibres. The same connections are also used to link the MCM to the SCMs and the MCM to the CTM, allowing time slips to occur from multiple sources, dependent on the type of data being analysed. Such “time slips” as they are known, are an obvious problem for effective timing calibration on an individual sub-detector, but also between sub-detectors, and as such their detection and the application of an associated correction became a priority for the ND280 calibration group.

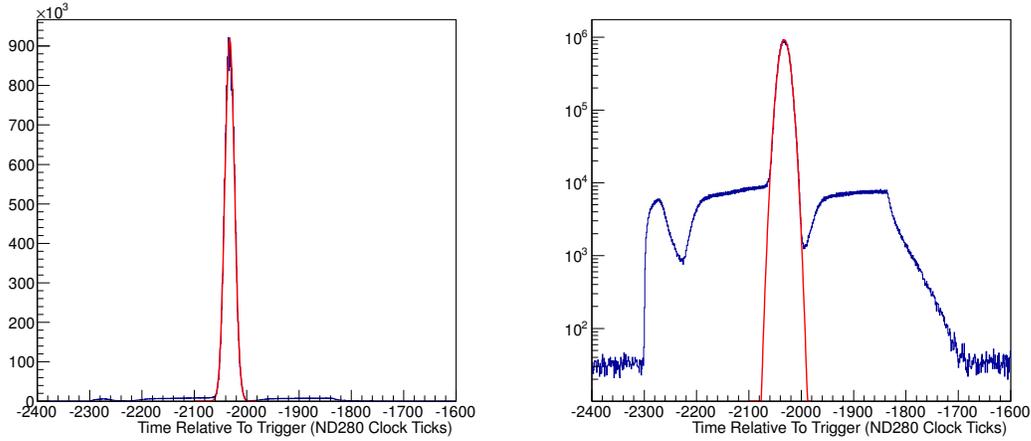
4.3.2 Considerations For RMM Time Slip Detection

Whilst developing a technique to locate and account for the RMM time slips a range of factors were considered.

- The RMM time slips do not only occur only within the ECal, but are an issue for data collected by the PØD and SMRD. It would therefore be best to detect them in a holistic, rather than sub-detector specific manner.
- The calibration techniques used by each sub-detector are already well established. It would therefore be appropriate to make any changes to the calibration chain upstream of the sub-detector level calibration, but in a manner that does not disrupt those subsequent calibration techniques.
- The spread in uncalibrated hit times on each sub-detector varies quite significantly, and on no sub-detector is it narrow enough to detect time slips on an event by event basis. Therefore data needs to be collected and combined efficiently such that time slips can be detected at a good resolution on all sub-detectors.
- In an ideal case, the collection, processing and identification would be nearly fully automated to prevent any significant increase in the workload of the calibration group.

4.3.3 Detection of RMM Time Slips

Given the considerations outlined above a technique for locating the RMM time slips was developed using the raw data output from the ND280. First an automated job is run on the ND280 semi-offline computers which produces a flattened data tree of useful variables for the Trip-T detectors from cosmic ray triggered events. For each hit with a high gain read out greater than 900 ADC counts, the sub-detector, cosmic ray trigger type, electronics read out boards, absolute hit time and time relative to trigger are stored. The use of a 900 ADC high gain cut aims to remove noise and pedestal level hits from the data sample. Once the hits



(a) Time relative to trigger distribution, linear scale. (b) Time relative to trigger distribution, logarithmic scale.

Figure 4.6: Time relative to trigger distribution for ECal RMM 5 during part of the third T2K experimental running period. By fitting a Gaussian function to the data sample it is easy to reject the noise, most easily seen in 4.6b, from the time slip analysis samples. One ND280 Clock Tick equates to 2.5 ns.

have been collected they can then be passed through the RMM time slip finding algorithm as described below.

Constrain Hit Times

The first stage is to constrain the time range of hits accepted for the time slip search. First histograms are filled with the time relative to trigger for all events in the flattened data tree files provided. This is done on a sub-detector and RMM basis, with a specified cosmic trigger, either TripT or FGD. Each timing distribution is then fitted with a Gaussian function to find the mean and RMS of the time relative to trigger distributions as seen in Figure 4.6.

Hits that reside within plus or minus three times the RMS of the Gaussian fit, from the Gaussian fit central value, $(\bar{t} \pm 3 \times \bar{t}_{RMS})$, are accepted for the next stage of the time slip search. This cut ensures that much of the noise that was not removed by the 900 ADC high gain cut is now removed from the samples.

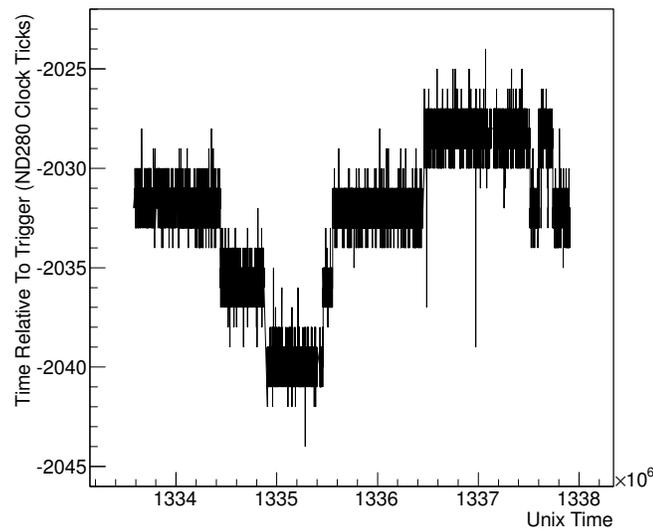


Figure 4.7: Mean time relative to trigger for five minute data taking periods as a function of Unix time. Data for ECal RMM 5 during the third T2K experimental running period.

Calculate Hit Time Averages

Next the hits are partitioned, on each sub-detector and RMM basis, into consecutive five minute periods. For each period the time relative to trigger for all hits that occur within each five minute window are averaged. Figure 4.7 shows an example distribution of the resultant mean time relative to trigger as a function of Unix time; eight time slips can be observed during the time period shown. Five minute averaging periods were used as this was deemed a good compromise between ease of observing RMM time slips out of the standard variation in hit times relative trigger whilst still finding the time slips with a high resolution in Unix time. To put this in context, the eight time slips shown for RMM 5 (which displays typical behaviour) occurred within a 52 day period; this is an average rate of 0.15 time slips per day, equivalent to approximately $5 \times 10^{-4}\%$ chance of a time slip occurring in each five minute period.

Smooth Spectrum

After the mean time relative to trigger has been calculated the resultant spectrum is smoothed to make it easier for the later algorithm to finally pick out the discontinuities in the spectrum due to the presence of time slips. The smoothing procedure uses a boxcar averaging method, as follows:

- Each mean time relative to trigger from the spectrum is selected in turn; \bar{t}_n .
- The three previous values from the spectrum are then selected and their mean value calculated; $\bar{t}_{pre} = \frac{\sum_{i=n-3}^{n-1} \bar{t}_i}{3}$.
- The three subsequent values are also selected, and likewise a mean value calculated between them; $\bar{t}_{post} = \frac{\sum_{i=n+1}^{n+3} \bar{t}_i}{3}$.
- If the mean time relative to trigger lies within the range produced by the mean from the three previous and subsequent values the mean time relative to trigger is retained; if $\bar{t}_{pre} \leq (\geq) \bar{t}_n \leq (\geq) \bar{t}_{post}$, then $\bar{t}_n = \bar{t}_n$, where $\bar{t}_{pre} \leq \bar{t}_{post}$ ($\bar{t}_{pre} \geq \bar{t}_{post}$).
- If the mean time relative to trigger lies outside of the range produced by the mean from the three previous and subsequent values, then the original value is rejected and is set equal the average of the mean values from the three previous and subsequent points; if $\bar{t}_n < \bar{t}_{pre}$ (\bar{t}_{post}) or $\bar{t}_n > \bar{t}_{post}$ (\bar{t}_{pre}), then $\bar{t}_n = \frac{\bar{t}_{pre} + \bar{t}_{post}}{2}$, where $\bar{t}_{pre} < \bar{t}_{post}$ ($\bar{t}_{post} < \bar{t}_{pre}$).

This boxcar-averaging smoothing algorithm is run over the spectrum three times to repeatedly smooth the distribution as can be seen in Figure 4.8. Special consideration is given at the start and end of the spectrum where there are less than three previous or subsequent mean time relative to trigger values to use. As can be seen, the smoothing of the spectrum reduces the noise present in the spectrum without affecting the time slips. The smoothing algorithm is run three times over the spectrum as further iterations beyond this level are of ever reducing benefit in cleaning up the noise. This reduction in benefit can be seen by

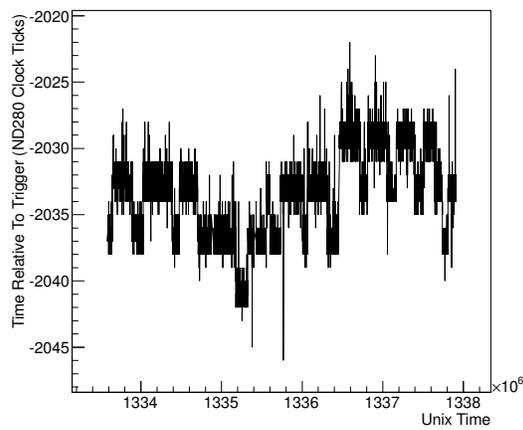
comparing the very limited change in noise rate on the spectrum between Figures 4.8d and 4.8c, compared to the benefit gained by the previous smoothing passes.

Moving Standard Deviation

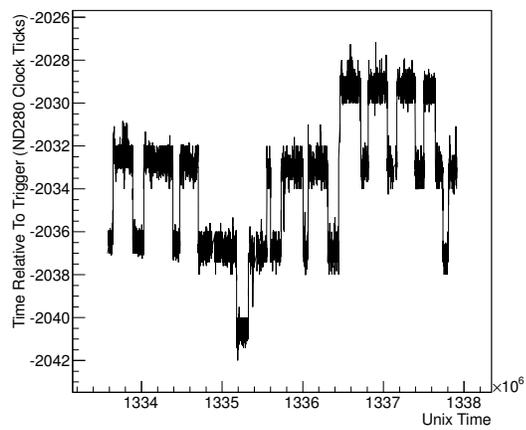
Once the mean time relative to trigger data spectrum has been smoothed attention turns to detecting the time slips. The first stage in this process is to run a moving standard deviation across the spectrum to highlight the location of the discontinuities due to the RMM time slips. This moving standard deviation proceeds as follows:

- Each mean time relative to trigger from the smoothed spectrum is selected in turn; \bar{t}_n .
- The six previous and six subsequent values are selected from the spectrum and a mean value is obtained from all thirteen values; $\bar{t}_{range} = \frac{\sum_{i=n-6}^{n+6} \bar{t}_i}{13}$.
- Using the mean over the thirteen point range the standard deviation is calculated; $\sigma_n = \sqrt{\frac{\sum_{i=n-6}^{n+6} (\bar{t}_i - \bar{t}_{range})^2}{13}}$.

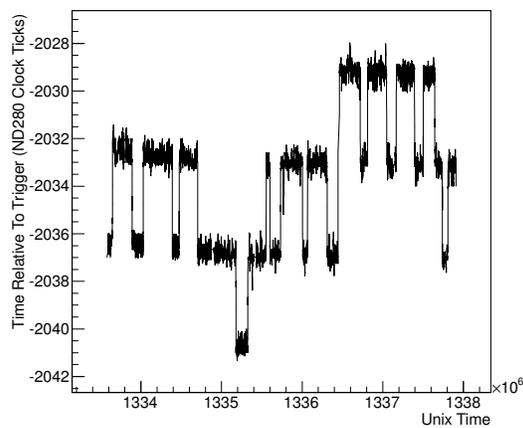
The resultant moving standard deviation spectrum features peaks which align with the position of the time slips as can be seen in Figure 4.9. Special consideration is given at the start and end of the spectrum where there are less than six previous or subsequent mean time relative to trigger values to use. The range over which the standard deviation calculation is applied can be increased or reduced, but there are obvious compromises to be taken into account in doing so. By reducing the standard deviation calculation range, random fluctuations in the noise from the smoothed mean time relative to trigger spectrum produce a greater number of spikes in the standard deviation, leading to a potential increase in false positive time slip identifications. Increasing the standard deviation calculation range reduces the noise in the standard deviation spectrum, but impinges upon the ability to identify time slips that occur close together in Unix time. Therefore



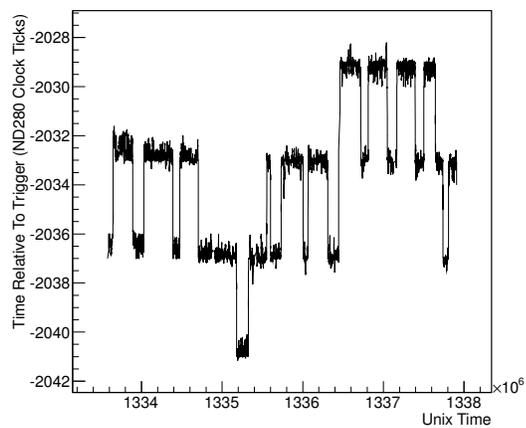
(a) Unsmoothed distribution.



(b) First pass smoothed distribution.



(c) Second pass smoothed distribution.



(d) Third pass smoothed distribution.

Figure 4.8: Progressive smoothing of mean time relative to trigger for five minute data taking periods as a function of Unix time. Data for ECal RMM 11 during the third T2K experimental running period.

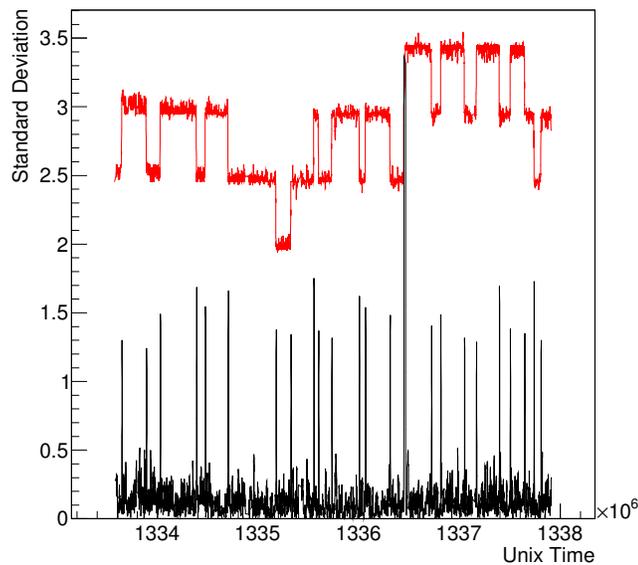


Figure 4.9: Standard deviation spectrum obtained from Figure 4.8d shown in black, Figure 4.8d is shown superimposed above in red.

a range of plus and minus six values was found to be a good compromise between these two considerations.

Standard Deviation Peak Search

After the standard deviation spectrum has been extracted from the data, a peak finding algorithm is applied to highlight the likely locations of time slips within the data set. This peak search proceeds as follows:

- Each standard deviation value is selected in turn; σ_n .
- If the standard deviation exceeds a value of 1, then that point is registered as the rising limb of a peak; if $\sigma_n > 1$ then $\sigma_{rise} = \sigma_n$.
- The standard deviation spectrum is then scanned forward in time to find the last position at which the standard deviation is still greater than 1, this is then registered as the falling limb of a peak; while $\sigma_{n+i} > 1$, $\sigma_{fall} = \sigma_{n+i}$.

- The Unix time of these two values are then taken as the upper and lower bound of the period in which a time slip is suspected to have occurred, with the central value of the period assumed to be the true slip time; $t_{slip} = \frac{t_{rise} + t_{fall}}{2}$.

The peak search runs across the entire standard deviation spectrum, extracting all the likely time slip positions for them to be finally confirmed, and their direction (slip up or slip down in time) registered.

Select Time Slips

The final selection of valid time slips then occurs based on the probable location of slips selected by the standard deviation peak search. The selection of valid slips proceeds as follows:

- First the point in the mean time relative to trigger spectrum at which the standard deviation exceeds 1 is selected; \bar{t}_n^{rise} .
- The six previous values are selected from the spectrum and a mean values is obtained; $\bar{t}_{rise} = \frac{\sum_{i=n-6}^{n-1} \bar{t}_i}{6}$.
- The point in the mean time relative to trigger spectrum at which the standard deviation falls back below 1 is also selected; \bar{t}_n^{fall} .
- The six subsequent values are then selected, and likewise a mean values is obtained; $\bar{t}_{fall} = \frac{\sum_{i=n+1}^{n+6} \bar{t}_i}{6}$.
- The difference in mean values prior to and after the suspected slip is then calculated time slip; $\Delta\bar{t}_{slip} = \bar{t}_{fall} - \bar{t}_{rise}$.
- If the absolute difference in mean values prior to and after the suspected slip is found to be greater than 2.4 ND280 clock ticks (6 ns), then the slip is deemed to be real and saved to a file; $|\Delta\bar{t}_{slip}| > 2.4$.
- The direction of the selected time slips is then obtained from the sign of the difference in mean values prior to and after the slip.

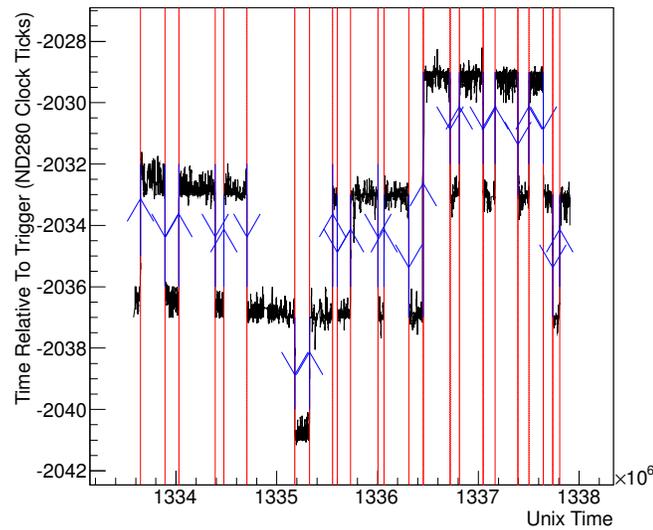


Figure 4.10: Selected time slips superimposed upon Figure 4.8d. The red bands indicate the lower and upper edges of the time slip regions as located by the standard deviation peak finding algorithm. The blue lines with arrow indicate the selected time slip position at the centre of the red bands, with the arrow indicating the direction of each time slip.

Given that the inter-slip noise in the spectrum shown in Figure 4.8d has a magnitude of about ± 0.5 ND280 Clock Ticks, it is not unreasonable for the algorithm to calculate a pre to post-slip difference of about 3 ND280 Clock Ticks. This is particularly likely if multiple time slips are closely coincident in time, allowing one slip to pull down the pre or post-slip average of an adjacent slip. As such, it was found best to place a reasonably low threshold of 2.4 ND280 Clock Ticks on the difference in mean values prior to and after the suspected slip, as this is then very unlikely to miss true time slips, but without including false positive slip identifications. The final time slip selection procedure runs over all locations highlighted by the standard deviation peak finding algorithm. Figure 4.10 shows the final output from the time slip finding algorithm, with the position and direction of all the slips clearly correctly determined.

4.3.4 Validation of RMM Time Slips

Once the time slips have been selected by the automated system described in Section 4.3.3 some validation checks are made to ensure that the timing distribution have been correctly flattened and that all the corrections are appropriate for the data.

Flattened Timing Distribution

Once all the time slips have been selected, if those time slips are applied to the data the resultant mean time relative to trigger distributions should be flat. This is not necessarily always the case though, because although the time slip algorithm has an accuracy of $> 95\%$, infrequently time slips can be missed or falsely identified. The algorithm also assumes all time slips are 10 ns in magnitude, however it is possible, particularly when the system is power cycled, for multiple time slips to occur simultaneously across the differing electronics connections. This can lead to time slips with a combined magnitude of greater than 10 ns, an example of which can be seen at Unix time 1336.45×10^6 s on Figure 4.10.

Therefore once all time slips have been selected, the calibrator can quickly apply the time slip corrections onto the data. Checks are then made to ensure the resultant distribution is flat, and if not make any necessary adjustments to the final data file containing the time slip corrections. Once this process has been repeated as necessary, a near flat timing distribution is produced. An example of this is shown in Figure 4.11.

Due to the choice of five minute averaging periods for the initial time relative to trigger distributions, and subsequent ranges used in the moving standard deviation and peak finding algorithms, the absolute time for selecting time slips is not perfect. This is evident in the spikes above and below the flattened time relative to trigger spectrum shown in Figure 4.11. However slips are normally corrected for within about 10-15 minutes of their occurrence.

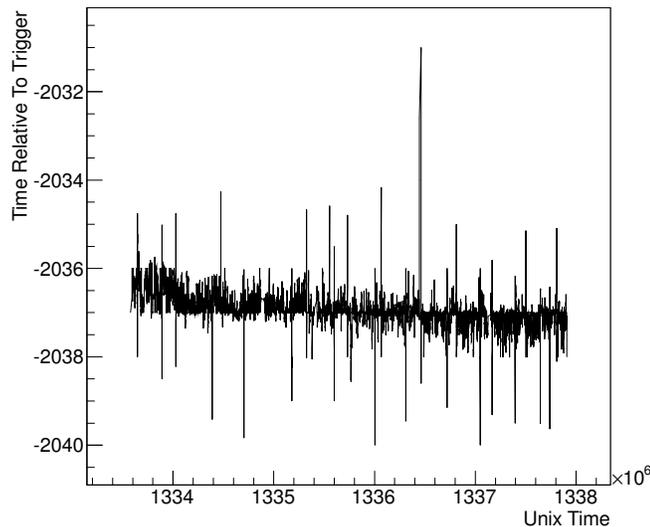


Figure 4.11: Resultant mean time relative to trigger timing distribution once the RMM time slip corrections, as identified in Figure 4.10 have applied to the data.

Removing the CTM Time Slip Correction

The final stage of the calibration is to remove any time slip corrections that apply to the link between the CTM and MCM. Like the other optical connections, the connection between the CTM and MCM exhibits 10 ns time slips, however unlike the other connections, slips to the CTM to MCM connection only appear when examining events from the cosmic ray trigger. This is because only events triggered by the cosmic ray trigger receive any timing information from the CTM.

Because the CTM to MCM connection is common to all the Trip-T sub-detectors, when a time slip occurs on this connection it manifests as a simultaneous 10 ns shift in the same direction across all channels on all the Trip-T detectors for cosmic ray triggered events (as used by this calibration). Beam triggered events (and all other non-cosmic ray triggered events) do not shift in time as they receive no timing from the CTM.

Were a time slip correction to be applied when a CTM to MCM time slips had occurred, the timing of the Trip-T detectors would all uniformly shift their timing alignment in the non-cosmic ray triggered events with respect to the other

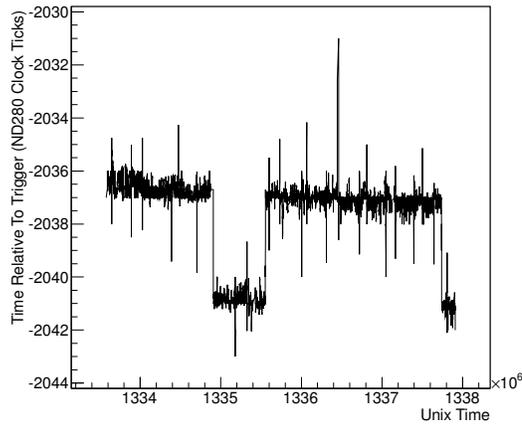
sub-detectors (FGD and TPC). Therefore applying the correction would be detrimental to the inter-detector timing calibration, manifesting as new time slips in the inter-detector timing distributions, as these are calculated through the timing of beam spill events.

Fortunately, because such time slips occur simultaneously across all the Trip-T sub-detectors, a simple algorithm which locates near simultaneous corrections on all sub-detectors can be used to find and remove these corrections as seen in Figure 4.12. Once this is done the RMM time slip calibration constants are fully prepared and can be applied to the data for all Trip-T sub-detectors prior to any other calibrations being calculated or applied.

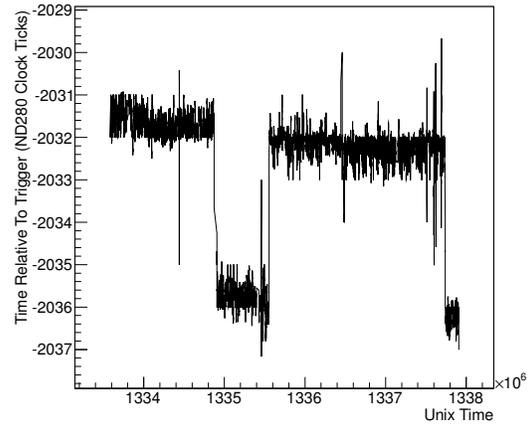
4.4 Application Of RMM Time Slip Correction

RMM time slip corrections have been calculated for all Trip-T sub-detectors for the first through fourth T2K experimental data taking periods. This data has been included in the ND280 calibration database for use by the different sub-detectors. Overall the application of the time slip corrections have been shown to improve the timing stability of the Trip-T sub-detectors.

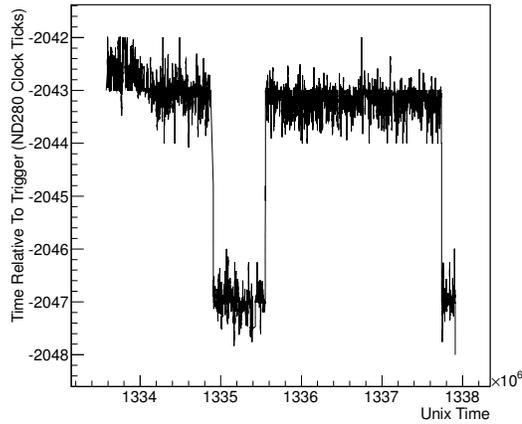
Figure 4.13 shows the new RMM calibration constants as a function of ND280 run number, obtained after the RMM time slip corrections have been applied. The RMM time slip corrections were designed to stabilise the data within the ranges between ND280 runs 8300 to 8390, 8390 to 8480, 8480 to 8510, and 8510 to 8760. Each range represents a distinct period of beam taking by the experiment. Within these periods, it is clear to see a significant improvement in the stability of the RMM calibration constants when compared to those seen in Figure 4.4. The variation that is left in offset is generally of the order less than ± 5 ns, often much less, which is comparable with variation in the TFB offsets for the same period (as seen in Figures 4.2 and 4.3 which are unaffected by the application of RMM time slip corrections).



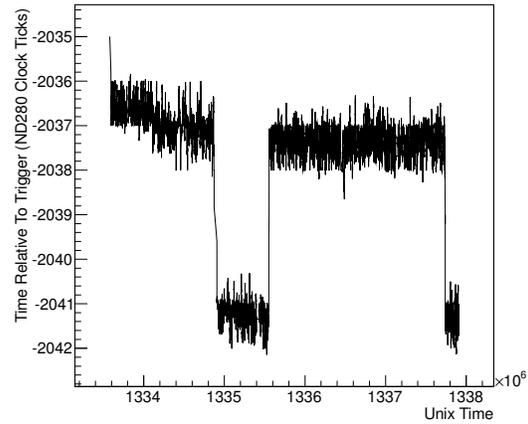
(a) ECal RMM 11.



(b) ECal RMM 5.



(c) PØD RMM 0.



(d) SMRD RMM 2.

Figure 4.12: The final mean time relative to trigger data spectra, once all relevant time slip corrections have been applied, for four different RMMs across three different Trip-T detectors. All the data spectra are now flat with time, except for the presence of the CTM to MCM time slips which occur simultaneously for all RMMs in all Trip-T sub-detectors.

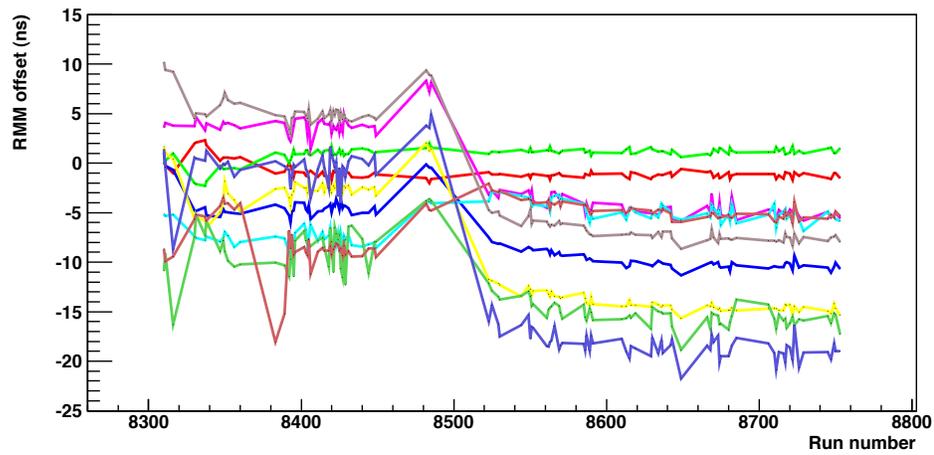


Figure 4.13: RMM calibration constants, with pre-applied RMM time slip corrections, as a function of ND280 run number for the third T2K experimental running period.

Chapter 5

Selecting Charged-Current Single π^0 Interactions in the FGDs

The analysis described within this chapter aims to select charged-current muon-neutrino interactions within the fiducial volume of the ND280 Fine Grain Detectors, with an associated single neutral-pion being emitted from the interaction nucleus (CC1 π^0). Any number of nucleons and photons may also be emitted from the recoiling interaction nucleus. Furthermore, the event topology of particular interest requires that the photons produced by the decay of the π^0 both propagate to, and subsequently shower within, the Downstream and/or Barrel ECals.

This is complementary to a similar analysis being developed at The Henryk Niewodniczański Institute of Nuclear Physics of the Polish Academy of Sciences. Their analysis requires at least one of the π^0 decay photons to pair-produce in the tracking detectors (FGDs and TPCs) from which the interaction may be reconstructed [90].

5.1 Analysis Motivation

The T2K oscillation analyses have significant systematic uncertainty contributions due to the uncertainty on the neutrino flux and cross-section models. To

constrain these uncertainties, T2K produces fits to external data sets and dedicated ND280 final state samples. Currently the ND280 provides CCQE, CC $1\pi^+$ and CC Other enhanced samples to these fits which help significantly constrain those systematic uncertainties. As more data is collected by T2K it is intended to further constrain these systematic uncertainties with additional ND280 final state samples and associated cross-section measurements. The analysis presented within this thesis will in future contribute to a dedicated CC $1\pi^0$ enhanced sample and cross-section measurement, and may also help in the development of equivalent neutral-current interaction measurements and may contribute to a π^0 veto for those studying other dedicated final states.

5.2 Analysis Samples

The analysis described in this chapter uses data from the second, third and fourth T2K experimental running periods, Run 2, Run 3 and Run 4 respectively. Each data taking period has a MC sample associated to it, flux-weighted to the beam operation during the data taking period. The different running periods are also further broken down based on changes in the water-in/out status of the PØD and accelerator operation.

5.2.1 Data Samples

Table 5.1 specifies the operational parameters of the different beam taking periods. All data used is required to pass spill quality and ND280 data taking quality cuts as described below.

It is important to note that the data obtained during the first T2K experimental running period, Run 1, is not used by this analysis as the BrECal modules were not installed into the detector at that time. Given the importance of these ECal modules to this analysis, and the low POT obtained during Run 1, it was decided not to analyse that data sample.

The Run 3a data is also not used by this, or any other ND280 analysis, as the magnetic focusing horns of the beamline were temporarily inoperable during that

T2K Run	Run Period	Good POT	PØD Water Status
Run 2 (Water)	Nov. 2010 - Feb. 2011	4.29×10^{19}	Water In
Run 2 (Air)	Feb. 2011 - Mar. 2011	3.55×10^{19}	Water Out
Run 3b	Feb. 2012 - Mar. 2012	2.15×10^{19}	Water Out
Run 3c	Apr. 2012 - Jun. 2012	13.48×10^{19}	Water Out
Run 4 (Water)	Oct. 2012 - Feb. 2013	16.25×10^{19}	Water In
Run 4 (Air)	Feb. 2013 - May 2013	17.62×10^{19}	Water Out
Total	Nov. 2010 - May 2013	57.34×10^{19}	-

Table 5.1: Experimental operation periods being used for this analysis showing PØD water status and the good POT collected. Good POT requirements are described in Section 5.2.1.

very short data taking period.

The beam power has increased with time, from about 50 kW to 220 kW, as upgrades have been made to the accelerator complex. This is reflected in the MC simulations and can lead to an increase in event pile up, but does not affect the total number of neutrino interactions which is dictated by the number of POT.

Data Quality Requirements

During the data taking periods used 60.95×10^{19} POT were delivered, of which 57.34×10^{19} POT passed the good POT data quality requirements, a 94% data taking efficiency. The data quality requirements are split into two categories, good spill delivery and good data taking at the ND280. Figure 5.1 shows the total and good POT collected by the T2K experiment as a function of time.

For data to pass the good spill data quality requirements four criteria must be met, as described in [55], these are:

- 1 All beamline hardware components are working correctly.
- 2 Horn currents are within ± 5 kA of the mean value.
- 3 Beam angle, as measured by MUMON, is within 1 mrad of the mean value.

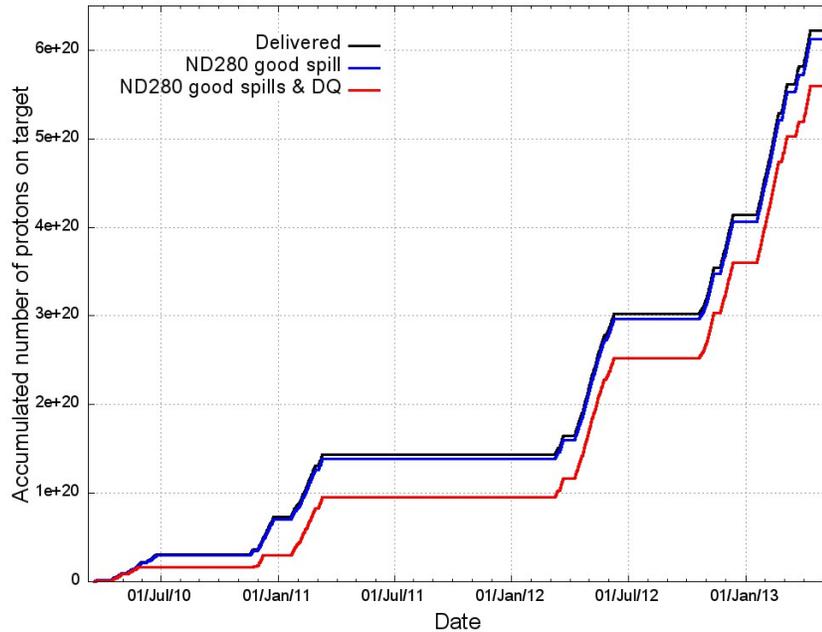


Figure 5.1: Accumulated total POT, good spill data quality POT, and good spill and ND280 data quality POT as a function of time.

4 Muon yield, as measured by MUMON, is within $\pm 5\%$ of the mean value.

For data to pass the good ND280 data quality criteria a dedicated team of analysts study the data on a spill by spill basis within each sub-detector. The checks made by each sub-detector are too numerous to describe individually here, but all checks can generally be described as checking on one of three key requirements:

- 1 Sub-detector hardware components are working correctly.
- 2 Sub-detector hits are being recorded correctly in time with the beam triggers.
- 3 Sub-detector calibrations are applying an appropriate correction to the data.

The results of the spill and ND280 data quality checks are stored in oaAnalysis files as Boolean flags on a spill by spill basis. The analysis described here requires a good spill flag from both the beam spill and from all ND280 sub-detectors.

T2K Run	MC POT	Multiple of Data POT	PØD Water Status	Beam Power
Run 2 (Water)	2.05×10^{20}	4.78	Water In	120 kW
Run 2 (Air)	1.80×10^{20}	5.05	Water Out	120 kW
Run 3b	1.04×10^{20}	4.82	Water Out	178 kW
Run 3c	6.45×10^{20}	4.78	Water Out	178 kW
Run 4 (Water)	7.87×10^{20}	4.84	Water In	178 kW
Run 4 (Air)	8.30×10^{20}	4.71	Water Out	178 kW
Total	27.51×10^{20}	4.80	-	-

Table 5.2: MC POT used to compare with each data taking period, along with the multiple of good data POT that it is equates to, the PØD water status and the simulated beam power.

5.2.2 Monte Carlo Samples

Each data running period is compared against a sample of NEUT [91] *magnet* Monte Carlo data. The *magnet* MC samples simulate the whole detector, including the PØD water status, the beam spill structure and the average beam power for each running period. Table 5.2 specifies the details of MC samples used to compare with data for each running period. Each MC sample uses approximately 5 times data POT, and the MC is always carefully scaled down to the true data POT of each running period when being compared.

Additional MC samples are employed during the training of the Boosted Decision Tree (BDT), these are described in Section 5.5.3.

5.3 Signal Definition

Interactions are considered to be signal events if, within the NEUT MC, they have an interaction code greater than 0 and less than 30 which denote charged-current neutrino interactions. Additionally, the interacting neutrino requires a PDG MC identification code of 14 (ν_μ) [33], the interaction had to occur within the fiducial volume of either FGD, and the final state of the interaction required a muon and

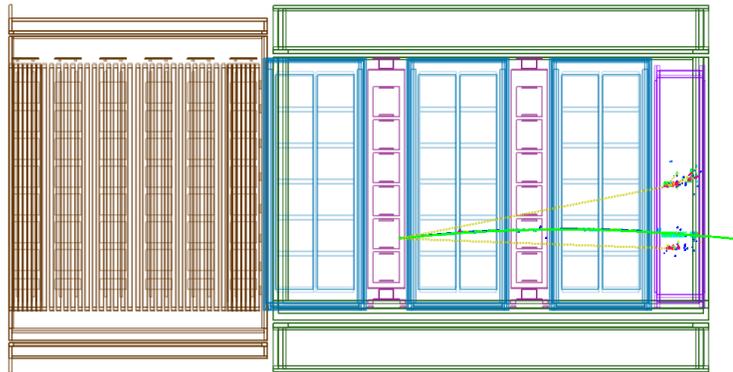
Topology Code	Interaction FGD	1 st Shower Location	2 nd Shower Location
FGD1DsDs	1	DsECal	DsECal
FGD1DsBr	1	DsECal	BrECal
FGD1BrBr	1	BrECal	BrECal
FGD2DsDs	2	DsECal	DsECal
FGD2DsBr	2	DsECal	BrECal
FGD2BrBr	2	BrECal	BrECal

Table 5.3: Codes used to describe the six event topologies of interested in this analysis based upon the interaction FGD and where the π^0 decay photons shower.

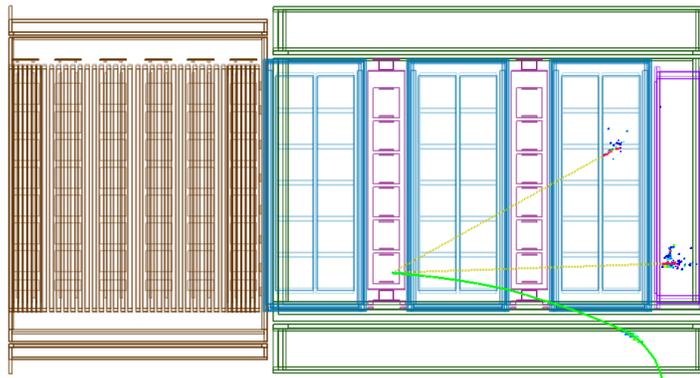
a single neutral-pion, along with any number of nucleons and photons, to emerge from the nuclear environment.

Of these signal events, this analysis is specifically designed to select the events whereby the π^0 decays to a pair of photons which both propagate to and shower within the Downstream and/or Barrel ECals. This leads to six different event topologies to reconstruct depending upon which FGD the interaction occurred within and where the pair of decay photons shower, these are described by Table 5.3. Figures 5.2 and 5.3 show example event displays for each of the three topologies of interest for FGDs 1 and 2 respectively. Within those figures, photon trajectories which appear to terminate within the TPC detector regions are showering within a side ECal module. Additionally simulation of the muon and π^0 was decoupled for these illustrative figures and so momentum is not conserved for the interactions shown.

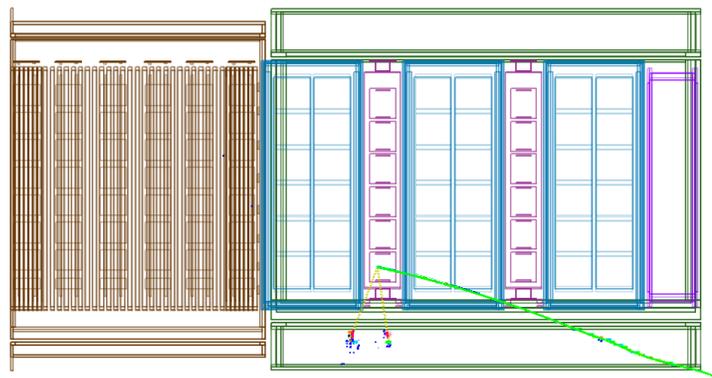
Table 5.4 shows the breakdown of final states for neutrino interaction within the fiducial volumes of the FGDs. Table 5.5 further shows the breakdown in event topology for the $CC1\pi^0$ final states. The majority, 82.4%, of $CC1\pi^0$ final states do not produce an interaction which results in a topology of interest to the analysis described in this chapter, these are denoted “Non-ECalECal” topologies. The Non-ECalECal topology final states are dominated by events where at least one of the π^0 decay photons has showered or pair-produced in material prior to the ECal modules, normally the FGD itself. A small contribution to Non-ECalECal final states comes from events where the path of at least one of the



(a) FGD1DsDs.

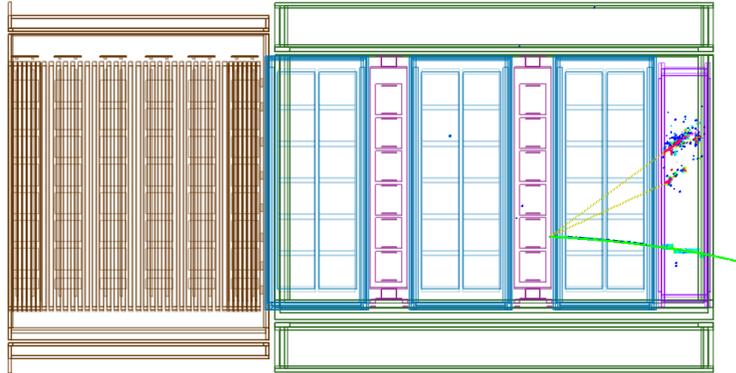


(b) FGD1DsBr.

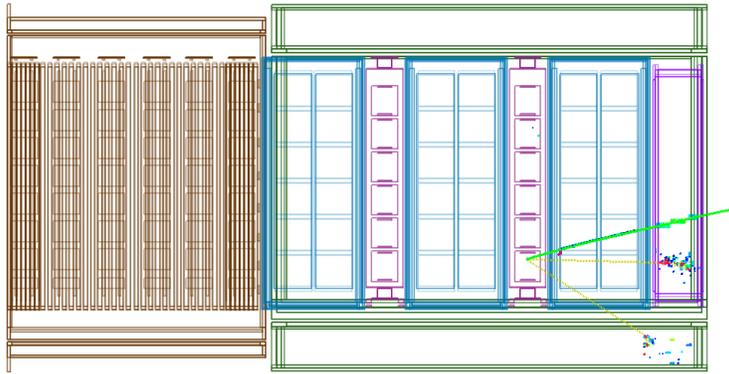


(c) FGD1BrBr.

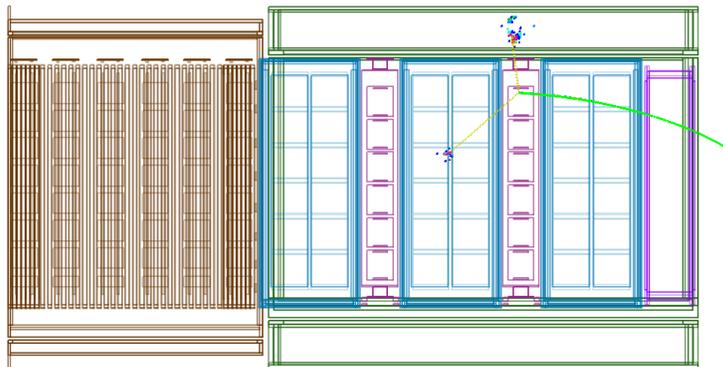
Figure 5.2: Example event displays for simulated interactions where both π^0 decay photons shower within the DsECal and/or BrECal modules for FGD1. Detector hits are shown with true particle trajectories overlaid in green for the muon and yellow for the decay photons.



(a) FGD2DsDs.



(b) FGD2DsBr.



(c) FGD2BrBr.

Figure 5.3: Example event displays for simulated interactions where both π^0 decay photons shower within the DsECal and/or BrECal modules for FGD2. Detector hits are shown with true particle trajectories overlaid in green for the muon and yellow for the decay photons.

Final State	Contribution (%)	Final State Particles (excluding nucleons and photons)
CCQE	61.8	μ
CC1 π^0	4.8	μ + single π^0
CCN π^0	1.0	μ + multiple π^0 's
CC π^0 +	9.5	μ + π^0 + any other particles (π^\pm , etc.)
CC Other	22.3	μ + any other particles excluding π^0 's
Other	0.7	All other final states

Table 5.4: Final state breakdown for FGD fiducial volume charged-current neutrino interactions.

Event Topology	Fraction of CC1 π^0 (%)	Fraction of All Final States (%)
FGD1DsDs	0.32	0.015
FGD1DsBr	0.92	0.044
FGD1BrBr	4.70	0.224
FGD2DsDs	3.40	0.162
FGD2DsBr	3.95	0.189
FGD2BrBr	4.28	0.204
FGD1 Non-ECalECal	43.9	2.10
FGD2 Non-ECalECal	38.5	1.84

Table 5.5: Further breakdown of CC1 π^0 final states by interaction vertex and π^0 decay topology.

photons has taken it into an uninstrumented region of the ND280.

The analysis is split into two sections, first a muon is selected within the fiducial volume of either FGD as described in Section 5.4. This is followed by the selection of π^0 decay photons in the ECals, this is where the analysis splits based upon the specific topology of the event and is described in Section 5.5.

5.4 Muon Selection

The first stage of the analysis aims to select muons from either FGD fiducial volume. The analysis is based upon the official ND280 ν_μ inclusive selection [92], but is implemented for vertices in either FGD fiducial volume (rather than just FGD1) and with looser selection criteria such that a larger sample of CC1 π^0 events propagate through to the π^0 selection phase of the analysis.

5.4.1 Muon Selection Criteria

Eight separate steps are applied in the selection of a muon candidate track, these are:

1 Event Quality,

Data events must have an associated good spill flag as previously described in Section 5.2.1. This cut has no effect in MC. All efficiencies and purities are measured relative to the number of events passing this cut.

2 Spill Bunching,

Reconstructed objects are separated in time into the eight individual bunches that are present in each spill. Subsequent cuts are then applied to each separate bunch in turn.

3 Negative FGD Fiducial Volume,

A minimum of one negative track must be present with both FGD and TPC components that originates from the fiducial volume of either FGD.

4 Track Quality,

At least one track passing the negative FGD fiducial volume cut must have produced a minimum of 18 vertical clusters in each TPC the track traverses. This requirement rejects short tracks for which the TPC, and therefore PID and momentum measurement, would otherwise be poor. Figure 5.4 shows the number of vertical clusters produced (nodes) by tracks crossing each TPC in a *magnet* MC sample after the previous cuts have been applied.

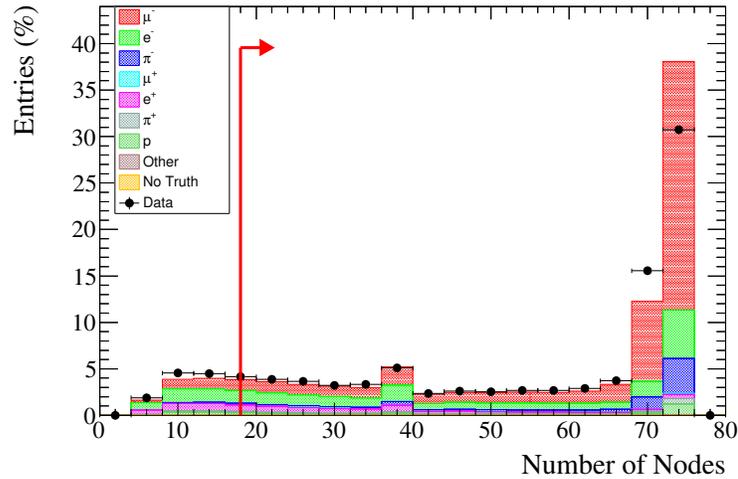


Figure 5.4: Number of TPC vertical clusters (nodes) crossed by tracks passing the Negative FGD Fiducial Volume cut. The red arrow indicates the tracks which pass the Track Quality cut criteria.

As can be seen such a cut mainly removes electron tracks, along with 6% of the muon tracks passing the previous cut.

5 Forward Going,

At least one track passing the track quality cut must reconstruct as forward going; its end point must be downstream of the track start position. This requirement is implemented as tracks which appear to be backwards going are often forward going positive tracks that have had their direction misreconstructed. 6% of muon tracks passing the previous cut are rejected by this cut.

6 TPC PID Quality,

At least one track passing the forward going track cut must have a good PID pull (as described in Section 3.3.1) in every TPC crossed. A good TPC pull is defined as an absolute pull value of less than four for any of the five particle hypotheses, $|pull_\alpha| < 4$ where $\alpha = \mu, e, p, \pi$ or K . Similar to the earlier track quality cut, this ensures that the track has been sufficiently well reconstructed to give a sensible PID in all TPCs crossed. Figure 5.5 shows each of the TPC pull hypothesis distributions in a *magnet* MC sample after

the previous cuts have been applied.

7 TPC Muon PID,

At least one track passing the TPC PID quality cut must simultaneously have a muon-like and not electron-like TPC PID in one of the TPCs crossed. A muon-like PID is defined as absolute muon pull of less than three, $|pull_\mu| < 3$, whilst a not electron-like PID is defined as an electron pull of less than minus one, $pull_e < -1$. Figure 5.6 shows the TPC muon and electron pull hypothesis distributions in a *magnet* MC sample after the previous cuts have been applied.

8 Muon Multiplicity,

After those cuts have been applied, in the majority of cases, about 95%, only one negatively charged track remains as the muon candidate and that track is selected as such. However, events will be rejected if multiple muon candidates remain, as this is most likely due to charged pions being misidentified as muons. Given that this analysis aims to select $CC1\pi^0$ interactions, events that likely contain charged pions in the final state should be rejected.

In *magnet* MC, the above cuts produce a CC inclusive-like event selection with an efficiency of 50% and purity of 86% as is shown in Figure 5.7 and described in detail in Table 5.6. The track selected as the muon candidate is a true muon 87% of the time with the dominant background being from misidentified negative pions as shown in Table 5.7.

5.4.2 Muon Selection Parameter Space

The true muon momentum and direction distributions in MC, prior to any cuts are shown in Figure 5.8. Figure 5.9 then shows the same distributions as Figure 5.8 for true muons selected by the analysis cuts described in Section 5.4.1.

The efficiency for selecting muons as a function of their momentum and direction is shown in Figure 5.10. Given the presence of a forward going track cut, and that any tracks moving at high angles will pass directly between the FGDs and BrECals, or traverse only a small region of the TPCs, it is clear that the above

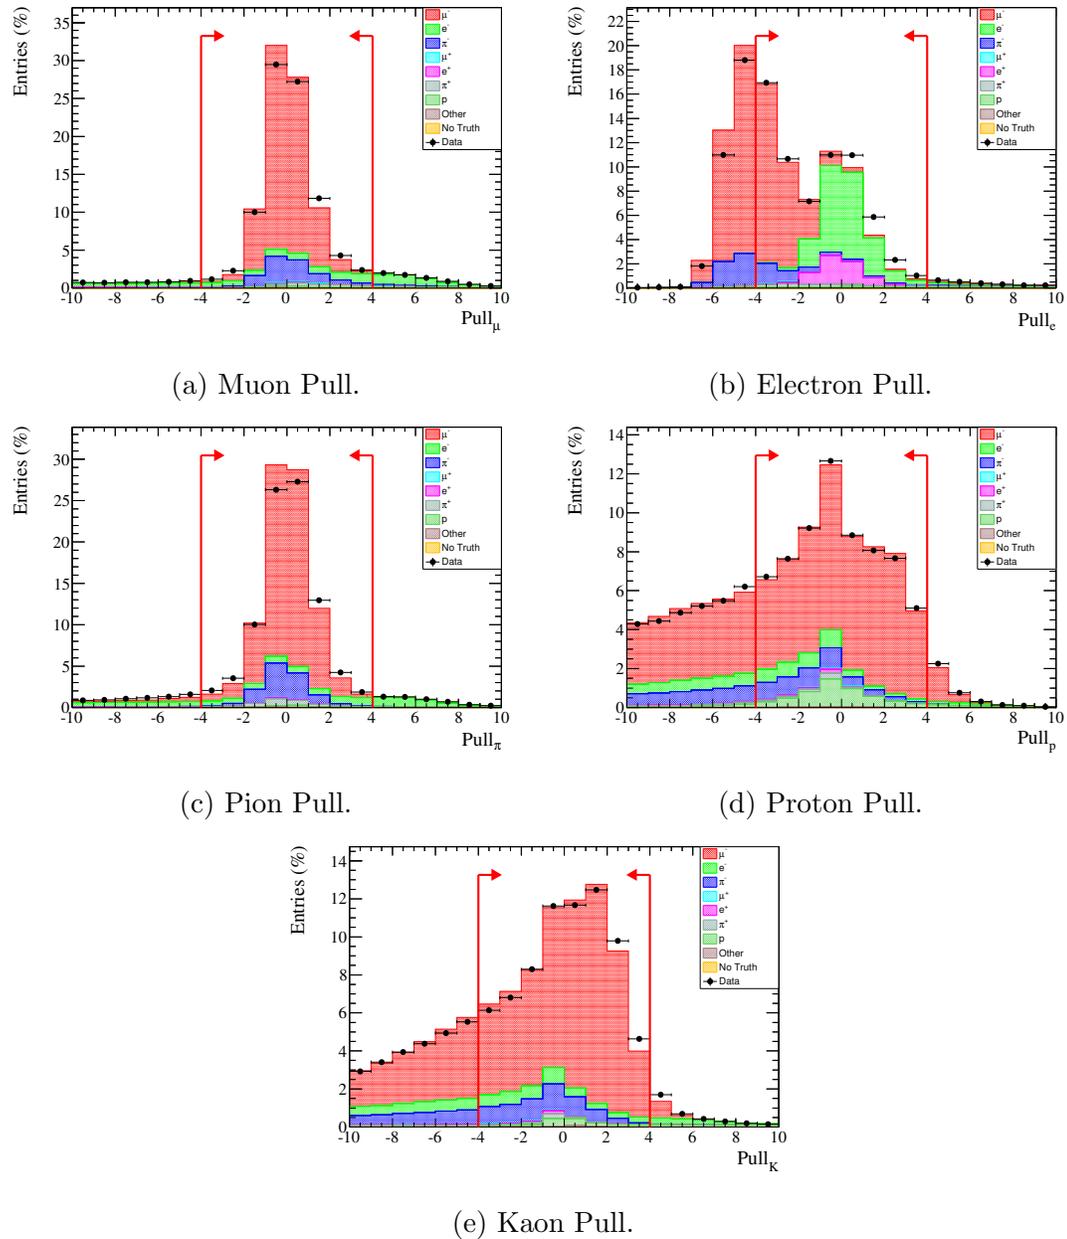
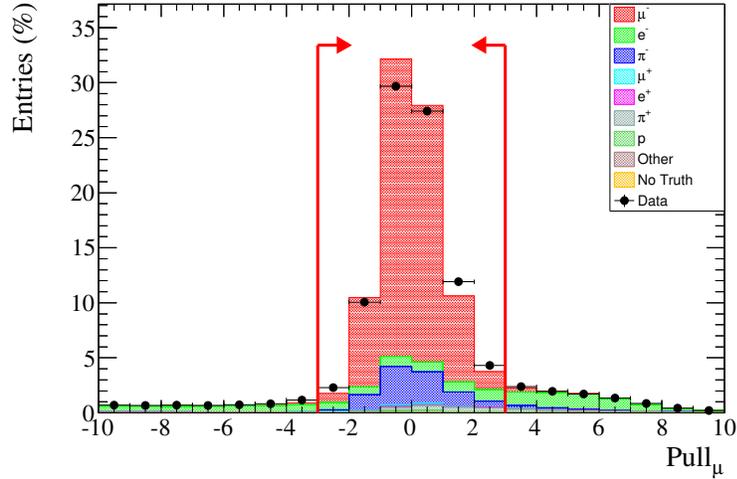
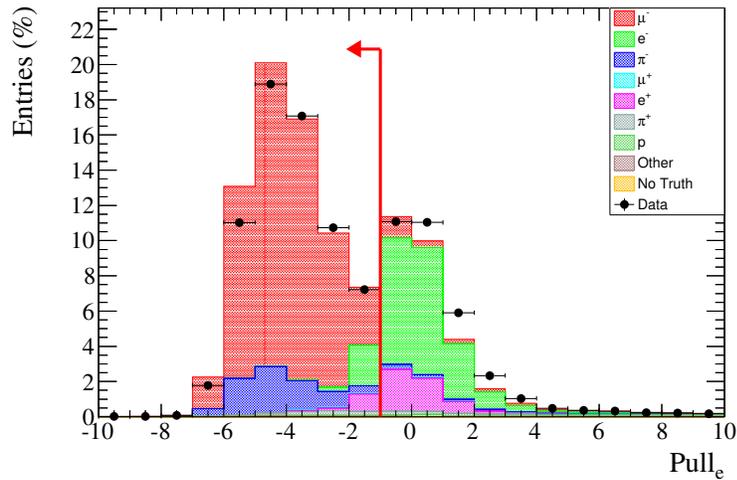


Figure 5.5: TPC pull hypothesis distributions for tracks passing the Forward Going track cut. Tracks must have a pull within the selection criteria, as indicated by the red arrows, for at least one of the TPC PID hypotheses in each TPC crossed.



(a) Muon Pull.



(b) Electron Pull.

Figure 5.6: TPC muon pull hypothesis, 5.6a, and electron pull hypothesis, 5.6b, distributions for tracks passing the Track PID Quality cut. Tracks must have a pull within the selection criteria, as indicated by the red arrows, for both the muon and electron TPC PID hypotheses simultaneously in at least one TPC that is crossed.

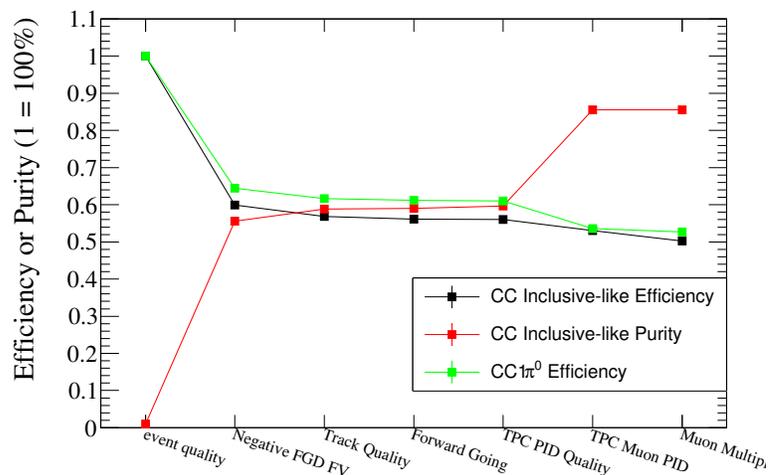
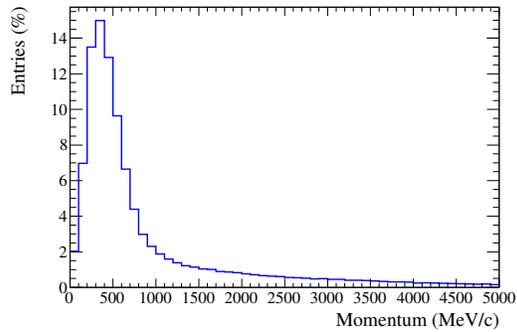


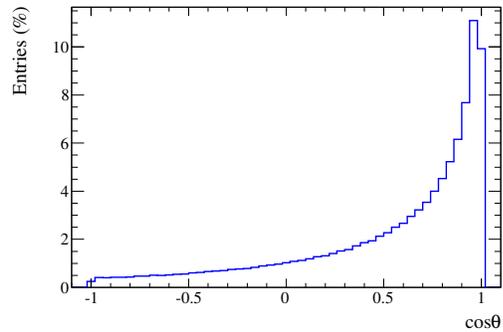
Figure 5.7: The efficiency and purity for selecting CC inclusive-like events, along with the efficiency for selecting CC1 π^0 , events as a function of cut.

Cut	Total Events	CC Inclusive Events	Efficiency (%)	Purity (%)
Event Quality	5317850.0	90434.73	100	1.7
Negative FGD FV	97460.79	54177.73	59.9	55.6
Track Quality	87357.34	51381.60	56.8	58.8
Forward Going	85964.20	50735.62	56.1	59.0
TPC PID Quality	84919.46	50639.63	56.0	59.6
TPC Muon PID	56081.30	47996.92	53.1	85.6
Muon Multiplicity	53111.01	45447.56	50.3	85.6

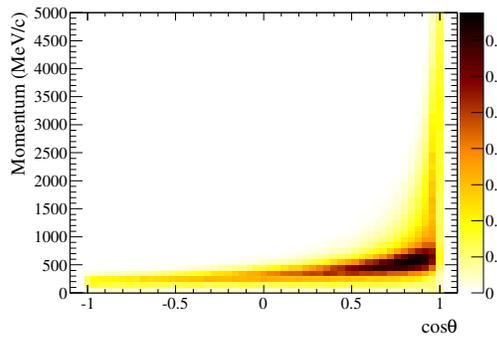
Table 5.6: MC prediction for the number of events selected as a function of cut, and associated CC inclusive selection efficiency and purity for all running periods to be analysed.



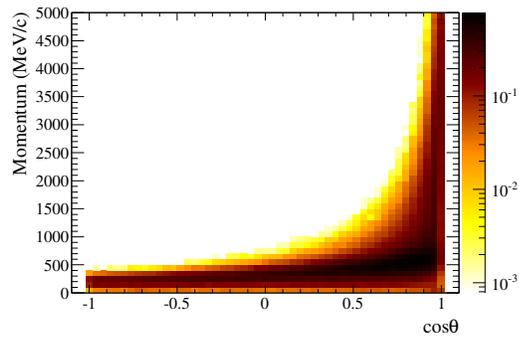
(a) True muon momentum.



(b) True muon direction.

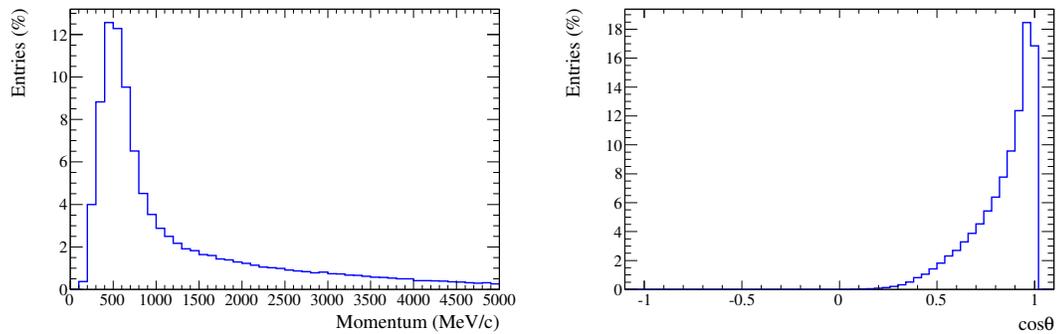


(c) True muon momentum as a function of direction.



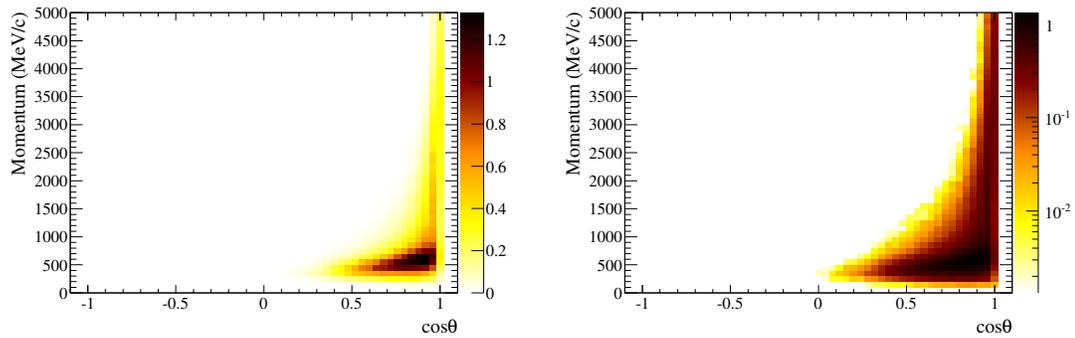
(d) True muon direction as a function of momentum, with logarithmic z axis.

Figure 5.8: The true momentum and direction of FGD FV muons prior to any selection criteria have been applied. 5.8c and 5.8d show the same information but with linear and logarithmic z axis scales, in both cases the z axis unit is the percentage of entries in each histogram bin.



(a) True muon momentum.

(b) True muon direction.



(c) True muon momentum as a function of direction.

(d) True muon direction as a function of direction, with logarithmic z axis.

Figure 5.9: The true momentum and direction of FGD FV muons selected by the criteria as described in Section 5.4. 5.9c and 5.9d show the same information but with linear and logarithmic z axis scales, in both case the z axis unit is the percentage of entries in each histogram bin.

True Particle Type	Selected Muon Candidate (%)
μ^-	86.9
e^-	1.1
π^-	9.7
μ^+	0.3
e^+	< 0.1
π^+	1.5
p	0.3
Other	0.1

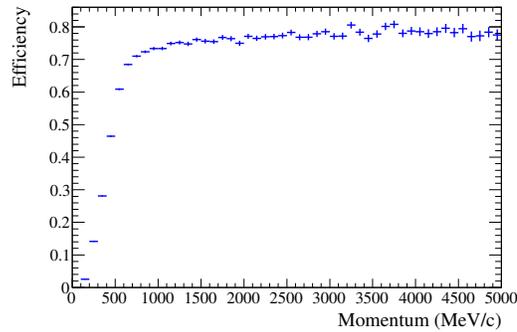
Table 5.7: Breakdown of true particle type for the track selected as the muon in the CC inclusive-like selection.

selection is somewhat restrictive in its muon kinematic sampling. This is best seen in Figure 5.10b showing the true muon selection efficiency as a function of direction, there is no selection efficiency in the backwards direction, and the efficiency then builds from zero at $\cos\theta = 0$, to nearly 80% for the most forward going tracks at $\cos\theta = 1$.

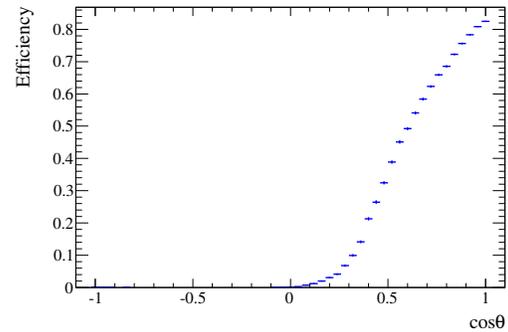
The high angle and backward going tracks are those with a lower momentum, < 500 MeV/c as can be seen in Figure 5.8c. Leading to a deficit in selection of muons with these lower momenta, this is seen in Figure 5.10a where the selection efficiency is seen to steadily climb from zero to 70%, between $0 \rightarrow 700$ MeV/c. Beyond about 600 MeV/c the selection efficiency then nearly plateaus as muons of greater momentum tend to be well reconstructed moving in the forward direction.

5.4.3 Muon Selection Results

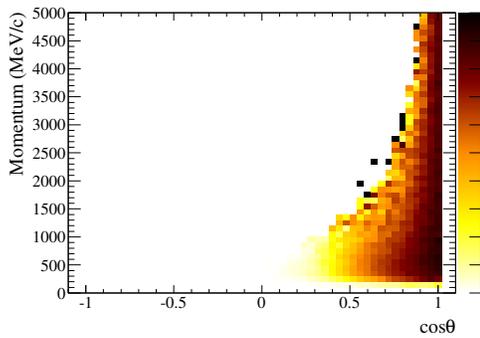
Applying the previously described selection criteria to each of the independent Run 2 to Run 4 data and MC sets produces muon candidate selection distributions as shown in Figure 5.11. Table 5.8 shows the POT scaled and flux weighted predicted number of events to pass each cut in *magnet* MC, along with the true number of events in data, broken down by running period. Table 5.9 then shows the data over MC ratio for the total number of events selected with the associated



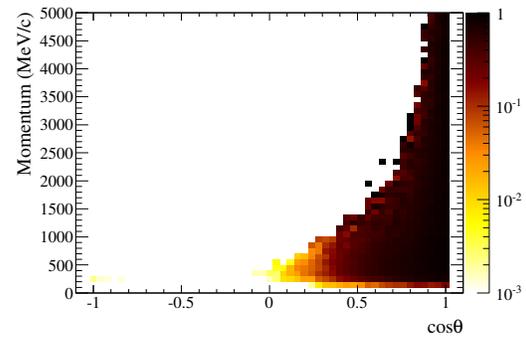
(a) True muon momentum.



(b) True muon direction.



(c) True muon momentum as a function of direction.



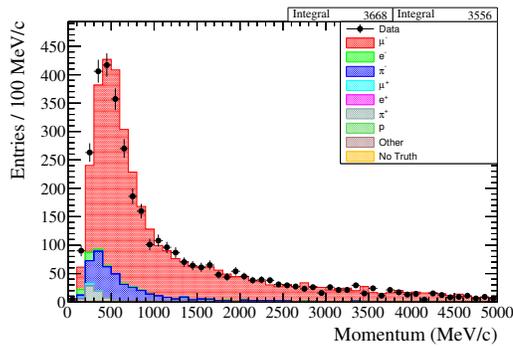
(d) True muon direction as a function of direction, with logarithmic z axis.

Figure 5.10: The efficiency for selecting FGD FV muons, using the criteria as described in Section 5.4, as a function of their true momentum and direction. 5.10c and 5.10d show the same information but with linear and logarithmic z axis scales, in both case the z axis unit is the selection efficiency in each histogram bin.

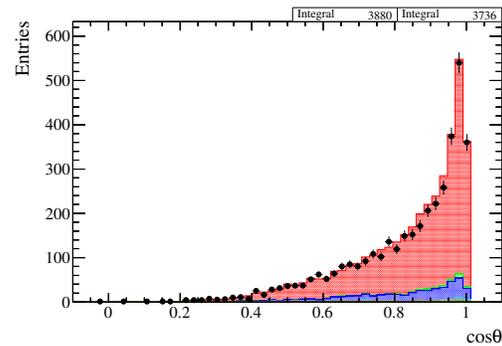
statistical error.

The selection shows an average 5% deficit in the number of selected data events over the MC prediction across all running periods. Figures 5.12 and 5.13 respectively show the momentum and direction of the track selected as the muon candidate for all running periods combined, along with the associated data over MC ratio. From Figure 5.13 it appears that the deficit is most acute for the more forward going tracks at $\cos\theta > 0.7$. Figure 5.12 shows that there is a significant excess in events with the lowest reconstructed momentum, < 400 MeV/c, which is immediately followed by a deficit in events between 400 and 1500 MeV/c, and again beyond 2500 MeV/c. These may point to an underestimate in the TPC momentum calculation in data, possibly due to calibration issues, or some unknown systematic effect, or may potentially be due to an overestimate in some aspect of the ND280 detector response in the MC simulation, or inaccuracies in the neutrino event generator. The observed muon deficit is not just isolated to this analysis, but is also observed within the official inclusive muon selection, which shows a 4.3% deficit, and related analyses.

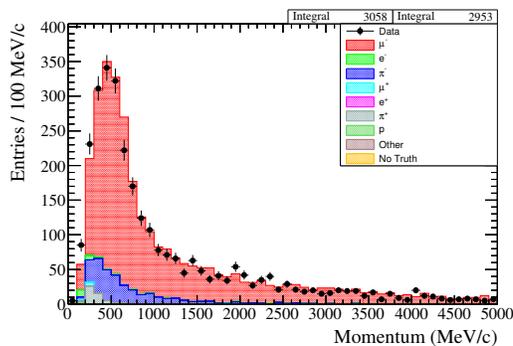
It should be noted at this time, that the significant excess in events passing the Event Quality cut in data are due to neutrino interactions in the Earth between the target and pit holding the ND280 detector; within T2K these are known as sand interactions. Sand interactions produce a flux of muons which propagate into the pit and pass through the ND280, and are reconstructed as through-going tracks (such interactions can be very useful for calibration purposes). These sand interactions are not present in the *magnet* MC sample used in this analysis, but appear in the data leading to the excess at the Event Quality cut. MC samples, known as “sand” MC, are available which model these interactions and brief studies by the author showed that almost all sand interactions failed to pass the Negative FGD FV cut, hence the return to expected data/MC ratios after the application of this cut, and that none propagated through the full π^0 selection, the main aim of this analysis. As such, the sand MC is not used in this analysis as it would be a waste of time and computing power, but does leave us with the data excess at the Event Quality stage.



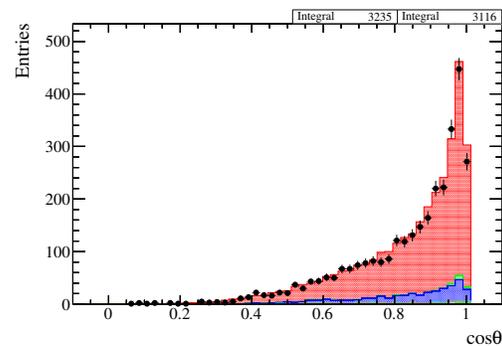
(a) Muon candidate momentum distribution for Run 2 Water.



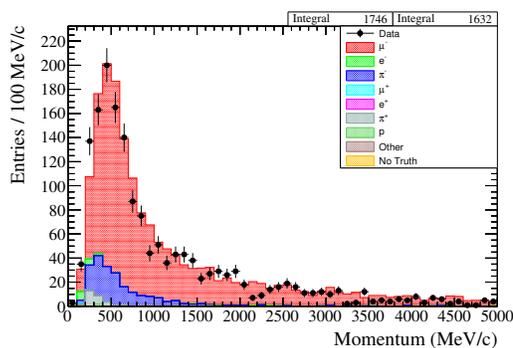
(b) Muon candidate direction distribution for Run 2 Water.



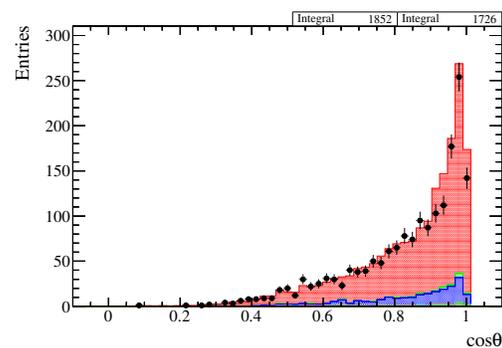
(c) Muon candidate momentum distribution for Run 2 Air.



(d) Muon candidate direction distribution for Run 2 Air.

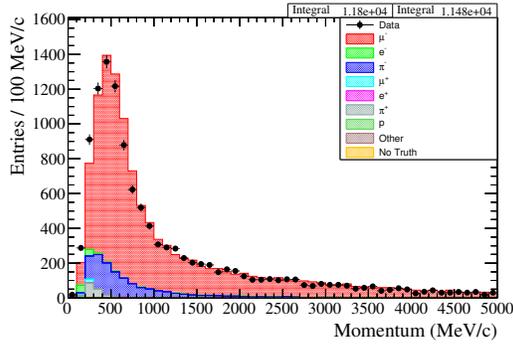


(e) Muon candidate momentum distribution for Run 3b (Air).

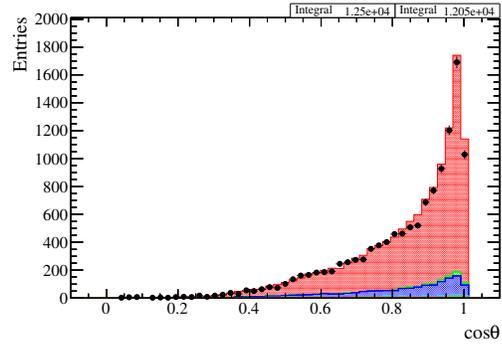


(f) Muon candidate direction distribution for Run 3b (Air).

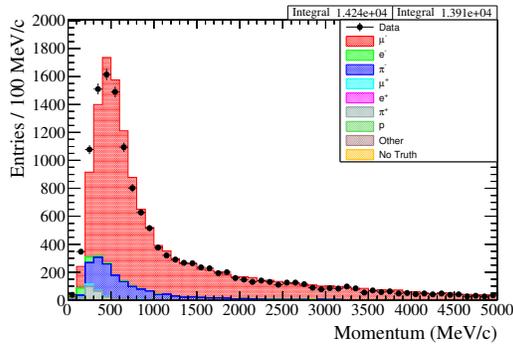
Figure 5.11: The measured momentum and direction for the track selected as the muon candidate for running periods Run 2 Air to Run 3b. The integral values at the top of each histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right.



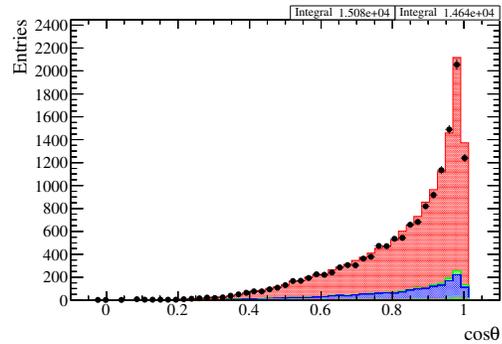
(g) Muon candidate momentum distribution for Run 3c (Air).



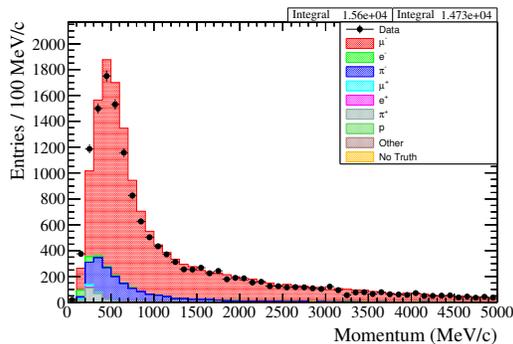
(h) Muon candidate direction distribution for Run 3c (Air).



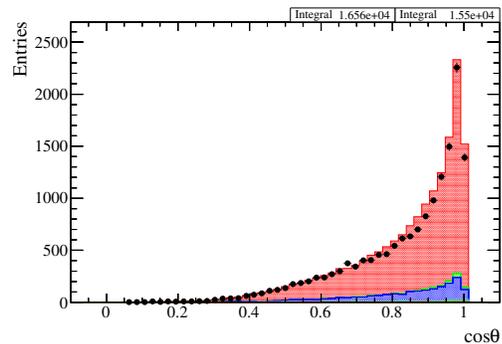
(i) Muon candidate momentum distribution for Run 4 Water.



(j) Muon candidate direction distribution for Run 4 Water.



(k) Muon candidate momentum distribution for Run 4 Air.

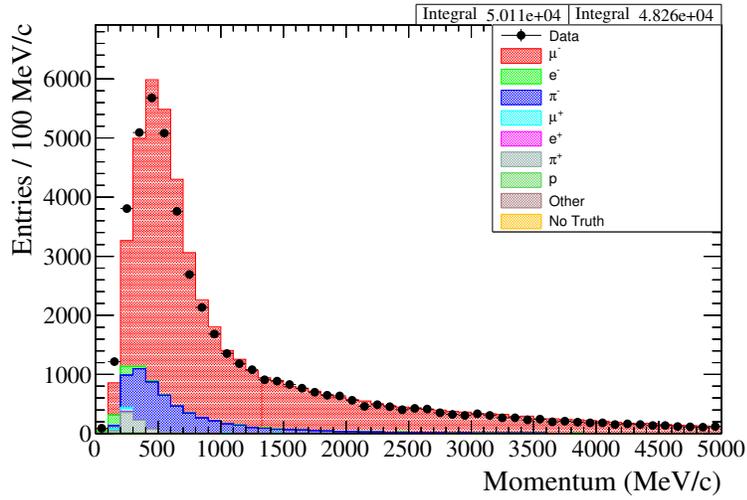


(l) Muon candidate direction distribution for Run 4 Air.

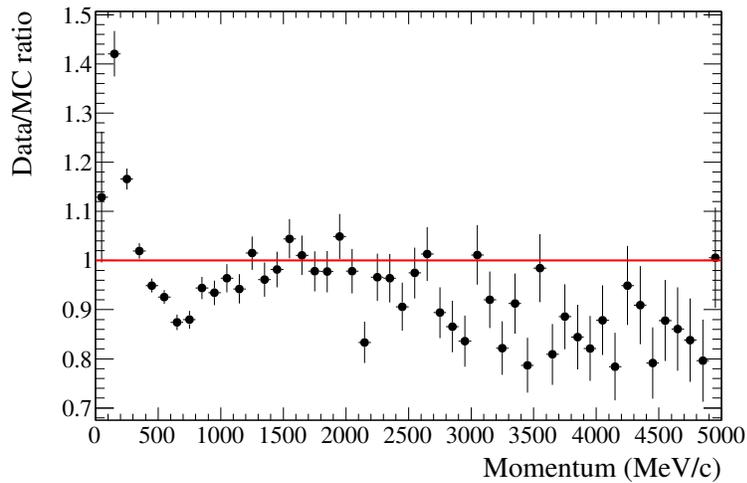
Figure 5.11: The measured momentum and direction for the track selected as the muon candidate for running periods Run 3c to Run 4 Air. The integral values at the top of each histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right.

Cut	Number of selected events in data and (MC)						<i>Total</i>
	Run 2 Water	Run 2 Air	Run 3b	Run 3c	Run 4 Water	Run 4 Air	
Event Quality	549089	(401680.94)	449988	(331051.5)	252299	(197414.33)	7221641 (5317850.0)
Negative FGD FV	7234	(7185.08)	5931	(5950.95)	3408	(3411.03)	97582 (97460.79)
Track Quality	6336	(6437.11)	5230	(5337.77)	3024	(3062.24)	85904 (87357.34)
Forward Going	6238	(6335.54)	5153	(5247.93)	2988	(3013.43)	84586 (85964.20)
TPC PID Quality	6139	(6263.36)	5082	(5186.67)	2933	(2976.77)	83281 (84919.46)
TPC Muon PID	3959	(4096.02)	3312	(3413.02)	1832	(1955.14)	53918 (56081.30)
Muon Multiplicity	3736	(3880.61)	3116	(3235.90)	1726	(1852.83)	50771 (53111.01)
Event Quality	1722363	(1245976.62)	2052889	(1510625.38)	2195013	(1631101.12)	7221641 (5317850.0)
Negative FGD FV	22952	(22973.30)	28184	(27788.76)	29873	(30151.67)	97582 (97460.79)
Track Quality	20239	(20578.36)	24797	(24875.02)	26278	(27066.85)	85904 (87357.34)
Forward Going	19920	(20241.93)	24392	(24480.48)	25877	(26644.89)	84586 (85964.20)
TPC PID Quality	19618	(19990.68)	24029	(24180.80)	25480	(26321.18)	83281 (84919.46)
TPC Muon PID	12821	(13213.31)	15486	(15911.41)	16508	(17492.40)	53918 (56081.30)
Muon Multiplicity	12049	(12503.10)	14641	(15078.43)	15503	(16560.15)	50771 (53111.01)

Table 5.8: The number of selected events in data and (MC) for each data taking period as a function of analysis cut.

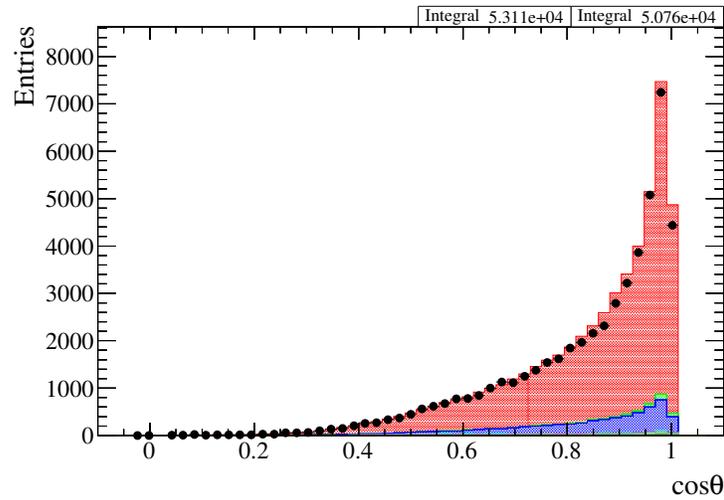


(a) Muon candidate momentum distribution for all running periods.

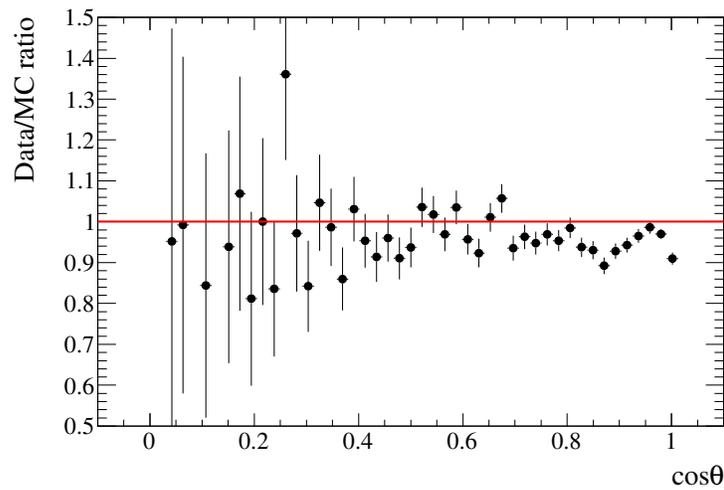


(b) Data over MC ratio for 5.12a.

Figure 5.12: 5.12a shows the measured momentum of the track selected as the muon candidate for all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.12b shows the data over MC ratio as a function of track momentum from 5.12a.



(a) Muon candidate direction distribution for all running periods.



(b) Data over MC ratio for 5.13a.

Figure 5.13: 5.13a shows the measured direction of the track selected as the muon candidate for all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.13b shows the data over MC ratio as a function of track momentum from 5.13a.

Cut	Total number of selected events		$\frac{data}{MC}$ ratio
	Data	(MC)	
Event Quality	7221641	(5317850.0)	1.358 ± 0.001
Negative FGD FV	97582	(97460.79)	1.001 ± 0.004
Track Quality	85904	(87357.34)	0.983 ± 0.004
Forward Going	84586	(85964.20)	0.984 ± 0.004
TPC PID Quality	83281	(84919.46)	0.981 ± 0.004
TPC Muon PID	53918	(56081.30)	0.961 ± 0.005
Muon Multiplicity	50771	(53111.01)	0.956 ± 0.005

Table 5.9: The total number of selected events in data and (MC) for each data taking period as a function of analysis cut. The data over MC ratio is along with its statistical error.

5.5 π^0 Selection

Upon the successful selection of a muon candidate, as described previously in Section 5.4, attempts are made to reconstruct a π^0 candidate with the Barrel and Downstream electromagnetic calorimeters. The basic requirements of the selection are described in the next section, with the details of the Multivariate Analysis (MVA) being described in detail subsequently in Section 5.5.2.

5.5.1 π^0 Selection Criteria

Six separate steps are applied in the selection of the π^0 candidate, these are:

1 ECal Multiplicity,

The first π^0 selection cut requires there to be a minimum of two isolated ECal clusters reconstructed in the bunch. Furthermore, it is exceedingly rare ($<2\%$) for a $CC1\pi^0$ interaction, where both decay photons entered the ECals and it has passed the muon selection, to contain more than three isolated ECal clusters. As such, bunches containing four or more isolated ECal clusters are also rejected. Figure 5.14 shows the number of isolated ECal clusters as a function of final state.

2 Select Pairwise Combinations,

If two or three isolated ECal clusters are present within the bunch, these are combined with each other in all possible pairwise combinations. One pairwise combination for the bunches with just two isolated ECal clusters, and three pairwise combinations for events with three isolated ECal clusters.

3 Reconstructed Mass,

At least one pair of isolated ECal clusters must combine to produce an invariant mass less than $500 \text{ MeV}/c^2$. Given that the true invariant mass of a π^0 is $135 \text{ MeV}/c^2$, a cut of $500 \text{ MeV}/c^2$ is sufficiently high as to not remove any true signal events, whilst slightly reducing the background pairs to be evaluated by the computationally intensive Multivariate Selection.

4 Multivariate Selection,

At least one pair of clusters passing the Reconstructed Mass cut must pass selection by a Boosted Decision Tree (BDT) multivariate analysis (MVA). Six MVA's are prepared, one for each of the six event topologies of interest, the details of which are described in the following section.

5 Pairs Multiplicity,

After the above cuts have been applied, in the majority of cases, about 92%, only one pair of isolated ECal clusters remains as the π^0 candidate and that pair is selected as such. However, bunches with multiple pairs of isolated ECal clusters passing the Multivariate Selection are rejected, as for any signal topology there should only be one pair of isolated ECal clusters that come from a true π^0 decay. Additionally, given the use of the earlier ECal Multiplicity cut, at least one of the isolated ECal clusters would be contributing to multiple multivariate selected pairs, something that should not occur for a well reconstructed signal π^0 decay.

6 Topology Selection,

If one pair of isolated ECal clusters has been selected as the candidate π^0 photon decay pair, the bunch is tagged as a signal event and partitioned based upon the reconstructed event topology.

5.5.2 Multivariate Analysis

To identify the pairs of isolated ECal clusters that are likely to originate from the decay of a π^0 a MVA selection is implemented using the TMVA software package [93] integrated within ROOT. The use of a MVA technique is motivated by the presence of minor differences between the signal and background distributions for a large range of ECal variables, as seen in Figure 5.22. The use of a straight cuts based approach with such poor discriminating variables would likely lead to a low efficiency and purity selection for signal events, and would also bias the selection towards the areas of the π^0 phase-space where the decay photon kinematics are most distinct from the background continuum. In these situations it is clear that building a discriminator which takes into account the correlations between multiple input variables would have the potential to produce a significantly improved signal and background separation; this is the MVA technique.

The TMVA software package provides a user friendly interface to train, optimise, apply and compare a large selection of MVA techniques. The MVA technique used by this analysis is a Boosted Decision Tree as this is found to be the most effective discriminator for this analysis, see Figure 5.15. BDT discriminators are becoming increasingly popular within the high energy physics community as they have the advantages of being simple to use, quick to train and rapidly produce very effective results. The Multilayer Perceptron (MLP), a type of feedforward artificial neural network, produces near identical discriminating power as the BDT, based on the ROC curve results. At one time MLPs were very popular MVA techniques, however they have fallen out of favour due to their complexity, requirement for very careful training and testing, and susceptibility to over-training.

Boosted Decision Tree Technique

A decision tree takes a set of discriminating variables and places binary branching cuts against them in turn. Each cut applied is the one that provides the best signal and background separation using the discriminating variable distributions available at that branching point. After the build up of branching cuts, the final

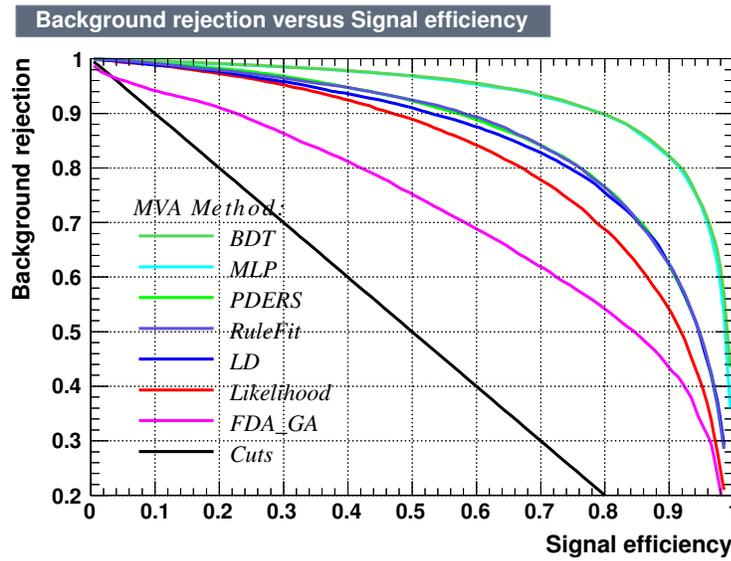


Figure 5.15: Receiver Operating Characteristic (ROC) curve for a range of MVA techniques. The better the technique performance, the closer to the top right hand corner ($signal\ efficiency = background\ rejection = 1$) of the plot the ROC curve will extend. The above curves are produced from training against the discriminating variables of the FGD2DsBr event topology training sample, described in detail in Section 5.5.3. The discriminating performance of the different MVA techniques for the FGD2DsBr event topology is typical of all event topologies.

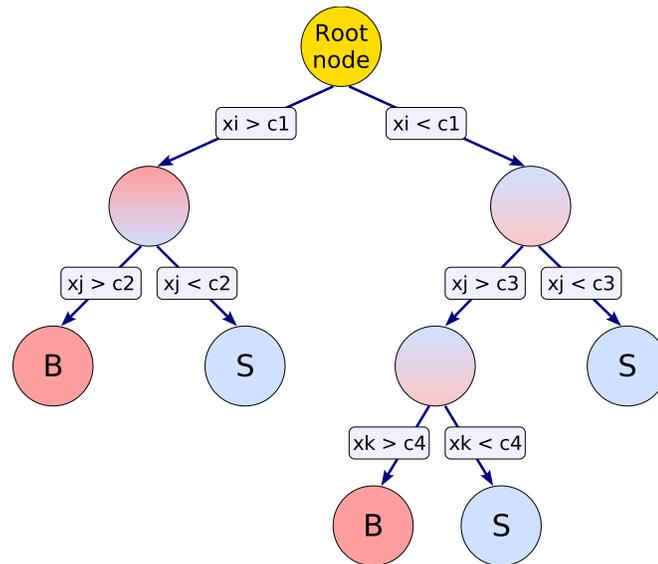


Figure 5.16: From the discriminating variables available at the root node, branching cuts are applied one at a time until samples are sufficiently segregated into signal-like and background-like leaf nodes [93].

leaf nodes at the end of the resultant branch structure are described as signal-like or background-like as shown in Figure 5.16. On their own, decision trees are poor multivariate discriminators, they are highly susceptible to fluctuations in the input training sample, can end up neglecting potentially useful variables, and be biased towards small regions of the signal phase-space. This is where it becomes useful to use a boosting method to improve the power of the decision tree method.

Training a Decision Tree The training, or “growth” of a decision tree begins by taking a sample of signal and background events with attributed discriminating variables; this is the training sample. In turn, each discriminating variable is considered, and a cut value is scanned across the full range of values the variable may take. After each cut is applied the number of signal and background events present either side of the cut is found and a separation parameter calculated. A range of different separation parameters are available (see the TMVA users guide for details [93]), but we shall consider the Gini Index, G , as used within this

analysis

$$G = p(1 - p), \quad (5.1)$$

where p is the sample purity, defined as the number of signal events, S , divided by the total number of signal and background, B , events passing the cut

$$p = \frac{S}{(S + B)}. \quad (5.2)$$

By finding the variable and cut value which minimises the Gini Index the optimal signal and background separation is achieved. Applying such a cut produces a branching node, with a signal-like sample to one side and background-like sample to the other. The branching process is then repeated for every resultant node, unless the resultant node either contains only signal or background events, contains too few events, or the number of cut branches has reached some pre-defined maximum depth. If one such criterion is reached then the resultant node is a final leaf node for the decision tree, and these are characterised as signal-like or background-like based upon their signal purity.

Boosting a Decision Tree

Once a decision tree has been trained, a potentially significant number of signal sample events are likely to have been misidentified as background-like, and vice versa for the background sample events. These misidentified events are misclassified, and it would be ideal to attempt to reclassify them into the correct samples. To do this the original samples are passed to another decision tree, but an additional weighting is given to events which were misclassified by the previous iteration of the decision tree training. There are several ways available to apply this event reweighting (boosting) process (see the TMVA users guide for details [93]), but we shall consider the Adaptive Boost (AdaBoost) [94] technique as used within this analysis.

The error rate, E_t , of a given decision tree, t , is defined as the sum of the weights,

w , of all misclassified, mc , events divided by the total sum weight of all events

$$E_t = \frac{\sum_{i=1}^{(S+B)_{mc}^{Total}} w_i}{\sum_{j=1}^{(S+B)^{Total}} w_j}. \quad (5.3)$$

Using the error rate it is then possible to define a boost weight, α , to be applied to all previously misclassified events from the previous decision tree training

$$\alpha_t = \frac{1 - E_t}{E_t}. \quad (5.4)$$

The boost weight is applied to each misclassified event such that

$$w_i^{t+1} = w_i^t \alpha^\beta, \quad (5.5)$$

where β is a user defined learning rate factor.

By definition $E_t \leq 0.5$ and the learning rate factor is set such that $\beta > 0$, ensuring that the misclassified event weights always increase $\alpha^\beta > 1$. The learning rate factor controls how rapidly the misclassified events increase in importance, allowing a slowing in the rate at which they become a priority to correctly identify.

All of the events are then reweighted, such that the total weight before and after application of the boost weight are equal. A new decision tree can then be trained with the new event weights. Looking back at the definition of the separation parameter in Equation 5.1, it is clear the reweighting of events would have no effect unless the separation parameter took this into account. Therefore the purity in Equation 5.2 needs to be redefined to account for the event weighting

$$p = \frac{\sum_{i=1}^S w_i}{\sum_{j=1}^{S+B} w_j}. \quad (5.6)$$

After repeatedly training decision trees and reweighting their misclassified events a forest of decision trees is built up, each of which should be optimised for discriminating between a different subset of signal and background events through the use of different discriminating variables and cut values.

To finally use a BDT classifier to evaluate whether an event is signal-like or background-like, it is tested with every single decision tree in the forest that has been built up, and a response $R(\mathbf{x})$ calculated

$$R(\mathbf{x}) = \frac{1}{N_{tree}} \sum_{t=1}^{N_{tree}} \ln(\alpha_t) R_t(\mathbf{x}), \quad (5.7)$$

where \mathbf{x} is an ordered list of values for the discriminating variables used by the BDT for the event being tested.

$R_t(\mathbf{x})$ is defined such that it equals +1 if an event is classified as signal-like, or -1 if classified as background-like by the decision tree $R_t(\mathbf{x})$. A cut can then be placed on the classification response to decide whether an event should be accepted by an analysis or not.

This boosting approach has the effect of increasing the separating power of the classifier against a stand alone decision tree, whilst also reducing its vulnerability to fluctuations that may be present within the input training samples.

Figure 5.17 shows the effect on BDT performance by varying the number of trees in the BDT forest and beta learning rate factor. It can be seen in general that increased separation is achieved with an increase in both parameters, however the significance of the classifier response generally decreases with an increase in both parameters.

The separation, S , of an MVA classifier, y , is defined as the integral

$$\langle S^2 \rangle = \frac{1}{2} \int \frac{(\hat{y}_S(y) - \hat{y}_B(y))^2}{\hat{y}_S(y) + \hat{y}_B(y)} dy, \quad (5.8)$$

where \hat{y}_S and \hat{y}_B are respectively the signal and background response distributions, such as those shown in Figure 5.18, for the classifier y .

The significance, \mathcal{S} , of an MVA classifier, y , is defined as

$$\mathcal{S} = \frac{|\bar{y}_S - \bar{y}_B|}{\sqrt{\frac{\sum_{i=1}^S y_S^2(i)}{S} + \frac{\sum_{j=1}^B y_B^2(j)}{B}}}, \quad (5.9)$$

where $|\bar{y}_S - \bar{y}_B|$ is the difference between the classifier means for signal, \bar{y}_S , and

background, \bar{y}_B . $\sqrt{\frac{\sum_{i=1}^S y_S^2(i)}{S} + \frac{\sum_{j=1}^B y_B^2(j)}{B}}$ is the quadratic sum of their root-mean-squares.

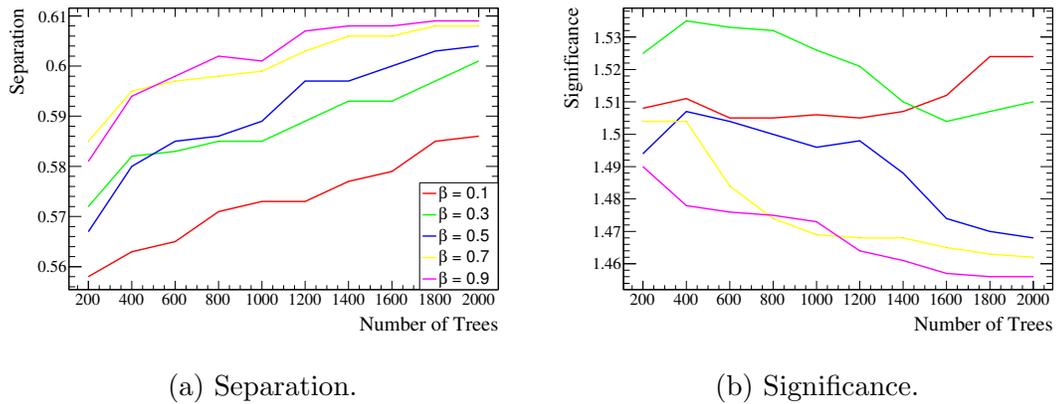


Figure 5.17: Separation, 5.17a, and Significance, 5.17b, achieved from training multiple BDTs with varied numbers of trees in the forest and beta learning rate factors for the FGD2DsBr topology.

Testing a Boosted Decision Tree

When training a classifier, not all signal and background events are used in growing the trees, a percentage of events are held back to subsequently test the performance of the resultant classifier. This is to ensure the performance of the classifier is as expected, an important test to check the classifier has not been over-trained, a potentially common pitfall when using certain classification methods; or when too small, and therefore an unrepresentative training sample has been used. Over-training most clearly manifests when a classifier is very good at discriminating between signal and background events within the training sample it has been provided with, but performs badly on an equally representative sample of signal and background events.

The easiest way to check for over-training is to measure the response, $R(\mathbf{x})$, of the classifier for both the training event sample and the testing events that were held back at the training stage. A statistical test, such as the Kolmogorov-Smirnov test [95], can then be used to ensure the resultant response distributions match one-another, and therefore the classifier hasn't been biased towards selecting events from only the training sample, see Figure 5.18.

Figure 5.19 shows the over-training test results as the number of trees in the BDT forest and beta learning rate factor are varied. It can be seen, that in general, the

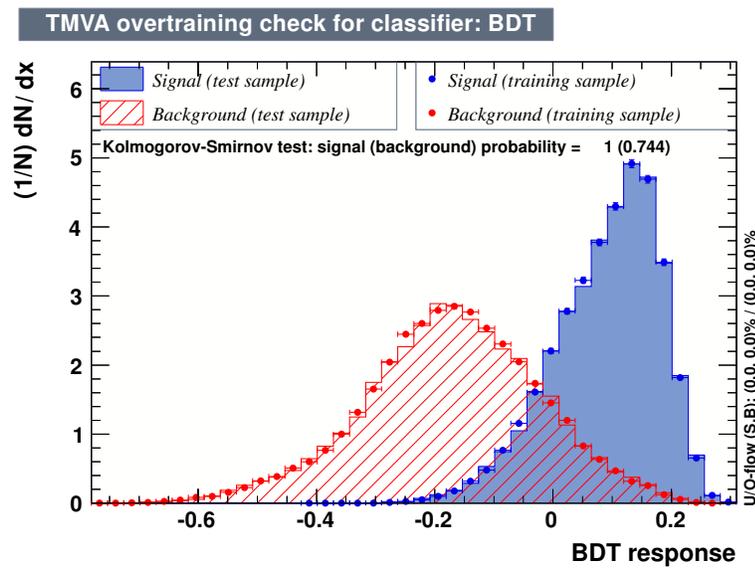


Figure 5.18: BDT response for FGD2DsBr training and testing samples with $\beta = 0.5$ and 1200 trees in the training forest. The results of a KS test to compare the training and testing sample show the two signal distributions are in near perfect agreement, but some minor disagreement is found for the background samples.

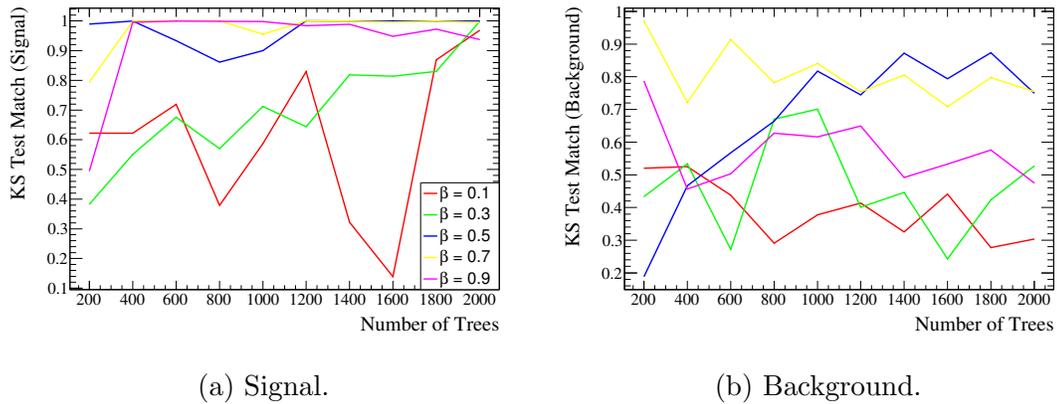


Figure 5.19: Strength of the matching between the training and testing sample response distributions for signal, 5.19a, and background, 5.19b, as obtained with a KS test from training multiple BDTs with varied numbers of trees in the forest and beta learning rate factors for the FGD2DsBr topology.

matching KS test shows poorer matching performance for the background sample compared to the signal sample. It is also seen that for both signal and background, generally higher beta learning factors produce better matched results, except for the highest tested factor, $\beta = 0.9$, which shows surprisingly poor performance for the background sample.

Once a BDT has been trained and tested, an appropriate BDT response cut value can be selected, such that only events exceeding some $R(\mathbf{x})$ value are selected by the analysis. The optimal cut value is found measuring the integrated number of signal and background events selected within the testing sample as function of BDT response value. A metric such as $Efficiency \times Purity$ or $S/\sqrt{S+B}$ can then be maximised and used as the cut value within an analysis, see Figure 5.20.

5.5.3 Boosted Decision Tree Analysis Implementation

Within this analysis the Boosted Decision Tree implemented for each event topology will use 1200 trees, a beta learning rate factor of 0.5 and a maximum of eight leaf nodes per decision tree in the forest. The use of eight leaf nodes protects against over-training, a learning rate factor of 0.5 provides a good compromise

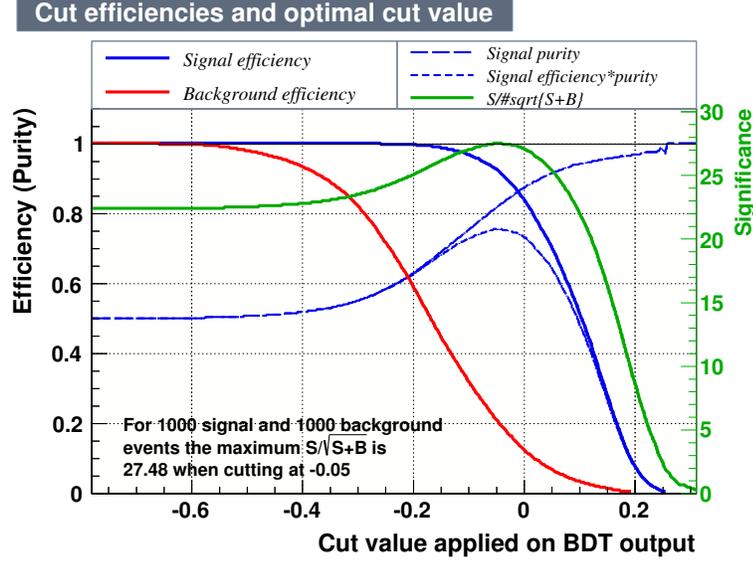


Figure 5.20: BDT cut tuning for FGD2DsBr training and testing samples with $\beta = 0.5$ and 1200 trees in the training forest. The optimal performance is achieved when applying a cut of -0.05, yielding a significance, $S/\sqrt{S+B}$, of 27.48. Note, this is a different significance definition to that described by Equation 5.9.

between maximising separation and significance, whilst 1200 trees keeps the data processing time to a reasonable length, but still maintains the compromise between significance and separation. 80% of signal and background events are used in the BDT training, with the remaining 20% reserved for testing.

Monte-Carlo Training Samples for the Boosted Decision Trees

In an ideal analysis scenario, the MC sample used to train a MVA discriminator would be a large sample of full scale experimental simulation, *magnet* MC in the case of the ND280, and all analysis cuts prior to the MVA implementation would be applied to the MVA training sample. However, providing such large MC samples is computationally very expensive, and T2K only has a limited GRID storage space for the data samples it produces. This means the quantity of *magnet* MC samples available to this analysis does not provide enough signal or background MC for an effective BDT training, with sufficient signal MC being a particular issue. To deal with this *magnet* MC deficit, a combination of *magnet*

and specialised, non-*magnet* ND280 MC samples are used to provide signal and background events, and bespoke *particle gun* MC was produced to bolster the number of signal events.

ND280 Signal and Background Monte-Carlo Training Samples

Table 5.10 shows the breakdown in the type of ND280 MC files used to extract signal and background events for training the BDT.

- The *magnet* MC samples were all that remained after setting aside nearly 5 times the data POT for the MC data analysis sample and were previously described in Section 5.2.2.
- The *basket* MC samples are simplified versions of the *magnet* MC whereby there is no spill and bunch structure, and the interaction vertex is constrained to be within the ND280 basket region.
- The $CC1\pi^0$ MC samples are the same as the *basket* MC, but only contain events with $CC1\pi^0$ final state.
- Finally the *FGD CC π^0 Inc.* MC samples are the same as the *basket* MC, but only contain events with an FGD interaction vertex and inclusive $CC\pi^0$ final state.

For all the *basket*, $CC1\pi^0$ and *FGD CC π^0 Inc.* MC samples the number of POT quoted is the number of POT worth of *magnet* files you would have to use to obtain an equivalent subset of interactions. Only events from these MC training samples that pass all of the muon selection, and pre-MVA π^0 selection criteria are accepted into the final BDT training sample.

Particle Gun Signal Monte-Carlo Training Sample

The signal events extracted from the above MC samples did not provide a sufficiently large event sample to train the BDT. It was therefore decided to simulate a large sample of π^0 decays within the FGD detectors and use the subset of

MC type	POT	PØD Water Status	Beam Power
Magnet	9.11×10^{20}	Water In	120 kW
Magnet	6.51×10^{20}	Water Out	120 kW
Magnet	4.23×10^{21}	Water In	178 kW
Magnet	1.12×10^{21}	Water Out	178 kW
Basket	1×10^{21}	Water In	-
Basket	1×10^{21}	Water Out	-
CC1 π^0	1.99×10^{22}	Water In	-
CC1 π^0	2.05×10^{22}	Water Out	-
FGD CC π^0 Inc.	1.91×10^{22}	Water In	-
FGD CC π^0 Inc.	2.08×10^{22}	Water Out	-

Table 5.10: Breakdown in the MC samples used to select signal and background events for the BDT training.

those where both decay photons showered in the ECals as the signal training events. To ensure the simulated interactions were an accurate representation of the expectation first the true energy and direction of π^0 s from CC1 π^0 final states were extracted from the CC1 π^0 MC samples as described in Table 5.10. Those interactions provided the energy and direction of π^0 s from 2.1×10^6 simulated events as shown in Figure 5.21a. Subsequently the ND280 particle gun was used to simulate 1.01×10^7 π^0 s in the FGDs with energy and direction characteristics drawn from the previously extracted CC1 π^0 MC sample distribution. The resultant energy and direction of the particle gun simulated events are shown in Figure 5.21b and are a near identical match to the originally extracted distribution. From the *particle gun* MC sample, only events that were simulated with an event vertex in the FGD fiducial volume were used in the BDT training signal sample.

The total number of signal and background pairwise ECal combinations extracted from the ND280 MC samples and *particle gun* samples, for each event topology of interest, are shown in Table 5.11.

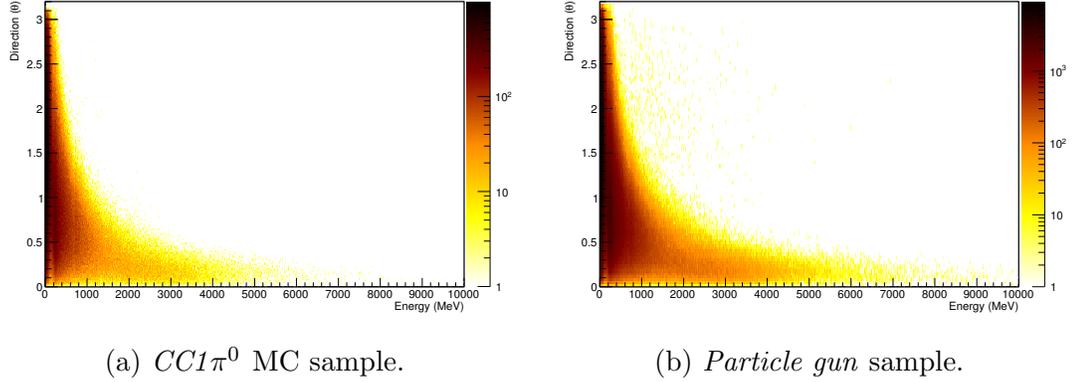


Figure 5.21: Energy and Direction of π^0 's from the $CC1\pi^0$ MC sample, 5.21a, and *particle gun* simulation, 5.21b.

Event Topology	Number of Pairwise ECal Combinations			
	ND280 MC		Particle Gun	Total
	Background	Signal	Signal	Signal
FGD1DsDs	83080	3977	3722	7699
FGD1DsBr	504830	6864	17463	24327
FGD1BrBr	758285	23326	84797	108123
FGD2DsDs	104371	9335	28448	37783
FGD2DsBr	451446	12418	62841	75259
FGD2BrBr	547686	21062	86827	107889
Total	2449698	76982	284098	361080

Table 5.11: The number of signal and background pairwise ECal combinations that were used to train the BDT discriminators for each signal topology.

Discriminating Variables for the Boosted Decision Trees

Each topology uses the same set of 17 discriminating variables derived from 9 different analysis quantities. The selected analysis quantities were derived from an initial input list of 14, from which the most highly correlated and poorly performing discriminating variables were rejected. Those used by this analysis are:

(a) **First Layer Hit,**

The most internal (closest to the ND280 tracking detectors) ECal layer hit by an ECal cluster.

(b) **Last Layer Hit,**

The most external (furthest from the ND280 tracking detectors) ECal layer hit by an ECal cluster.

(c) **Average z Position,**

The average position of an ECal cluster along the ND280 z-axis.

(d) **Thrust,**

The thrust value calculated for an ECal cluster (calculated as described in Section 3.3.2).

(e) **Zenith Angle,**

The angle between the reconstructed direction of an ECal cluster (from the major axis of a PCA fit to the cluster hits) and the local z-axis of the ECal module (the direction pointing straight out of the module from the most internal layer towards the most external layer).

(f) **Thrust to Trajectory Angle,**

The angle between the thrust axis (calculated as described in Section 3.3.2) and the hypothetical decay photon trajectory (the direction between the selected muon candidate's FGD vertex position and the thrust origin of the ECal cluster).

(g) **EM Energy,**

The reconstructed electromagnetic energy deposition of the ECal cluster.

(h) **EmHip,**

The log likelihood ratio result of the ECal electromagnetic-hadronic shower discriminator for an ECal cluster. This one of several particle identification variables which use the circularity of ECal clusters, along with a truncated maximum ratio, RMS, and front over back positional ratio of their constituent hit charges, to form a likelihood ratio using a TMVA likelihood implementation [96].

(i) **Energy Asymmetry,**

The energy asymmetry calculated for the pair of ECal clusters, $\frac{E_h - E_l}{E_h + E_l}$, where E_h is the energy of the higher energy cluster, and E_l is that of the lower.

The rejected analysis quantities included the number of ECal hits which was highly correlated with the EM Energy; the track-shower, MIP-pion and MIP-electromagnetic log likelihood ratio discriminators which were all highly correlated with the electromagnetic-hadronic shower discriminator; and the angle between the pair of ECal clusters as this feeds directly into the reconstructed π^0 invariant mass.

The analysis quantities (a) through (h) are calculated for the higher, E_h , and lower energy, E_l , ECal cluster separately, and in combination with the Energy Asymmetry, (i), give the 17 discriminating variables used by the BDT for each topology. Figure 5.22 shows the resultant signal and background discriminating variable distributions for the FGD2DsBr topology. See Appendix A for the discriminating variable distributions for all other topologies.

Figure 5.23 shows the correlation matrices for the signal and background discriminating variable distributions from Figure 5.22. The linear correlation is calculated by plotting all discriminating variables against one another and performing a linear fit through each pair individually. The gradient of this fit is then used to calculate the linear correlation between each pair of variables. It is seen that generally the correlation between discriminating variables is very low, thus maximising their discriminating potential. A notable exception is in the correlation between the high and low energy cluster average z position variables, which for this topology are highly anti-correlated, but this as expected as if one cluster

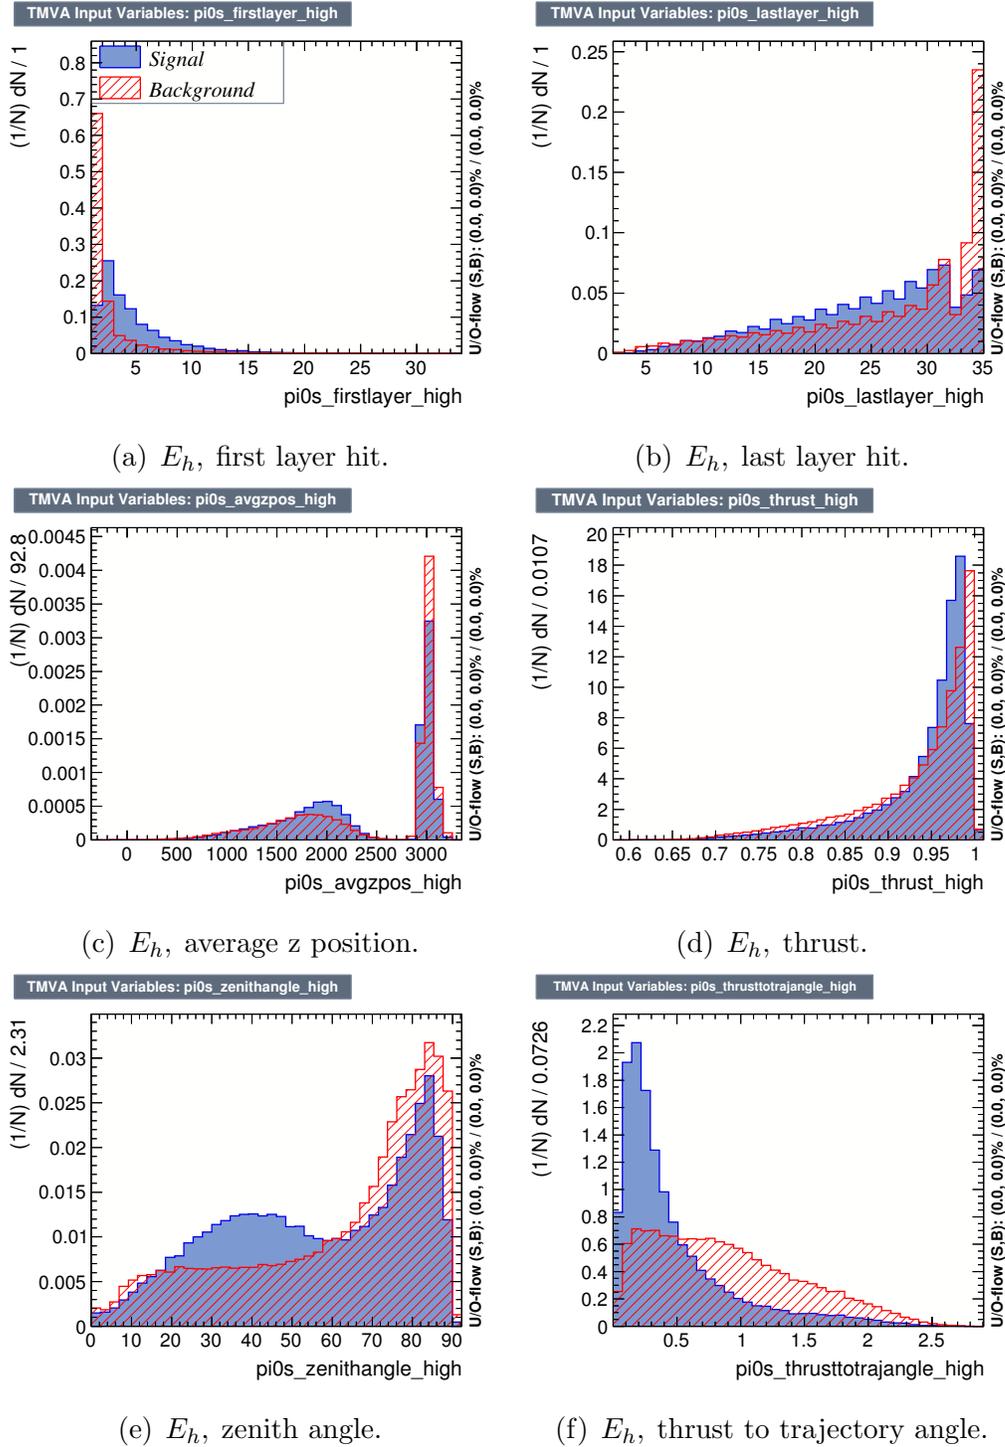


Figure 5.22: Discriminating variables used in the training of the FGD2DsBr topology BDT.

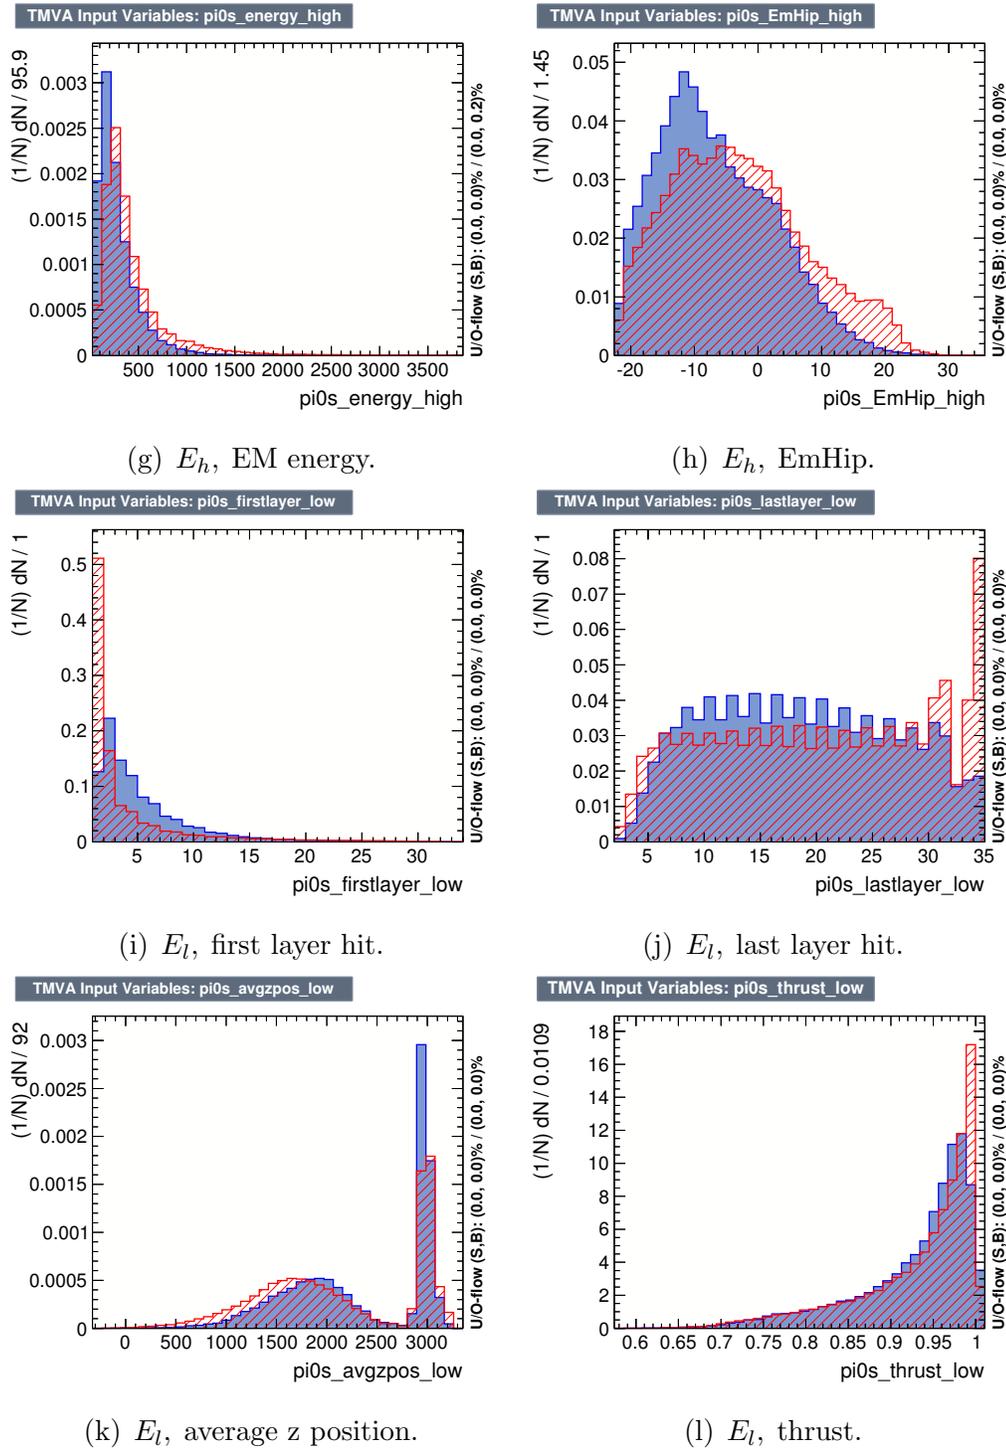


Figure 5.22: Discriminating variables used in the training of the FGD2DsBr topology BDT.

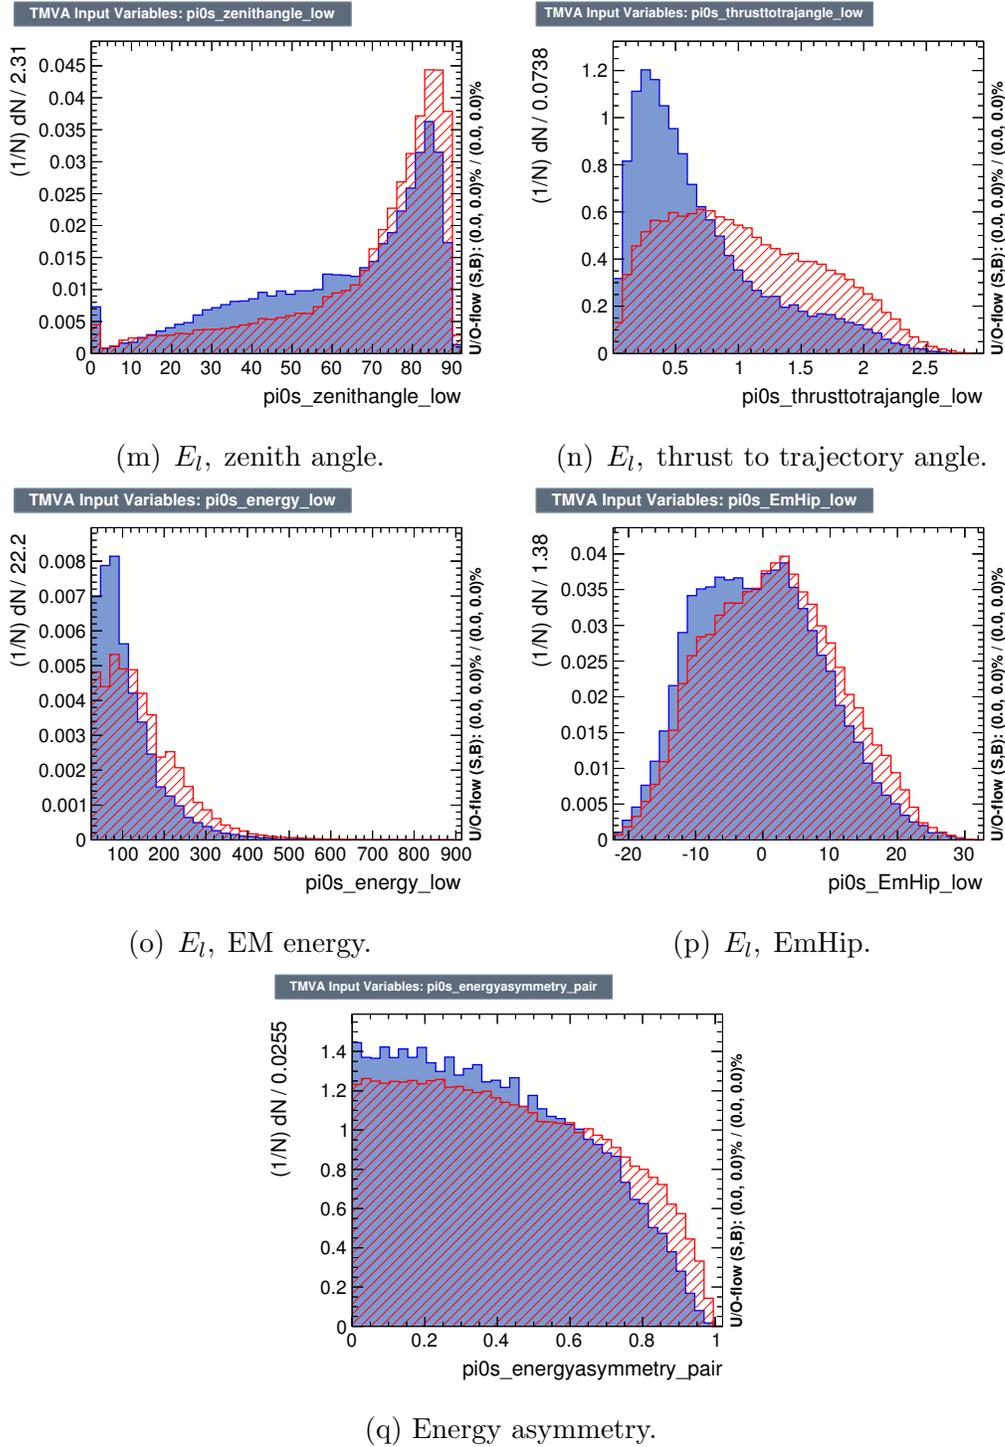


Figure 5.22: Discriminating variables used in the training of the FGD2DsBr topology BDT.

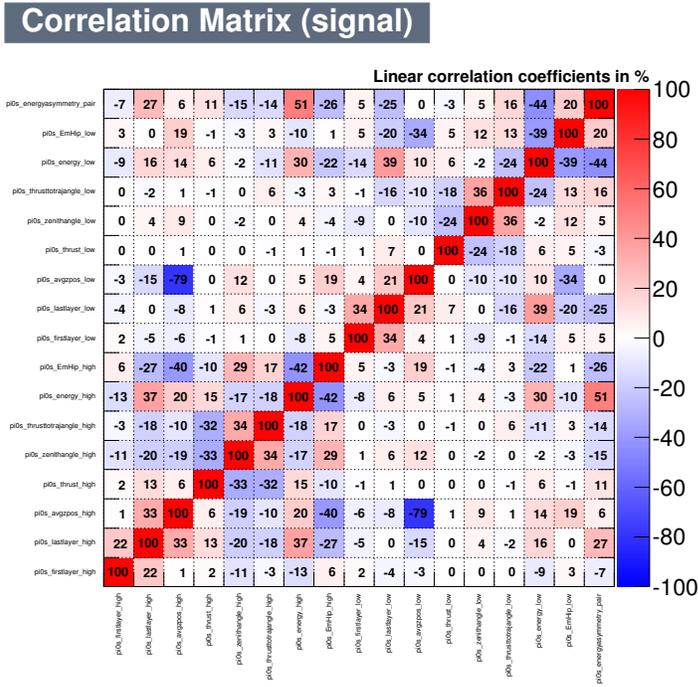
is in a BrECal module the other has to be in the DsECal. See Appendix A for the discriminating variable distributions for all other topologies.

Tables 5.12 and 5.13 show the resultant importance ranking of variables for each of the six event topologies being considered, separated by the event vertex being in FGD1 and FGD2 respectively. It is interesting to note that regardless of the topology of interest, the EM energy, first layer hit and EmHip analysis quantities, for both the higher and lower energy ECal clusters, tend to be the most powerful discriminators. Equally, the zenith angle, last layer hit and thrust, for both the higher and lower energy ECal clusters, along with the energy asymmetry, are consistently among the least powerful discriminating analysis quantities.

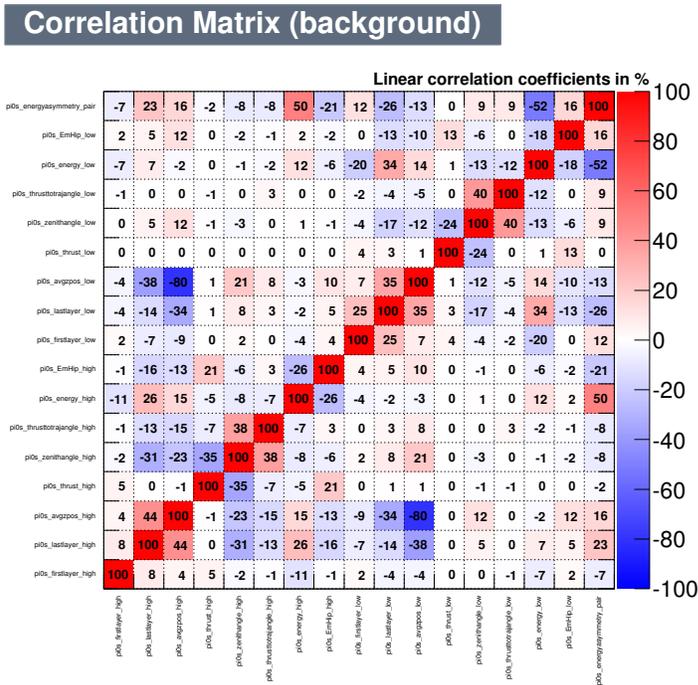
5.5.4 Boosted Decision Tree Training Performance

Once the BDT for each topology was produced, the standard over-training checks and cut optimisation were performed to ensure everything was as expected, as shown in Figure 5.24. From the over-training plots it is clear to see that in all cases the agreement between the BDT response distributions for the training and testing samples is fairly good, with the worst KS test matching probability performance being 16.8% and 39.4% for the signal, Figure 5.24i (FGD2DsBr), and background, Figure 5.24e (FGD1BrBr), samples respectively. From the cut optimisation plots it is seen that the significance, $S/\sqrt{S+B}$, is rather consistent across all the topologies, with the poorest performance reported by the FGD1DsDs topology, Figure 5.24b ($S/\sqrt{S+B} = 26.21$), with the best reported by the FGD2DsBr topology, Figure 5.24j ($S/\sqrt{S+B} = 27.57$). Given the relatively small training sample size for the FGD1DsDs topology (see Table 5.11) it is no surprise that the curves showing the cut performance, Figure 5.24b, are not as smooth as one would expect for a fully optimised training. This would be overcome by the production of a larger MC training sample for this topology, but given the very small number of events in the *magnet* MC which conform to this topology (see Table 5.5), it was deemed not worth the additional effort to further enhance this training sample.

The discrepancy in the BDT response for the FGD2DsBr topology between the



(a) Signal.



(b) Background.

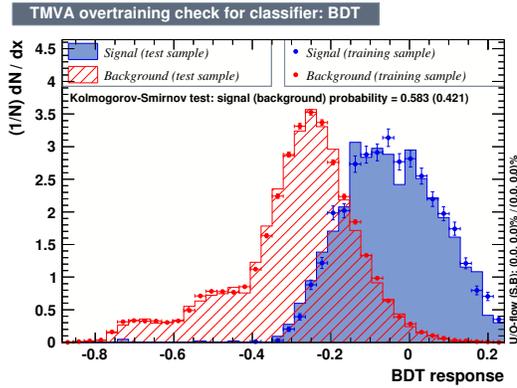
Figure 5.23: Correlation between the signal, 5.23a, and background, 5.23b, discriminating variables used in the training of the FGD2DsBr topology BDT.

Variable Ranking	Topology		
	FGD1DsDs	FGD1DsBr	FGD1BrBr
1st	E_h , EmHip	E_h , first layer hit	E_h , first layer hit
2nd	E_l , EM energy	E_l , first layer hit	E_h , EM energy
3rd	E_h , average z position	E_l , EM energy	E_l , first layer hit
4th	E_l , EmHip	E_h , EmHip	E_h , EmHip
5th	E_h , first layer hit	E_l , average z position	E_l , EM energy
6th	E_l , first layer hit	E_h , EM energy	E_l , average z position
7th	E_l , average z position	E_h , thrust to trajectory angle	E_l , EmHip
8th	Energy asymmetry	E_l , EmHip	E_h , thrust to trajectory angle
9th	E_h , thrust to trajectory angle	E_h , average z position	E_h , average z position
10th	E_l , thrust to trajectory angle	E_h , last layer	E_l , thrust to trajectory angle
11th	E_h , EM energy	E_l , thrust to trajectory angle	E_h , thrust
12th	E_l , zenith angle	E_l , last layer	E_h , last layer
13th	E_h , thrust	E_h , zenith angle	E_l , thrust
14th	E_h , zenith angle	E_l , zenith angle	E_h , zenith angle
15th	E_l , last layer	E_h , thrust	E_l , zenith angle
16th	E_h , last layer	E_l , thrust	E_l , last layer
17th	E_l , thrust	Energy asymmetry	Energy asymmetry

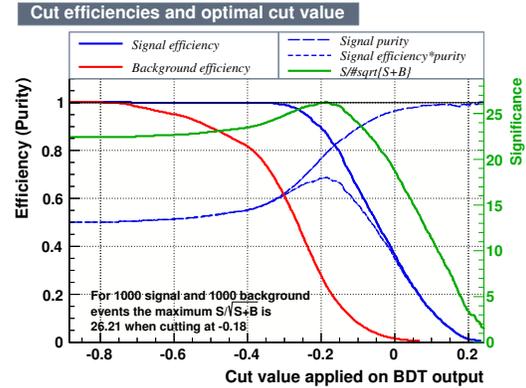
Table 5.12: Variable importance ranking from the BDT training for each of the event topologies with an FGD1 vertex. The number in brackets next to each variable represents its MVA method independent ranking.

Variable Ranking	Topology			
	FGD2DsDs	FGD2DsBr	FGD2BrBr	
1st	E_h , EmHip (4th)	E_l , EM energy (7th)	E_h , first layer hit (1st)	E_h , first layer hit (1st)
2nd	E_l , first layer hit (3rd)	E_h , first layer hit (1st)	E_h , EM energy (5th)	E_h , EM energy (4th)
3rd	E_l , EM energy (16th)	E_h , EM energy (5th)	E_l , first layer hit (2nd)	E_l , first layer hit (2nd)
4th	E_h , first layer hit (1st)	E_l , average z position (9th)	E_l , EM energy (9th)	E_l , EM energy (9th)
5th	E_l , EmHip (7th)	E_h , thrust to trajectory angle (3rd)	E_h , EmHip (13th)	E_h , EmHip (13th)
6th	E_h , average z position (14th)	E_l , first layer hit (2nd)	E_h , average z position (16th)	E_h , average z position (16th)
7th	E_h , thrust to trajectory angle (2nd)	E_l , EmHip (16th)	E_h , thrust to trajectory angle (3rd)	E_h , thrust to trajectory angle (3rd)
8th	E_h , zenith angle (8th)	E_h , average z position (14th)	E_l , EmHip (17th)	E_l , EmHip (17th)
9th	E_h , thrust (6th)	E_h , EmHip (10th)	E_l , average z position (15th)	E_l , average z position (15th)
10th	E_h , EM energy (15th)	E_h , last layer (6th)	E_h , thrust (8th)	E_h , thrust (8th)
11th	E_l , thrust to trajectory angle (5th)	E_l , last layer (11th)	E_l , thrust to trajectory angle (5th)	E_l , thrust to trajectory angle (5th)
12th	E_h , last layer (10th)	E_l , thrust to trajectory angle (4th)	E_h , last layer (12th)	E_h , last layer (12th)
13th	E_l , average z position (13th)	E_h , thrust (8th)	E_h , zenith angle (6th)	E_h , zenith angle (6th)
14th	E_l , last layer (9th)	E_l , zenith angle (13th)	E_l , zenith angle (7th)	E_l , zenith angle (7th)
15th	Energy asymmetry (17th)	E_h , zenith angle (12th)	Energy asymmetry (11th)	Energy asymmetry (11th)
16th	E_l , thrust (12th)	Energy asymmetry (17th)	E_l , last layer (14th)	E_l , last layer (14th)
17th	E_l , zenith angle (11th)	E_l , thrust (15th)	E_l , thrust (10th)	E_l , thrust (10th)

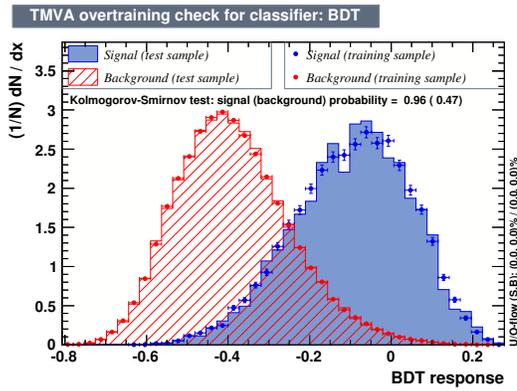
Table 5.13: Variable importance ranking from the BDT training for each of the event topologies with an FGD2 vertex. The number in brackets next to each variable represents its MVA method independent ranking.



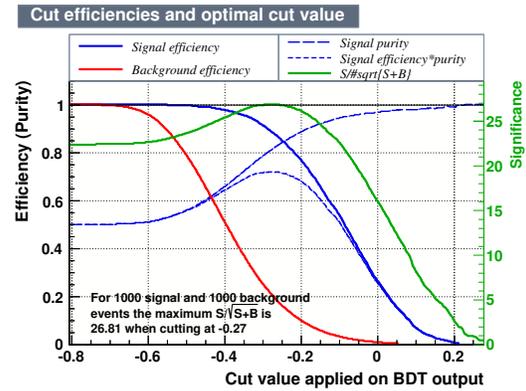
(a) BDT signal and background sample response for FGD1DsDs topology.



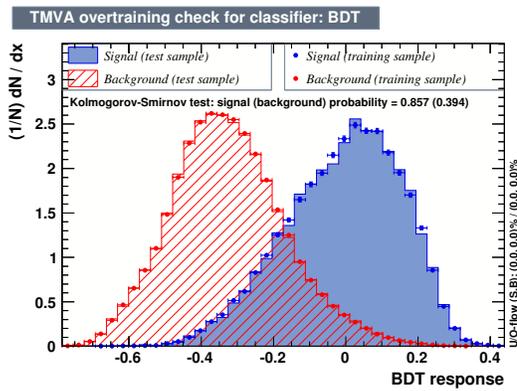
(b) BDT cut optimisation for FGD1DsDs topology.



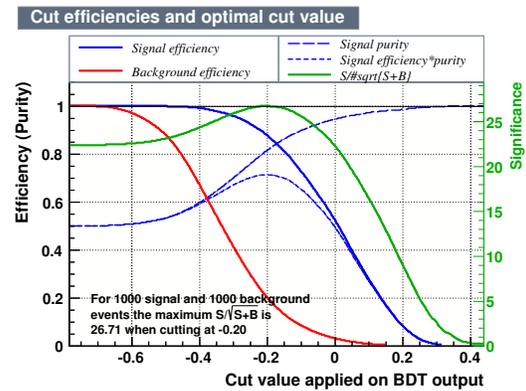
(c) BDT signal and background sample response for FGD1DsBr topology.



(d) BDT cut optimisation for FGD1DsBr topology.

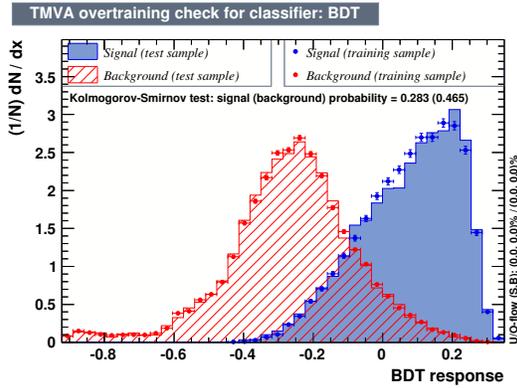


(e) BDT signal and background sample response for FGD1BrBr topology.

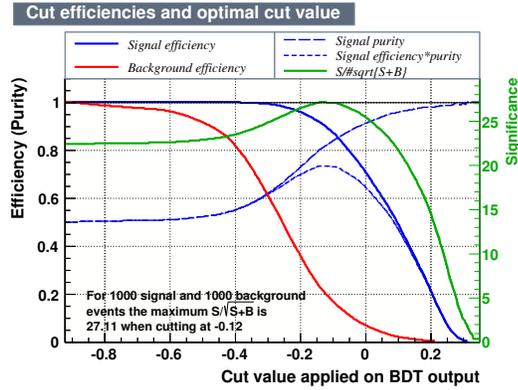


(f) BDT cut optimisation for FGD1BrBr topology.

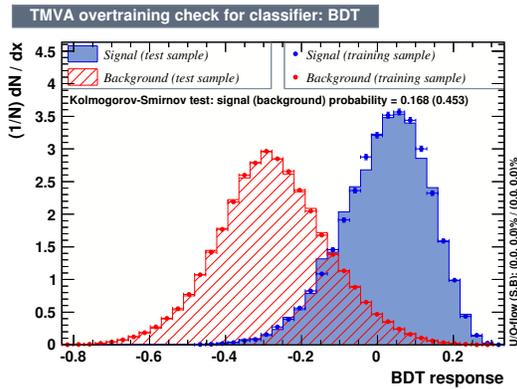
Figure 5.24: BDT response and cut optimisation for the FGD1 vertex analysis topologies.



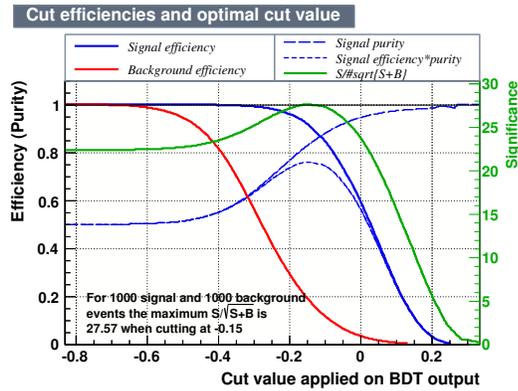
(g) BDT signal and background sample response for FGD2DsDs topology.



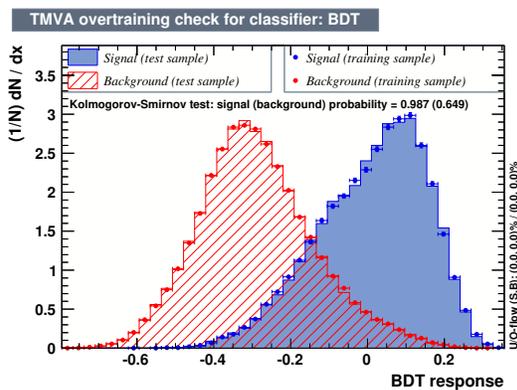
(h) BDT cut optimisation for FGD2DsDs topology.



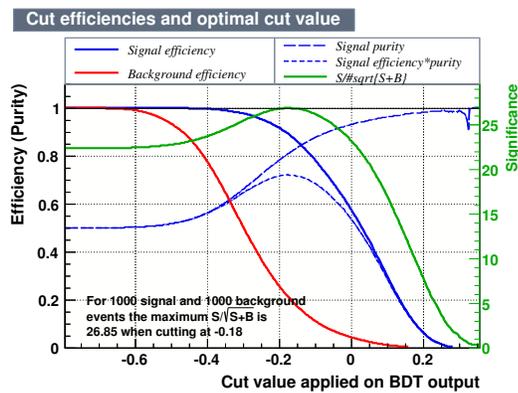
(i) BDT signal and background sample response for FGD2DsBr topology.



(j) BDT cut optimisation for FGD2DsBr topology.



(k) BDT signal and background sample response for FGD2BrBr topology.



(l) BDT cut optimisation for FGD2BrBr topology.

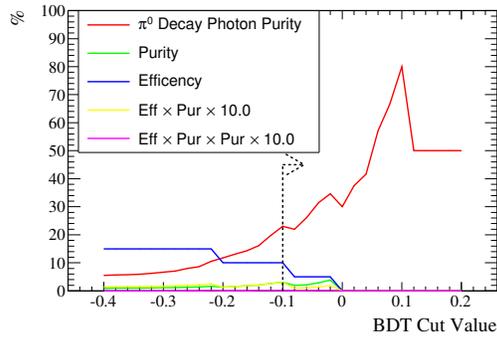
Figure 5.24: BDT response and cut optimisation for the FGD2 vertex analysis topologies.

final training shown here, Figure 5.24i, and the earlier BDT input parameter optimisation, see Figure 5.18, is due to slight changes in the quantity of signal and background MC samples used in the separate trainings. Although there were some changes in this regard the BDT response has remained rather consistent between the two separate trainings, $S/\sqrt{S+B} = 27.57$ for the final training, Figure 5.24j, and $S/\sqrt{S+B} = 27.48$ for the earlier optimisation, Figure 5.20, suggesting that the selected BDT input parameters were appropriate.

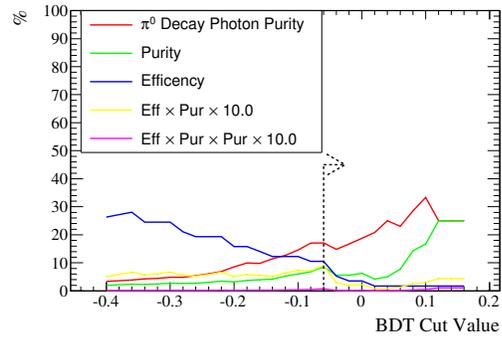
Final Boosted Decision Tree Cut Optimisation

Since the MC samples used to train the BDT discriminators for the event selection are not a true representation of the *magnet* MC over which the analysis is otherwise refined it was necessary to optimise the BDT cut value for each topology. This was done by running all of the analysis cuts over the Run 4 Air MC sample, as this has the highest POT of any running period, whilst varying the BDT response cut for each of the six topologies in turn between -0.4 and 0.2 at intervals of 0.02. The efficiency and purity for selecting events from each signal topology was then calculated, along with the total purity for correctly selecting π^0 decay photons from any final state. Figure 5.25 shows how these metrics vary as a function of cut value.

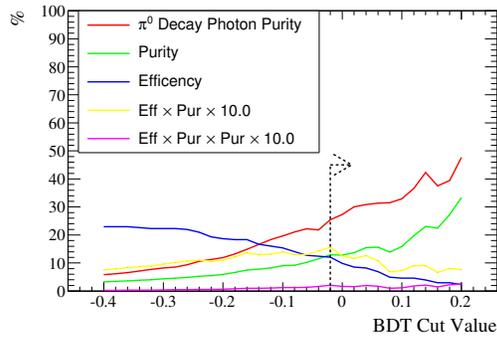
For all topologies the BDT is behaving as expected, increasing purity and decreasing efficiency as the cut value is incrementally increased. This, in conjunction with the increasing π^0 decay photon purity, shows that in all cases the BDT discriminators have been effectively trained to select pairs of ECal clusters arising from the decay of π^0 s to two photons, which have both propagated to, and showered within, the barrel and/or downstream ECals. The final BDT cut values to be used for each topology are those where the maximum *Efficiency* \times *Purity* is achieved, these are shown in Table 5.14, and are significantly more exclusionary than those suggested by the internal MVA optimisation.



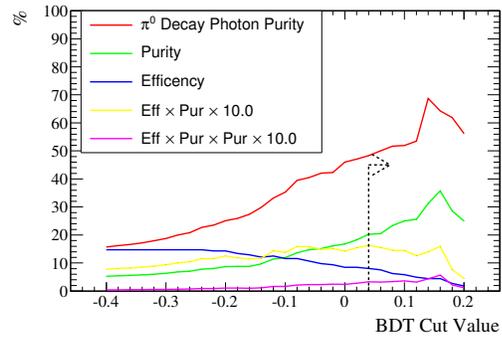
(a) Final BDT cut optimisation for FGD1DsDs topology.



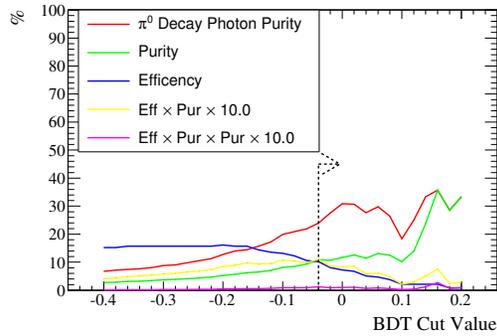
(b) Final BDT cut optimisation for FGD1DsBr topology.



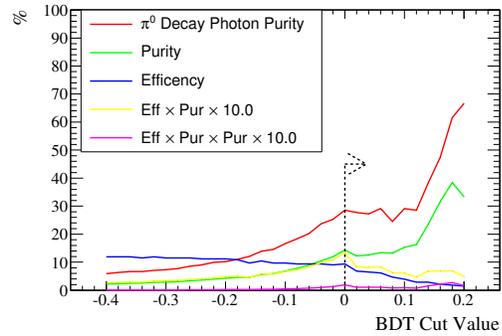
(c) Final BDT cut optimisation for FGD1BrBr topology.



(d) Final BDT cut optimisation for FGD2DsDs topology.



(e) Final BDT cut optimisation for FGD2DsBr topology.



(f) Final BDT cut optimisation for FGD2BrBr topology.

Figure 5.25: Efficiency and Purity of final selection sample for each of the six signal topologies of interest, as a function of BDT response cuts. The dashed black lines and arrows, indicate the position of the optimised cut values.

Event Topology	MVA Suggested Cut Value	Optimised Cut Value
FGD1DsDs	-0.18	-0.10
FGD1DsBr	-0.27	-0.06
FGD1BrBr	-0.20	-0.02
FGD2DsDs	-0.12	0.04
FGD2DsBr	-0.15	-0.04
FGD2BrBr	-0.18	0.00

Table 5.14: Optimal BDT cut values as suggested by the MVA training and the newly optimised cut vales that are applied to the final analysis.

5.5.5 Data and Monte-Carlo Comparisons for the Boosted Decision Tree

Given that the BDT training is purely based on MC samples, it is important that the distributions for the discriminating variables are well matched between the data and MC samples. Figure 5.26 shows area normalised data and MC comparisons for all 17 discriminating variables used by the BDT, after all pre-MVA implementation cuts have been applied. The particles identified in the legend for each figure represent the true particle type which produced the most hits within each ECal cluster. The results from all five ECal modules are aggregated together, for a module by module break down see Appendix B.

The distributions are area normalised as the BDT places cuts on each variable based on what proportion of the signal and background lie each side of the selected cut position, not the absolute number of events. If there was a data excess or deficit, and it was proportionally spread across the distributions for each discriminating variable, then the relative number of the signal and background events lying each side of the cut position would remain constant and the predicted efficiency and purity of the selection would remain valid. However, if the shape of the distribution of data compared to MC for some discriminating variables differed, then the relative number of the signal and background events lying each side of the selected cut position would change and the efficiency and purity of the selection would no longer be accurately predicted. Such a situations is easiest to

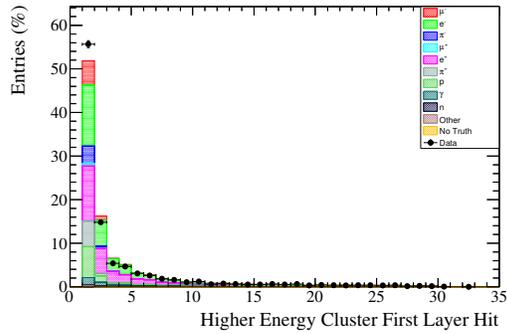
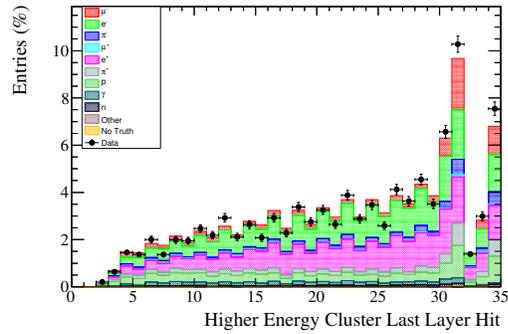
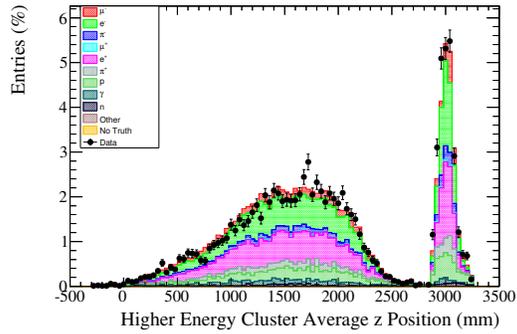
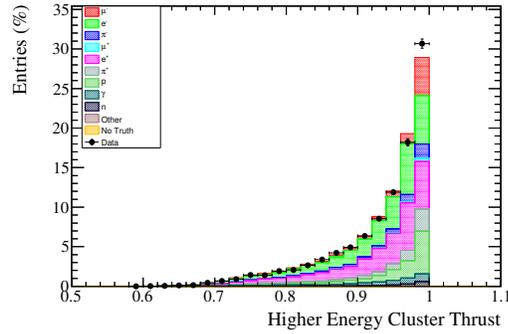
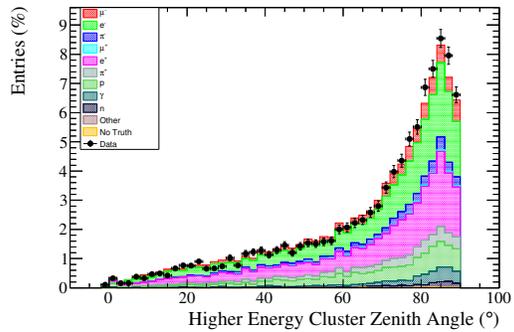
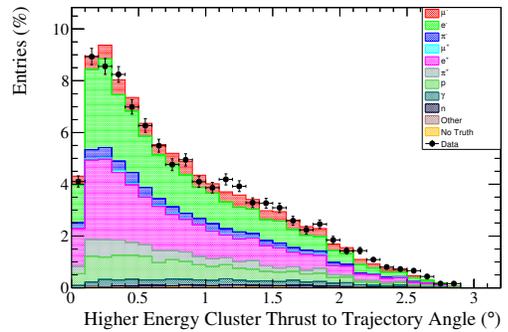
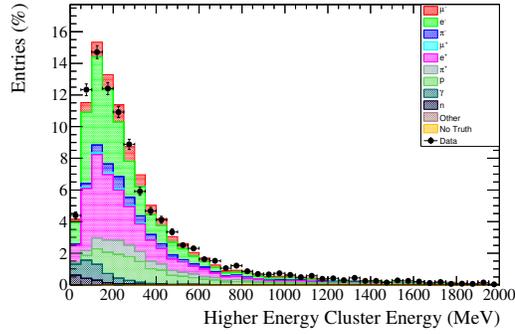
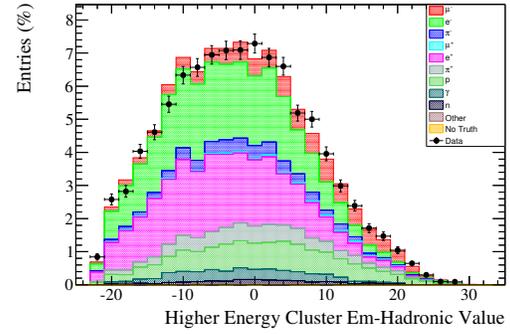
(a) E_h , first layer hit.(b) E_h , last layer hit.(c) E_h , average z position.(d) E_h , thrust.(e) E_h , zenith angle.(f) E_h , thrust to trajectory angle.

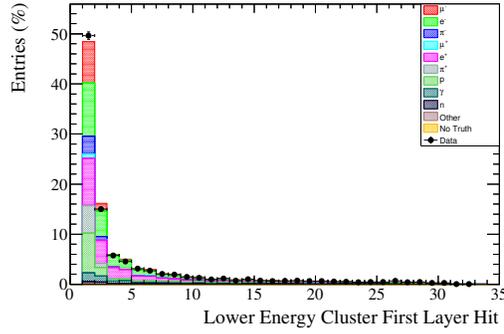
Figure 5.26: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection.



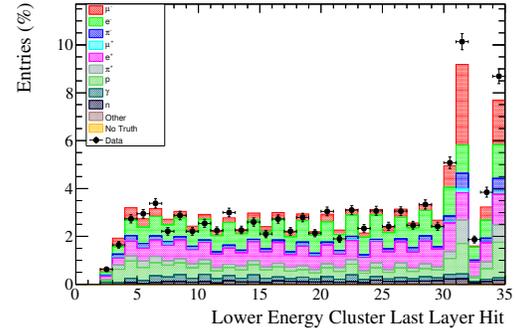
(g) E_h , EM energy.



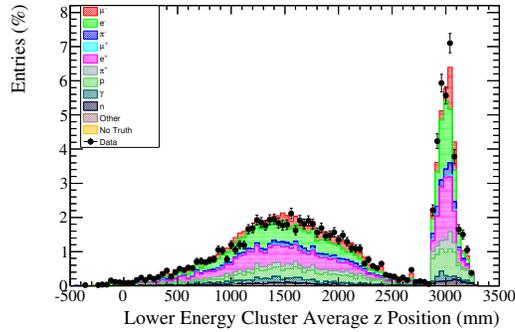
(h) E_h , EmHip.



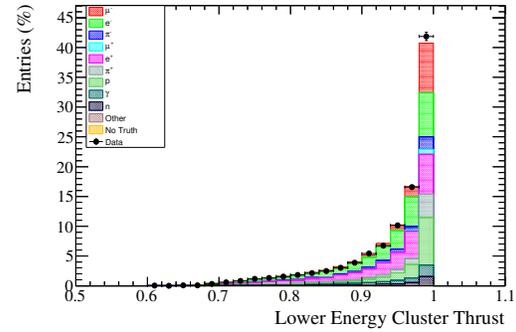
(i) E_l , first layer hit.



(j) E_l , last layer hit.

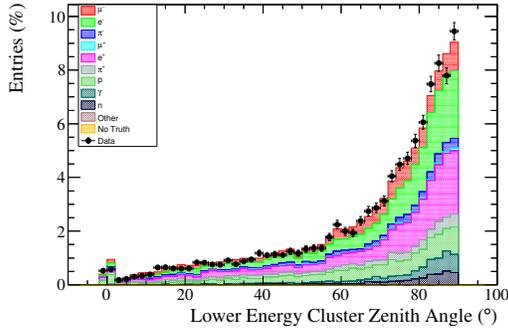
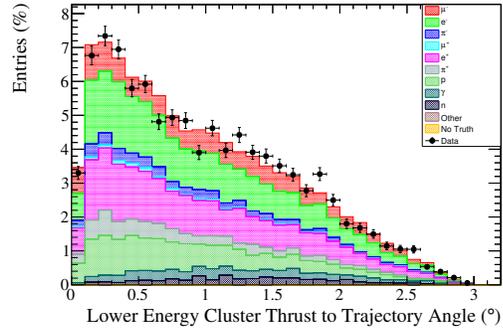
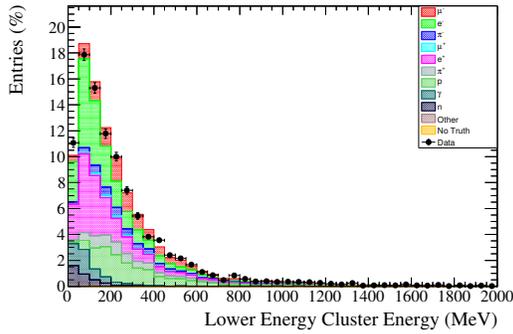
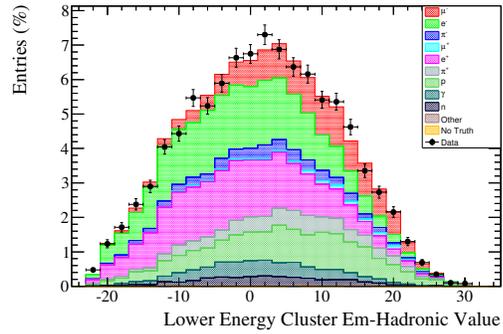
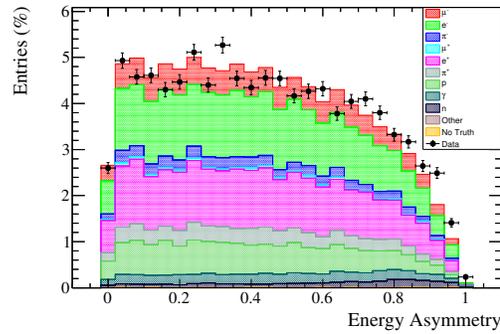


(k) E_l , average z position.



(l) E_l , thrust.

Figure 5.26: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection.

(m) E_l , zenith angle.(n) E_l , thrust to trajectory angle.(o) E_l , EM energy.(p) E_l , EmHip.

(q) Energy asymmetry.

Figure 5.26: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection.

observe within area normalised data and MC distributions.

As can be seen, the matching between data and MC for the distributions shown is generally very good, providing confidence in the BDT discriminator cut when applied to the data samples. Areas where some minor disagreement can be seen include at higher values of the last layer hit for both the high and low energy clusters, as seen in Figures 5.26b and 5.26j respectively, and at values of the energy asymmetry approaching unity, Figure 5.26q. However minor disagreement within the distribution of these variables is not of significant concern because, as previously stated in Section 5.5.3, these are some of the worst performing discriminators across all topologies and so will minimally contribute to the final result.

5.5.6 π^0 Selection Performance

Applying the selection criteria as described in Section 5.5.1, with the optimised BDT cut values, to the *magnet* MC event sample passing the muon selection criteria, produces a total selection efficiency, across all $CC1\pi^0$ topologies of interest, of 9.61%, with 12.1% purity as shown in Figure 5.27 and described in detail in Table 5.15. The efficiency and purity are 2.5% and 18.1% respectively when considering any $CC1\pi^0$ final state. The final state dominating the background is predominantly $CC\pi^0+$, along with near equal contributions from each of $CCN\pi^0$, $CCQE$, CC Other, Other and Out of FV events as detailed in Table 5.16.

5.5.7 π^0 Selection Results

Applying the $CC1\pi^0$ selection criteria to the Run 2 to Run 4 data and MC sets produced the following π^0 invariant mass distributions from the final selected pair of ECal clusters. Figures 5.28 through 5.33 show the reconstructed mass peak for each of the six $CC1\pi^0$ signal topologies individually, broken down by true selection topology. Figures 5.34 and 5.35 show the same information, but broken down by final state and π^0 decay photon reconstruction respectively. In all cases the vertical error bar shown for each data point is statistical. Table 5.17

Cut	Total Events	CC1 π^0 Signal Topology			π^0 Decay Photon	
		Events	Efficiency (%)	Purity (%)	Events	Purity (%)
Event Quality	5317850.0	763.01	100	8.5×10^{-3}	-	-
Negative FGD FV	97460.79	451.70	59.2	0.46	-	-
Track Quality	87357.34	435.10	57.0	0.50	-	-
Forward Going	85964.20	432.55	56.6	0.50	-	-
TPC PID Quality	84919.46	431.89	56.6	0.51	-	-
TPC Muon PID	56081.30	404.87	53.0	0.72	-	-
Muon Multiplicity	53111.01	399.45	52.3	0.75	-	-
ECal Multiplicity	6877.52	177.01	23.2	2.57	-	-
Reconstructed Mass	6117.73	164.88	21.6	2.70	-	-
Multivariate Selection	659.85	75.93	10.0	11.5	-	-
Pairs Multiplicity	605.09	73.51	9.61	12.1	171.55	28.4
Topology Selection						
FGD1DsDs	39.72	1.29	7.1	3.2	10.28	25.9
FGD1DsBr	48.14	3.57	10.0	7.4	9.17	19.0
FGD1BrBr	191.2	20.32	10.2	10.6	49.23	25.7
FGD2DsDs	58.54	13.21	9.4	22.6	29.08	49.7
FGD2DsBr	145.9	16.82	10.4	11.5	35.10	24.1
FGD2BrBr	121.7	16.08	9.0	13.2	38.69	31.8

Table 5.15: MC prediction for the number of events selected as a function of cut, CC1 π^0 selection efficiency and purity, and π^0 decay photon purity across all running periods to be analysed.

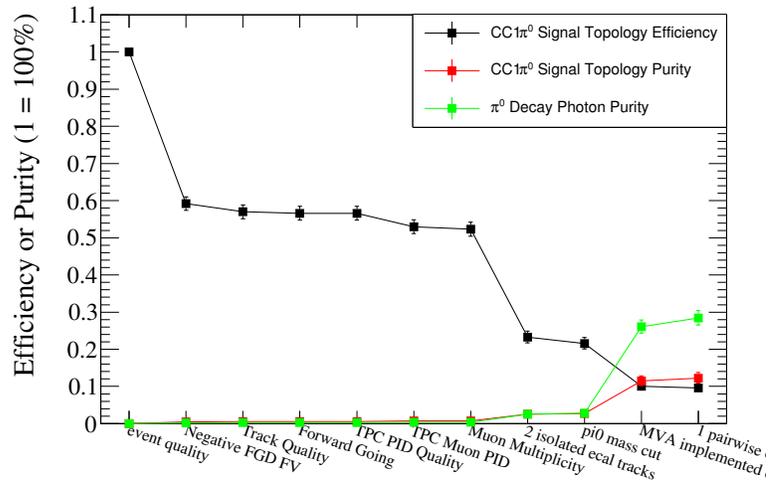


Figure 5.27: The efficiency and purity for selecting CC1 π^0 signal topology events, along with the purity for selecting π^0 decay photon events as a function cut.

Final State	Contribution (%)
CC1 π^0	18.1
CCN π^0	8.4
CC π^0 +	40.4
CCQE	6.4
CC Other	9.1
Other	9.0
Out of FV	8.5

Table 5.16: Final state contributions to final event selection.

shows the POT scaled and flux weighted predicted number of events to pass each cut in magnet MC, along with the true number of events in data, broken down by running period.

The reconstructed π^0 invariant mass is calculated as (see Appendix D for a derivation)

$$m_{\pi^0} = \sqrt{2E_h E_l (1 - \cos \theta_{hl})}, \quad (5.10)$$

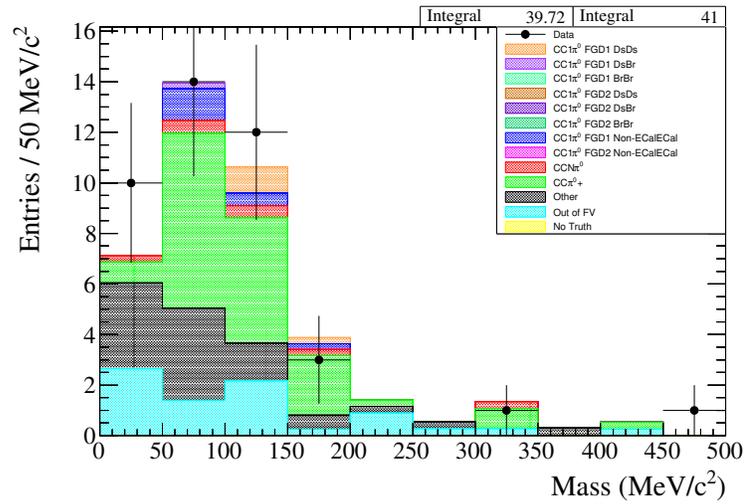
where E_h and E_l are the energy of the ECal clusters, and $\cos \theta_{hl}$ is the angle between the ECal clusters anchored to the selected muon vertex.

The ECal Multiplicity and Reconstructed Mass cuts continue to maintain the 5% deficit that was observed during the muon selection, the subsequent multivariate implementation then further drops the data against the MC leading to a final 13% deficit. However, when looking at the breakdown by event topology it is clear that the data deficit is predominantly due to a FGD1BrBr topology, whereby a 40% deficit has occurred. All other topologies show excellent agreement with the MC prediction, or a very minor deficit, consistent with the drops seen over the previous selection cuts.

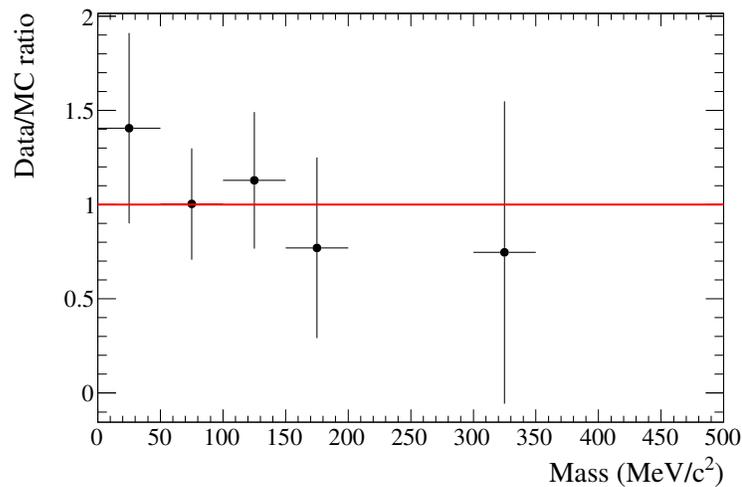
Given that the BDT used for each topology were trained and optimised in the same manner, with the same training and testing samples it seems very odd that the results would be so deviant from the MC expectation for one topology compared to all the others. At this time there is no clear explanation for this behaviour, but some initial investigations into this behaviour is described below.

Investigating the FGD1BrBr Data Deficit

Having made careful checks of each step within the data processing chain, and ensuring that the BDT cut was reading in the correct weights files and applying the correct cut value, it is believed that the FGD1BrBr topology deficit is a real effect, rather than a procedural error. Looking at the event rates, broken down by running period as shown in Table 5.17, it is clear that the data deficit is not constrained to any particular data taking period, with all periods showing a final deficit of between 25% and 57%. Therefore, we can rule out the issue being due to one or more bad data runs, which is to be expected as bad data taking would

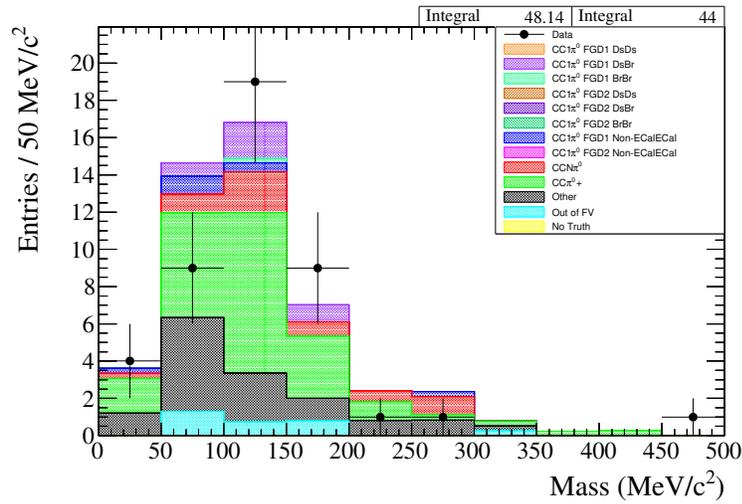


(a) Reconstructed π^0 invariant mass for FGD1DsDs event topology.

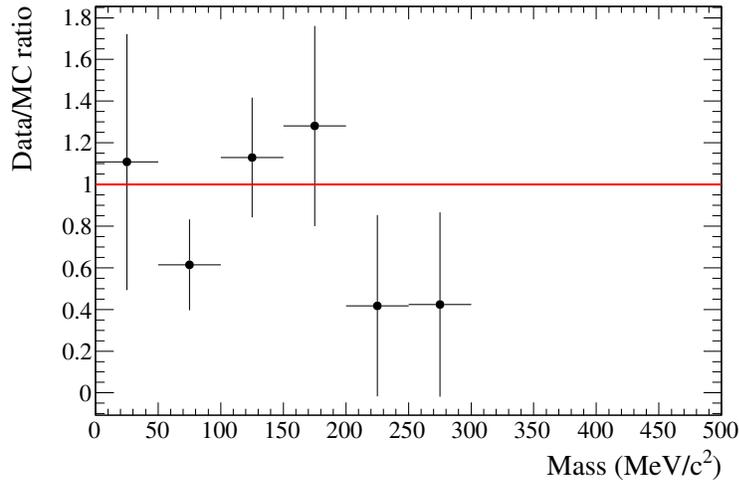


(b) Data over MC ratio for 5.28a.

Figure 5.28: 5.28a shows the reconstructed π^0 invariant mass distribution for the FGD1DsDs event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.28b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.28a.

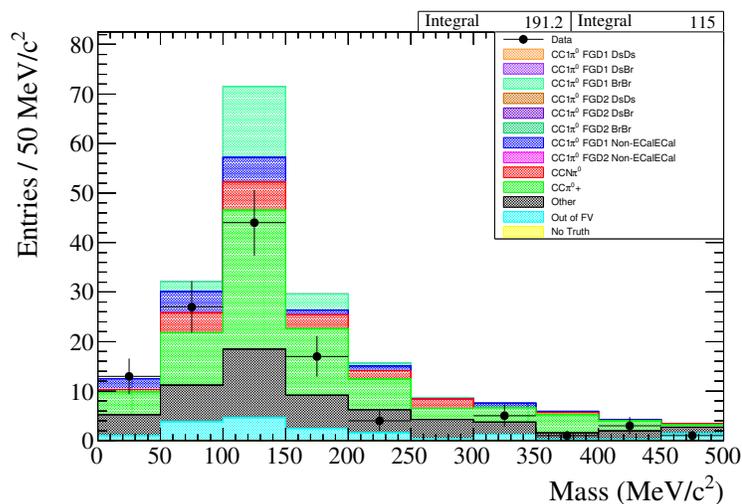


(a) Reconstructed π^0 invariant mass for FGD1DsBr event topology.

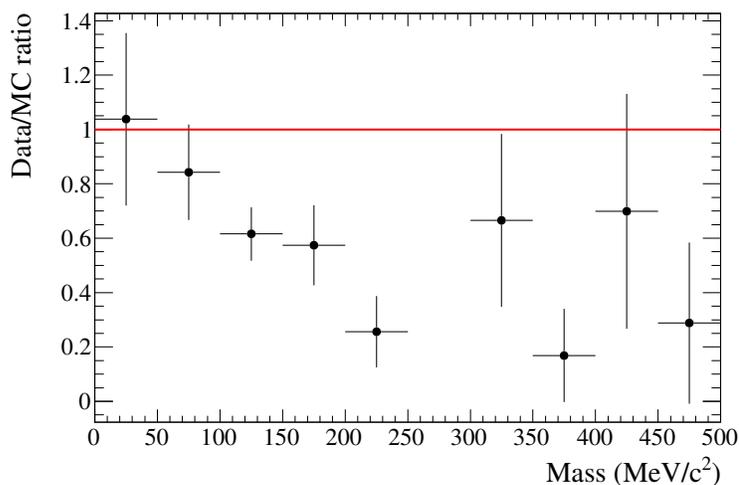


(b) Data over MC ratio for 5.29a.

Figure 5.29: 5.29a shows the reconstructed π^0 invariant mass distribution for the FGD1DsBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.29b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.29a.

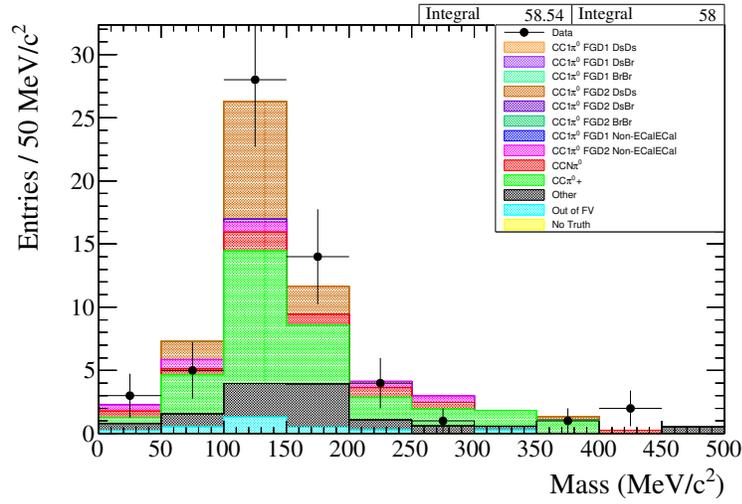


(a) Reconstructed π^0 invariant mass for FGD1BrBr event topology.

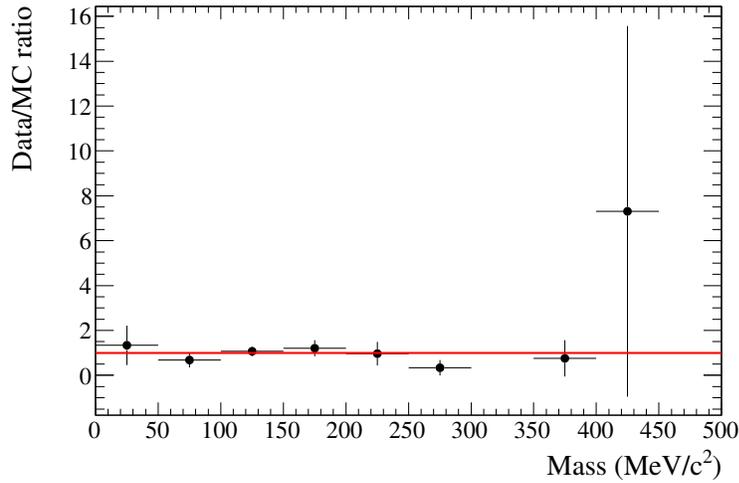


(b) Data over MC ratio for 5.30a.

Figure 5.30: 5.30a shows the reconstructed π^0 invariant mass distribution for the FGD1BrBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.30b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.30a.

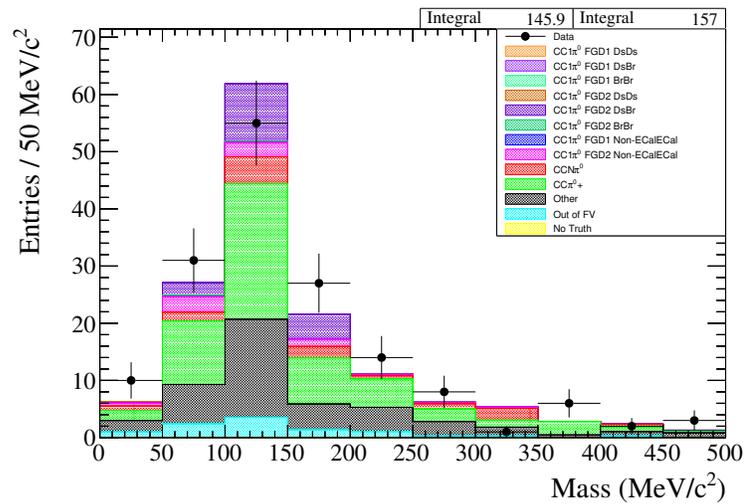


(a) Reconstructed π^0 invariant mass for FGD2DsDs event topology.

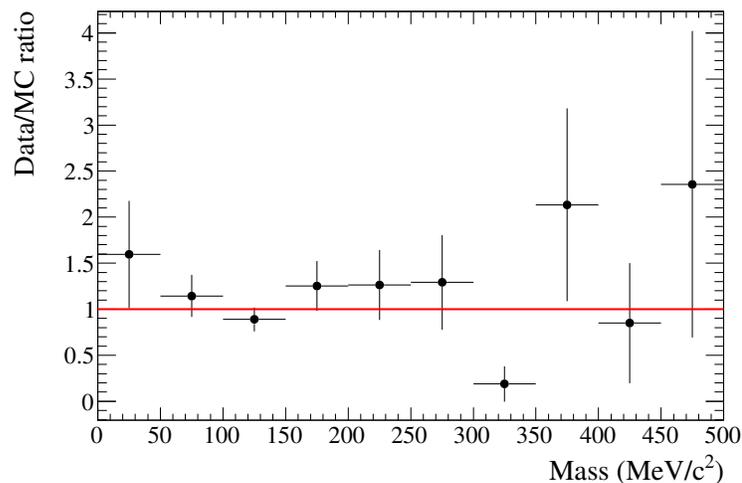


(b) Data over MC ratio for 5.31a.

Figure 5.31: 5.31a shows the reconstructed π^0 invariant mass distribution for the FGD2DsDs event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.31b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.31a.

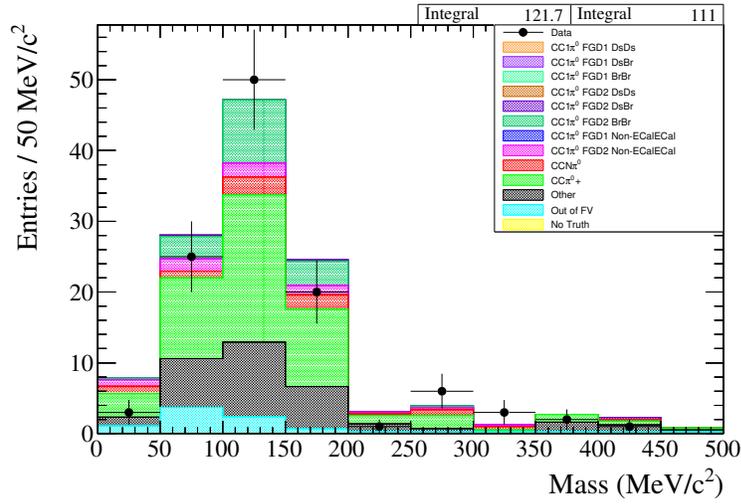


(a) Reconstructed π^0 invariant mass for FGD2DsBr event topology.

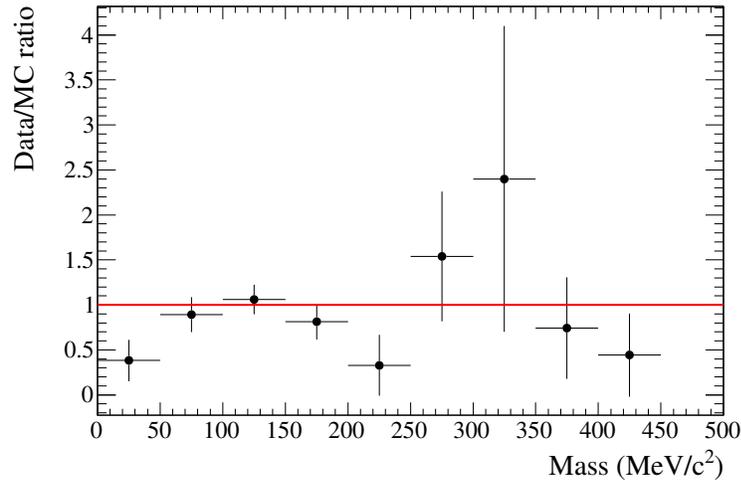


(b) Data over MC ratio for 5.32a.

Figure 5.32: 5.32a shows the reconstructed π^0 invariant mass distribution for the FGD2DsBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.32b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.32a.



(a) Reconstructed π^0 invariant mass for FGD2BrBr event topology.



(b) Data over MC ratio for 5.33a.

Figure 5.33: 5.33a shows the reconstructed π^0 invariant mass distribution for the FGD2BrBr event topology over all running periods. The integral values at the top of the histogram shows the integrated number of entries in the displayed histogram range for POT scaled and flux weighted MC on the left, and the data on the right. 5.33b shows the data over MC ratio as a function of reconstructed π^0 invariant mass from 5.33a.

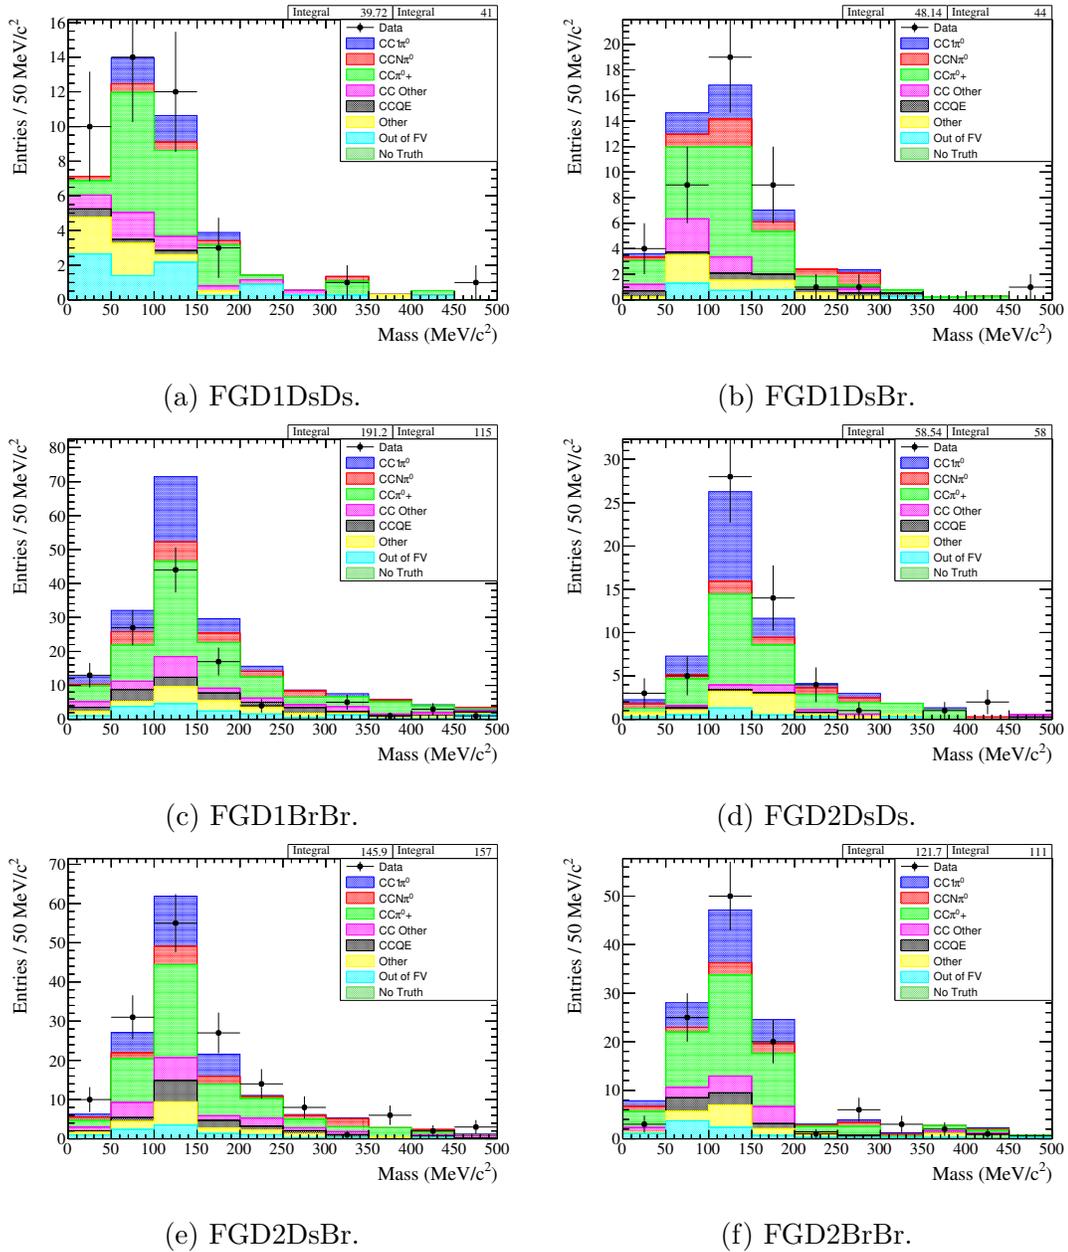
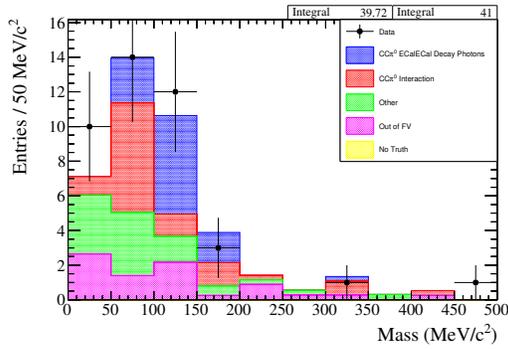
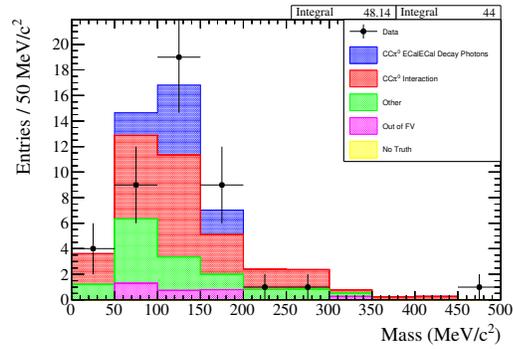


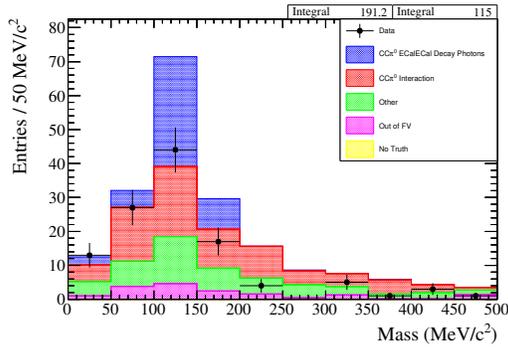
Figure 5.34: Reconstructed π^0 invariant mass distribution for each event topology over all running periods, broken down by final state.



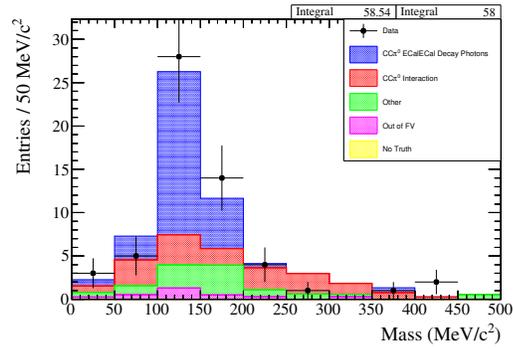
(a) FGD1DsDs.



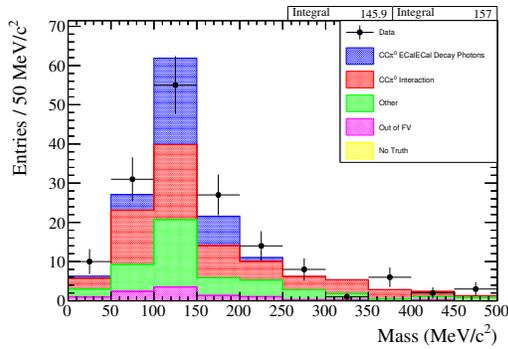
(b) FGD1DsBr.



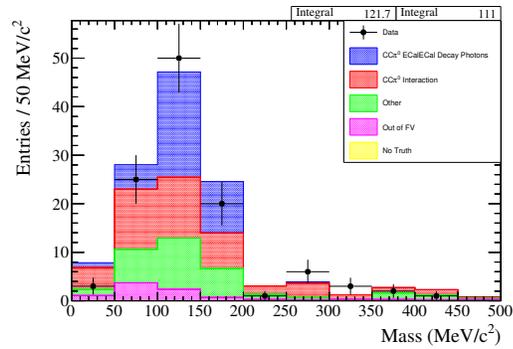
(c) FGD1BrBr.



(d) FGD2DsDs.



(e) FGD2DsBr.



(f) FGD2BrBr.

Figure 5.35: Reconstructed π^0 invariant mass distribution for each event topology over all running periods, broken down by π^0 decay photon reconstruction.

Cut	Number of selected events in data and (MC)		
	Run 2 Water	Run 2 Air	Run 3b
Ecal Multiplicity	433 (497.31)	394 (407.56)	191 (242.62)
Reconstructed Mass	383 (443.97)	345 (366.32)	173 (218.85)
Multivariate Selection	42 (51.28)	45 (41.41)	16 (25.63)
Pairs Multiplicity	38 (45.62)	40 (37.88)	14 (23.83)
Topology Selection			
FGD1DsDs	0 (2.70)	3 (2.70)	2 (1.79)
FGD1DsBr	7 (4.13)	2 (3.16)	1 (1.89)
FGD1BrBr	8 (14.67)	5 (11.60)	3 (6.67)
FGD2DsDs	2 (4.40)	4 (3.33)	2 (3.69)
FGD2DsBr	11 (8.77)	13 (8.73)	1 (7.05)
FGD2BrBr	10 (10.86)	13 (8.36)	5 (2.74)

Table 5.17: The number of selected events in data and (MC) for each data taking period as a function of analysis cut. The total across all running periods is also shown along with the data over MC ratio with its statistical error.

	Number of selected events in data and (MC)				Total	$\frac{\text{data}}{\text{MC}}$ ratio
	Run 3c	Run 4 Water	Run 4 Air			
ECal Multiplicity	1561 (1632.47)	1966 (1971.18)	2032 (2126.39)	<i>6577</i> <i>(6877.52)</i>	<i>0.96</i> \pm <i>0.01</i>	
Reconstructed Mass	1401 (1452.77)	1748 (1757.03)	1813 (1878.79)	<i>5863</i> <i>(6117.73)</i>	<i>0.96</i> \pm <i>0.01</i>	
Multivariate Selection	122 (156.38)	161 (185.98)	187 (199.16)	<i>573</i> <i>(659.85)</i>	<i>0.87</i> \pm <i>0.04</i>	
Pairs Multiplicity	112 (143.26)	152 (171.83)	170 (182.67)	<i>526</i> <i>(605.09)</i>	<i>0.87</i> \pm <i>0.04</i>	
Topology Selection						
FGD1DsDs	13 (9.10)	9 (10.44)	14 (13.00)	<i>41</i> <i>(39.72)</i>	<i>1.03</i> \pm <i>0.18</i>	
FGD1DsBr	9 (12.76)	14 (12.28)	11 (13.92)	<i>44</i> <i>(48.14)</i>	<i>0.91</i> \pm <i>0.15</i>	
FGD1BrBr	21 (44.97)	41 (54.69)	37 (58.51)	<i>115</i> <i>(191.20)</i>	<i>0.60</i> \pm <i>0.06</i>	
FGD2DsDs	14 (16.44)	18 (13.39)	18 (17.29)	<i>58</i> <i>(58.54)</i>	<i>0.99</i> \pm <i>0.14</i>	
FGD2DsBr	32 (34.76)	41 (41.69)	59 (44.85)	<i>157</i> <i>(145.90)</i>	<i>1.08</i> \pm <i>0.10</i>	
FGD2BrBr	23 (25.24)	29 (39.35)	31 (35.10)	<i>111</i> <i>(121.70)</i>	<i>0.91</i> \pm <i>0.09</i>	

Table 5.17: The number of selected events in data and (MC) for each data taking period as a function of analysis cut. The total across all running periods is also shown along with the data over MC ratio with its statistical error.

ECal Module	$\frac{data}{MC}$ ratio							
	E_h				E_l			
	FGD1 Event Vertex				FGD2 Event Vertex			
BrECal left side	0.88	± 0.05	0.94	± 0.04	0.95	± 0.05	0.86	± 0.05
BrECal right side	1.02	± 0.05	0.89	± 0.04	0.96	± 0.05	0.90	± 0.05
BrECal top	0.88	± 0.04	0.91	± 0.04	0.94	± 0.04	0.91	± 0.05
BrECal bottom	0.94	± 0.04	0.95	± 0.04	0.94	± 0.04	1.00	± 0.05
DsECal	1.10	± 0.05	1.04	± 0.04	1.02	± 0.04	1.03	± 0.03

Table 5.18: Data over MC ratio, with statistical error, for the total number of higher and lower energy isolated ECal clusters reconstructed in each ECal module after the Reconstructed Mass cut has been applied, broken down by event vertex location.

have also affected the results for the other event topologies, and all bad data should have already been removed from the analysis by the data quality cuts.

Therefore to further investigate the issue, data over MC ratios were calculated for the numbers of isolated objects that were reconstructed in each module, and for the number of events that could possibly contribute to each event topology of interest, prior to the BDT cut being applied, as shown in Tables 5.18 and 5.19 respectively.

From Table 5.18 it is seen that all the BrECal modules tend to exhibit a deficit, around 10%, in both higher and lower energy isolated ECal clusters, whereas the DsECal shows a slight excess, about 5%. As seen in Table 5.19, this leads to an excess in contributions towards the DsDs event topologies and a deficit in the BrBr topologies, with the differences effectively cancelling out for the DsBr topologies, leading to good data-MC agreement in event rates for those topologies. Crucially, both FGDS behave similarly as there is no discrepancy between the event rates for the comparable FGD1 and FGD2 topologies, with both the FGD1BrBr and FGD2BrBr topologies both showing a 10% deficit in their pre-BDT cut event rate contributions. For the FGD2BrBr topology, this 10% deficit then continues through to the final selection as one might expect, and a similar effect is seen for all other topologies except FGD1BrBr which then exhibits the resultant 40%

Event Topology	$\frac{data}{MC}$	ratio
FGD1DsDs	1.18	± 0.09
FGD1DsBr	1.02	± 0.04
FGD1BrBr	0.90	± 0.02
FGD2DsDs	1.07	± 0.06
FGD2DsBr	1.00	± 0.03
FGD2BrBr	0.91	± 0.03

Table 5.19: Data over MC ratio, with statistical error, for the total number of events that could contribute to each topology after the Reconstructed Mass cut has been applied.

deficit.

Further possible explanations for the final FGD1BrBr deficit could be due to a data-MC discrepancy within the discriminating variable distributions used as inputs to the BDT discriminator, or a discrepancy between the final data and MC samples, and the initial BDT training sample distributions. Figure 5.36 shows the data and MC distributions for each BDT input to the FGD1BrBr topology after the Reconstructed Mass cut has been applied. As can be seen, no one variable shows a significant discrepancy within these distributions, and certainly no more than the other topologies as shown in Appendix C. Furthermore, comparing the shape of the distributions in Figure 5.36 with those used to train the FGD1BrBr BDT in Figure A.5, no variable looks likely to be inaccurately modelled within the training samples. Either possible explanation also seems unlikely, as it is doubtful that one variable alone is sufficiently important across the forest of all decision trees to significantly effect the final result to the extent observed. To fully test both possibilities each variable in turn would have to be removed from the BDT training, a new training and cut optimisation performed, and the results reproduced with each of the new, reduced BDTs. Unfortunately, there was not sufficient time to perform such a comprehensive examination of the FGD1BrBr deficit and so this remains an outstanding issue for future investigation.

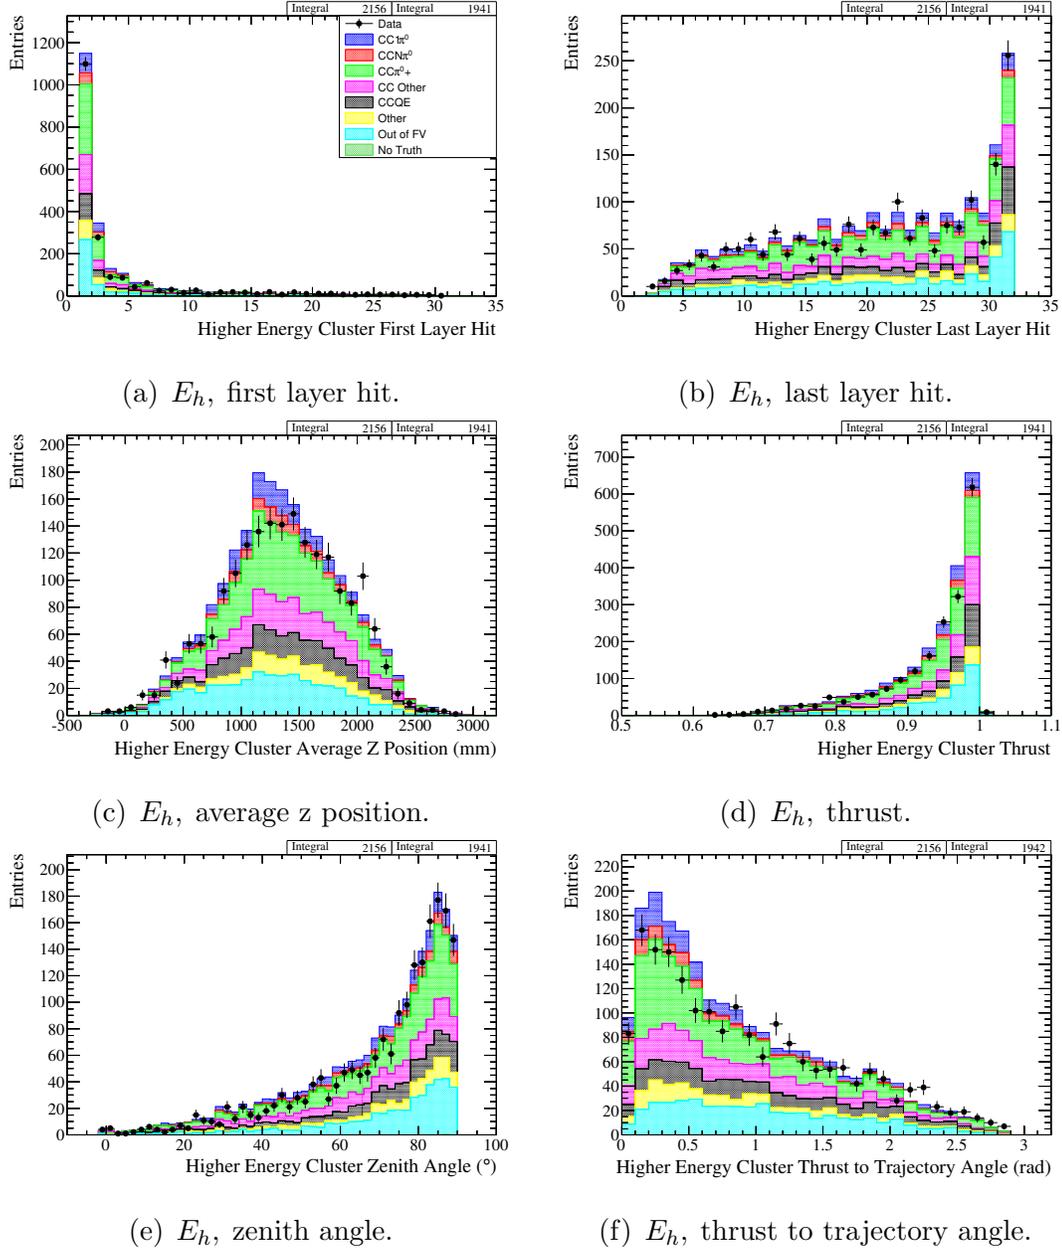


Figure 5.36: Data-MC comparison plots for events that could contribute to the final FGD1BrBr event topology, after the Reconstructed Mass cut has been applied.

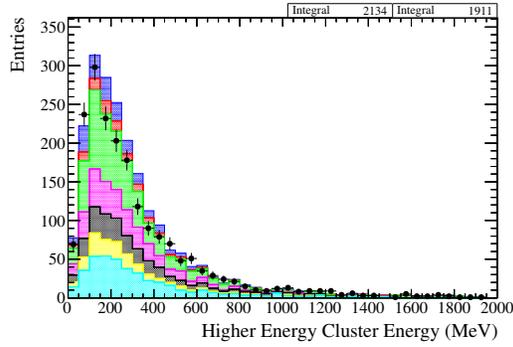
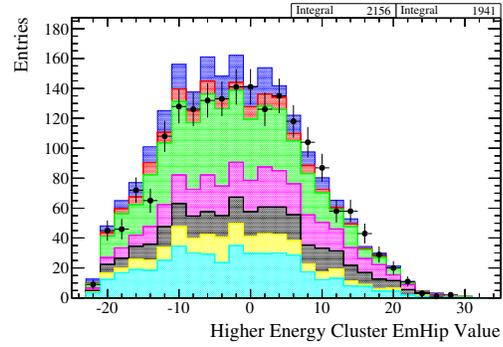
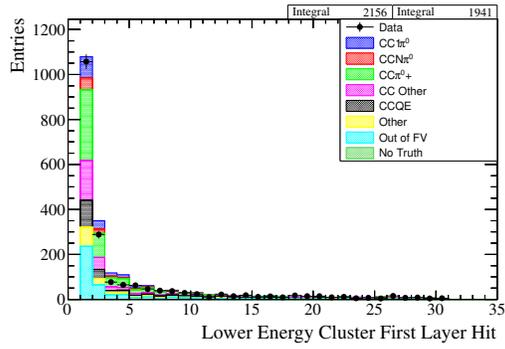
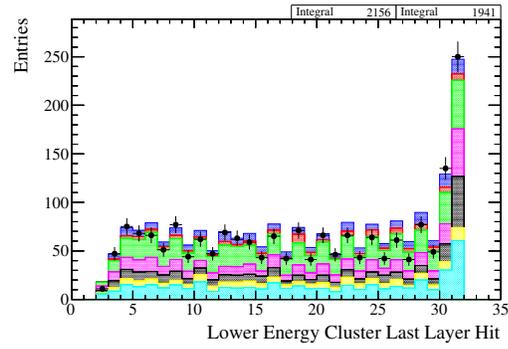
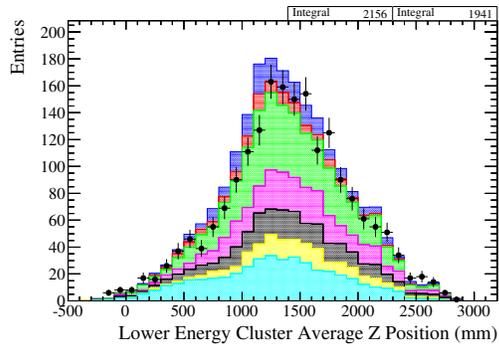
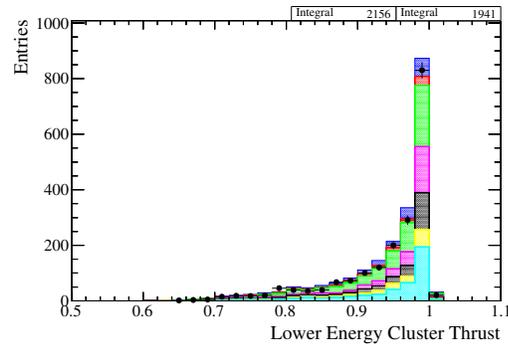
(g) E_h , EM energy.(h) E_h , EmHip.(i) E_l , first layer hit.(j) E_l , last layer hit.(k) E_l , average z position.(l) E_l , thrust.

Figure 5.36: Data-MC comparison plots for events that could contribute to the final FGD1BrBr event topology, after the Reconstructed Mass cut has been applied.

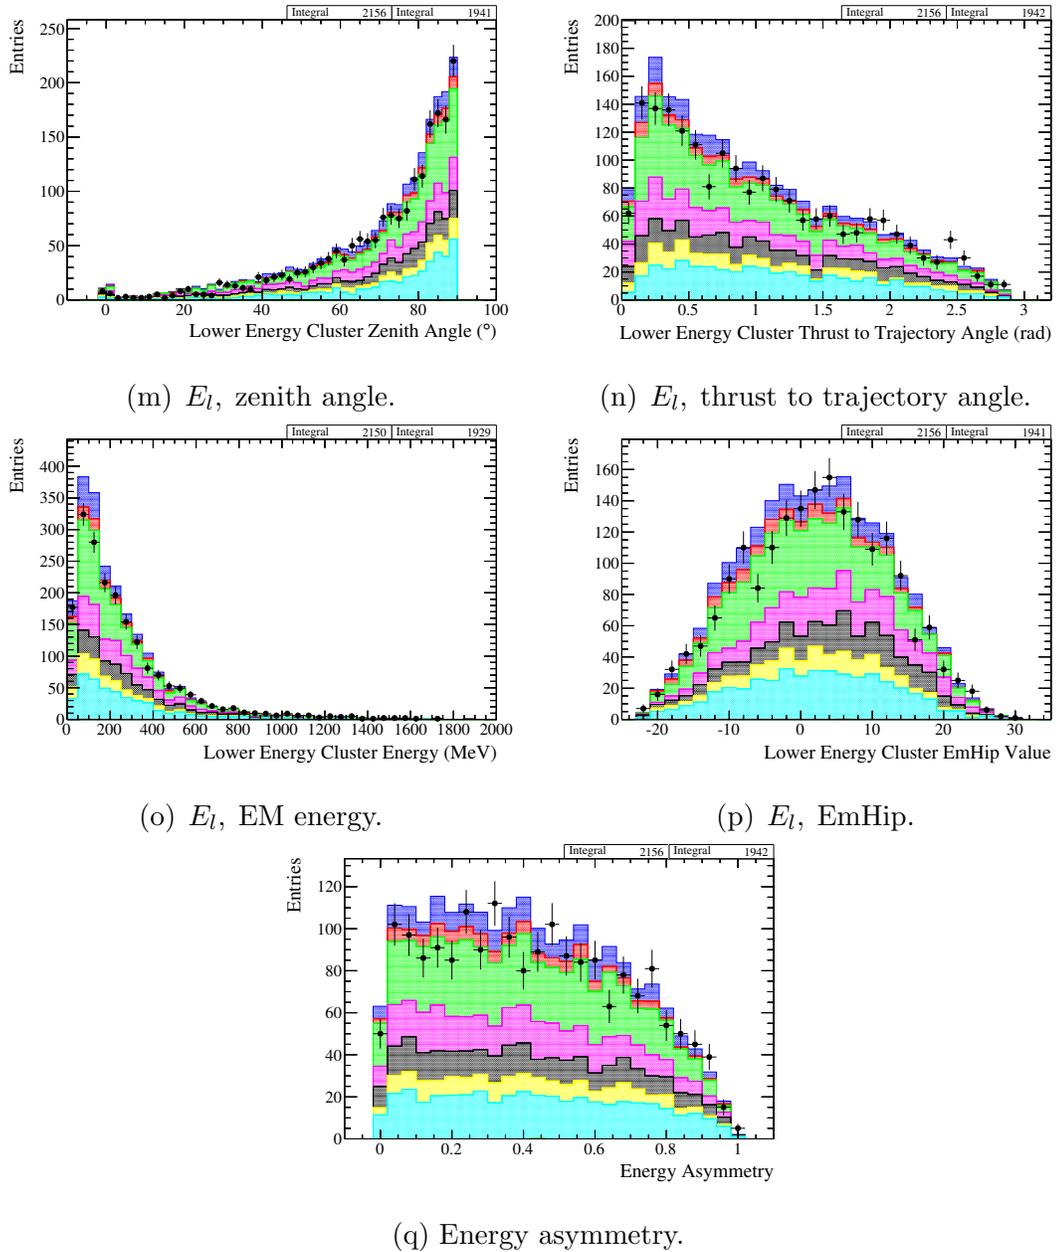


Figure 5.36: Data-MC comparison plots for events that could contribute to the final FGD1BrBr event topology, after the Reconstructed Mass cut has been applied.

5.6 Systematic Uncertainty

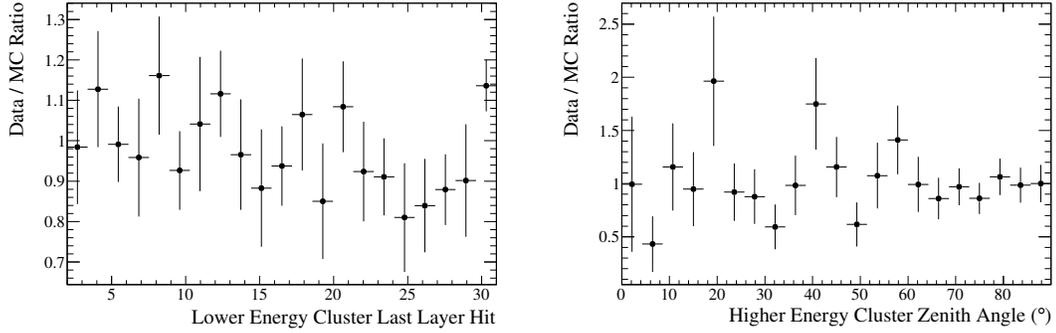
At this time, a full analysis of the systematic uncertainties associated with this analysis is not possible. However, a consideration of the systematic error associated with the use of the Multivariate Analysis is considered in Section 5.6.1, along with the flux uncertainty in Section 5.6.2. A qualitative discussion of the cross-section and detector systematics are then provided in Sections 5.6.3 and 5.6.4 respectively.

5.6.1 Multivariate Analysis Systematic Error

As shown in Section 5.5.5 the area normalised data and MC comparisons show good agreement. However, given that there are some disagreements within those distributions, disagreements that would systematically lead to different proportions of events passing the cut nodes in each Decision Tree, a measure of this effect is required. The method employed within this analysis is outlined below with a detailed description of each step:

1 Calculate Data-MC Ratio for BDT Inputs,

For each variable passed to the BDT, a bin by bin Data over MC ratio is calculated after all the pre-MVA cuts have been applied to the data and MC samples. This is done for each topology independently, such that pairwise combinations of ECal clusters that would be passed to a specific topology-dependent BDT, only contribute to that topology's data over MC ratio calculation. The binning used to calculate the data over MC ratio for each variable and for each topology, is matched to the binning used at the root node when the BDT training took place. This always employed 21 bins of equal width with a range between the minimum and maximum values found within the BDT training sample for each variable. Prior to the bin by bin ratio being calculated, the distributions in data and MC were area normalised across the entire range to ensure the ratio is a measure of shape difference, and does not convolve this systematic error calculation with those that may explain the overall data deficit prior to the MVA cut



(a) E_l , last layer hit data over MC ratio for clusters contributing to the FGD1BrBr topology. (b) E_h , zenith angle data over MC ratio for clusters contributing to the FGD2DsDs topology.

Figure 5.37: Example of binned data over MC ratios used in calculating the Multivariate Analysis systematic error. The error bar shown on each data point is the statistical error.

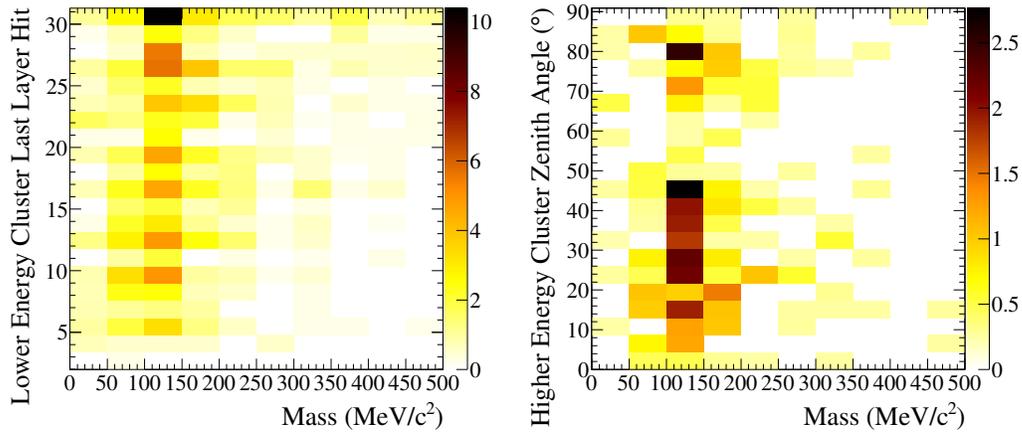
such as the flux or cross-section uncertainty. Figure 5.37 shows example data over MC ratios for a couple of the BDT input variables.

2 Obtain Reconstructed π^0 Mass as a Function of BDT Inputs,

For events passing all the analysis cuts, for each topology of interest, the reconstructed π^0 mass is calculated as a function of each BDT input variable. The binning for each variable continues to use the same binning as when the data over MC ratio was calculated. Figure 5.38 shows the reconstructed π^0 mass as a function of BDT input variable, for those data over MC ratio examples shown in Figure 5.37.

3 Multiply Through by Data-MC Ratio,

To then measure the effect of the BDT input shape discrepancy on the final result for each topology, the reconstructed π^0 mass as a function of BDT input variable is multiplied through by the binned data over MC ratio calculated at the first step. This reweights the final result by the data over MC discrepancy for each variable, examples of which are shown in Figure 5.39.



(a) E_l , last layer hit as a function of reconstructed π^0 mass for all events contributing to the final FGD1BrBr topology selection.

(b) E_h , zenith angle as a function of reconstructed π^0 mass for all events contributing to the final FGD2DsDs topology selection.

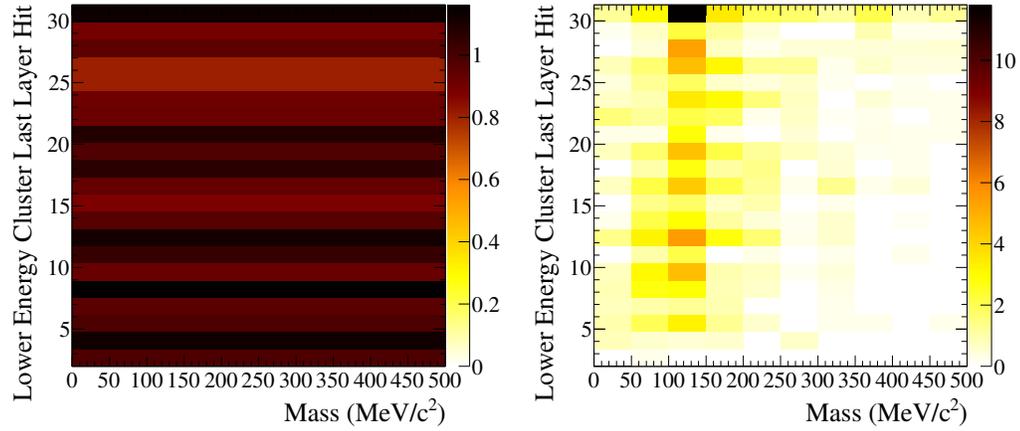
Figure 5.38: Example of reconstructed π^0 mass as a function of BDT inputs for all events contributing to the final selection within those topologies.

4 Integrate Data-MC Ratio Scaled Reconstructed π^0 Mass and Calculate Systematic Error,

Once the reconstructed π^0 mass has been scaled by the data over MC ratio, the integrated area of the resultant histogram is calculated. The difference between the scaled and nominal MC is then a measure of the systematic error attributed to that variable within the BDT implementation. By summing in quadrature the differences between the scaled and nominal MC for every BDT input variable, a total multivariate analysis implementation systematic error is obtained.

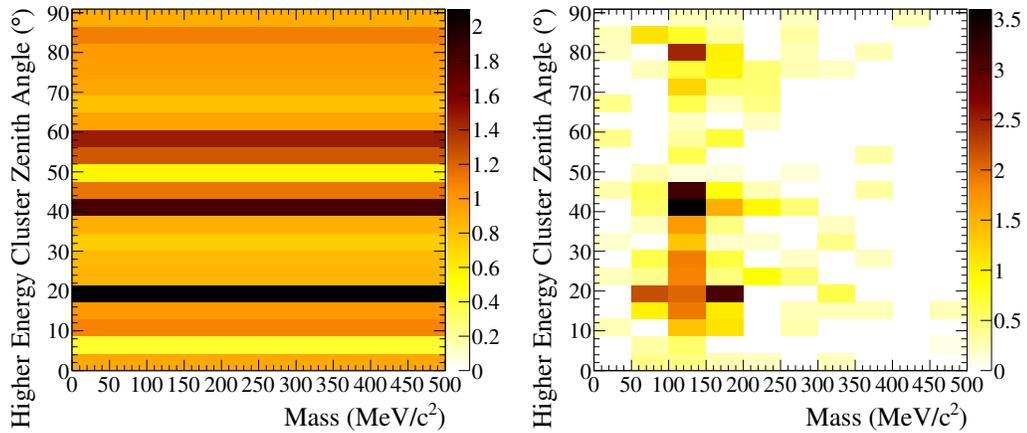
There are two clear issues with the methodology outlined above which, overall, lead it to being a conservative estimate of the systematic error attributed to each BDT implementation:

- Firstly, each variable is considered to be independent of all others, therefore failing to take into account the correlations that occur between some inputs as shown by Figure 5.23 and within Appendix A.



(a) E_l , last layer hit data over MC ratio from FGD1BrBr topology to multiply through Figure 5.38a with.

(b) E_l , last layer hit as a function of reconstructed π^0 mass from 5.38a once multiplied through by 5.39a.



(c) E_h , zenith angle data over MC ratio from FGD2DsDs topology to multiply through Figure 5.38b with.

(d) E_h , zenith angle as a function of reconstructed π^0 mass from 5.38b once multiplied through by 5.39c.

Figure 5.39: Data over MC ratio multiplying distributions from Figure 5.37 and the result of their application to Figure 5.38.

- Secondly, all variables are considered to carry equal weight in the way they influence the result produced by each BDT. This is clearly not true, as some variables will occur more frequently, and with greater discriminating power, in the forest of decision trees constructed for each BDT.

The first issue is difficult to deal with, and normally would lead to an underestimate of the systematic uncertainty, however it probably has minimal effect on the resultant systematic calculation as the variables used within the BDT implementation were chosen due to the low correlations between them (see Figure 5.23 and Appendix A). The second issue however, which leads to an overestimate of the systematic uncertainty, can be addressed by reweighting the initial data over MC ratios, used to scale the final MC selected events, by a factor proportional to each variable's importance in each BDT.

When the BDT for each topology was trained, a variable ranking was produced, as shown in Tables 5.12 and 5.13, which is based on the importance value calculated for each variable. The importance value *“is derived by counting how often the variables are used to split decision tree nodes, and by weighting each split occurrence by the separation gain-squared it has achieved and by the number of events in the node”* [93]. If the variable with the highest ranking in each BDT is defined such that it has a weight of 1 in the systematic error calculation, we can then scale the weight of the lower ranked variables by the ratio of their importance value divided by the importance of the highest ranked variable

$$w_i = \frac{\mathcal{I}_i}{\mathcal{I}_1}, \quad (5.11)$$

where w_i and \mathcal{I}_i are the systematic weight and importance value respectively of a BDT input variable i , and \mathcal{I}_1 is the importance value of the highest ranked BDT input variable.

This can then be used to reweight back towards unity, bin by bin, the data over MC ratios used to calculate the multivariate systematic error, such that

$$\mathcal{R}(x)_i^{reweigh} = 1 + ((\mathcal{R}(x)_i - 1) \times w_i), \quad (5.12)$$

where $\mathcal{R}(x)_i$ and $\mathcal{R}(x)_i^{reweigh}$ are the nominal and reweighted data over MC ratios in bin x for variable i .

Event Topology	Nominal MC Events Selected	Systematic Error			
		With		(Without)	
FGD1DsDs	39.7	± 7.6	$\pm 19.1\%$	(± 9.5)	($\pm 23.9\%$)
FGD1DsBr	48.1	± 2.7	$\pm 5.6\%$	(± 5.0)	($\pm 10.4\%$)
FGD1BrBr	191.2	± 17.0	$\pm 8.9\%$	(± 26.8)	($\pm 14.0\%$)
FGD2DsDs	58.5	± 7.8	$\pm 13.3\%$	(± 11.0)	($\pm 18.8\%$)
FGD2DsBr	145.9	± 9.0	$\pm 6.2\%$	(± 12.0)	($\pm 8.2\%$)
FGD2BrBr	121.7	± 11.0	$\pm 9.0\%$	(± 15.9)	($\pm 13.1\%$)

Table 5.20: Nominal MC event selection and associated multivariate analysis implementation systematic error with and (without) the variable importance reweighting.

Table 5.20 shows the resultant multivariate analysis implementation systematic, as calculated via the method outlined above, for each of the six topologies of interest. The mass peak distributions obtained when multiplying through by the data over MC ratios, with and without the variable importance reweighting are shown for comparison in Figures 5.40 and 5.41 respectively.

From Table 5.20 it is seen that the size of the systematic uncertainty, as a percentage of number of events selected for each topology, is comparable between the equivalent FGD1 and FGD2 event topologies. This suggests that the equivalent FGD1 and FGD2 topologies feature similar data and MC differences as would be expected. This is further supported by the fact that although the statistics contributing to each equivalent FGD1 and FGD2 topology can be significantly different, the results remain consistent, implying that it is the true shape difference between the data and MC that is being measured, not statistical fluctuations within each of the data samples.

It is also important to note that although the FGD1BrBr topology features the largest absolute MVA systematic uncertainty, as a percentage of the number of events selected it is consistent with the other topologies. This lends credence to the belief that there is nothing unusual about the scale of data to MC differences, or the BDT implementation, for this topology, and that the observed deficit is a genuine effect.

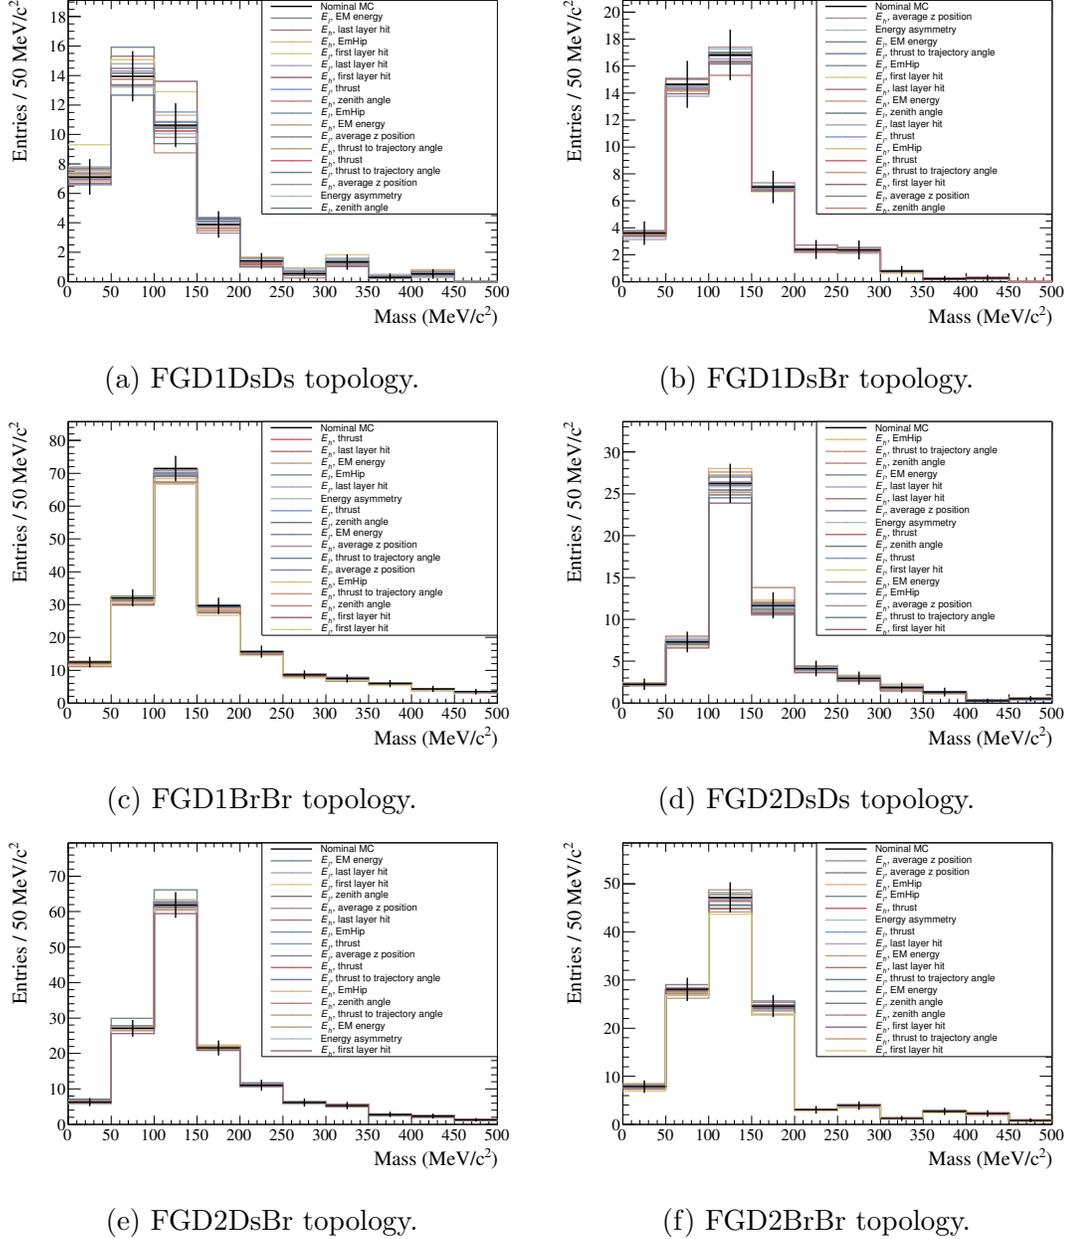


Figure 5.41: Reconstructed π^0 mass peak after multiplying through by the data over MC ratio for each of the six event topologies of interest. No variable importance reweighting is applied for these distributions.

5.6.2 Neutrino Flux Prediction Systematic Error

One potentially significant source of systematic error is that due to the uncertainty in the neutrino beam flux. As described in Section 2.1.6, the T2K experiment aims to accurately predict the neutrino flux, and a dedicated working group regularly provides updated beam fluxes and uncertainties for each neutrino flavour at both the ND280 and Super-Kamiokande. The flux uncertainties are provided to analysts in the form of a fractional covariance matrix which is binned by neutrino flavour, energy and detector as seen in Figure 5.42. The covariance matrix has 160 bins, 80 each for the ND280 (ND) and Super-Kamiokande (SK), and these are further broken down into groups of 20 bins for each of the four neutrino flavours present in the beam. The 20 bins for each neutrino flavour correspond to a different neutrino energy range; in GeV these are 0.0-0.1, 0.1-0.2, 0.2-0.3, 0.3-0.4, 0.4-0.5, 0.5-0.6, 0.6-0.7, 0.7-0.8, 0.8-1.0, 1.0-1.2, 1.2-1.5, 1.5-2.0, 2.0-2.5, 2.5-3.0, 3.0-3.5, 3.5-4.0, 4.0-5.0, 5.0-7.0, 7.0-10.0 and >10.0 .

The fractional covariance matrix is produced by evaluating the potential variation brought about by the uncertainty inherent in a comprehensive range of factors which contribute to the final neutrino flux observed at each detector, these include:

- The pion and kaon, primary and secondary nucleon production uncertainty, along with their associated interaction cross-sections; these are obtained from the results of the dedicated NA61/SHINE[57] target experiment, in combination with older experimental data from the work of by Allaby *et al.*[97], Eichten *et al.*[98], Bellettini *et al.*[99], Denisov *et al.*[100] and Carroll *et al.*[101].
- Off-axis angle uncertainty; this is derived from the variation in the measured beam direction by the MUMON and INGRID detectors.
- Horn and target alignment uncertainty; measured by dedicated survey.
- Proton beam uncertainty; derived from measurements made by the various beam monitors and their inherent uncertainties, along with beamline magnetic field strength and alignment uncertainties.

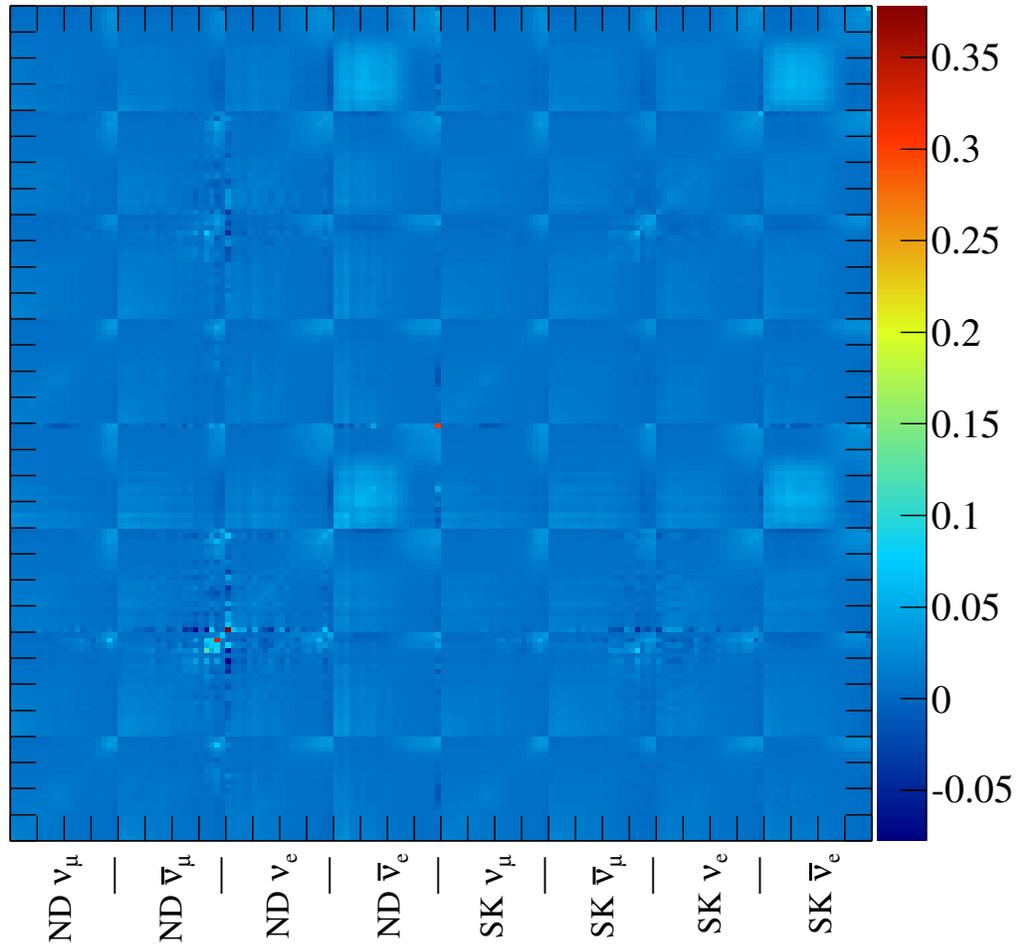


Figure 5.42: Fractional covariance matrix for the neutrino flux uncertainty as a function of neutrino flavour and detector.

- Absolute horn current uncertainty; measured by Rogowski coils.

To calculate the effect of the flux systematic uncertainty on the resultant π^0 mass distribution for each topology of interest, an event reweighting method is applied based upon random throws of the flux fractional covariance matrix as follows.

Each event selected in the MC samples corresponds to an original neutrino interaction of tuned event weight, Φ_i , which falls within a neutrino energy, detector and flavour flux bin, i , of the fractional covariance matrix. To calculate the effect of the uncertainty on that neutrino flux bin, the bin weight can be varied through a systematic parameter b_i , such that

$$\Phi'_i = b_i \Phi_i. \quad (5.13)$$

To generate a variation factor, first the fractional covariance matrix is decomposed, via Cholesky decomposition as shown in Figure 5.43, and then multiplied by a vector of random Gaussian distributed numbers with a mean of zero and sigma of one

$$b_i = \sum_{j=0}^{160} C_{ij} u_j, \quad (5.14)$$

where C_{ij} is the Cholesky decomposition of the fractional covariance matrix and u_j is the Gaussian throw corresponding to bin j of the covariance matrix.

Φ'_i is the fractional change in the neutrino flux within a specific detector for the specified neutrino flavour and energy, for one throw of the fractional covariance matrix, therefore

$$\Phi_i^{Throw} = \Phi_i + \Phi'_i. \quad (5.15)$$

After performing a throw of the flux fractional covariance matrix, a new π^0 mass distribution can be formed from the Φ_i^{Throw} values, and the difference between the thrown and nominal π^0 mass distribution calculated. By performing many independent throws, and taking the average of those throws, a true measure of the influence of the flux uncertainty can be obtained.

Because the different bins in the π^0 mass distribution are correlated, a covariance matrix is built up using each throw of the flux fractional covariance matrix such

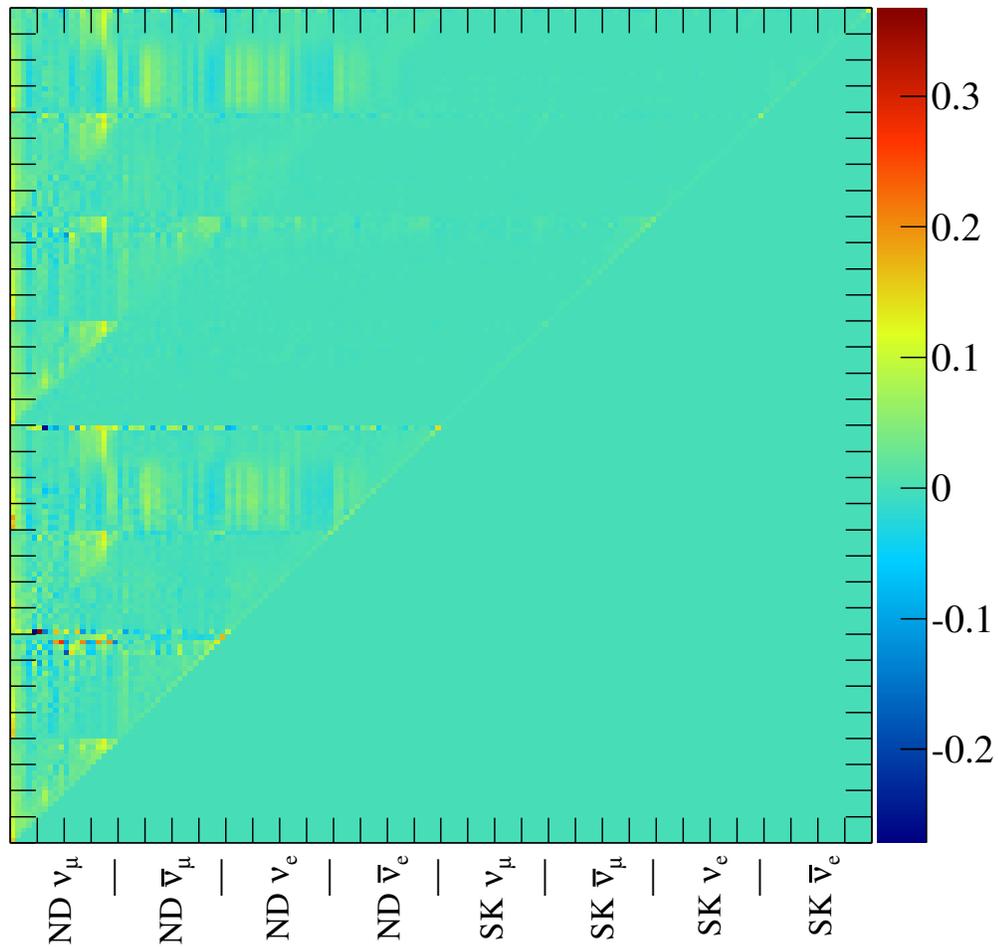


Figure 5.43: Cholesky decomposition of the fractional covariance matrix for the neutrino flux uncertainty as a function of neutrino flavour and detector.

that

$$v_{ij} = \frac{1}{N_t} \sum_{N_t} (x_i - n_i)(x_j - n_j), \quad (5.16)$$

where v_{ij} is the covariance between π^0 mass peak bins i and j , x and n are the number of entries in the corresponding π^0 mass peak bins for the thrown and nominal distributions respectively, and N_t is the number of throws performed.

The associated fractional covariance matrix is obtained by dividing by the product of nominal bin entries

$$v_{ij}^{\text{frac}} = \frac{v_{ij}}{n_i n_j}. \quad (5.17)$$

The correlation between the independent bins is given by

$$c_{ij} = \frac{v_{ij}}{\sigma_i \sigma_j}, \quad (5.18)$$

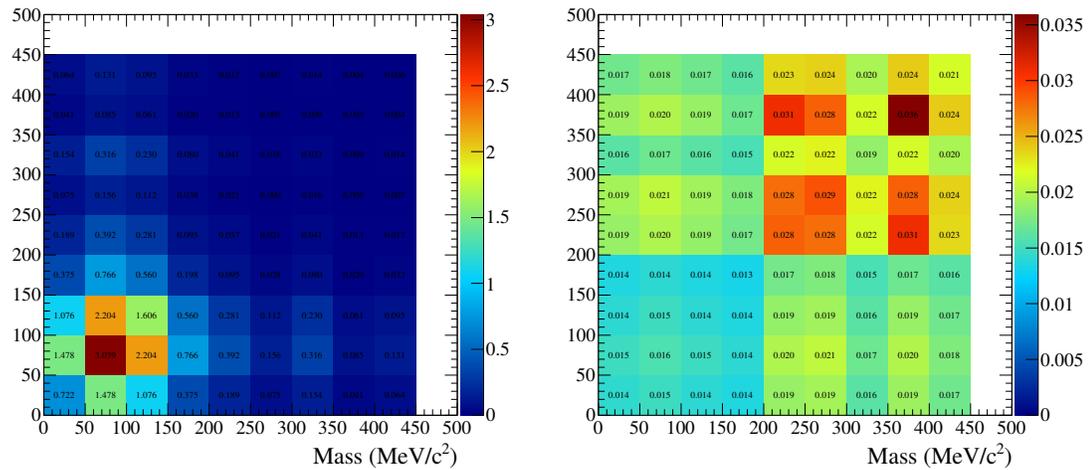
where c_{ij} is the correlation between π^0 mass peak bins i and j , and σ is the standard deviation (square root of the variance) for the corresponding π^0 mass peak bin.

For each event topology of interest, 5000 throws of the flux fractional covariance matrix were performed and correlation and covariance matrices constructed as shown for the FGD1DsDs topology in Figure 5.44, and for all other topologies in Appendix E.

The final error on each topology is obtained by taking the square root of the sum of variances plus twice the covariances, equivalent to the square root of the integral of the covariance matrix

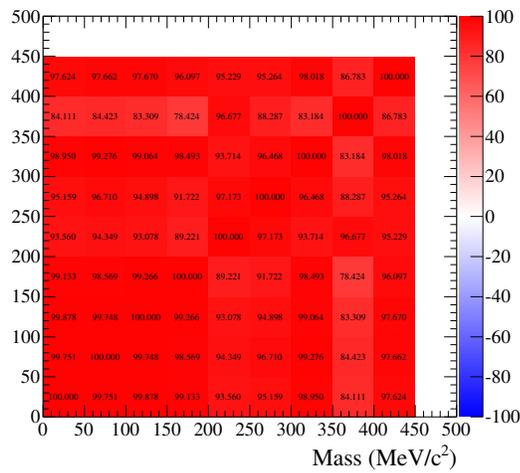
$$\sigma = \sqrt{\sum_i \sum_j v_{ij}}. \quad (5.19)$$

The results of which are shown in Table 5.21, with all event topologies evaluating to a flux systematic error of approximately 11%. This is consistent with the flux systematic uncertainty calculated for previous ND280 cross-section measurements, see for example [92]. The consistency across all topologies is also reassuring, as there should be no difference in the way the flux systematic contributes to each individual topology.



(a) Covariance matrix.

(b) Fractional covariance matrix.



(c) Correlation matrix (%).

Figure 5.44: Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD1DsDs event topology.

Event Topology	Nominal MC	Systematic Error	
	Events Selected	Events	(%)
FGD1DsDs	39.7	± 4.9	(± 12.4)
FGD1DsBr	48.1	± 5.4	(± 11.2)
FGD1BrBr	191.2	± 20.9	(± 10.9)
FGD2DsDs	58.5	± 6.3	(± 10.7)
FGD2DsBr	145.9	± 15.7	(± 10.8)
FGD2BrBr	121.7	± 13.5	(± 11.1)

Table 5.21: Nominal MC event selection and associated flux systematic error.

5.6.3 Cross-Section Systematic Error

The systematic error source that is likely to produce the largest uncertainty for this analysis is that associated with the neutrino-nucleus interaction as modelled by the NEUT MC generator. This systematic has two associated components, that associated with the neutrino interaction cross-section model parametrisation and that associated with the modelling of final state interactions within the nuclear environment.

It is important to consider the cross-section systematic uncertainty when comparing the obtained result to the MC prediction. In the future, this analysis should directly lead to, or contribute towards (in combination with other $CC\pi^0$ event topologies such as those described in [90]), a measurement of the the $CC1\pi^0$ interaction cross-section. Such a measurement would then be combined with previous results, such as those discussed in Section 1.4.2, to further constrain the MC model parameters, as described below, which directly feed into the systematic uncertainty calculation for the T2K oscillation analysis.

Cross-Section Model Parametrisation Systematic

The NEUT generator has 17 relevant input parameters when generating the cross-sections associated with different types of neutrino interactions. The values for these parameters, and their associated uncertainty, are calculated by the T2K

Parameter	E_ν range (GeV)	Nominal value	Error
M_A^{QE}	all	1.21 GeV/ c^2	0.45
M_A^{RES}	all	1.41 GeV/ c^2	0.11
p_F	all	274 MeV/ c	30
E_B	all	25 MeV	9
SF	all	0 (off)	1 (on)
<i>CC Other shape</i>	all	0.00	0.40
<i>Pion-less Δ decay</i>	all	0.00	0.20
<i>CCQE E1</i>	$0.0 < E_\nu < 1.5$	1.00	0.11
<i>CCQE E2</i>	$1.5 < E_\nu < 3.5$	1.00	0.30
<i>CCQE E3</i>	$E_\nu > 3.5$	1.00	0.30
<i>CC1π^+ E1</i>	$0.0 < E_\nu < 2.5$	1.15	0.43
<i>CC1π^+ E2</i>	$E_\nu > 2.5$	1.00	0.40
<i>CC Coh</i>	all	1.00	1.00
<i>NC1π^0</i>	all	0.96	0.43
<i>NC1π^\pm</i>	all	1.00	0.30
<i>NC Coh</i>	all	1.00	0.30
<i>NC Other</i>	all	1.00	0.30

Table 5.22: NIWG cross-section parametrisation and associated errors with applicable neutrino energy ranges for the NEUT MC generator.

Neutrino Interaction Working Group (NIWG) are detailed in Table 5.22. How these parametrisations and their uncertainties are obtained is briefly discussed as follows.

- M_A^{QE} and *CCQE E1*,

M_A^{QE} is the axial mass in the axial form factor for the CCQE cross-section calculation. This is used by NEUT in the Llewellyn-Smith [102] formalism of CCQE scattering on a single nucleon, and the Smith-Moniz [103] relativistic Fermi gas formalism for describing the nucleon as bound within a nuclear potential. The error on this value is calculated by first minimising the χ^2 value obtained when fitting NEUT CCQE MC generated interactions to the MiniBooNE CCQE-corrected double-differential cross-section data [104], as a function of M_A^{QE} and a normalisation factor. The difference between the χ^2 minimised value for M_A^{QE} obtained by the NIWG group,

$1.64 \pm 0.03 \text{ GeV}/c^2$, plus the fitted error, and the NEUT nominal value, $1.21 \text{ GeV}/c^2$, is then used as the error. *CCQE* E1 is the CCQE normalisation factor at low energies, and is assigned an uncertainty of 11% from the neutrino flux uncertainty in this MiniBooNE sample.

- **M_A^{RES} , $CC1\pi^+$ E1, CC *Coh*, $NC1\pi^0$, $NC1\pi^\pm$, NC *Coh* and NC *Other*,**

M_A^{RES} is the axial mass in the axial form factor for the resonance cross-section calculation. This is used by NEUT in the Rein-Seghal [105] model of CC and NC single pion production. Much like the M_A^{QE} error determination, the uncertainty on M_A^{RES} is obtained through fits to MiniBooNE data. Three MiniBooNE data samples are used, the $CC1\pi^0$ [45], $CC1\pi^+$ [106] and $NC1\pi^0$ [50] samples. Much like the M_A^{QE} uncertainty determination, a χ^2 minimisation is attempted simultaneously across all the data samples. The $CC1\pi^+$ E1 and $NC1\pi^0$ parameters are normalisation factors that may be simultaneously constrained by the fit, and are thus correlated with M_A^{RES} . The remaining normalisation parameters cannot be constrained independently by the fit, but appear as penalty terms in the χ^2 calculation as they effect the overall MC prediction. These normalisation parameters are then individually scaled against the relevant data samples to obtain an error value.

- **p_F and E_B ,**

p_F and E_B are the Fermi momentum and nuclear binding energy respectively. These are parameters used in the Relativistic Fermi Gas (RFG) model for the characterisation of the target nuclear potential. Their values and associated uncertainties have been measured through electron scattering experiments [107].

- **SF ,**

SF represents the spectral function, a more complex alternative to the RFG model of the nuclear potential, and one which better agrees with the electron scattering data. The fractional difference between the default RFG and SF models is taken as the model-dependent systematic uncertainty, after the total cross-sections have been renormalised to match one another at 1 GeV.

- ***CCQE E2, CCQE E3 and CC1 π^+ E2,***

These high energy normalisation factors cannot be constrained by MiniBooNE or the ND280 alone due to the low neutrino fluxes received by these detectors at those energies. The CCQE normalisations therefore come from the discrepancy seen between MiniBooNE and the higher energy NOMAD experiment [108], whilst the $CC1\pi^+$ E2 uncertainty comes from extrapolation of the discrepancy between the NEUT nominal and MiniBooNE cross-section fit at 2 GeV, to higher energies.

- ***CC Other shape,***

CC Other shape is an energy dependent uncertainty on the cross-section for CC multi-pion production, CC Deep Inelastic Scattering (DIS) and CC resonant production of $\eta/K/\gamma$. From MINOS data [109] this uncertainty is known to be 10% at 4 GeV, and from this the error is defined such that

$$\sigma_{\text{CC Other}} = \frac{0.4}{E_\nu (\text{GeV})}. \quad (5.20)$$

This formulation is prevented from going to infinity as E_ν approaches zero due to a required threshold energy of approximately 0.6 GeV.

- ***Pion-less Δ decay,***

Within NEUT it is possible for a Δ to decay without producing pions 20% of the time. These *Pion-less Δ decay* events are an irreducible background to CCQE fits, as it migrates $CC1\pi^+$ events to *CCQE*-like samples. This influences the previously discussed fits to single-pion data and was investigated by reducing the pion-less Δ decay fraction to zero, leading to a reduction in the $CC1\pi^+$ E1 and $NC1\pi^0$ normalisations by 20%.

In combination, the uncertainties on the cross-section model parametrisation have been shown to produce fractional systematic errors for ND280 tracker ν_μ samples of 17.7% for $CC0\pi$, 16.2% for $CC1\pi^+$, 10.4% for CC Other and 14.3% for CC Inclusive samples. The application of the cross-section model parametrisation directly to the $CC1\pi^0$ analysis described within this thesis is ongoing, but an error of approximately 15% to 20% would seem the likely outcome.

Final State Interactions Systematic

Once a neutrino-nucleon interaction has occurred, NEUT employs a microscopic nuclear cascade model, much like the one described by [110], to simulate the re-interaction of the resultant hadrons with the nuclear environment via the strong force. The re-interactions possible are QE scattering, particle absorption, charge exchange and particle production. Six parameters within NEUT scale the different contributing factors to the microscopic nuclear cascade model, these are:

- FSIABS, the pion absorption parameter.
- FSIQEL, the low energy QE scattering parameter, which includes single charge exchange.
- FSICXL, the low energy single charge exchange branching fraction parameter.
- FSIQEH, the high energy QE scattering parameter.
- FSICXH, the high energy single charge exchange branching fraction parameter.
- FSIINEL, the pion production parameter.

Each of these final state interaction (FSI) parameters are tuned through comparison between NEUT MC simulation of pion scattering, with varied FSI parametrisations, and 25 separate external pion scattering data sets (see [111] as an example). The simulated data has simple selection criteria applied to it, and the resultant selected events are compared to the relevant data samples. Certain combinations of FSI parametrisations provide good agreement to the data, constituting a tuned MC value, from which an error is extracted as the difference between the tuned and nominal NEUT FSI parametrisation values. Because different combinations of FSI parametrisation can provide good agreement to the data samples an error covariance is built up from the range of agreeable values for each FSI parametrisation.

To assess the effect of the FSI parametrisation systematic on the resultant event selection, events can be reweighted based on splines that are created for each FSI

parameter. The FSI parameters are considered factorisable, and the splines are evaluated for the desired throw of each FSI parameter and constrained via the covariance matrices produced from the tuned and nominal MC discrepancy.

The FSI systematic has been shown to produce fractional systematic errors for ND280 tracker ν_μ samples of 1.0% for CC0 π , 3.1% for CC1 π^+ , 2.6% for CC Other and 0.3% for CC Inclusive samples. Like the cross-section model parametrisation, direct application of the FSI parametrisation systematic uncertainty to the CC1 π^0 analysis described within this thesis is ongoing, but an error of approximately 3% seem the likely outcome.

5.6.4 Detector Systematic Error

The final sources of systematic uncertainty that will be briefly considered are those associated with the detector performance and how significant a difference between the MC simulation and real data samples this may introduce. This encompasses a very large number of systematic uncertainties and can be generally grouped into those relating to the muon selection in the tracking detectors, those relating to the π^0 selection in the ECals, and those relating to the event selection as a whole.

Many systematics relating to the selection of muons within the tracking detectors have been to some extent considered by the official ND280 ν_μ selection analysis. However, many of these systematics have only been considered for events with an event vertex in FGD1, and not yet assessed for FGD2, nor have they been fully integrated and validated within the highland analysis framework as used by the analysis described in this chapter. These include:

- FGD and TPC track reconstruction efficiency.
- FGD-TPC track matching efficiency.
- Track direction reconstruction efficiency.
- Event vertex reconstruction efficiency.
- TPC PID systematic uncertainty.

- TPC charge identification efficiency.
- TPC momentum resolution uncertainty.
- TPC E and B fields distortion.
- TPC track quality.
- TPC track merging uncertainty.
- FGD mass uncertainty.

Some of the systematics relating to the selection of events within the ECals have been considered by the official ND280 ν_e selection analysis, but like the ν_μ selection, most of these have neither been assessed for events with an FGD2 vertex, nor integrated and validated within the highland analysis framework. Relevant systematics include:

- ECal cluster reconstruction efficiency.
- FGD-ECal and TPC-ECal track matching efficiency.
- ECal energy reconstruction uncertainty.
- ECal mass uncertainty.
- ECal cluster merging uncertainty.
- ECal PID systematic uncertainty.

Finally there are some additional uncertainties that have implications for event selection as a whole as they may impact upon both the muon and π^0 components of the analysis.

- Detector alignment uncertainty.
- Event pile up uncertainty.
- Sand interaction uncertainty.
- Pion secondary interaction uncertainty.

Where considered in other, fully developed analyses, the systematic uncertainty introduced by the detector systematics have been found to be smaller ($< 10\%$) than the flux and cross-section uncertainties which always dominate. As such those systematics have been the primary consideration of this analysis, but detector systematics will be an important consideration during its future development.

5.7 Summary

The selection of $CC1\pi^0$ final states has been undertaken for events where both photons from the decay of the π^0 enter the Barrel and/or Downstream ECal modules. The selection is broken down into six topologies, dependent on the location of the muon vertex and position where the decay photons shower in the ECal. The MC prediction for the number of events in each topology, along with their associated efficiency and purity are summarised in Table 5.23. The dominant background to the analysis is from $CC\pi^0+$ final states, about 40%, whilst final states not containing a π^0 combine to produce around 30% of the background contamination.

When considering the systematic uncertainty from the use of the MVA technique and flux uncertainty (and 20% cross-section uncertainty), along with the statistical uncertainty the predicted event rates, and total found in data are:

- **FGD1DsDs,**
MC: 39.7 ± 2.9 (stat.) ± 9.0 (12.0) (syst.), Data: 41.
- **FGD1DsBr,**
MC: 48.1 ± 3.2 (stat.) ± 6.0 (11.4) (syst.), Data: 44.
- **FGD1BrBr,**
MC: 191.2 ± 6.3 (stat.) ± 26.9 (46.8) (syst.), Data: 115.
- **FGD2DsDs,**
MC: 58.5 ± 3.5 (stat.) ± 10.0 (15.4) (syst.), Data: 58.
- **FGD2DsBr,**
MC: 145.9 ± 5.5 (stat.) ± 18.1 (34.3) (syst.), Data: 157.

Topology	Total Events	CC1 π^0 Signal Topology			Any CC1 π^0 final state		π^0 Decay Photon	
		Events	Efficiency (%)	Purity (%)	Events	Purity (%)	Events	Purity (%)
FGD1DsDs	39.7	1.3	7.1	3.2	3.5	8.8	10.3	25.9
FGD1DsBr	48.1	3.6	10.0	7.4	5.8	12.0	9.2	19.0
FGD1BrBr	191.2	20.3	10.2	10.6	35.1	18.4	49.2	25.7
FGD2DsDs	58.5	13.2	9.4	22.6	16.4	28.0	29.1	49.7
FGD2DsBr	145.9	16.8	10.4	11.5	25.3	17.2	35.1	24.1
FGD2BrBr	121.7	16.1	9.0	13.2	23.4	19.2	38.7	31.8

Table 5.23: MC prediction for the number of events selected for each topology of interest, CC1 π^0 selection efficiency and purity for the specific event topology of interest, the selected purity from all CC1 π^0 final states, and π^0 decay photon purity, across all running periods.

Event Topology	p -value
FGD1DsDs	0.77
FGD1DsBr	0.62
FGD1BrBr	0.06
FGD1DsDs	0.30
FGD1DsBr	0.26
FGD1BrBr	0.42

Table 5.24: p -values calculated from the mass peak distributions for each event topology using the statistical uncertainties on each histogram bin as error input.

- **FGD2BrBr,**

MC: 121.7 ± 5.0 (stat.) ± 17.4 (29.9) (syst.), Data: 111.

The agreement, within the uncertainties quoted, between the data and MC is excellent except for the FGD1BrBr topology. As previously discussed at the end of Section 5.5.7, there is no clear explanation for the cause of this discrepancy and so further work is required to fully understand this issue. However, it is worth considering that if we combine the statistical and systematic uncertainties, including 20% cross-section uncertainty, together in quadrature, then the FGD1BrBr deficit is equivalent to a 1.6σ effect, something that should occur for about 10% of observations.

Furthermore, a χ^2 goodness-of-fit statistical test was used to calculate p -values [112] checking the compatibility of the data and MC distributions. Using only the statistical uncertainties on the data and MC samples for each topology, the p -values shown in Table 5.24 were obtained. These all evaluate to a value greater than the standard incomparability threshold of 0.05, including the FGD1BrBr topology which has a p -value of 0.06. This indicates that the data and MC event selections are behaving as expected and that any discrepancy could be due to statistical fluctuations, rather than any specific short coming within the analysis.

Chapter 6

Discussion and Conclusions

The T2K experiment is a long-baseline neutrino oscillation experiment based at the J-PARC facility in eastern Japan. The experiment features a high intensity ν_μ beam, 2.5° off-axis angle and 295 km baseline with both near and far detectors. It was the first experiment to report observation of $\nu_\mu \rightarrow \nu_e$ oscillations in 2011 [34], indicating a non-zero value for the final unknown mixing angle θ_{13} .

In the most recent T2K ν_e appearance analysis [113] the fractional systematic error on the expected event rate at Super-Kamiokande was about 10%, which includes important contributions from the beam flux, cross-section and final state interaction uncertainties. The systematic uncertainties associated with 17 independent neutrino interaction cross-section model parameters were constrained through fits to external data sets as described in Section 5.6.3, including the Mini-BooNE $\text{CC}1\pi^0$ cross-section measurement shown in Section 1.4.2. Additionally, constraints were placed on the energy spectrum of the neutrino beam and 5 cross-section parameters which are common to both the ND280 and SK, by dedicated fits to three ν_μ ND280 data sets; the CCQE, $\text{CC}1\pi^+$ and CC Other enhanced event samples. In the future, both the independent and common cross-section parameters, along with the beam energy spectrum uncertainties will be better constrained by use of native ND280 cross-section measurements and additional ND280 final state enhanced data samples such as $\text{CC}1\pi^0$.

As such, the analysis presented within Chapter 5 is the first step towards measuring $\text{CC}1\pi^0$ interactions within the FGD fiducial volume of the ND280, for

which both the π^0 decay photons shower within the Barrel and/or Downstream ECal modules. In combination with other $CC1\pi^0$ final states, such as those described in [90], it is expected that the analysis described will lead to a $CC1\pi^0$ cross-section measurement. Additionally the techniques being developed within this analysis may assist in the development of the equivalent NC interaction measurement. A native $NC1\pi^0$ cross-section measurement by the ND280 would be invaluable to reducing the background uncertainty for ν_e appearance searches with SK, as $NC1\pi^0$ interactions are the second largest background contribution to these searches after ν_e contamination in the ν_μ beam. The analysis presented may also contribute to the development of a π^0 veto for those studying other dedicated final states such as CCQE or $CC1\pi^\pm$. This should then improve the quality of the current and future ND280 final state enhanced data samples used to constrain the cross-section and beam energy spectrum uncertainties for the T2K oscillation analysis.

The analysis described was divided into two sections, firstly a search was undertaken to locate a muon-like track with a vertex within the fiducial volume of either FGD. This produced a CC inclusive-like event selection with 50% efficiency and 86% purity. The muon candidate is a true muon for 87% of selected events, with the background dominated by misidentified negatively charged pions.

If a muon candidate is successfully selected, then isolated ECal clusters are identified, and if two or three are present these can be used in an attempt to identify a π^0 candidate pair. A Boosted Decision Tree Multivariate Analysis technique is employed to select π^0 candidates using the ROOT TMVA software package. Six separate BDTs are used to select the π^0 candidate based on the specific event topology presented.

Across all topologies of interest, a total selection efficiency of 10% and purity of 12% is obtained. The purity increased to 18% when any $CC1\pi^0$ final state is considered, and the selection produces a purity of 28% when considering the correct identification of π^0 s from any final state. $CC1\pi^0+$ final states make up the majority of the background at 40%. Given that the analysis is effective at correctly identifying π^0 interactions, in future it would probably be worth further examining the other tracks emerging from the event vertex in the hope of better

partitioning the resultant samples by final state and to veto background events. The selection was performed over the T2K Run 2 to Run 4 data sets, which after data quality cuts, constitute 57.34×10^{19} POT. After all cuts 526 events were selected in data against a MC expectation of 605. The deficit can almost entirely be attributed to the FGD1BrBr event topology which displayed a 40% deficit. When broken down by each topology of interest the results were as follows,

- **FGD1DsDs,**
MC: 39.7 ± 2.9 (stat.) ± 9.0 (12.0) (syst.), Data: 41.
- **FGD1DsBr,**
MC: 48.1 ± 3.2 (stat.) ± 6.0 (11.4) (syst.), Data: 44.
- **FGD1BrBr,**
MC: 191.2 ± 6.3 (stat.) ± 26.9 (46.8) (syst.), Data: 115.
- **FGD2DsDs,**
MC: 58.5 ± 3.5 (stat.) ± 10.0 (15.4) (syst.), Data: 58.
- **FGD2DsBr,**
MC: 145.9 ± 5.5 (stat.) ± 18.1 (34.3) (syst.), Data: 157.
- **FGD2BrBr,**
MC: 121.7 ± 5.0 (stat.) ± 17.4 (29.9) (syst.), Data: 111.

which are consistent with the MC expectation within the calculated errors for all topologies except FGD1BrBr.

The analysis presented shows promise in selecting $CC1\pi^0$ final states through the identification of the muon in the ND280 tracking detectors and the showers from the decay of the π^0 within the surrounding ECals. However, the resultant efficiency and purity are low, with large systematic uncertainties, but it should be feasible to improve this through future updates to the presented selection. For instance, further examination of the charged tracks exiting the event vertex may allow the identification and rejection of charged pions, particularly if they can be tracked into the surrounding ECals and found to begin showering. Additionally, a larger *magnet* MC sample would help remove the requirement to use

the broad range of MC samples for training the BDTs, hopefully leading to an improvement in the accuracy for selecting π^0 s. One of the main challenges this analysis faces is the difficulty presented by π^0 decay photons with energies below 150 MeV failing to be reconstructed within the ECals. This immediately causes the potential selection efficiency to be significantly reduced, and so low energy reconstruction improvements could provide great benefit to future iterations of this, and similar analyses.

To eventually make a cross-section measurement competitive with the MiniBooNE result, (see Section 1.4.2 and Reference [45]), it will also be crucial to reduce the size of the systematic uncertainty associated with the result shown. The MiniBooNE analysis features a systematic uncertainty of 18.7%, and this has contributions from the flux uncertainty at 10.5%, cross-section and final state interaction uncertainty at about 14% and detector systematics at nearly 6.5%. The analysis presented within Chapter 5 has a comparable flux uncertainty at about 11% across all event topologies considered. The cross-section and final state interaction uncertainty is anticipated to be larger than that of MiniBooNE, at around 20%, as are the detector uncertainties at near 10%. Additionally, the use of the Boosted Decision Tree introduces a further systematic uncertainty, which although varying significantly between topology, shall be considered to be about 10% for the considerations of this discussion. Together these uncertainties combine to give a total uncertainty of 26.9% and would therefore require reducing to become directly competitive with the published MiniBooNE result.

It is not infeasible that the BDT systematic could be reduced to 5% or less across all topologies, but the flux uncertainty is unlikely to change, and given the complexity of the ND280 and the number of components used within this analysis, it is unlikely that a significant reduction in the detector systematic can be made either. Even if major reductions were possible for those systematic uncertainties, the total systematic is still dominated by the cross-section and final state interaction uncertainty, which alone is comparable in size to the total MiniBooNE systematic uncertainty. As it stands reducing this uncertainty is difficult as multiple external cross-section measurements contribute to its constraint, particularly those measurements made by MiniBooNE. As seen in Figure 1.9, these

measurements can show considerable disagreement with the interaction models currently implemented. As such, T2K takes a conservative approach, which ensures the calculated cross-section uncertainties account for the full discrepancy between the current model predictions and available cross-section measurements, including those from MiniBooNE.

There are two clear ways in which the cross-section uncertainty could be reduced. Firstly, the interaction models implemented within the NEUT MC generator used by T2K can be improved to show reduced disagreement with data measurements. This is a process that is ongoing, with significant improvements being made by the T2K neutrino interaction working group. Recently the NEUT MC generator was updated by the group to include models that consider Multi Nucleon Interactions (MEC) [114][115], Spectral Functions (SF) [116] and the Random Phase Approximation (RPA) [117]. Secondly, it is crucial to improve the purity of the $CC1\pi^0$ enhanced sample that is obtained. This would reduce the cross-section uncertainty attributed to the significant range of final states within the background contamination which are in the majority for the sample obtained. If these were reduced, the total cross-section uncertainty would also reduce, and so this should be a primary aim when improving the analysis in the future.

As it stands, the analysis would not provide a cross-section measurement that could be considered competitive with the MiniBooNE result. But given the intense work being done to improve the cross-section uncertainties at T2K, and the planned improvements for the analysis presented, it should be possible to become competitive in the not too distant future.

Bibliography

- [1] W. Pauli, “Pauli letter collection: letter to Lise Meitner.” Typed copy, December 1930.
- [2] C. D. Ellis and W. A. Wooster, “The average energy of disintegration of radium E,” *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 117, no. 776, pp. 109–123, 1927.
- [3] J. Chadwick, “The existence of a neutron,” *Proceedings of the Royal Society A*, vol. 136, no. 830, pp. 692–708, 1932.
- [4] E. Fermi, “Versuch einer theorie der β -strahlen. i,” *Zeitschrift für Physik*, vol. 88, pp. 161–177, 1934.
- [5] F. Reines and C. L. Cowan, “Detection of the free neutrino,” *Phys. Rev.*, vol. 92, pp. 830–831, Nov 1953.
- [6] F. Reines and C. L. Cowan, “The neutrino,” *Nature*, vol. 178, pp. 446–449, September 1956.
- [7] B. Pontecorvo, “Chalk River Laboratory PD-205 report.” 1946.
- [8] J. N. Bahcall, *Neutrino Astrophysics*. Cambridge University Press, 1989.
- [9] R. Davis, “Attempt to detect the anti-neutrinos from a nuclear reactor by the $\text{Cl}^{37}(\bar{\nu}, e^-)\text{A}^{37}$ reaction,” *Phys. Rev.*, vol. 97, pp. 766–769, Feb 1955.
- [10] G. Danby *et al.*, “Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos,” *Phys. Rev. Lett.*, vol. 9, pp. 36–44, Jul 1962.

- [11] M. L. Perl *et al.*, “Evidence for anomalous lepton production in $e^+ - e^-$ annihilation,” *Phys. Rev. Lett.*, vol. 35, pp. 1489–1492, Dec 1975.
- [12] J. Burmester *et al.*, “Anomalous muon production in $e^+ - e^-$ annihilations as evidence for heavy leptons,” *Physics Letters B*, vol. 68, no. 3, pp. 297–300, 1977.
- [13] B. Adeva *et al.*, “A determination of the properties of the neutral intermediate vector boson Z^0 ,” *Physics Letters B*, vol. 231, no. 4, pp. 509–518, 1989.
- [14] D. DeCamp *et al.*, “Determination of the number of light neutrino species,” *Physics Letters B*, vol. 231, no. 4, pp. 519–529, 1989.
- [15] M. Akrawy *et al.*, “Measurement of the Z^0 mass and width with the opal detector at LEP,” *Physics Letters B*, vol. 231, no. 4, pp. 530–538, 1989.
- [16] P. Aarnio *et al.*, “Measurement of the mass and width of the Z^0 -particle from multi-hadronic final states produced in $e^+ e^-$ annihilations,” *Physics Letters B*, vol. 231, no. 4, pp. 539–547, 1989.
- [17] G. S. Abrams *et al.*, “Measurements of Z -boson resonance parameters in $e^+ e^-$ annihilation,” *Phys. Rev. Lett.*, vol. 63, pp. 2173–2176, Nov 1989.
- [18] The ALEPH Collaboration, The DELPHI Collaboration, The L3 Collaboration, The OPAL Collaboration and The SLD Collaboration, “Precision electroweak measurements on the Z resonance,” *Physics Reports*, vol. 427, no. 5-6, pp. 257–454, 2006.
- [19] K. Kodama *et al.*, “Observation of tau neutrino interactions,” *Physics Letters B*, vol. 504, no. 3, pp. 218–224, 2001.
- [20] B. T. Cleveland *et al.*, “Measurement of the solar electron neutrino flux with the Homestake chlorine detector,” *The Astrophysical Journal*, vol. 496, pp. 505–526, March 1998.
- [21] J. N. Bahcall, M. H. Pinsonneault, and G. J. Wasserburg, “Solar models with helium and heavy-element diffusion,” *Rev. Mod. Phys.*, vol. 67, pp. 781–808, Oct 1995.

- [22] W. Hampel *et al.*, “GALLEX solar neutrino observations: results for GALLEX IV,” *Physics Letters B*, vol. 447, no. 1-2, pp. 127–133, 1999.
- [23] J. N. Abdurashitov *et al.*, “Measurement of the solar neutrino capture rate with gallium metal. iii. results for the 2002-2007 data-taking period,” *Phys. Rev. C*, vol. 80, p. 015807, Jul 2009.
- [24] J. N. Bahcall, A. M. Serenelli, and S. Basu, “New solar opacities, abundances, helioseismology, and neutrino fluxes,” *The Astrophysical Journal*, vol. 621, pp. 85–88, March 2005.
- [25] D. Casper *et al.*, “Measurement of atmospheric neutrino composition with the IMB-3 detector,” *Phys. Rev. Lett.*, vol. 66, pp. 2561–2564, May 1991.
- [26] K. Hirata *et al.*, “Observation of a small atmospheric ν_μ/ν_e ratio in Kamiokande,” *Physics Letters B*, vol. 280, no. 1-2, pp. 146–152, 1992.
- [27] Y. Fukuda *et al.*, “Atmospheric ν_μ/ν_e ratio in the multi-GeV energy range,” *Physics Letters B*, vol. 335, no. 2, pp. 237–245, 1994.
- [28] Y. Fukuda *et al.*, “Study of the atmospheric neutrino flux in the multi-GeV energy range,” *Physics Letters B*, vol. 436, no. 1-2, pp. 33–41, 1998.
- [29] Y. Fukuda *et al.*, “Evidence for oscillation of atmospheric neutrinos,” *Phys. Rev. Lett.*, vol. 81, pp. 1562–1567, Aug 1998.
- [30] Q. R. Ahmad *et al.*, “Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury neutrino observatory,” *Phys. Rev. Lett.*, vol. 89, p. 011301, Jun 2002.
- [31] B. Pontecorvo, “Mesonium and anti-mesonium,” *J. Exptl. Theoret. Phys.*, vol. 33, pp. 549–551, 1957.
- [32] Z. Maik, M. Nakagawa, and S. Sakata, “Remarks on the unified model of elementary particles,” *Progress of Theoretical Physics*, vol. 28, pp. 870–880, Nov 1962.
- [33] J. Beringer *et al.*, “2012 review of particle physics,” *Phys. Rev. D*, vol. 86, 2012.

- [34] K. Abe *et al.*, “Indication of electron neutrino appearance from an accelerator-produced off-axis muon neutrino beam,” *Phys. Rev. Lett.*, vol. 107, p. 041801, Jul 2011.
- [35] Y. Abe *et al.*, “Indication of reactor $\bar{\nu}_e$ disappearance in the Double Chooz experiment,” *Phys. Rev. Lett.*, vol. 108, p. 131801, Mar 2012.
- [36] J. K. Ahn *et al.*, “Observation of reactor electron antineutrinos disappearance in the RENO experiment,” *Phys. Rev. Lett.*, vol. 108, p. 191802, May 2012.
- [37] F. P. An *et al.*, “Observation of electron-antineutrino disappearance at Daya Bay,” *Phys. Rev. Lett.*, vol. 108, p. 171803, Apr 2012.
- [38] K. A. Olive *et al.*, “2014 review of particle physics,” *Chin. Phys. C*, vol. 38, 2014.
- [39] M. Hitoshi, “Oscillation parameter plot, 2014 PDG update,” 2014. Available at <http://hitoshi.berkeley.edu/neutrino/>.
- [40] D. Ayres *et al.*, “NO ν A proposal to build a 30 kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline,” 2005. Available at hep-ex/0503053v1.
- [41] S. Davidson, E. Nardi, and Y. Nir, “Leptogenesis,” *Physics Reports*, vol. 466, no. 4-5, pp. 105–177, 2008.
- [42] S. F. King and C. Luhn, “Neutrino mass and mixing with discrete symmetry,” *Reports on Progress in Physics*, vol. 76, no. 5, p. 056201, 2013.
- [43] J. A. Formaggio and G. P. Zeller, “From eV to EeV: Neutrino cross sections across energy scales,” *Rev. Mod. Phys.*, vol. 84, pp. 1307–1341, Sep 2012.
- [44] D. Casper, “The nuance neutrino physics simulation, and the future,” *Nuclear Physics B - Proceedings Supplements*, vol. 112, pp. 161–170, 2002.
- [45] A. A. Aguilar-Arevalo *et al.*, “Measurement of ν_μ -induced charged-current neutral pion production cross sections on mineral oil at $E_\nu \in 0.5\text{--}2.0$ GeV,” *Phys. Rev. D*, vol. 83, p. 052009, Mar 2011.

- [46] J. Catala-Perez, *Measurement of neutrino induced charged current neutral pion production cross section at SciBooNE*. PhD thesis, Universitat de València, 2014.
- [47] C. Mariani *et al.*, “Measurement of inclusive π^0 production in the Charged-Current Interactions of Neutrinos in a 1.3-GeV wide band beam,” *Phys. Rev. D*, vol. 83, p. 054023, 2011.
- [48] S. Nakayama *et al.*, “Measurement of single π^0 production in neutral current neutrino interactions with water by a 1.3-GeV wide band muon neutrino beam,” *Phys. Lett. B*, vol. 619, pp. 255–262, 2005.
- [49] Y. Kurimoto *et al.*, “Measurement of Inclusive Neutral Current Neutral π^0 Production on Carbon in a Few-GeV Neutrino Beam,” *Phys. Rev. D*, vol. 81, p. 033004, 2010.
- [50] A. A. Aguilar-Arevalo *et al.*, “Measurement of ν_μ and $\bar{\nu}_\mu$ induced neutral current single π^0 production cross sections on mineral oil at $E_\nu \sim \mathcal{O}(1 \text{ GeV})$,” *Phys. Rev. D*, vol. 81, p. 013005, Jan 2010.
- [51] K. Abe *et al.*, “The T2K experiment,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 659, no. 1, pp. 106–135, 2011.
- [52] Y. Hayato *et al.*, “Letter of intent: Neutrino oscillation experiment at JHF,” 2003. Available at <http://neutrino.kek.jp/jhfnu/loi/loi.v2.030528.pdf>.
- [53] K. Matsuoka *et al.*, “Design and performance of the muon monitor for the T2K neutrino oscillation experiment,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 624, no. 3, pp. 591–600, 2010.
- [54] S. Bhadra *et al.*, “Optical transition radiation monitor for the T2K experiment,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 703, pp. 45–58, 2013.

- [55] K. Abe *et al.*, “T2K neutrino flux prediction,” *Phys. Rev. D*, vol. 87, p. 012001, Jan 2013.
- [56] D. Beavis, A. Carroll, and I. Chiang, “Long baseline neutrino oscillation experiment at the AGS,” tech. rep., Apr. 1995. Available at <http://adsabs.harvard.edu/abs/1995bnl...rept.....B>.
- [57] N. Abgrall *et al.*, “Pion emission from the T2K replica target: Method, results and application,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 701, pp. 99–114, 2013.
- [58] R. Brun *et al.*, “GEANT3,” *CERN-DD-EE-84-1*, 1987.
- [59] A. Ferrari *et al.*, “FLUKA: a multi-particle transport code,” *CERN-2005-10*, 2005.
- [60] C. Zeitnitz and T. Gabriel, “The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 349, no. 1, pp. 106–111, 1994.
- [61] K. Abe *et al.*, “Measurements of the T2K neutrino beam properties using the INGRID on-axis near detector,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 694, pp. 211–223, 2012.
- [62] A. Pla-Dalmau, A. D. Bross, and V. V. Rykalin, “Extruding plastic scintillator at Fermilab,” *FERMILAB-Conf-03-318-E*, 2003.
- [63] G. Arnison *et al.*, “Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $s^{*(1/2)} = 540\text{-GeV}$,” *Phys.Lett.*, vol. B122, pp. 103–116, 1983.
- [64] J. Altegoer *et al.*, “The NOMAD experiment at the CERN SPS,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 404, no. 1, pp. 96–128, 1998.

- [65] S. Aoki *et al.*, “The T2K side muon range detector (SMRD),” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 698, pp. 135–146, 2013.
- [66] S. Assylbekov *et al.*, “The T2K ND280 off-axis pi-zero detector,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 686, pp. 48–63, 2012.
- [67] P.-A. Amaudruz *et al.*, “The T2K fine-grained detectors,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 696, pp. 1–31, 2012.
- [68] N. Abgrall *et al.*, “Time projection chambers for the T2K near detectors,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 637, no. 1, pp. 25–46, 2011.
- [69] D. Allan *et al.*, “The electromagnetic calorimeter for the T2K near detector ND280,” *Journal of Instrumentation*, vol. 8, October 2013.
- [70] B. Still, *T2K ND280 π^0 Electromagnetic Calorimeter*. PhD thesis, The University of Sheffield, 2009.
- [71] J. Estrada, C. Garcia, B. Hoeneisen, and P. Rubinov, “MCM II and the Trip chip,” *FERMILAB-TM-2226 (DØ note 4009)*, 2002.
- [72] A. Vacheret *et al.*, “Characterization and simulation of the response of multi-pixel photon counters to low light levels,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 656, no. 1, pp. 69–83, 2011.
- [73] Kuraray Co. Ltd., “Wavelength shifting fibers: Formulations and technical data.” Available at <http://kuraraypsf.jp/psf/ws.html>.

- [74] S. Fukuda *et al.*, “The Super-Kamiokande detector,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 501, no. 2-3, pp. 418–462, 2003.
- [75] R. Wendell *et al.*, “Atmospheric neutrino oscillation analysis with subleading effects in Super-Kamiokande I, II, and III,” *Phys. Rev. D*, vol. 81, p. 092004, May 2010.
- [76] K. Abe *et al.*, “Solar neutrino results in Super-Kamiokande-III,” *Phys. Rev. D*, vol. 83, p. 052010, Mar 2011.
- [77] R. Brun and F. Rademakers, “ROOT - an object oriented data analysis framework,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 389, no. 1-2, pp. 81–86, 1997. New Computing Techniques in Physics Research V.
- [78] S. Agostinelli *et al.*, “Geant4, - a simulation toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.
- [79] S. Ritt, P. Amaudruz, and K. Olchanski, “MIDAS - Maximum Integration Data Acquisition System.” Available at <https://midas.triumf.ca>.
- [80] A. Cervera-Villanueva, J. Gómez-Cadenas, and J. Hernando, “RecPack, - a reconstruction toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 534, no. 1-2, pp. 180–183, 2004. Proceedings of the IXth International Workshop on Advanced Computing and Analysis Techniques in Physics Research.
- [81] C. Andreopoulos *et al.*, “The GENIE neutrino monte carlo generator,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 614, no. 1, pp. 87–104, 2010.
- [82] Y. Hayato, “A neutrino interaction simulation program library NEUT,” *Acta Phys.Polon.*, vol. B40, pp. 2477–2489, 2009.

- [83] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Journal of Basic Engineering*, vol. 82, pp. 35–45, 1960.
- [84] K. Pearson, “On lines and planes of closest fit to systems of points in space,” *Philosophical Magazine*, vol. 2, pp. 559–572, 1901.
- [85] P. Guzowski, *Reconstruction of neutrino induced neutral current neutral pion events with the T2K ND280 Tracker and ECAL*. PhD thesis, Imperial College London, 2011.
- [86] R. C. Prim, “Shortest connection networks and some generalizations,” *The Bell System Technical Journal*, vol. 36, pp. 1389–1401, Nov 1957.
- [87] S. Short, *Study of Neutrino-Induced Neutral Current Neutral Pion Production in the T2K Near Detector*. PhD thesis, Imperial College London, 2013.
- [88] P. Adamson *et al.*, “First observations of separated atmospheric ν and $\bar{\nu}$ events in the MINOS detector,” *Phys. Rev. D*, vol. 73, p. 072002, Apr 2006.
- [89] Xilinx, “RocketIO™ Transceiver User Guide,” 2007. Available at http://www.xilinx.com/support/documentation/user_guides/ug024.pdf.
- [90] M. Batkiewicz, for the T2K collaboration, “ ν_μ CC π^0 reaction in the Tracker of the ND280 detector in the T2K experiment,” *NuInt12: Eighth International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region*, 2012. Available at <https://indico.fnal.gov/contributionDisplay.py?contribId=47&sessionId=16&confId=5361>.
- [91] Y. Hayato, “NEUT,” *Nucl.Phys.Proc.Suppl.*, vol. 112, pp. 171–176, 2002.
- [92] K. Abe *et al.*, “Measurement of the inclusive ν_μ charged current cross section on carbon in the near detector of the T2K experiment,” *Phys. Rev. D*, vol. 87, p. 092003, May 2013.
- [93] A. Hocker *et al.*, “TMVA - Toolkit for Multivariate Data Analysis,” *Proceedings, 11th International Workshop on Advanced Computing and Analysis Techniques in Physics Research*, 2007.

- [94] Y. Freund and R. E. Schapire, “A decision-theoretic generalization of on-line learning and an application to boosting,” *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119–139, 1997.
- [95] F. J. Massey, Jr., “The Kolmogorov-Smirnov Test for Goodness of Fit,” *Journal of the American Statistical Association*, vol. 46, pp. 68–78, March 1951.
- [96] G. Barker *et al.*, “Implementation of the second generation PID for the ND280 tracker ECals.” T2K Technical Note 111, August 2012.
- [97] J. Allaby *et al.*, “High-energy particle spectra from proton interactions at 19.2 GeV/c,” 1970.
- [98] T. Eichten *et al.*, “Particle production in proton interactions in nuclei at 24 GeV/c,” *Nuclear Physics B*, vol. 44, no. 2, pp. 333–343, 1972.
- [99] G. Bellettini *et al.*, “Proton-nuclei cross sections at 20 GeV,” *Nuclear Physics*, vol. 79, no. 3, pp. 609–624, 1966.
- [100] S. Denisov *et al.*, “Absorption cross sections for pions, kaons, protons and antiprotons on complex nuclei in the 6 to 60 GeV/c momentum range,” *Nuclear Physics B*, vol. 61, pp. 62–76, 1973.
- [101] A. Carroll *et al.*, “Absorption cross section of π^\pm , K^\pm , p and \bar{p} on nuclei between 60 and 280 GeV/c,” *Physics Letters B*, vol. 80, no. 3, pp. 319–322, 1979.
- [102] C. H. Llewellyn-Smith, “Neutrino reactions at accelerator energies,” *Physics Reports*, vol. 3, no. 5, pp. 261–379, 1972.
- [103] R. Smith and E. Moniz, “Neutrino reactions on nuclear targets,” *Nuclear Physics B*, vol. 43, pp. 605–622, 1972.
- [104] A. A. Aguilar-Arevalo *et al.*, “First measurement of the muon neutrino charged current quasielastic double differential cross section,” *Phys. Rev. D*, vol. 81, p. 092005, May 2010.

- [105] D. Rein and L. M. Sehgal, “Neutrino-excitation of baryon resonances and single pion production,” *Annals of Physics*, vol. 133, no. 1, pp. 79–153, 1981.
- [106] A. A. Aguilar-Arevalo *et al.*, “Measurement of neutrino-induced charged-current charged pion production cross sections on mineral oil at $E_\nu \sim 1$ GeV,” *Phys. Rev. D*, vol. 83, p. 052007, Mar 2011.
- [107] E. J. Moniz, I. Sick, R. R. Whitney, J. R. Ficenec, R. D. Kephart, and W. P. Trower, “Nuclear fermi momenta from quasielastic electron scattering,” *Phys. Rev. Lett.*, vol. 26, pp. 445–448, Feb 1971.
- [108] V. Lyubushkin *et al.*, “A Study of quasi-elastic muon neutrino and antineutrino scattering in the NOMAD experiment,” *Eur.Phys.J.*, vol. C63, pp. 355–381, 2009.
- [109] P. Adamson *et al.*, “Neutrino and antineutrino inclusive charged-current cross section measurements with the MINOS near detector,” *Phys. Rev. D*, vol. 81, p. 072002, Apr 2010.
- [110] L. Salcedo, E. Oset, M. Vicente-Vacas, and C. Garcia-Recio, “Computer simulation of inclusive pion nuclear reactions,” *Nuclear Physics A*, vol. 484, no. 3-4, pp. 557–592, 1988.
- [111] D. Ashery, I. Navon, G. Azuelos, H. K. Walter, H. J. Pfeiffer, and F. W. Schlepütz, “True absorption and scattering of pions on nuclei,” *Phys. Rev. C*, vol. 23, pp. 2173–2185, May 1981.
- [112] N. D. Gagunashvili, “Comparison of weighted and unweighted histograms,” *ArXiv Physics e-prints*, May 2006.
- [113] K. Abe *et al.*, “Observation of electron neutrino appearance in a muon neutrino beam,” *Phys. Rev. Lett.*, vol. 112, p. 061802, Feb 2014.
- [114] J. Nieves, I. R. Simo, and M. J. V. Vacas, “Inclusive charged-current neutrino-nucleus reactions,” *Phys. Rev. C*, vol. 83, p. 045501, Apr 2011.

- [115] R. Gran, J. Nieves, F. Sanchez, and M. J. V. Vacas, “Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV,” *Phys. Rev. D*, vol. 88, p. 113007, Dec 2013.
- [116] O. Benhar, A. Fabrocini, S. Fantoni, and I. Sick, “Spectral function of finite nuclei and scattering of GeV electrons,” *Nuclear Physics A*, vol. 579, no. 3-4, pp. 493–517, 1994.
- [117] J. Nieves, J. E. Amaro, and M. Valverde, “Inclusive quasielastic charged-current neutrino-nucleus reactions,” *Phys. Rev. C*, vol. 70, p. 055503, Nov 2004.

Appendix A

Discriminating Variable

Distributions and Correlations from BDT Training Samples

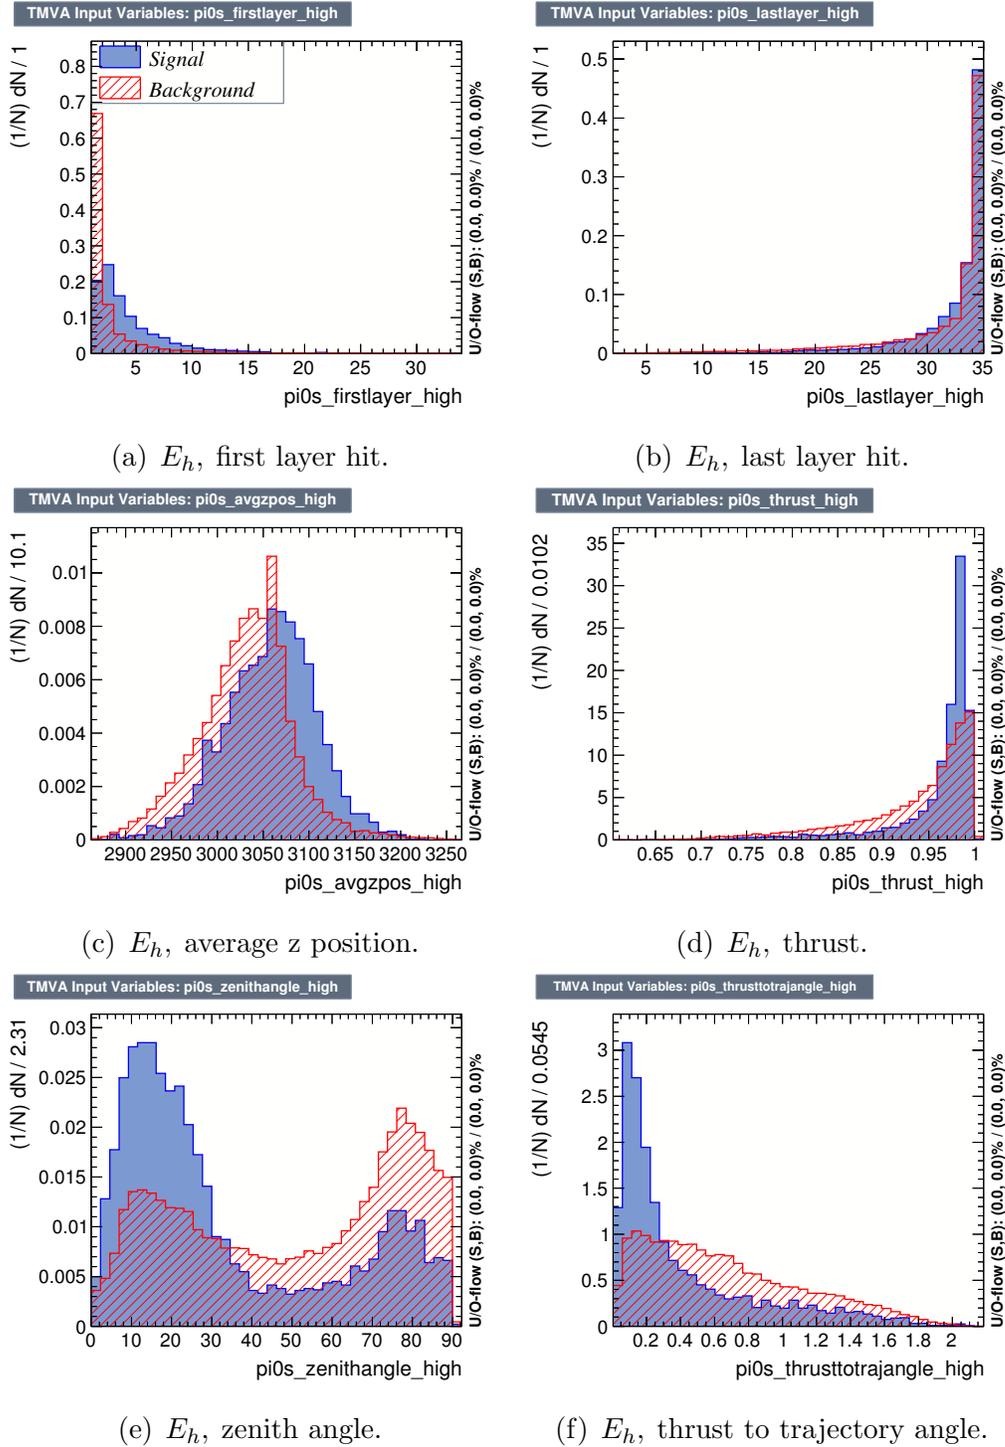
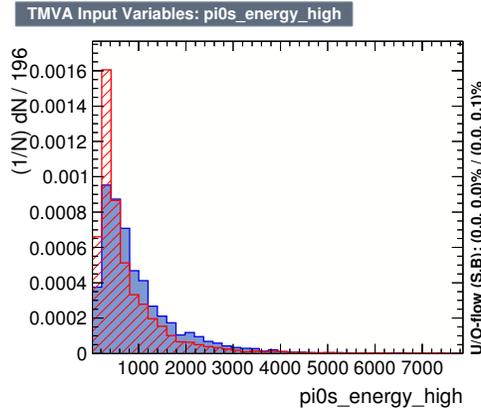
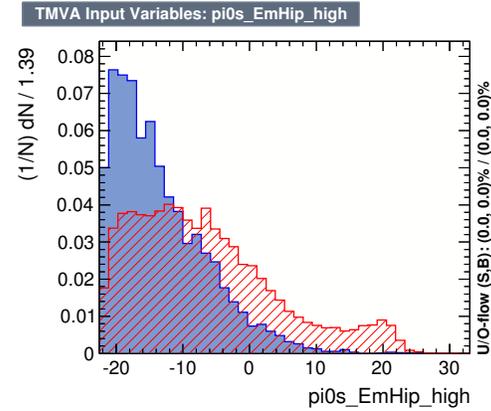


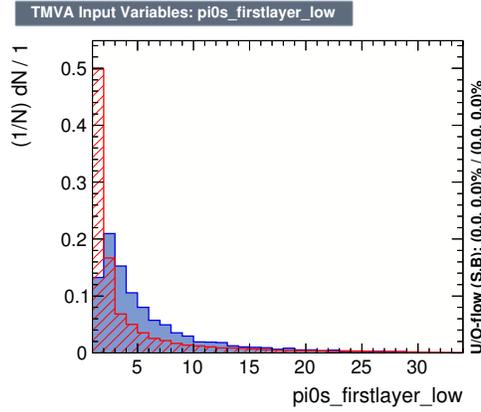
Figure A.1: Discriminating variables used in the training of the FGD1DsDs topology BDT.



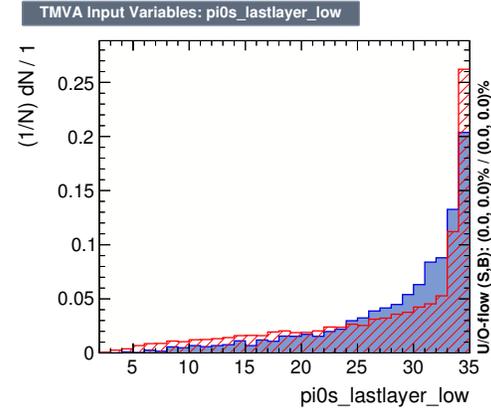
(g) E_h , EM energy.



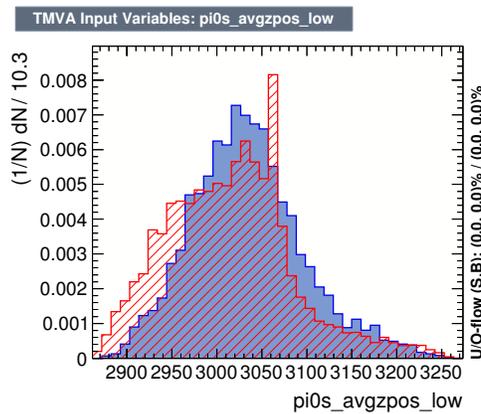
(h) E_h , EmHip.



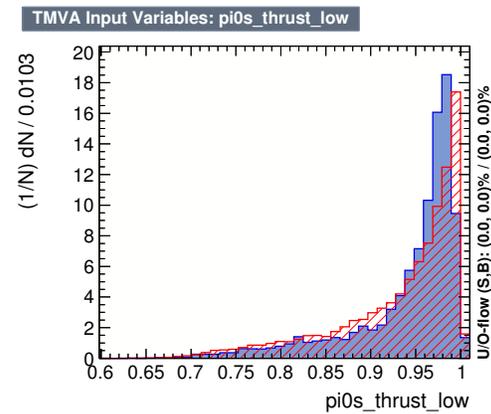
(i) E_l , first layer hit.



(j) E_l , last layer hit.



(k) E_l , average z position.



(l) E_l , thrust.

Figure A.1: Discriminating variables used in the training of the FGD1DsDs topology BDT.

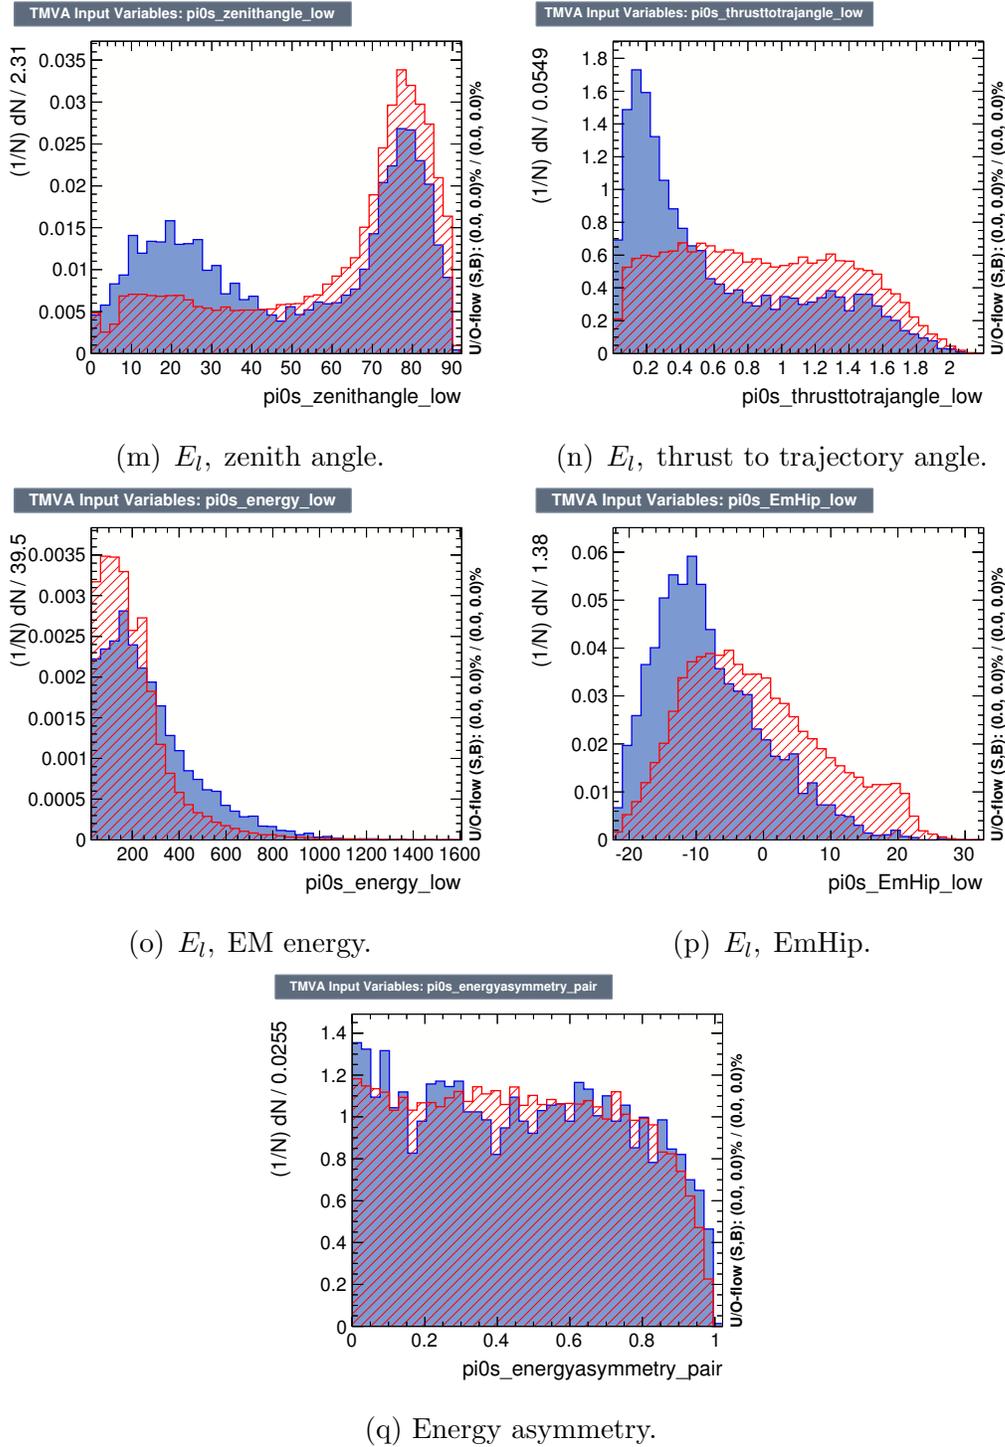
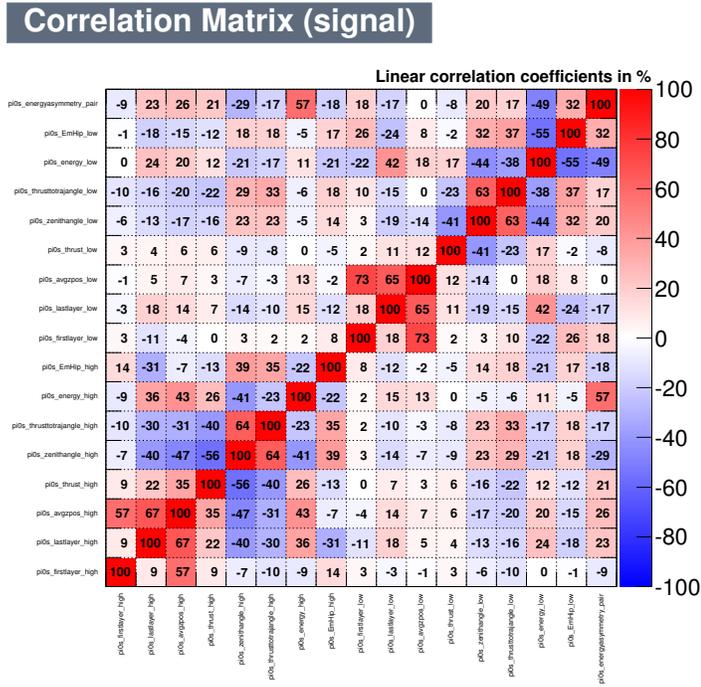
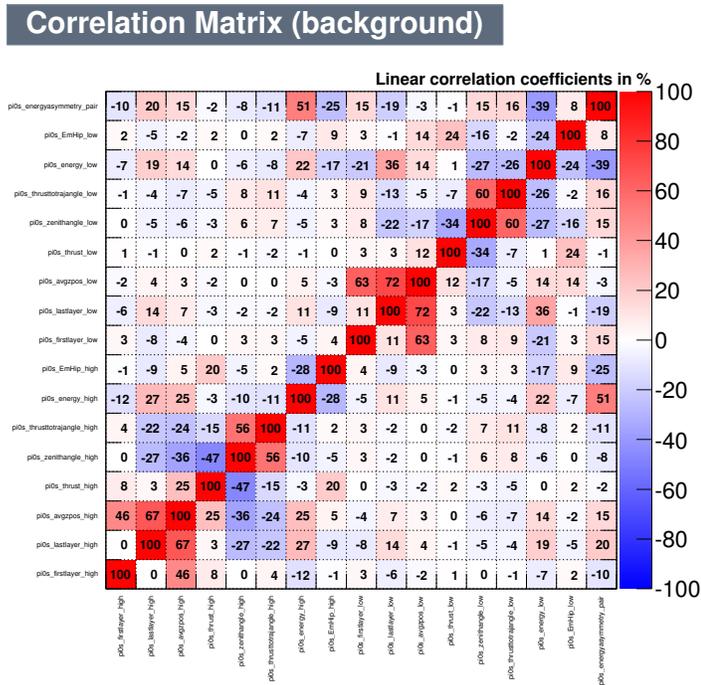


Figure A.1: Discriminating variables used in the training of the FGD1DsDs topology BDT.



(a) Signal.



(b) Background.

Figure A.2: Correlation between the signal, A.2a, and background, A.2b, discriminating variables used in the training of the FGD1DsDs topology BDT.

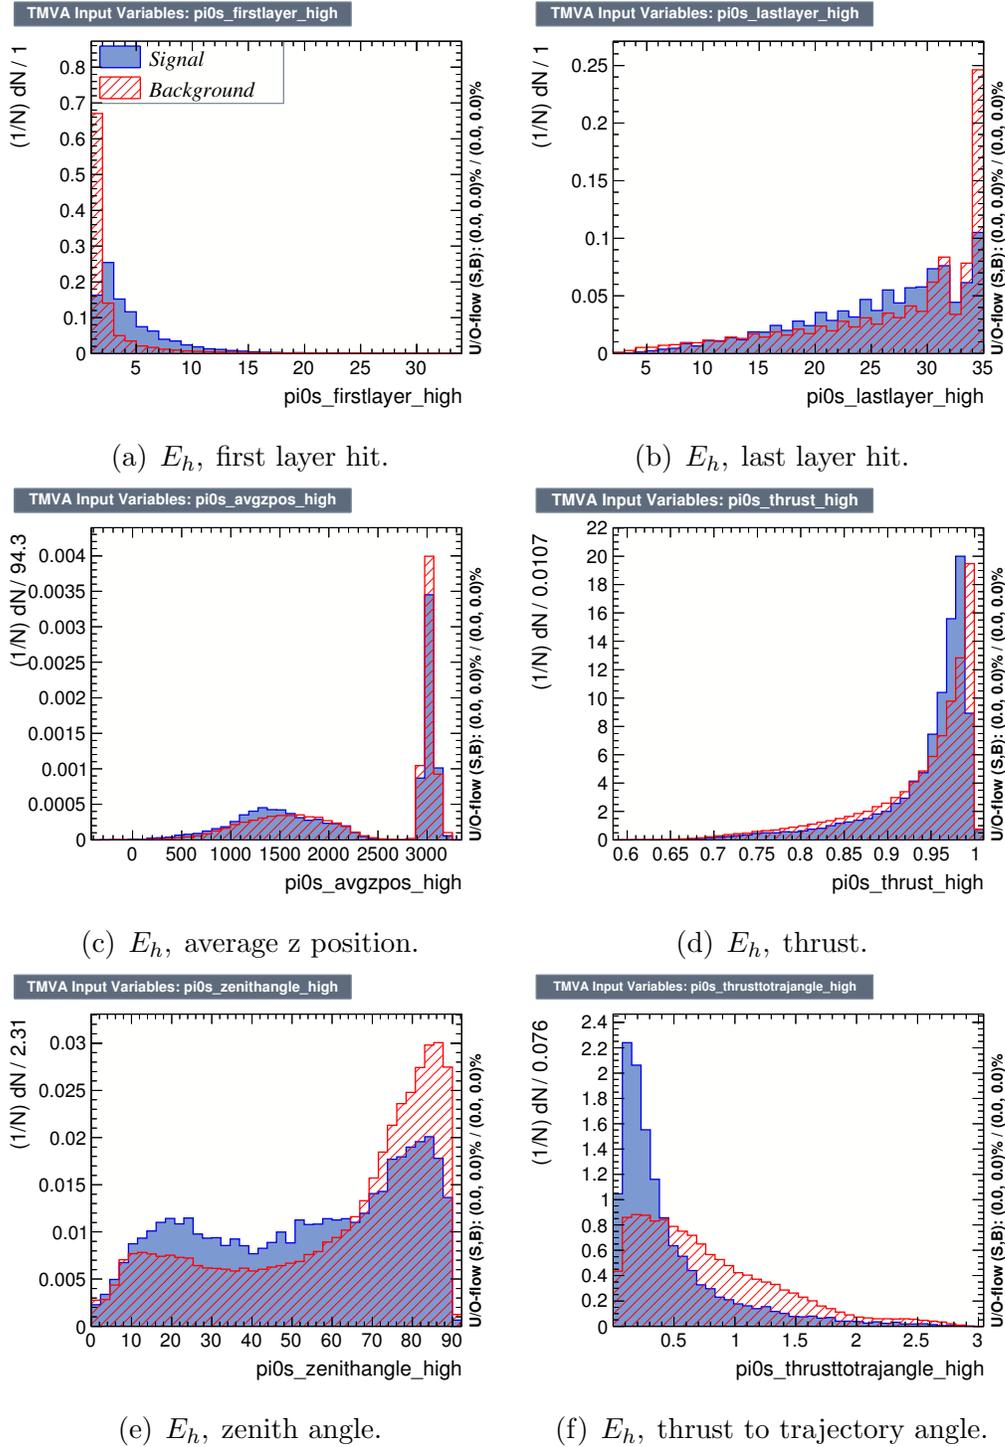
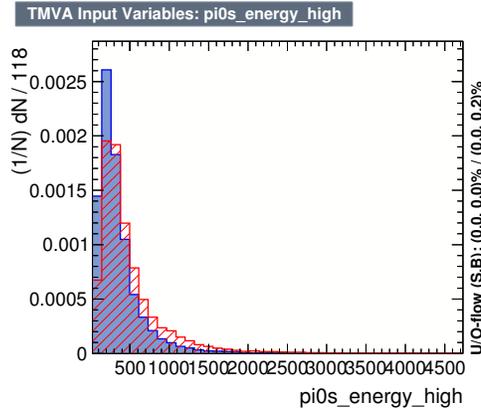
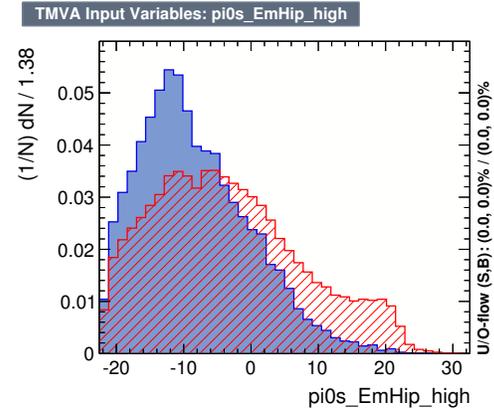


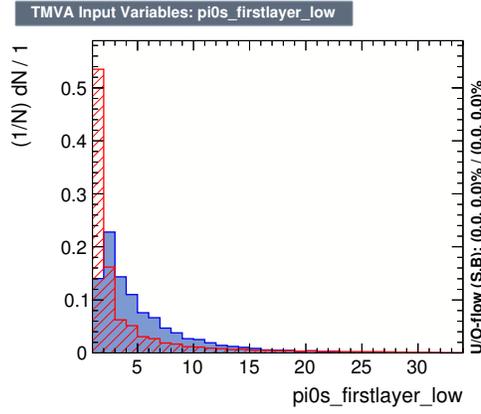
Figure A.3: Discriminating variables used in the training of the FGD1DsBr topology BDT.



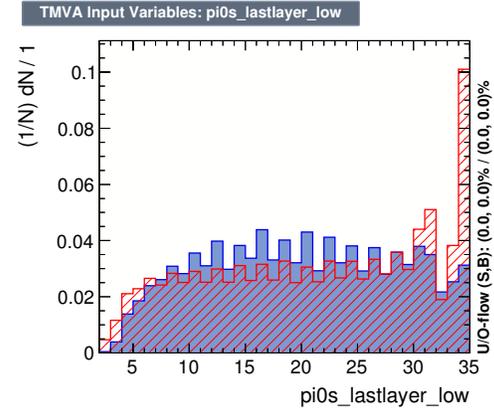
(g) E_h , EM energy.



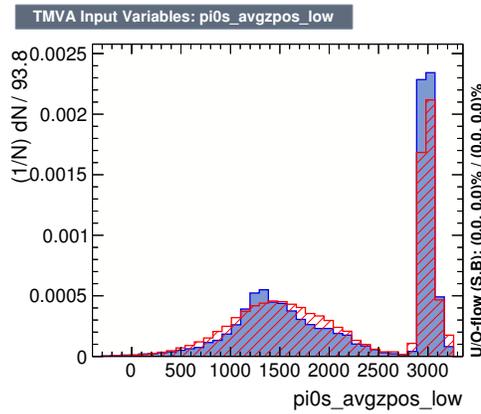
(h) E_h , EmHip.



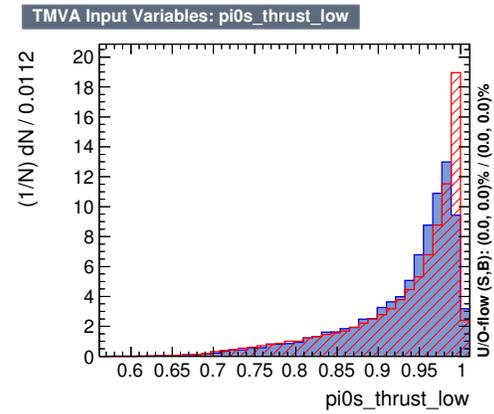
(i) E_l , first layer hit.



(j) E_l , last layer hit.



(k) E_l , average z position.



(l) E_l , thrust.

Figure A.3: Discriminating variables used in the training of the FGD1DsBr topology BDT.

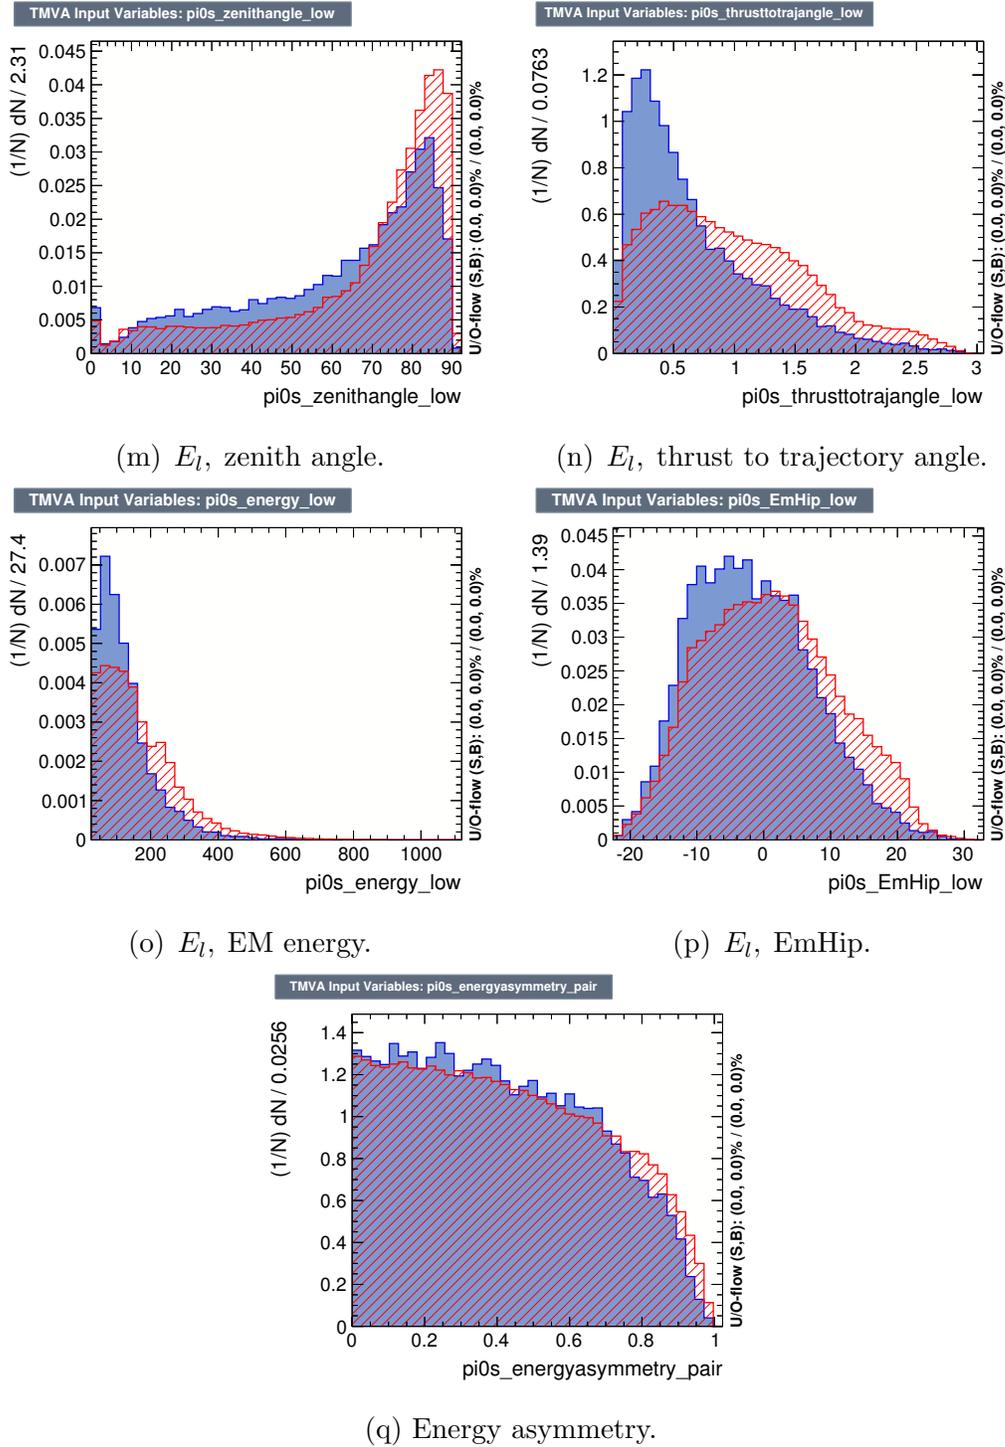
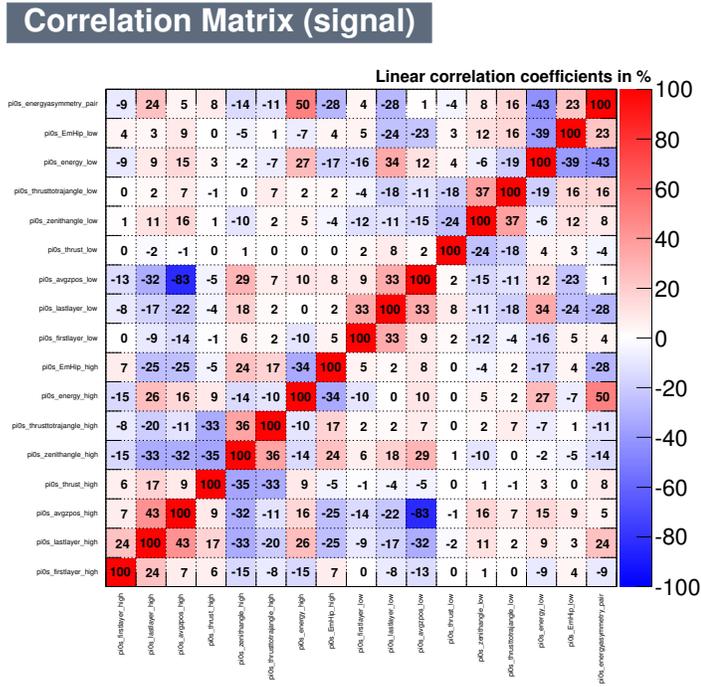
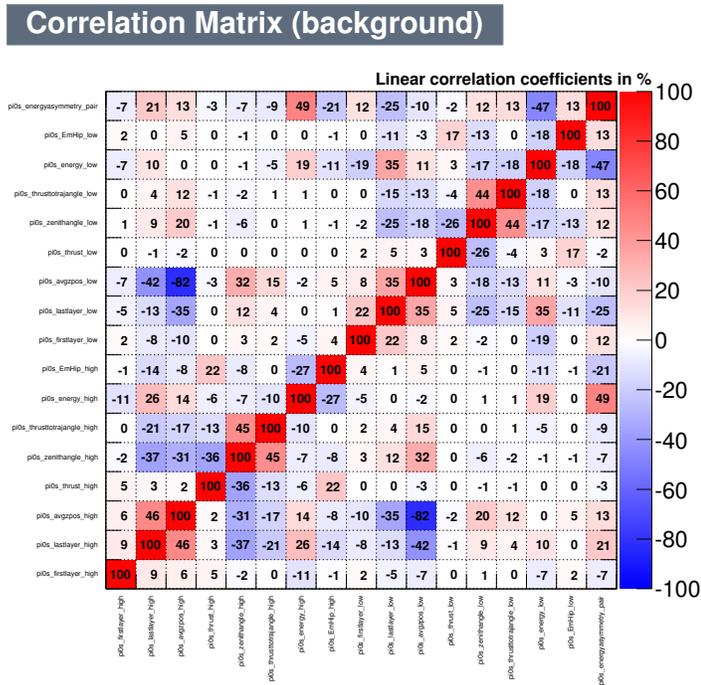


Figure A.3: Discriminating variables used in the training of the FGD1DsBr topology BDT.



(a) Signal.



(b) Background.

Figure A.4: Correlation between the signal, A.4a, and background, A.4b, discriminating variables used in the training of the FGD1DsBr topology BDT.

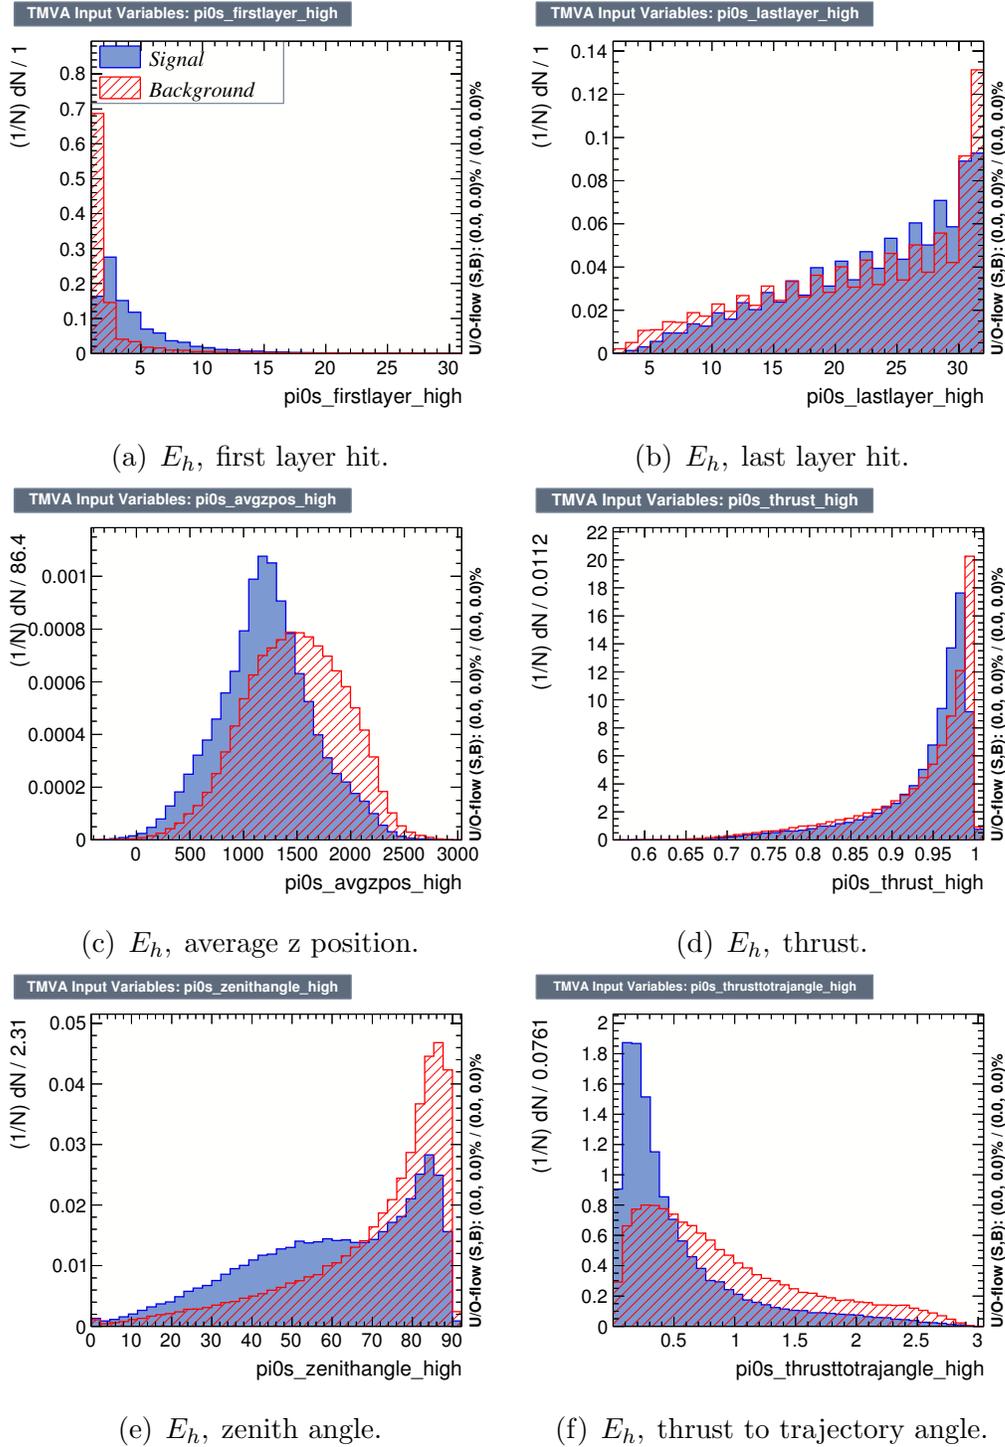
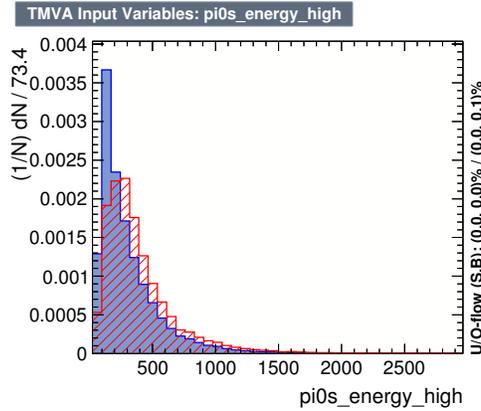
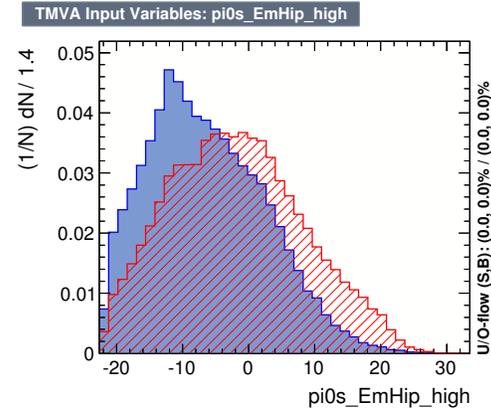


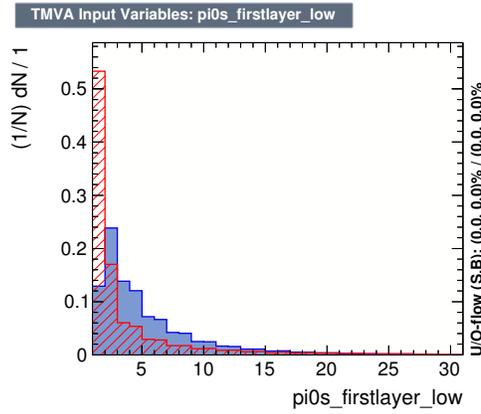
Figure A.5: Discriminating variables used in the training of the FGD1BrBr topology BDT.



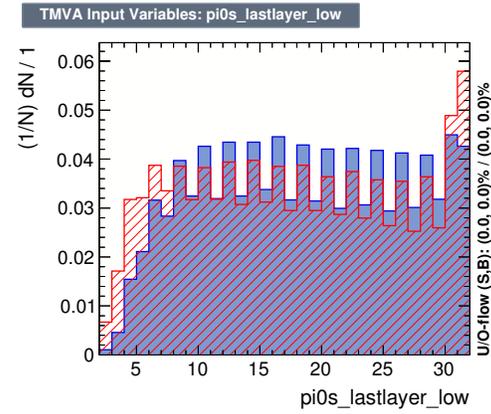
(g) E_h , EM energy.



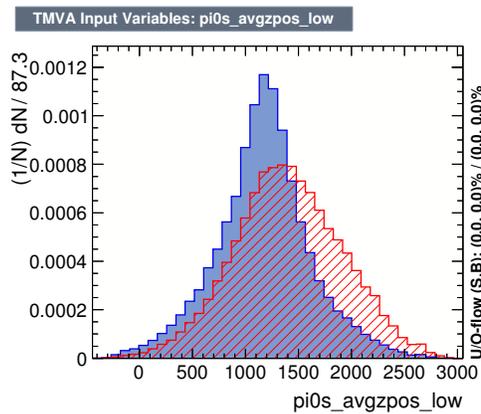
(h) E_h , EmHip.



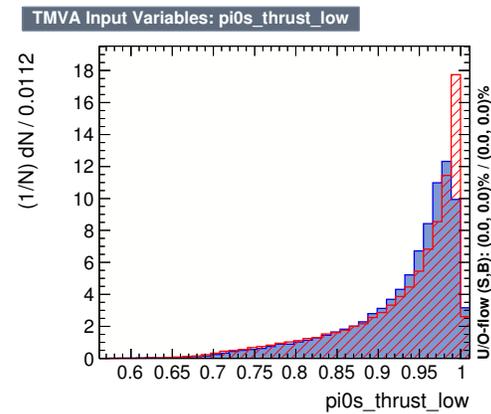
(i) E_l , first layer hit.



(j) E_l , last layer hit.



(k) E_l , average z position.



(l) E_l , thrust.

Figure A.5: Discriminating variables used in the training of the FGD1BrBr topology BDT.

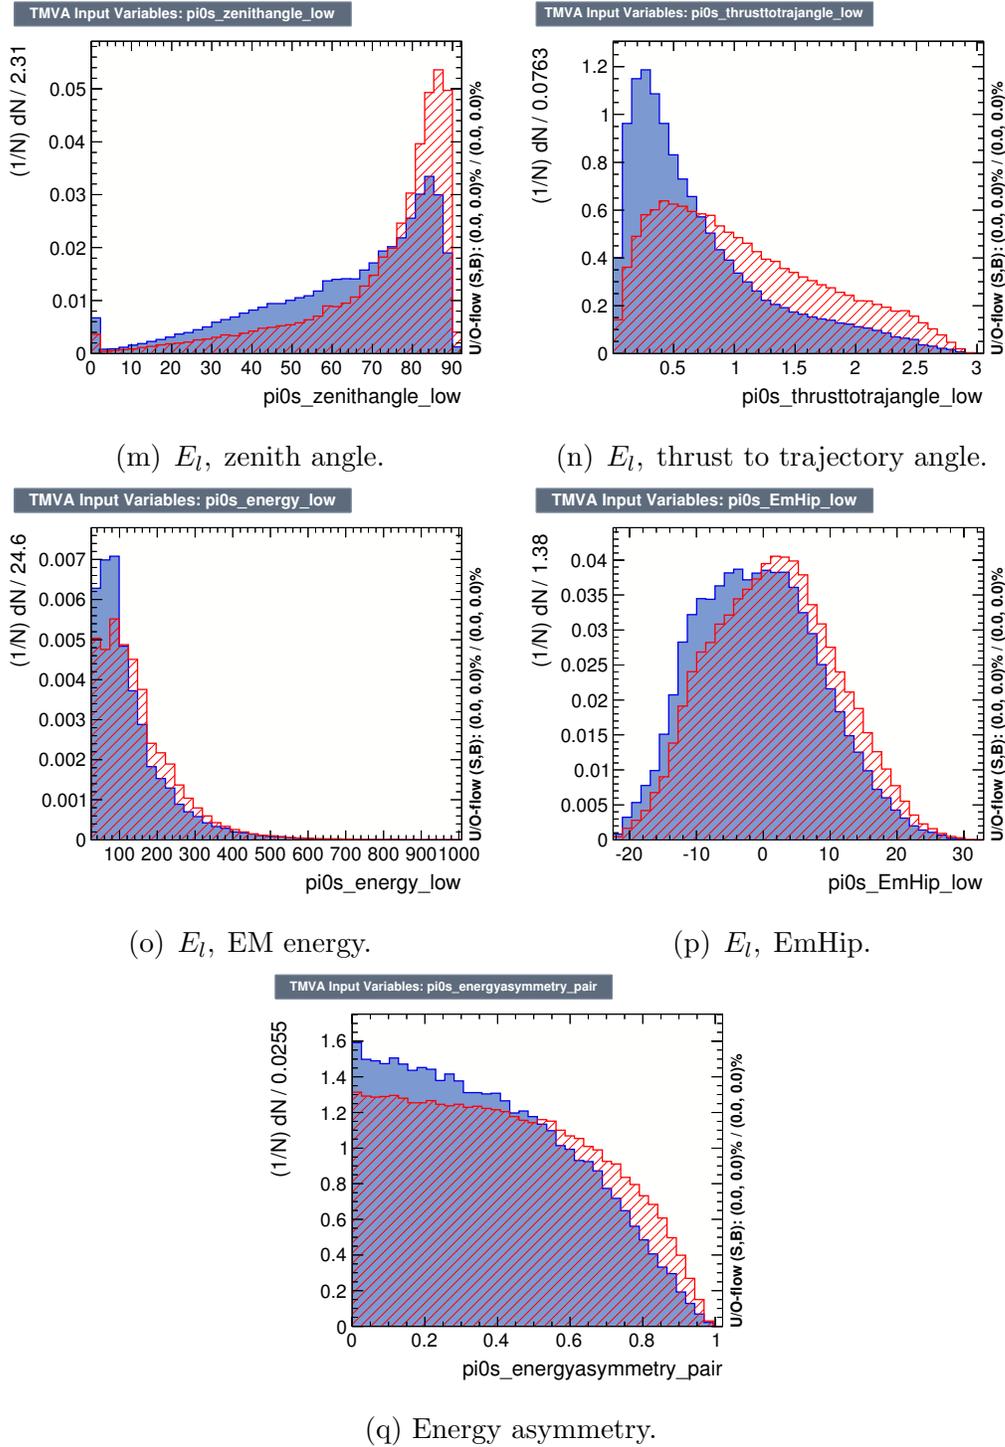
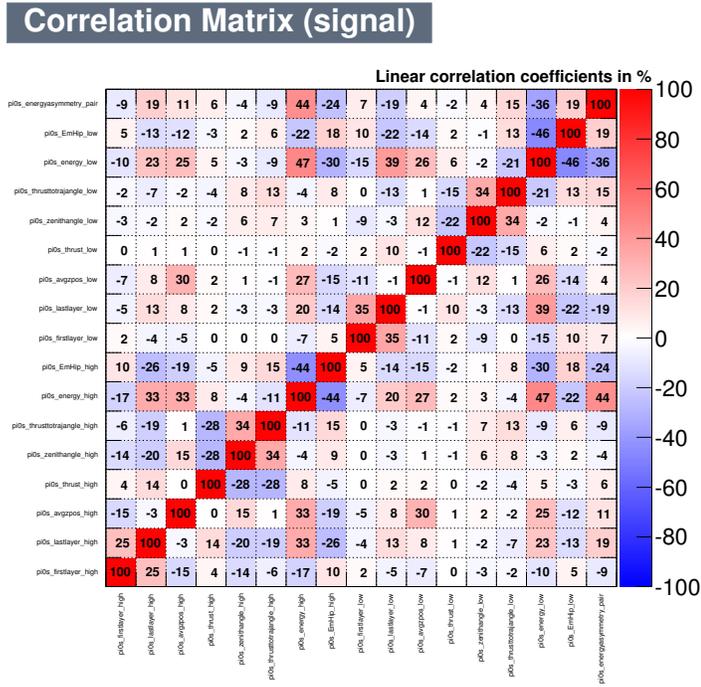
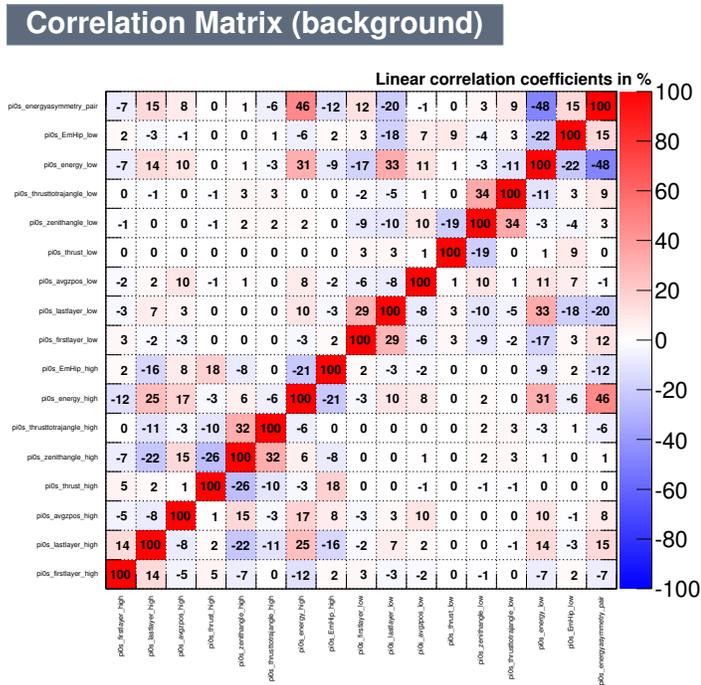


Figure A.5: Discriminating variables used in the training of the FGD1BrBr topology BDT.



(a) Signal.



(b) Background.

Figure A.6: Correlation between the signal, A.6a, and background, A.6b, discriminating variables used in the training of the FGD1BrBr topology BDT.

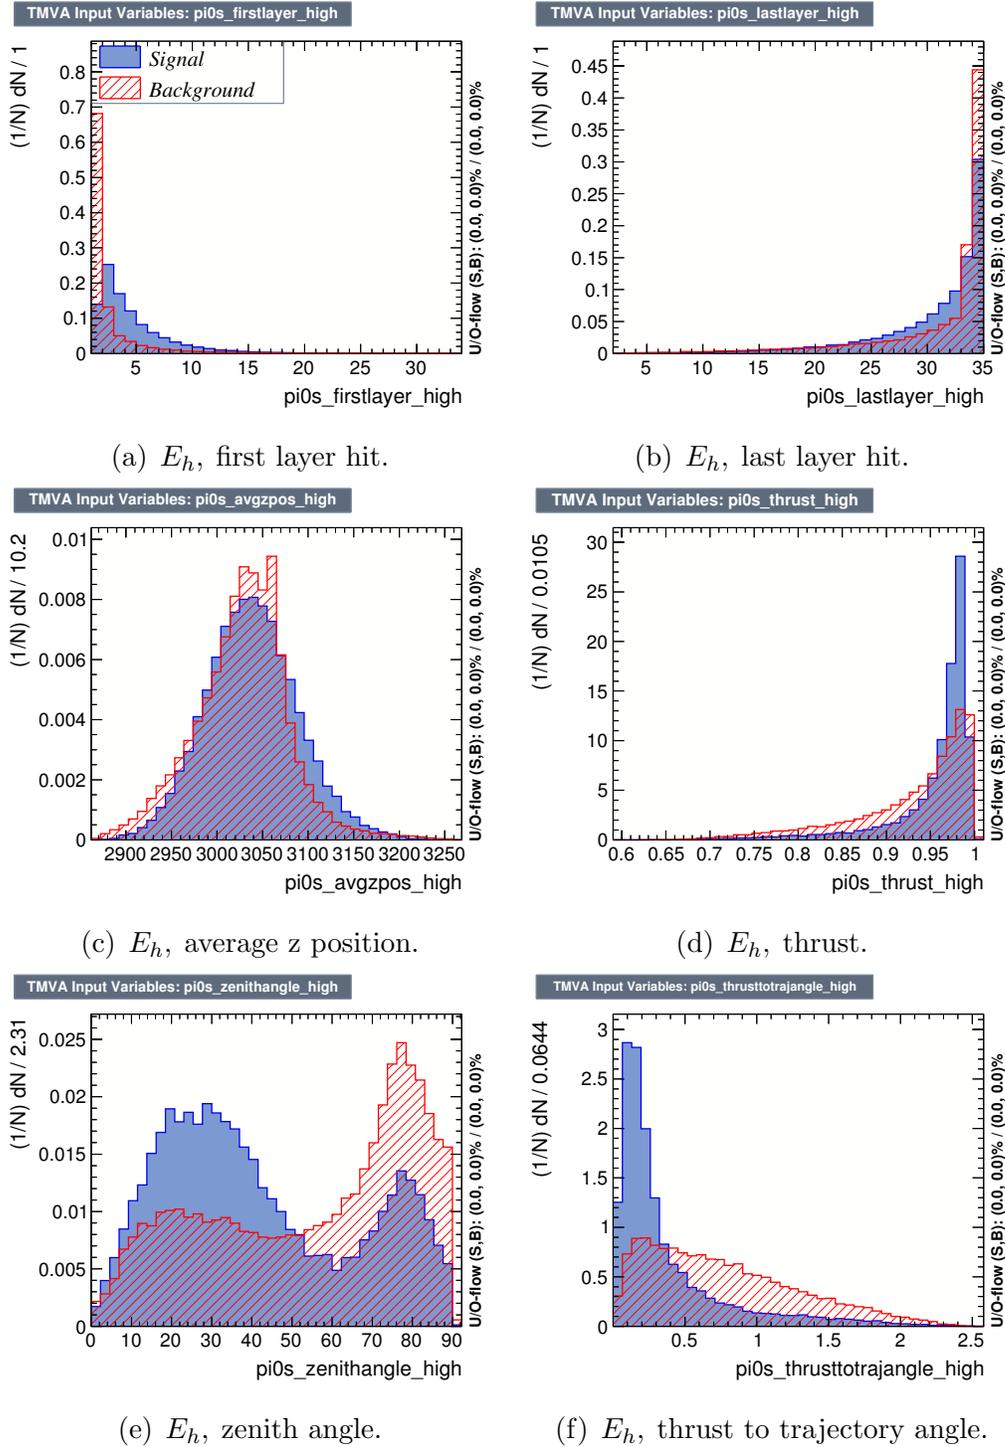


Figure A.7: Discriminating variables used in the training of the FGD2DsDs topology BDT.

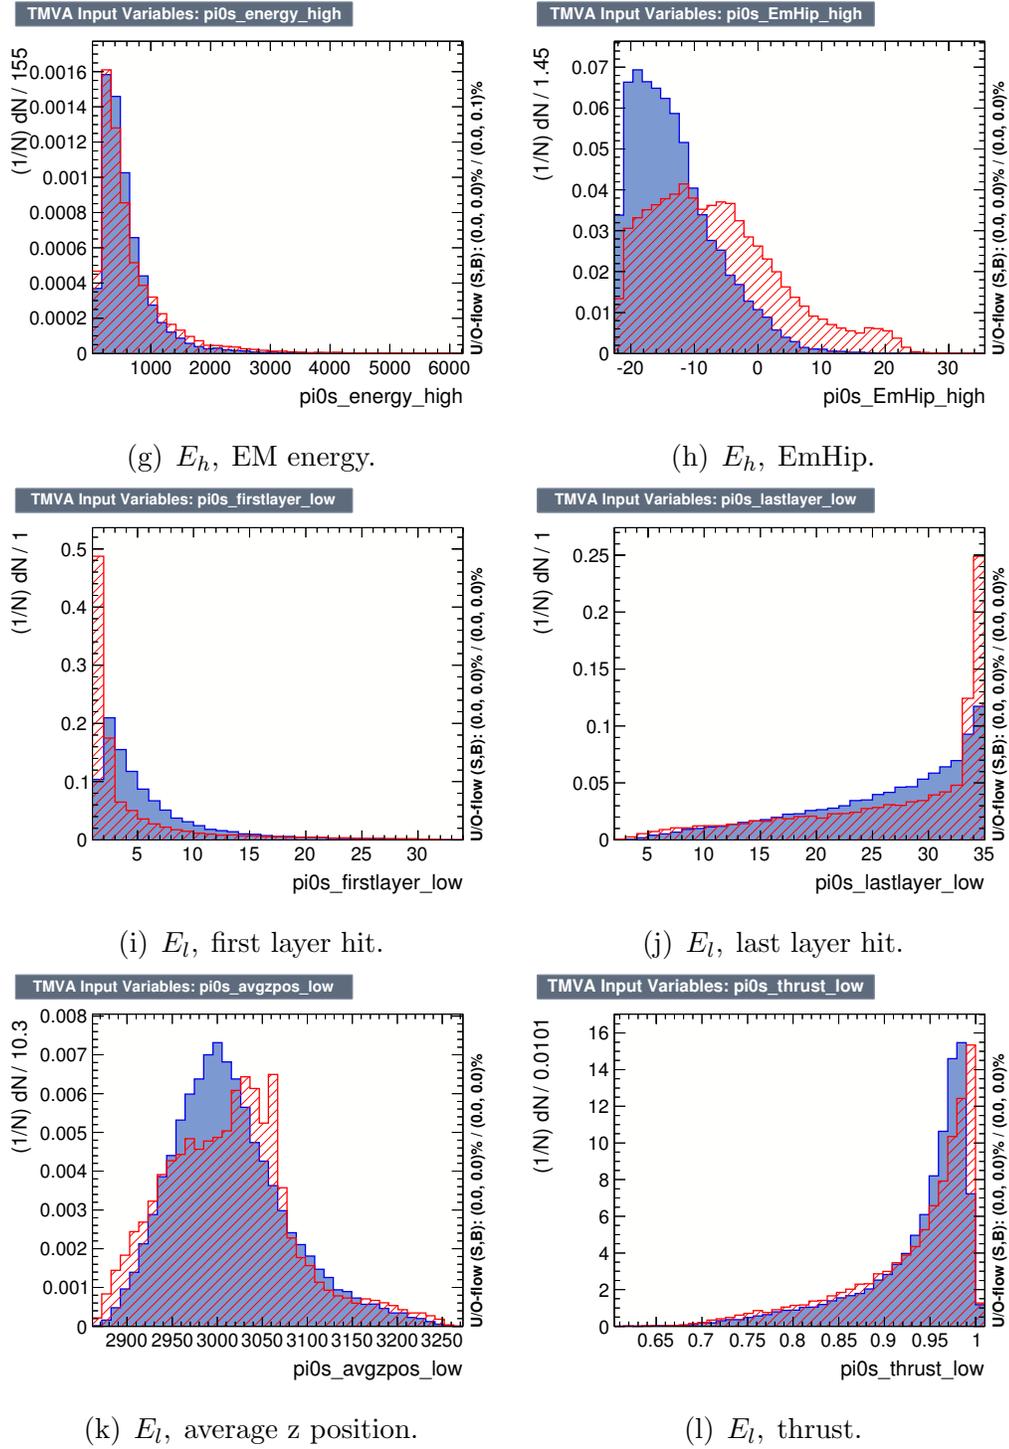


Figure A.7: Discriminating variables used in the training of the FGD2DsDs topology BDT.

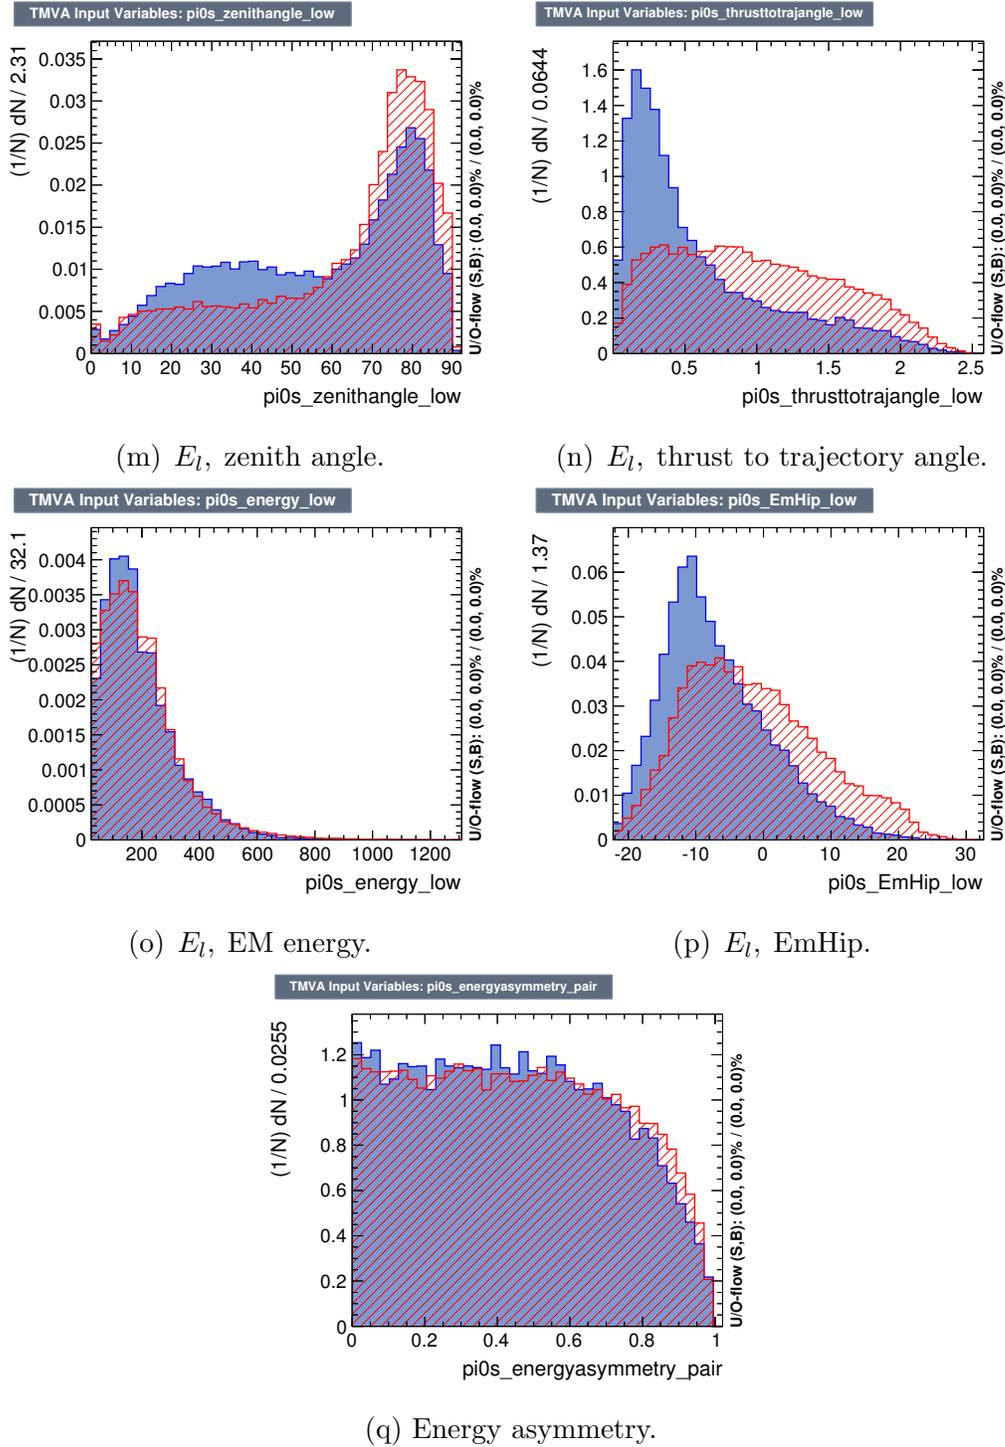
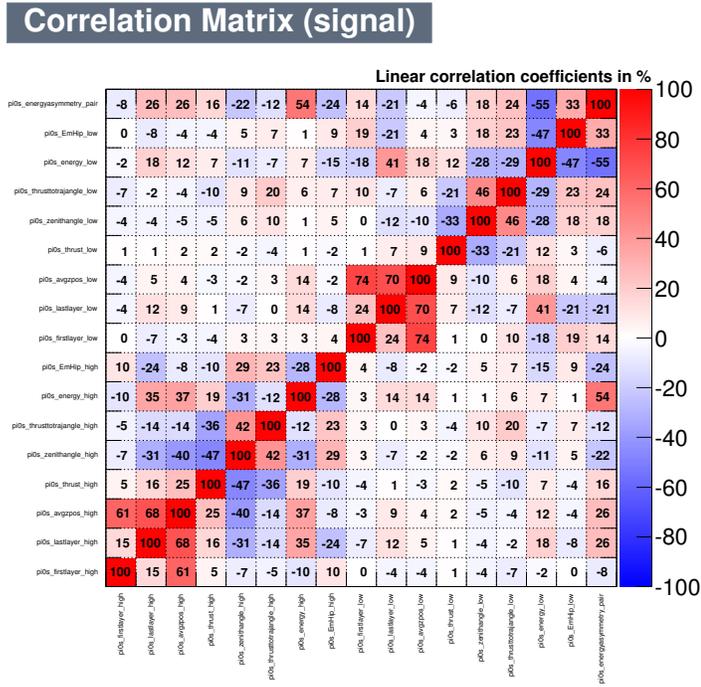
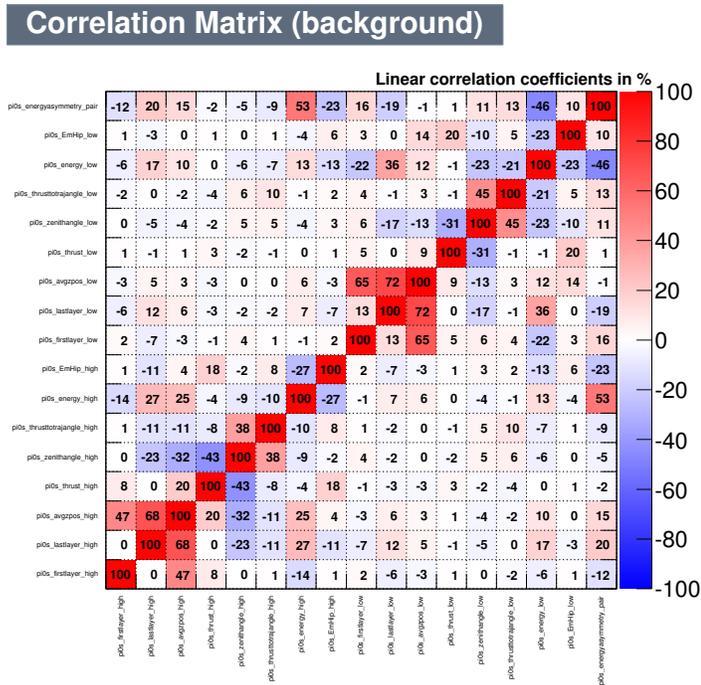


Figure A.7: Discriminating variables used in the training of the FGD2DsDs topology BDT.



(a) Signal.



(b) Background.

Figure A.8: Correlation between the signal, A.8a, and background, A.8b, discriminating variables used in the training of the FGD2Ds topology BDT.

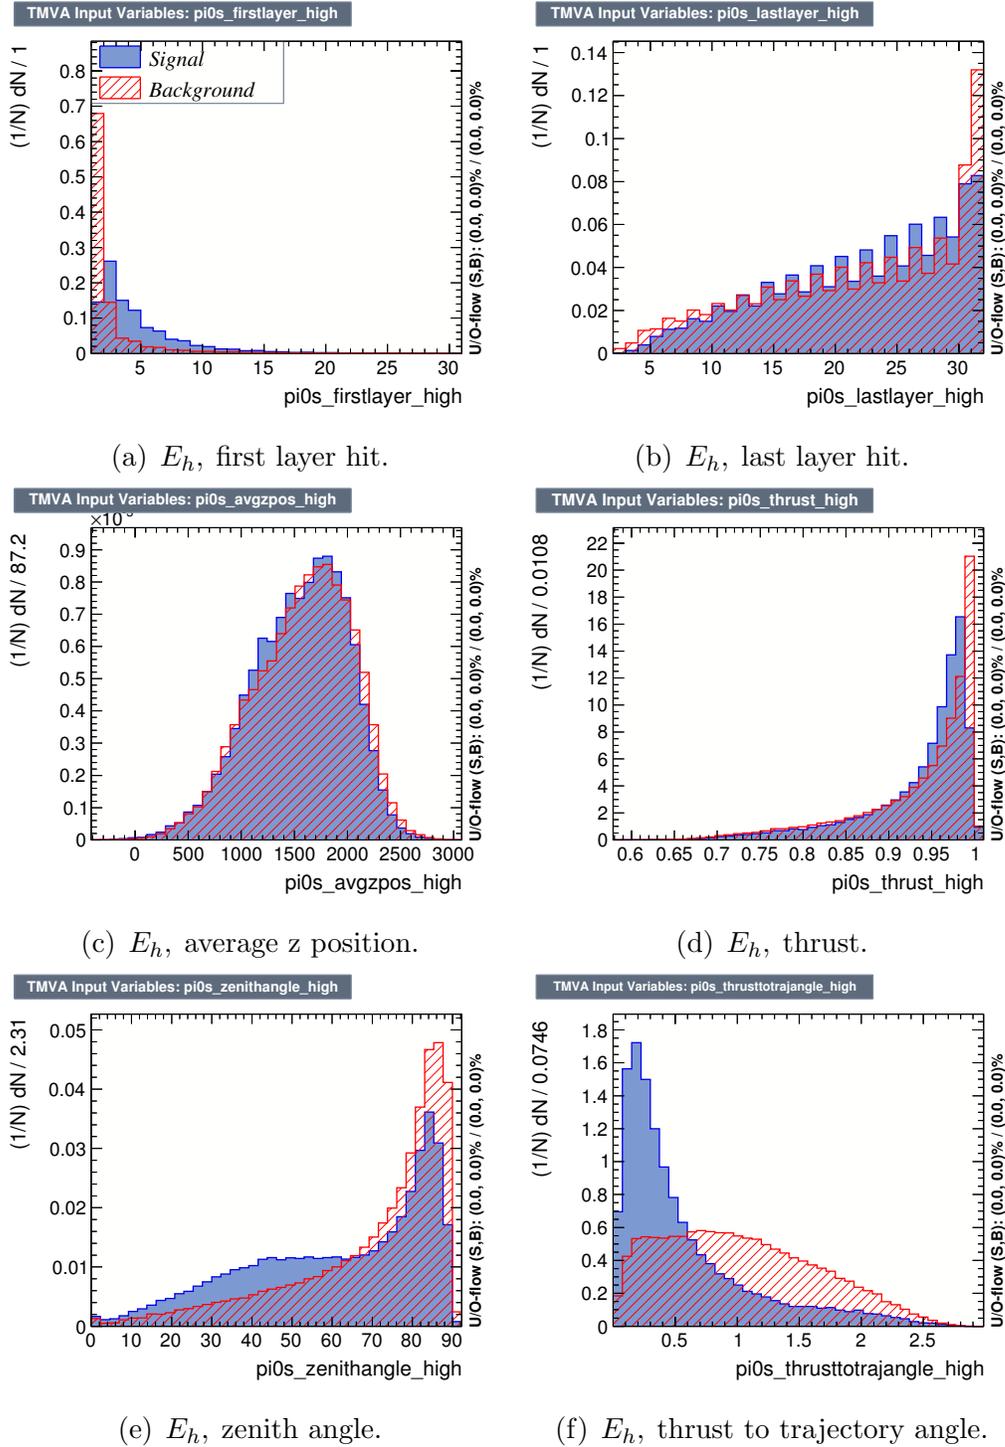


Figure A.9: Discriminating variables used in the training of the FGD2BrBr topology BDT.

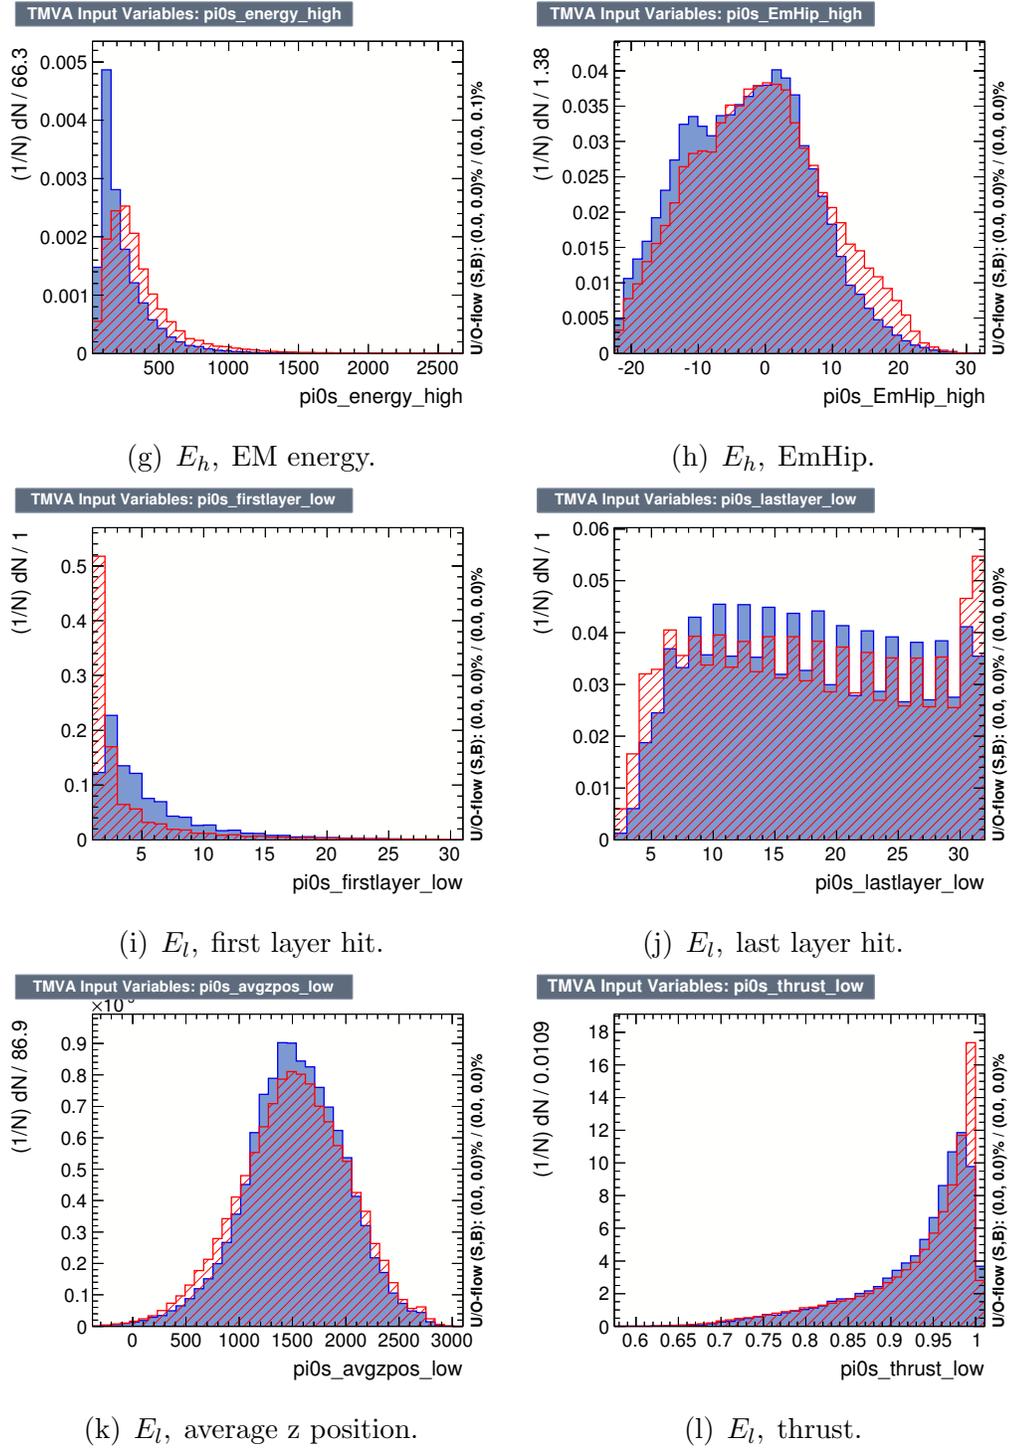


Figure A.9: Discriminating variables used in the training of the FGD2BrBr topology BDT.

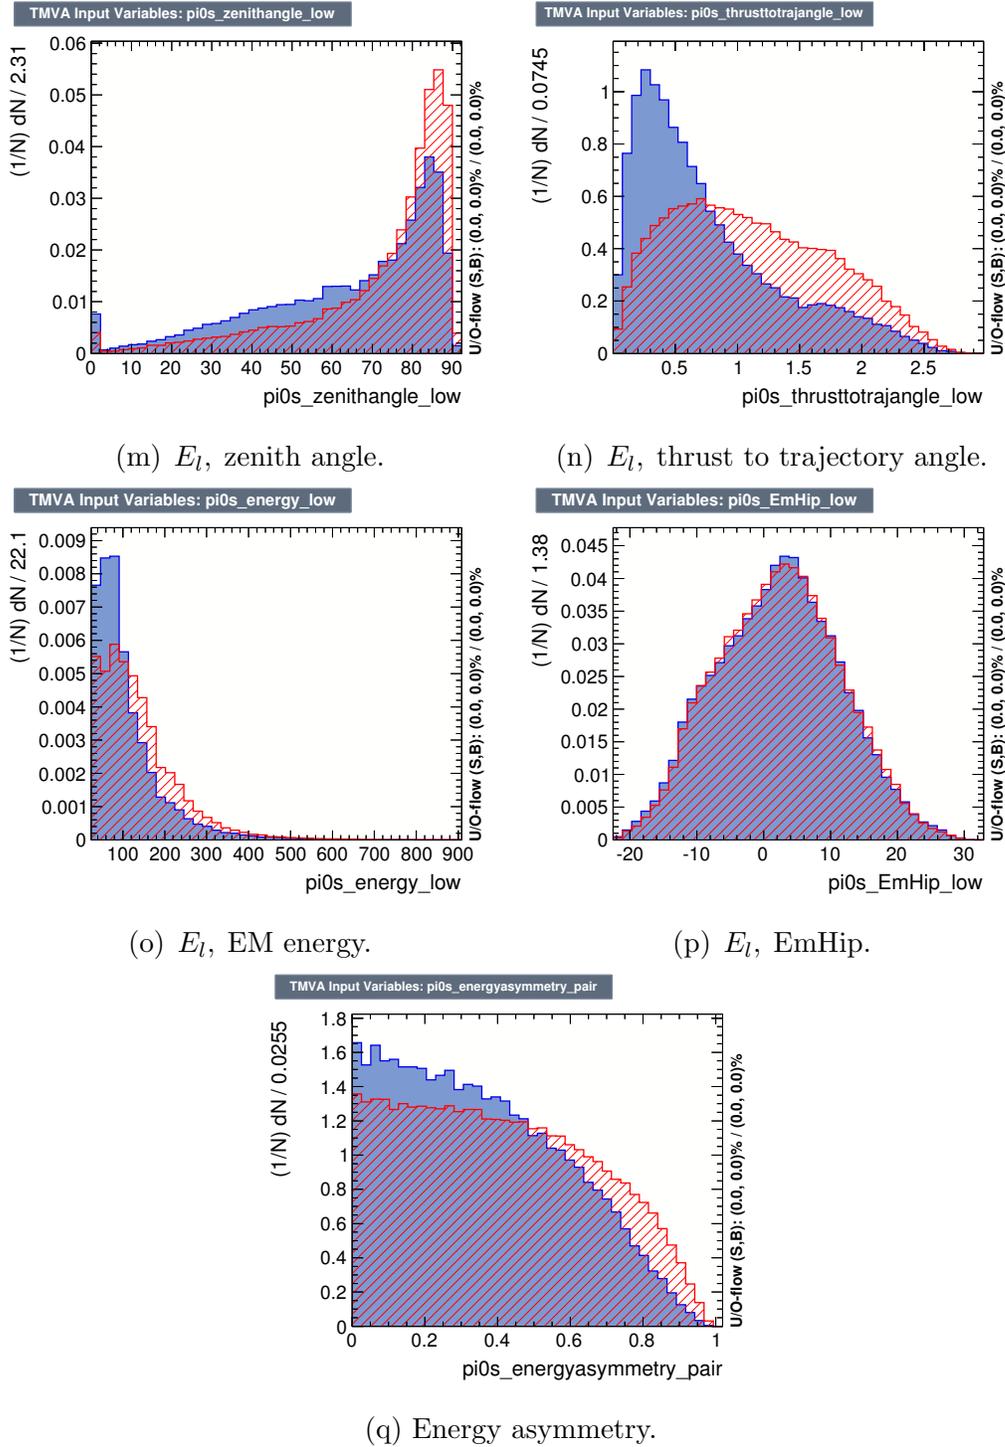


Figure A.9: Discriminating variables used in the training of the FGD2BrBr topology BDT.

Appendix B

Data and Monte-Carlo Comparisons for BDT Discriminating Variables

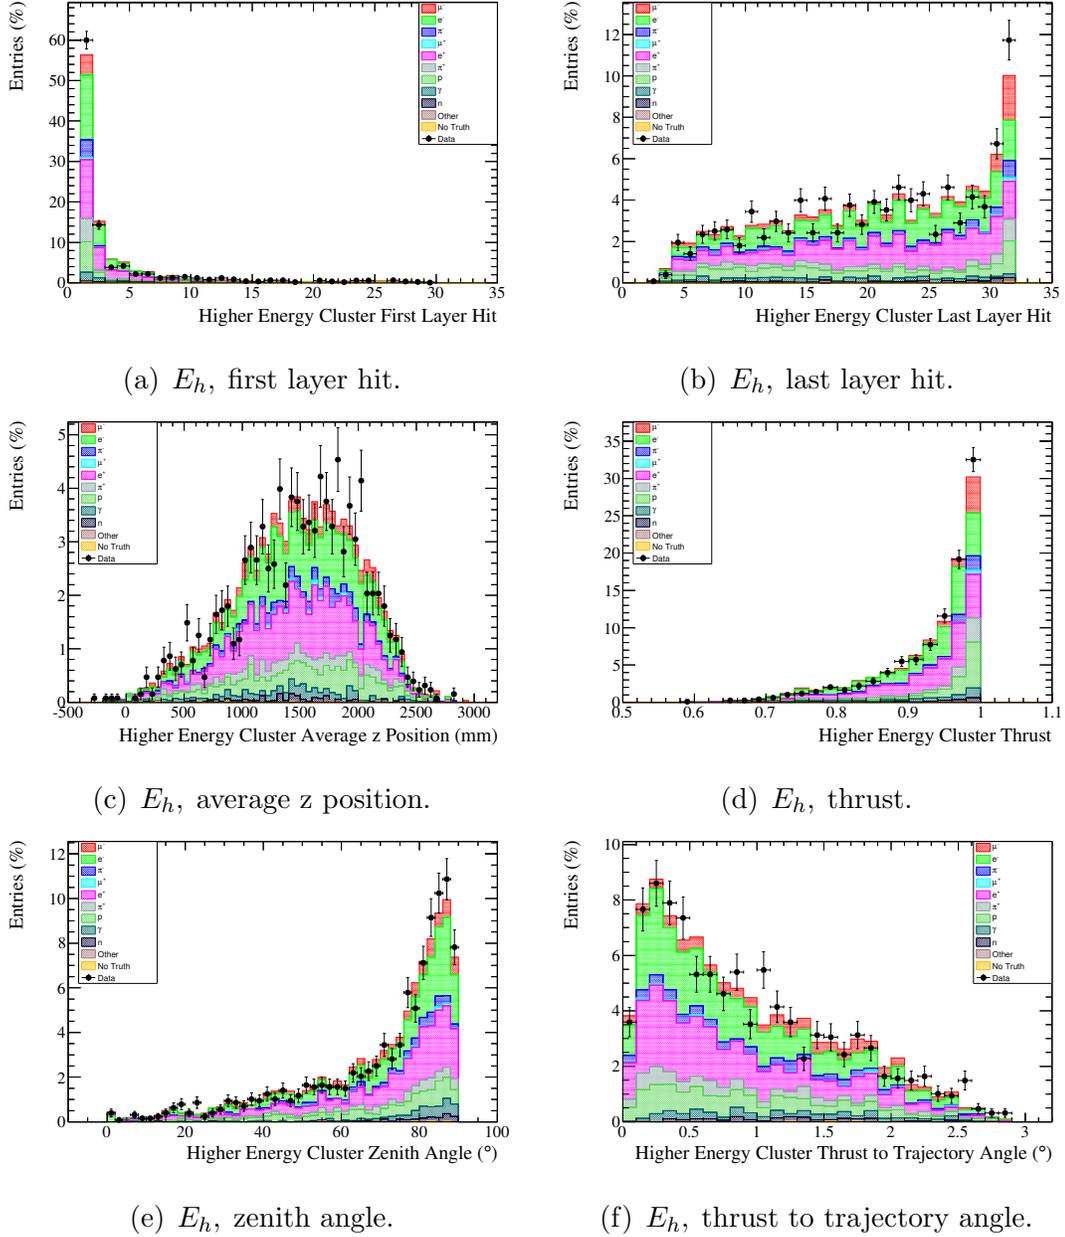


Figure B.1: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the left side BrECal.

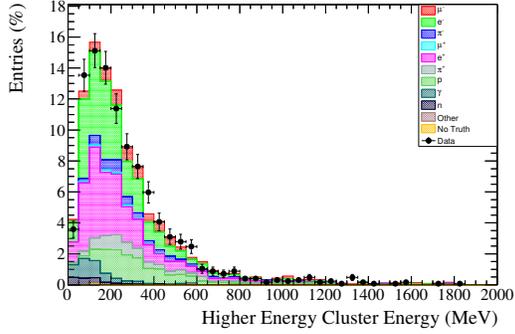
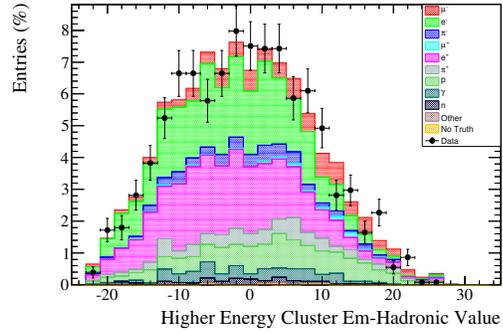
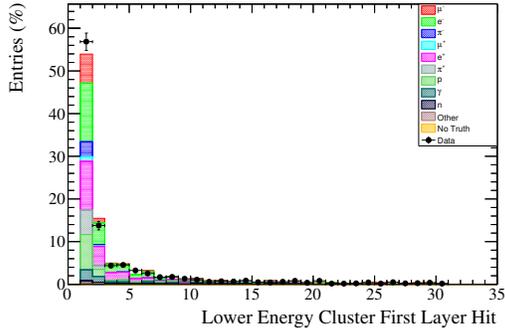
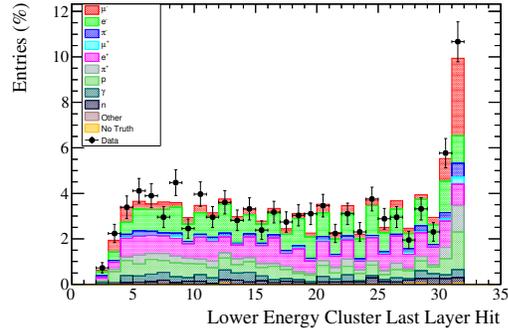
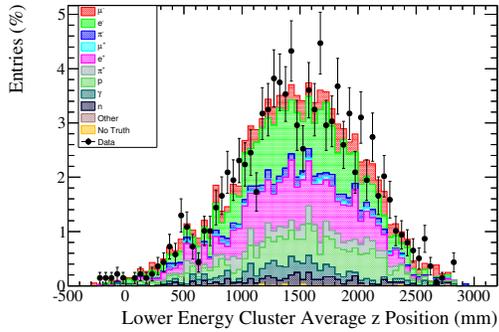
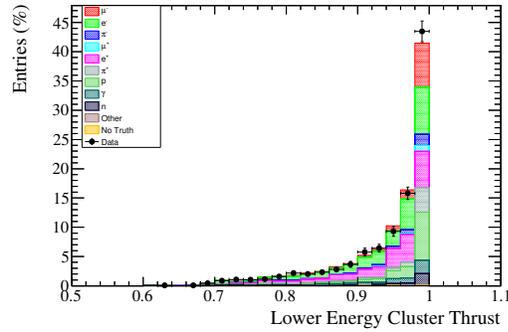
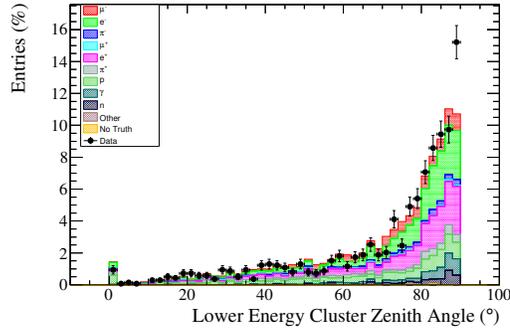
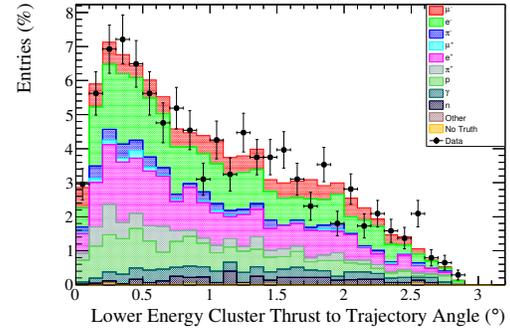
(g) E_h , EM energy.(h) E_h , EmHip.(i) E_l , first layer hit.(j) E_l , last layer hit.(k) E_l , average z position.(l) E_l , thrust.

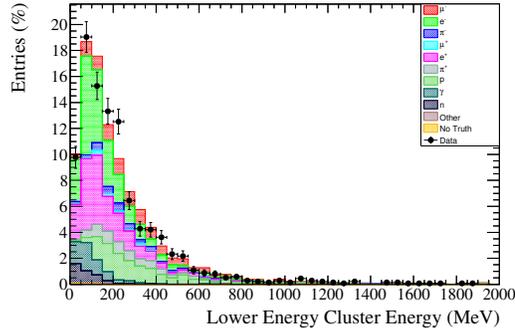
Figure B.1: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the left side BrECal.



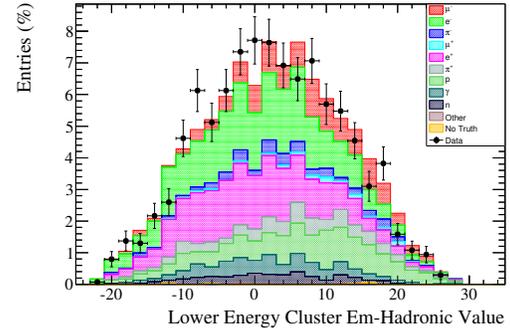
(m) E_l , zenith angle.



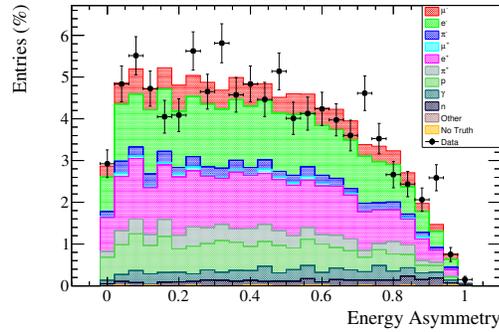
(n) E_l , thrust to trajectory angle.



(o) E_l , EM energy.



(p) E_l , EmHip.



(q) Energy asymmetry.

Figure B.1: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the left side BrECal.

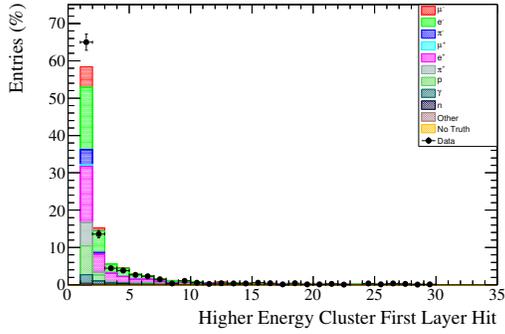
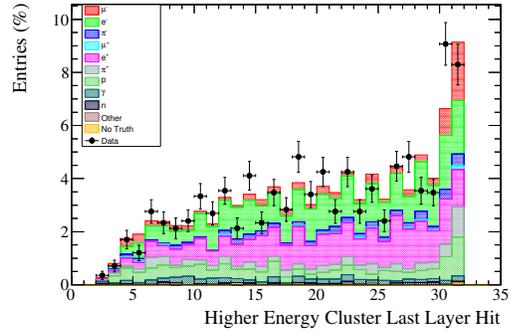
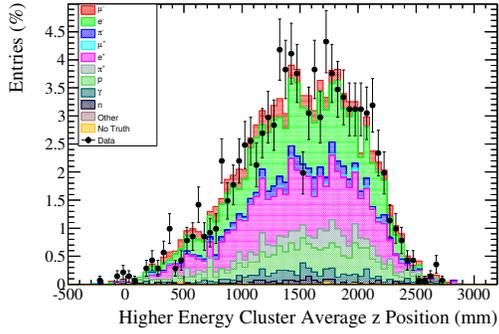
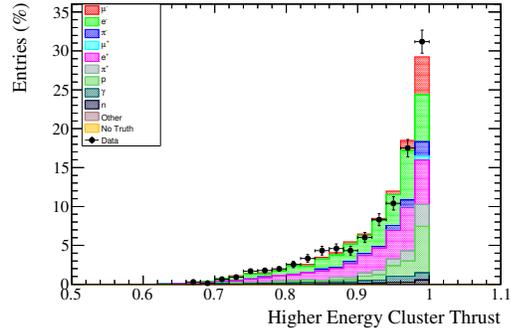
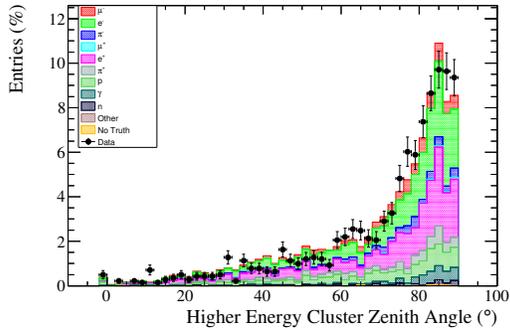
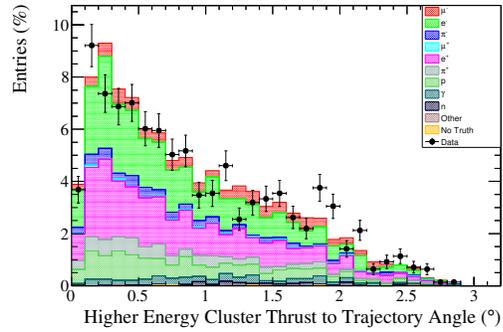
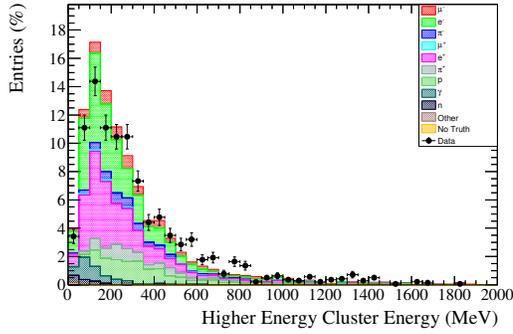
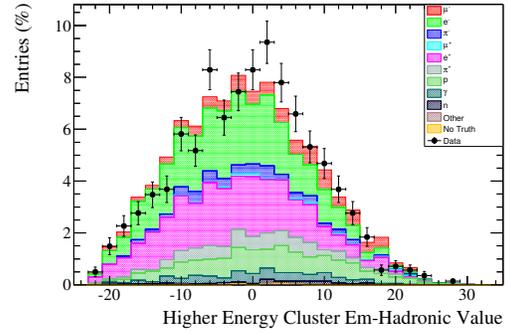
(a) E_h , first layer hit.(b) E_h , last layer hit.(c) E_h , average z position.(d) E_h , thrust.(e) E_h , zenith angle.(f) E_h , thrust to trajectory angle.

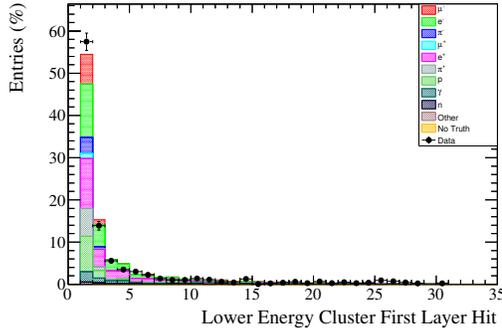
Figure B.2: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the right side BrECal.



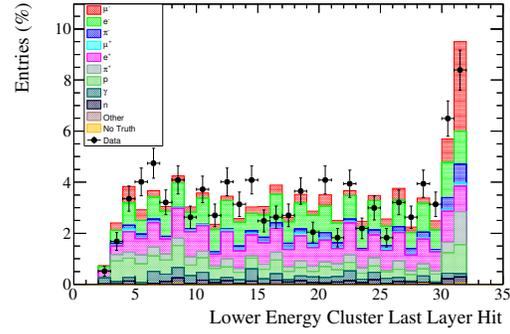
(g) E_h , EM energy.



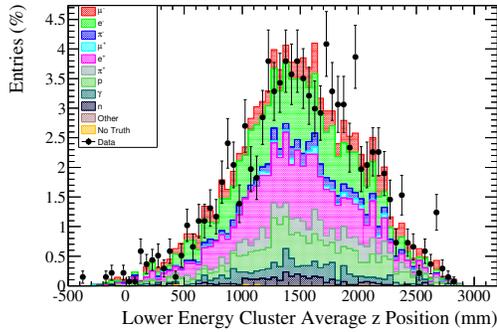
(h) E_h , EmHip.



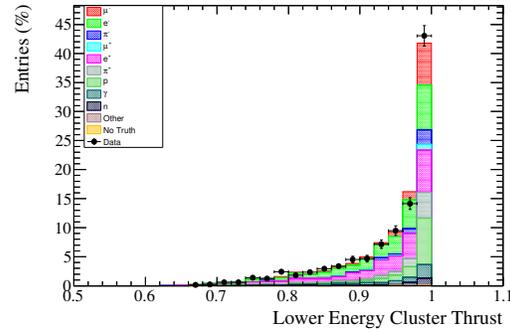
(i) E_l , first layer hit.



(j) E_l , last layer hit.

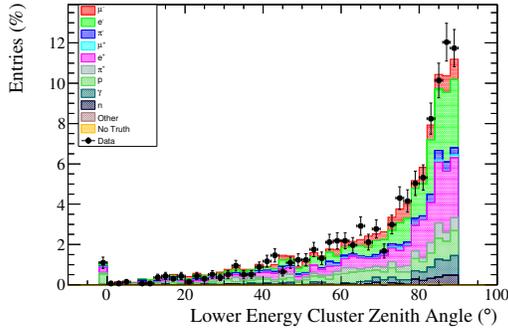
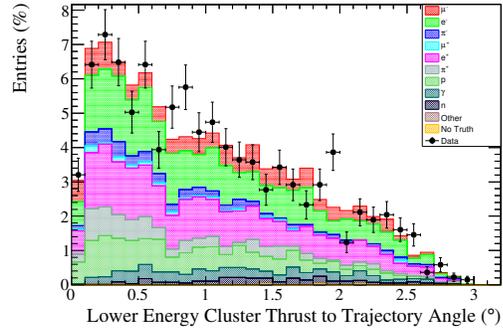
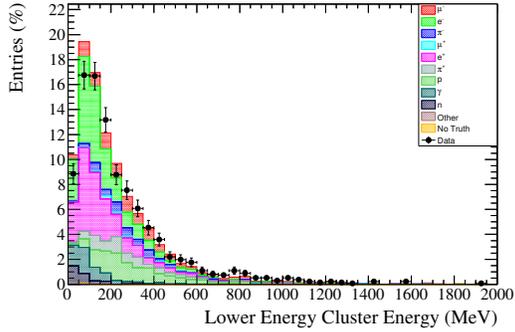
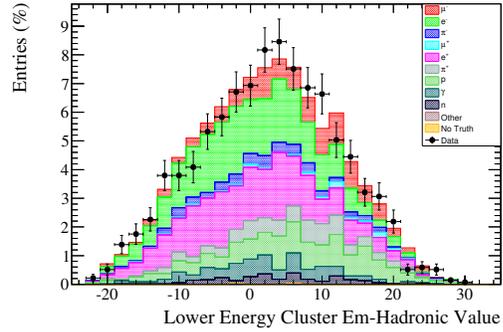
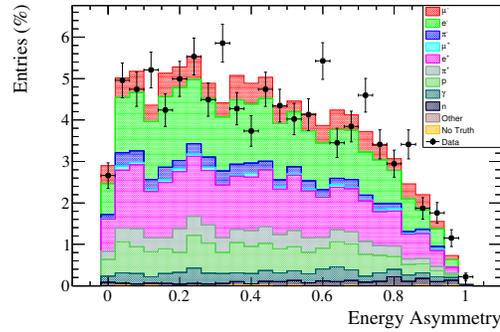


(k) E_l , average z position.



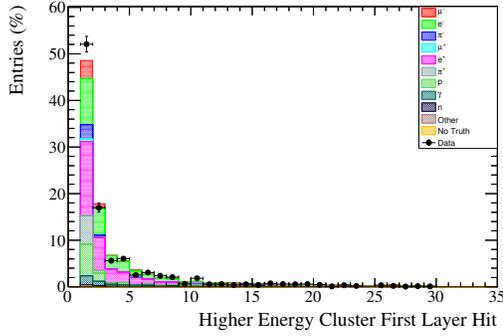
(l) E_l , thrust.

Figure B.2: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the right side BrECal.

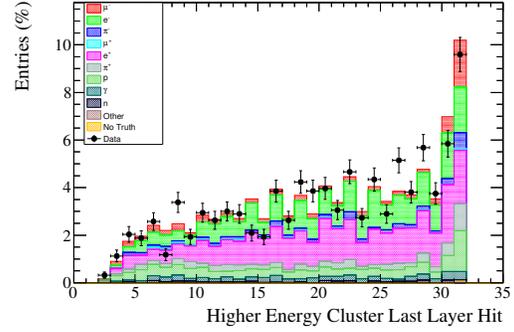
(m) E_l , zenith angle.(n) E_l , thrust to trajectory angle.(o) E_l , EM energy.(p) E_l , EmHip.

(q) Energy asymmetry.

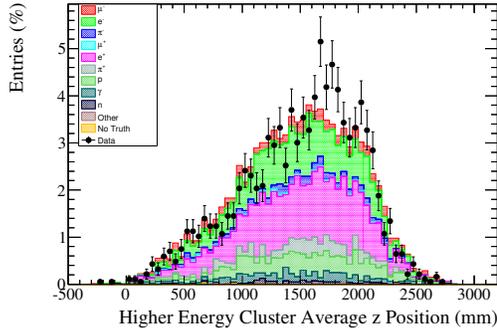
Figure B.2: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the right side BrECal.



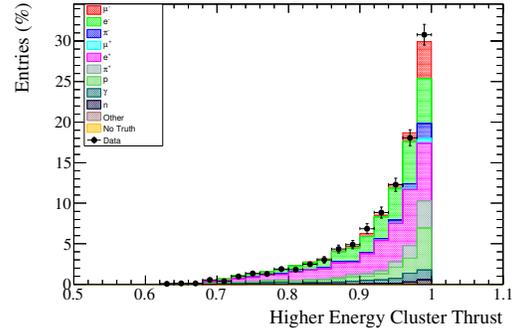
(a) E_h , first layer hit.



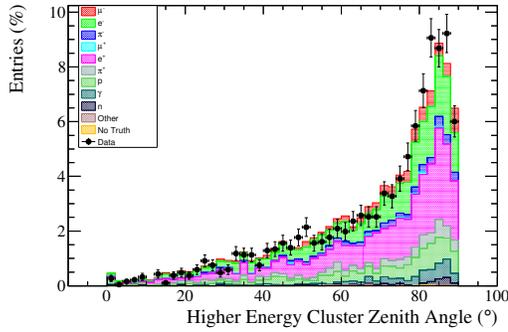
(b) E_h , last layer hit.



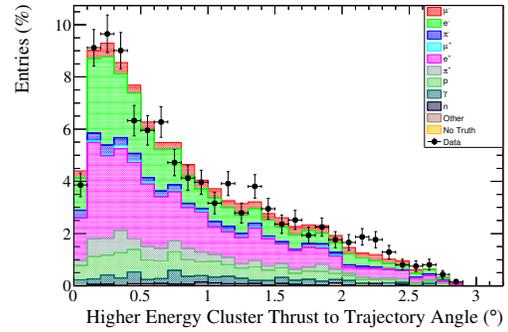
(c) E_h , average z position.



(d) E_h , thrust.



(e) E_h , zenith angle.



(f) E_h , thrust to trajectory angle.

Figure B.3: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the top BrECal.

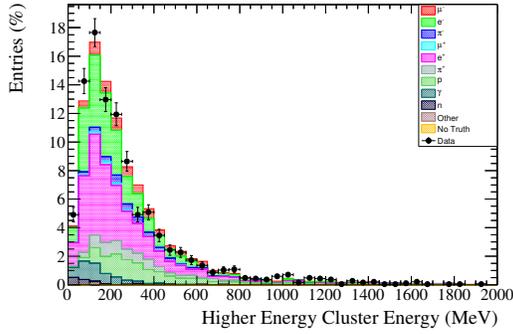
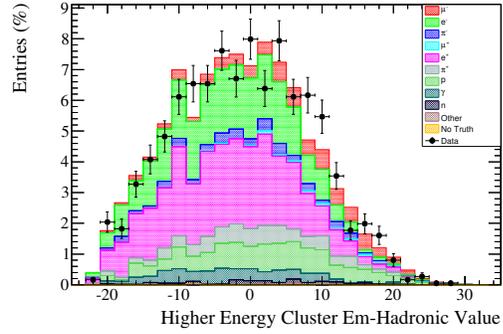
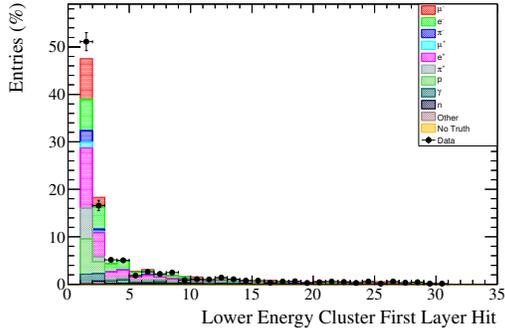
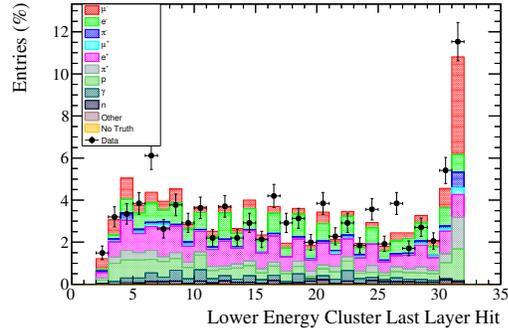
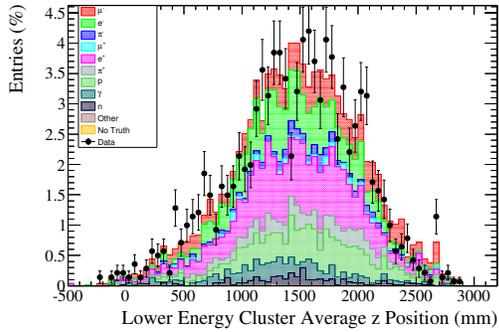
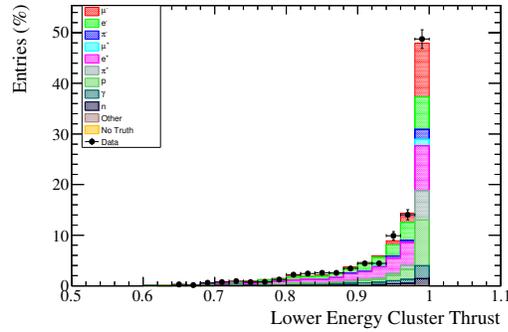
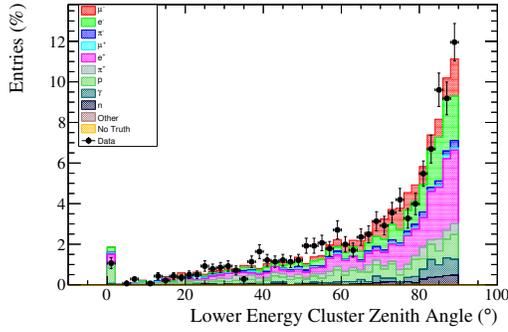
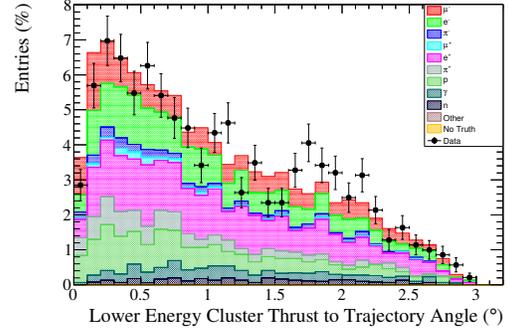
(g) E_h , EM energy.(h) E_h , EmHip.(i) E_l , first layer hit.(j) E_l , last layer hit.(k) E_l , average z position.(l) E_l , thrust.

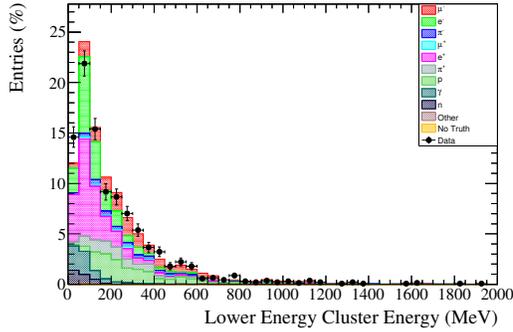
Figure B.3: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the top BrECal.



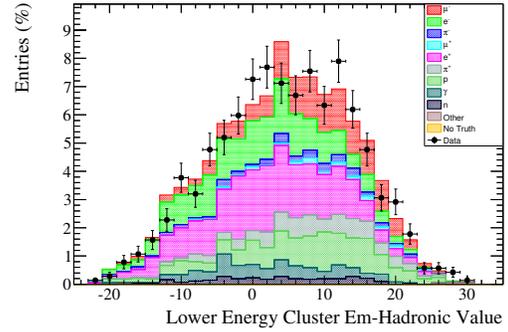
(m) E_l , zenith angle.



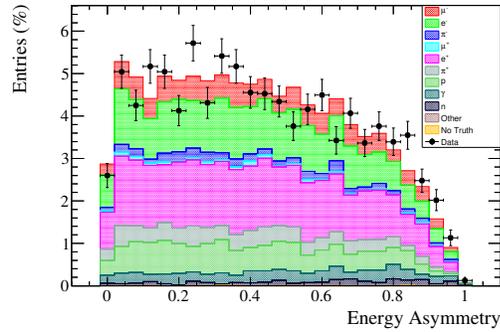
(n) E_l , thrust to trajectory angle.



(o) E_l , EM energy.



(p) E_l , EmHip.



(q) Energy asymmetry.

Figure B.3: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the top BrECal.

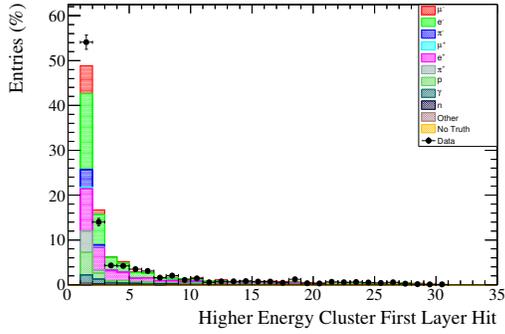
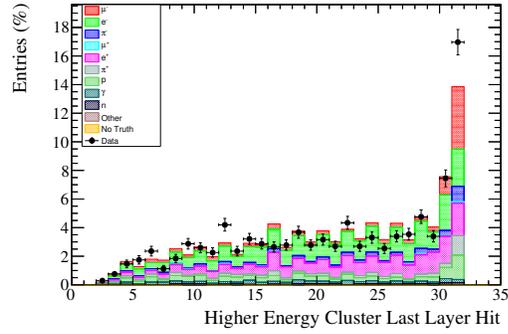
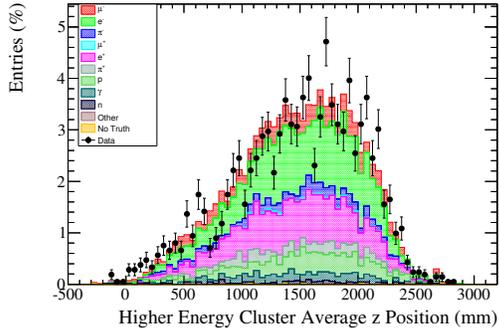
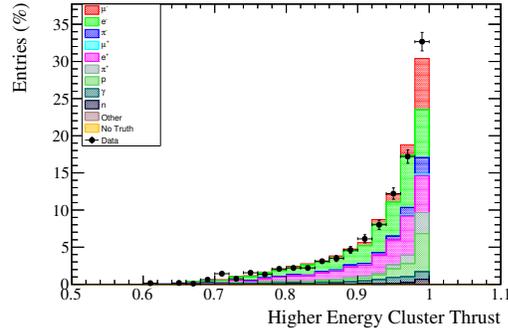
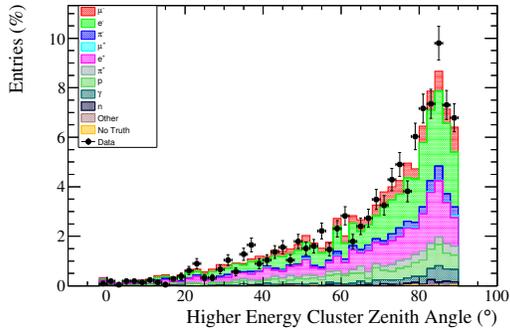
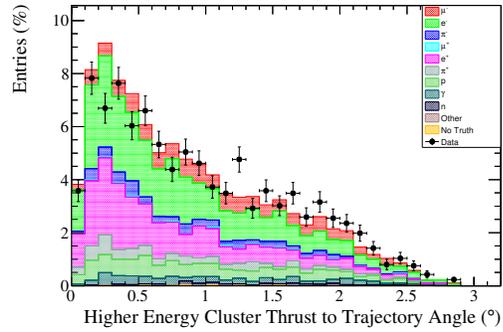
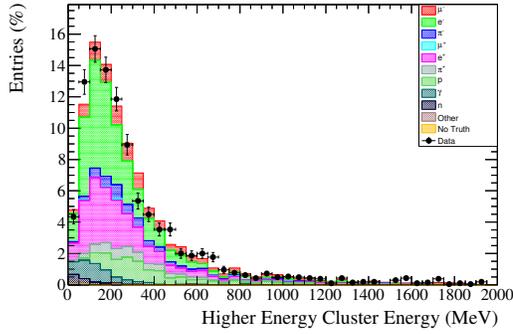
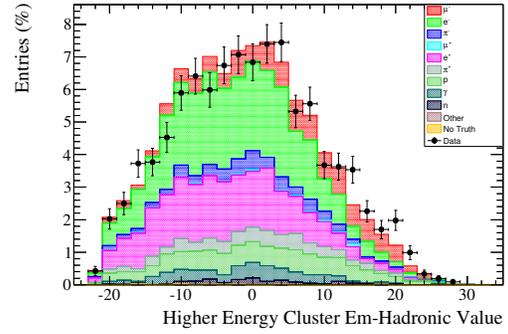
(a) E_h , first layer hit.(b) E_h , last layer hit.(c) E_h , average z position.(d) E_h , thrust.(e) E_h , zenith angle.(f) E_h , thrust to trajectory angle.

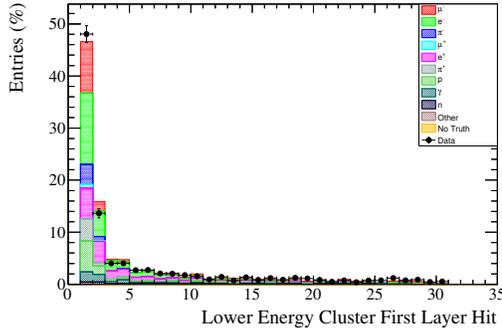
Figure B.4: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the bottom BrECal.



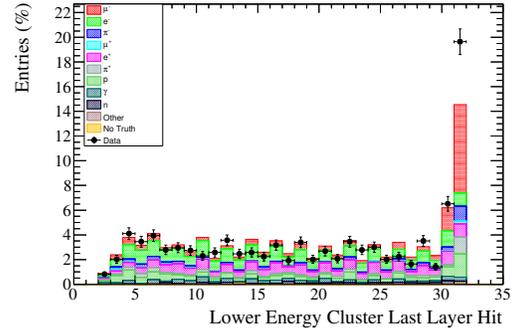
(g) E_h , EM energy.



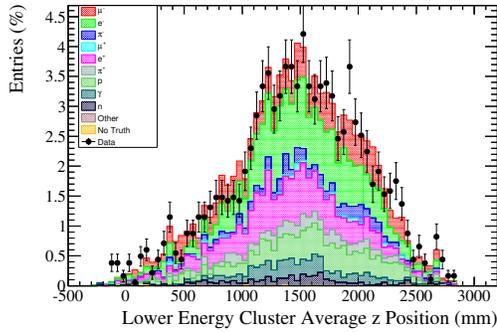
(h) E_h , EmHip.



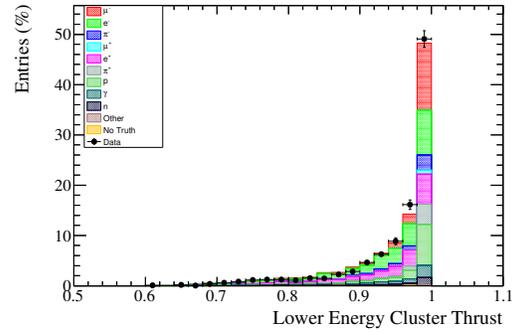
(i) E_l , first layer hit.



(j) E_l , last layer hit.

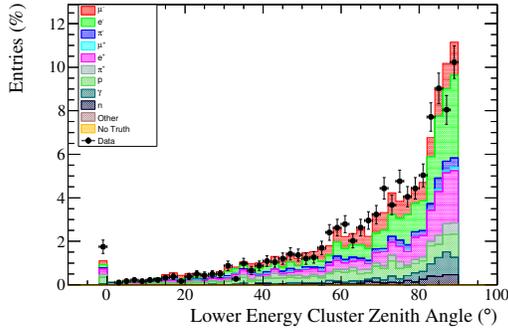
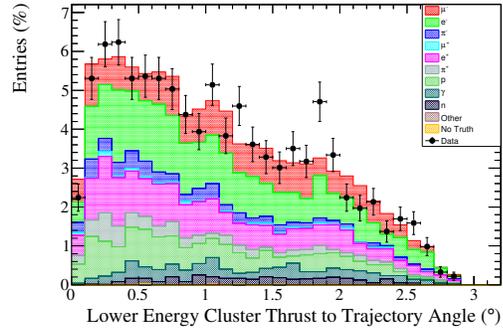
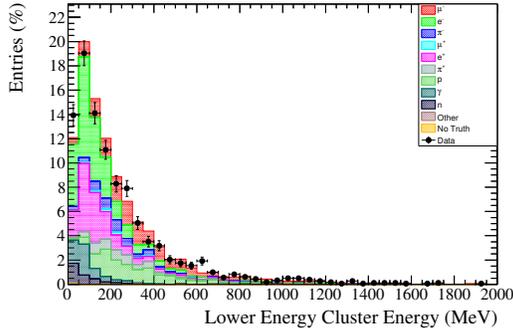
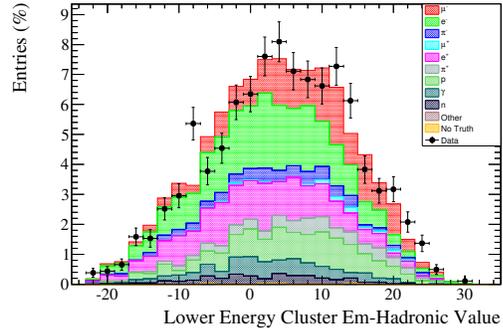
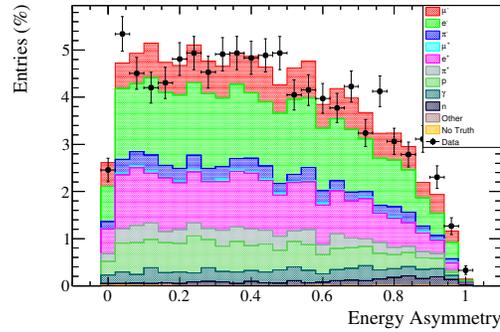


(k) E_l , average z position.



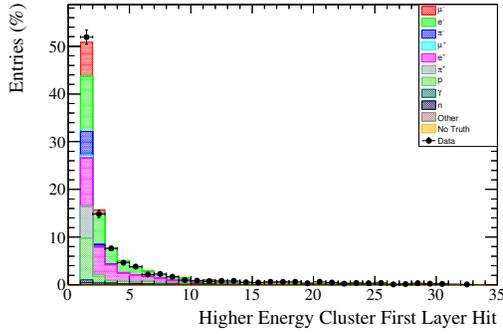
(l) E_l , thrust.

Figure B.4: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the bottom BrECal.

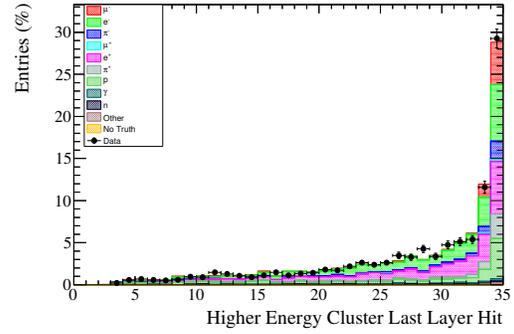
(m) E_l , zenith angle.(n) E_l , thrust to trajectory angle.(o) E_l , EM energy.(p) E_l , EmHip.

(q) Energy asymmetry.

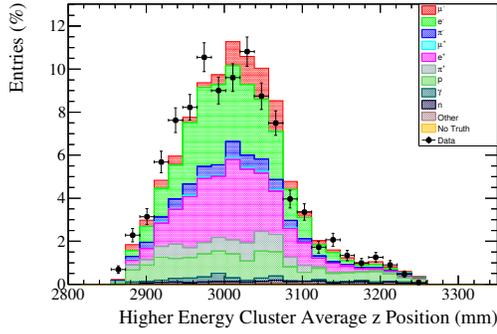
Figure B.4: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the bottom BrECal.



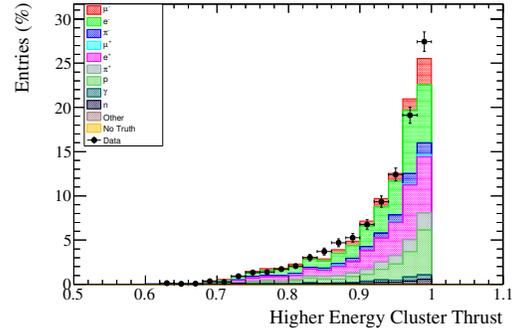
(a) E_h , first layer hit.



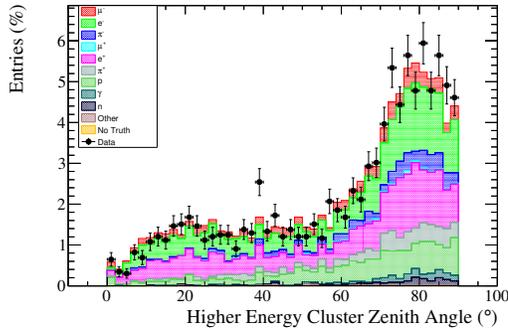
(b) E_h , last layer hit.



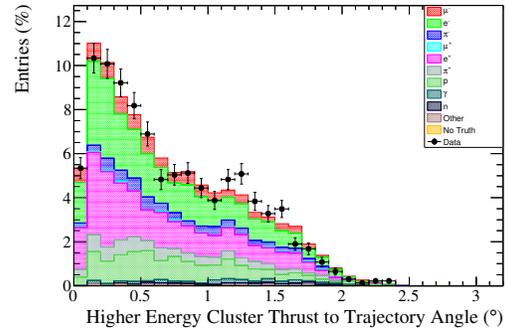
(c) E_h , average z position.



(d) E_h , thrust.



(e) E_h , zenith angle.



(f) E_h , thrust to trajectory angle.

Figure B.5: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the DsECal.

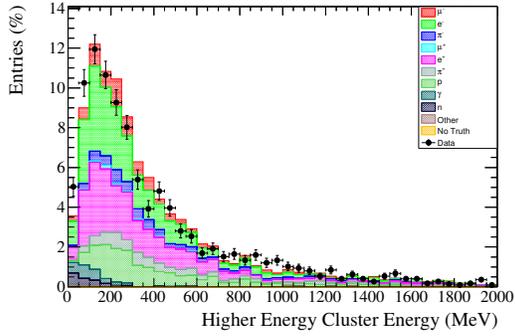
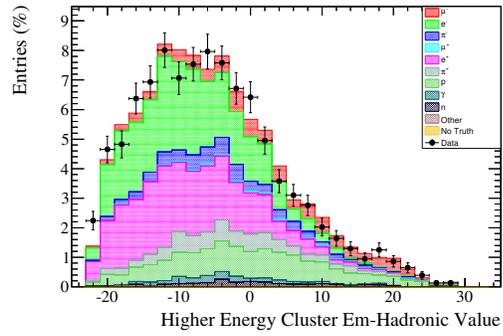
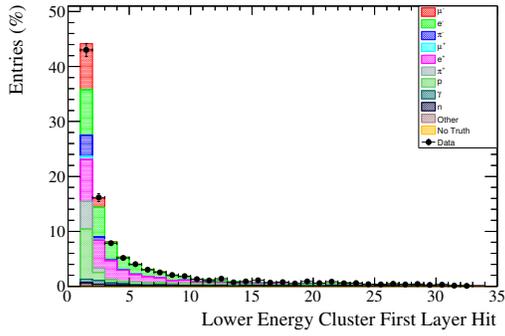
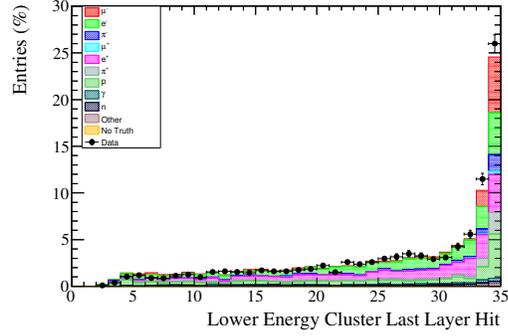
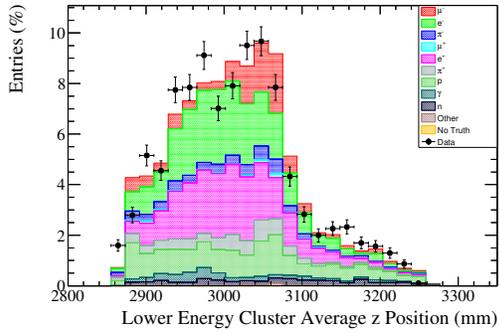
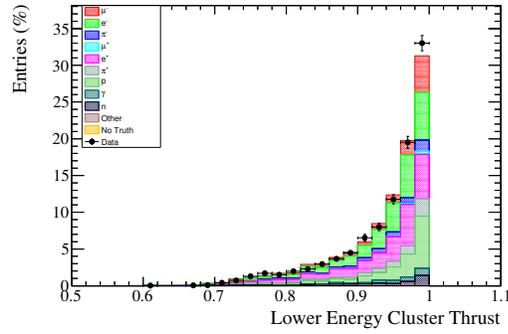
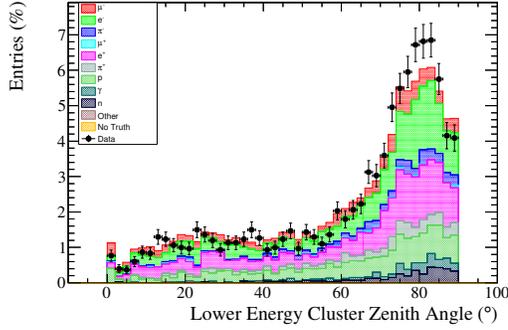
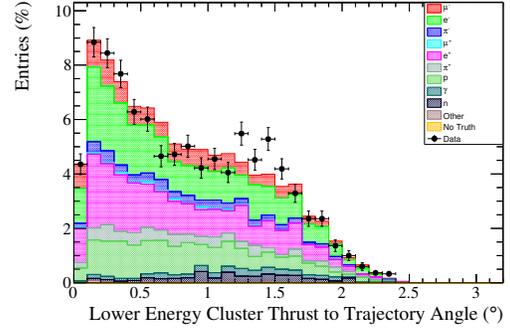
(g) E_h , EM energy.(h) E_h , EmHip.(i) E_l , first layer hit.(j) E_l , last layer hit.(k) E_l , average z position.(l) E_l , thrust.

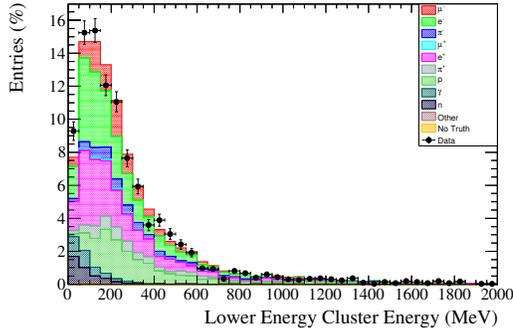
Figure B.5: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the DsECal.



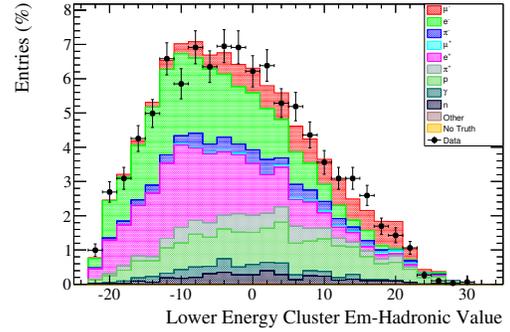
(m) E_l , zenith angle.



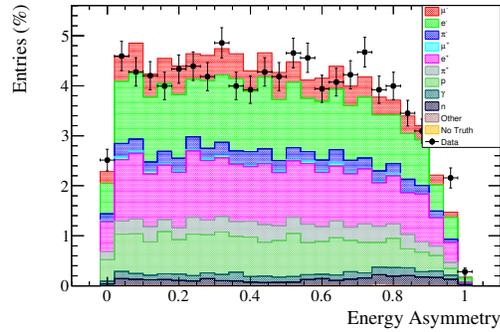
(n) E_l , thrust to trajectory angle.



(o) E_l , EM energy.



(p) E_l , EmHip.

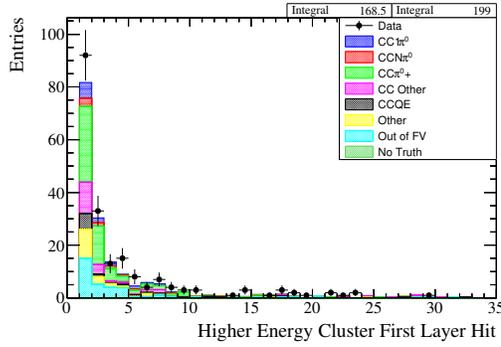


(q) Energy asymmetry.

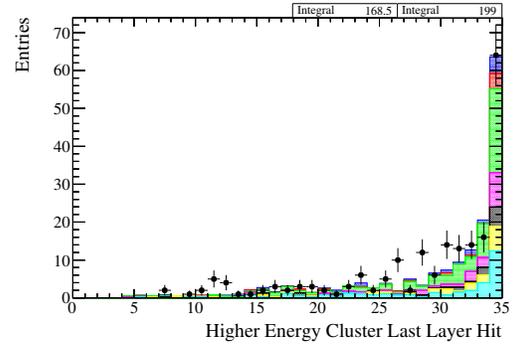
Figure B.5: Area normalised Data-MC comparison plots for all discriminating variables used by the MVA selection in the DsEcal.

Appendix C

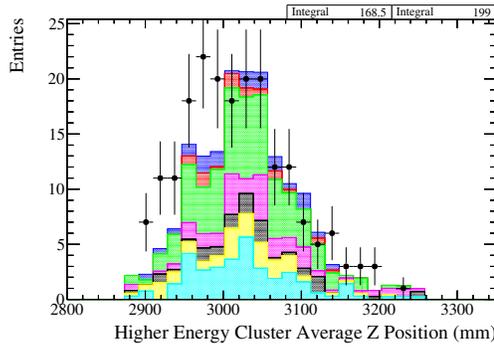
Data and Monte-Carlo BDT Discriminating Variable Comparisons for Each Event Topology



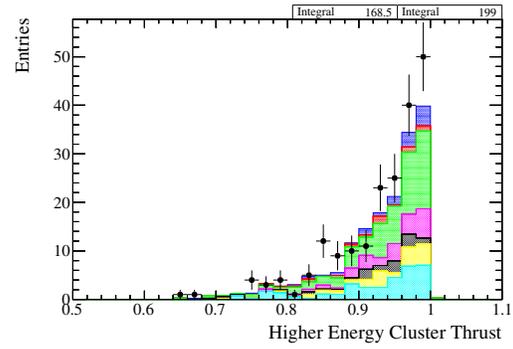
(a) E_h , first layer hit.



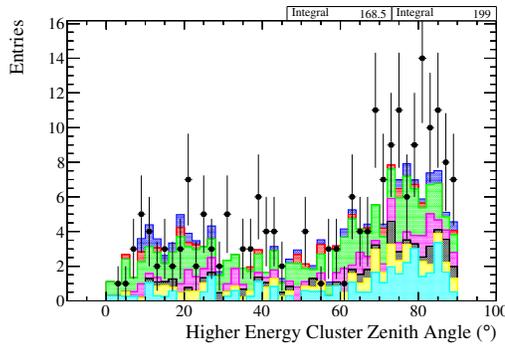
(b) E_h , last layer hit.



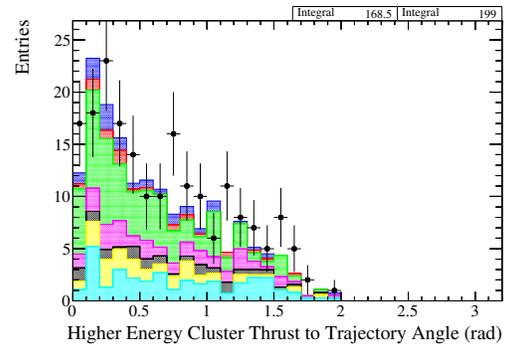
(c) E_h , average z position.



(d) E_h , thrust.



(e) E_h , zenith angle.



(f) E_h , thrust to trajectory angle.

Figure C.1: Data-MC comparison plots for events that could contribute to the final FGD1DsDs event topology, after the Reconstructed Mass cut has been applied.

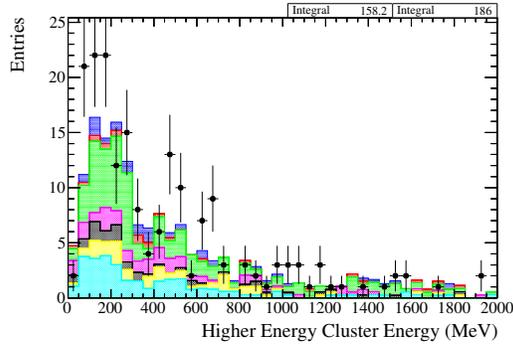
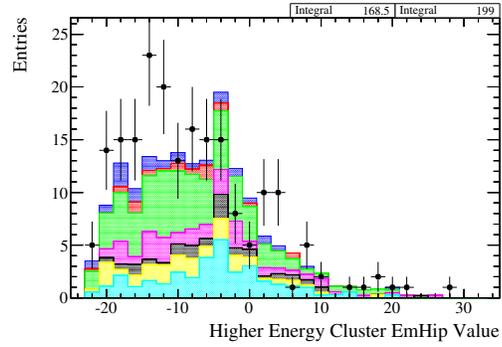
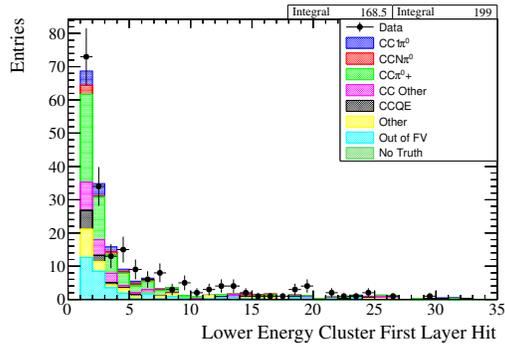
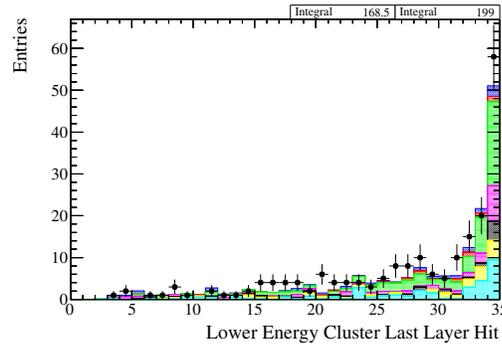
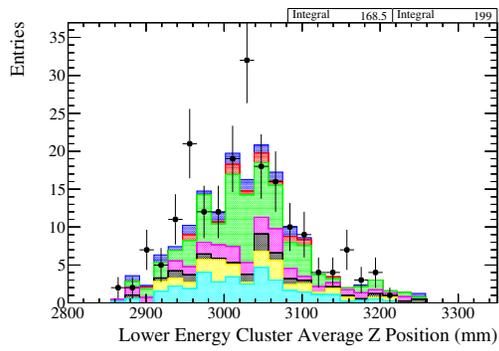
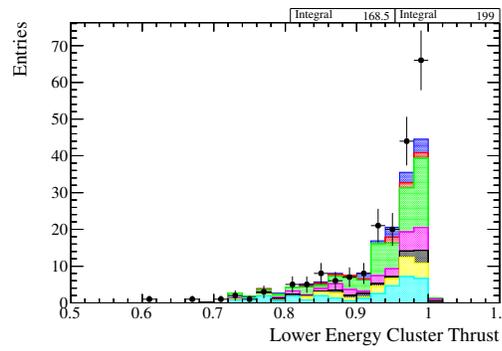
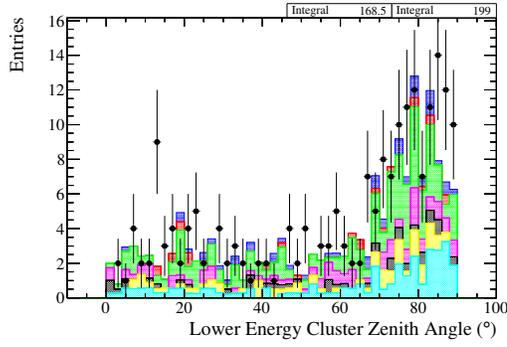
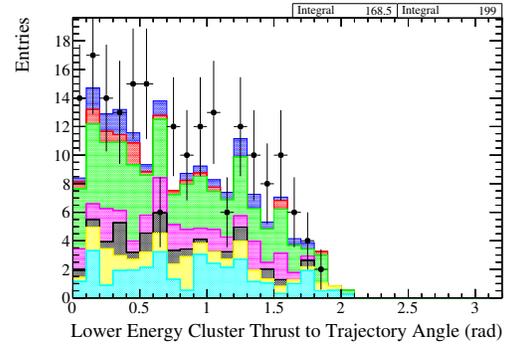
(g) E_h , EM energy.(h) E_h , EmHip.(i) E_l , first layer hit.(j) E_l , last layer hit.(k) E_l , average z position.(l) E_l , thrust.

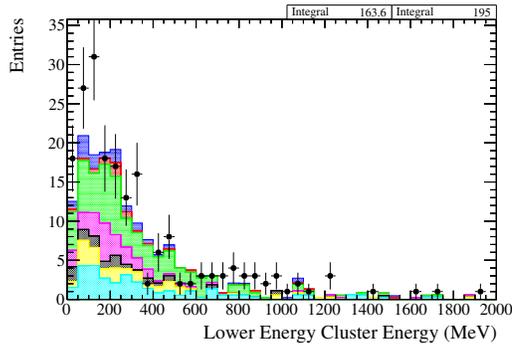
Figure C.1: Data-MC comparison plots for events that could contribute to the final FGD1DsDs event topology, after the Reconstructed Mass cut has been applied.



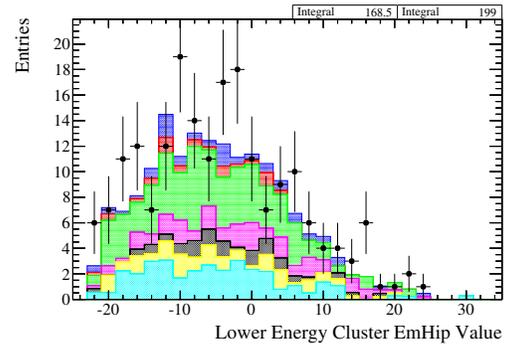
(m) E_l , zenith angle.



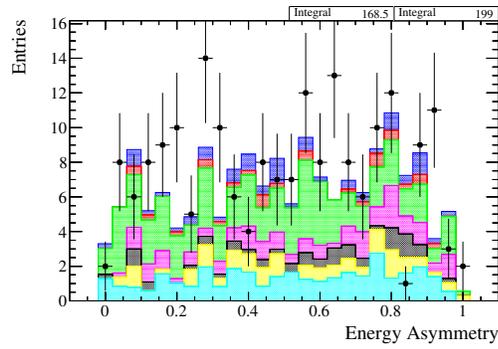
(n) E_l , thrust to trajectory angle.



(o) E_l , EM energy.



(p) E_l , EmHip.



(q) Energy asymmetry.

Figure C.1: Data-MC comparison plots for events that could contribute to the final FGD1DsDs event topology, after the Reconstructed Mass cut has been applied.

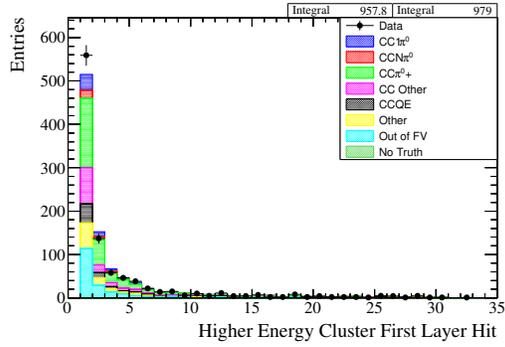
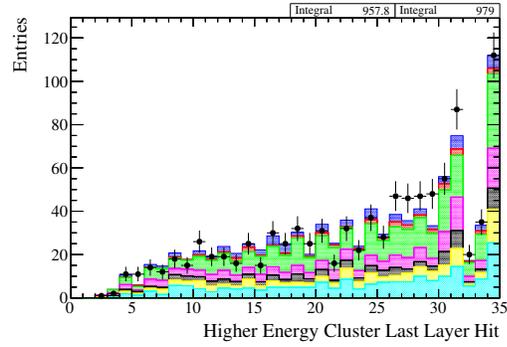
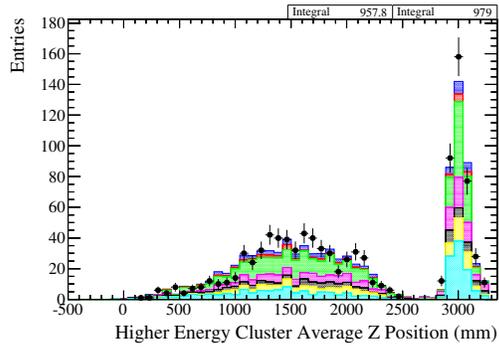
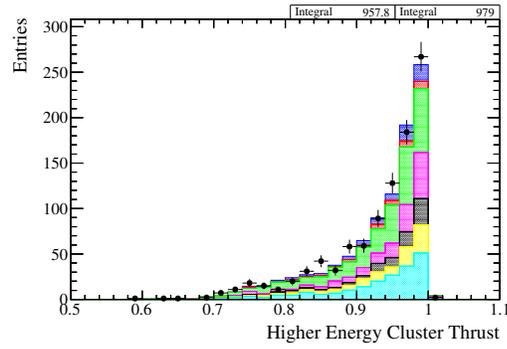
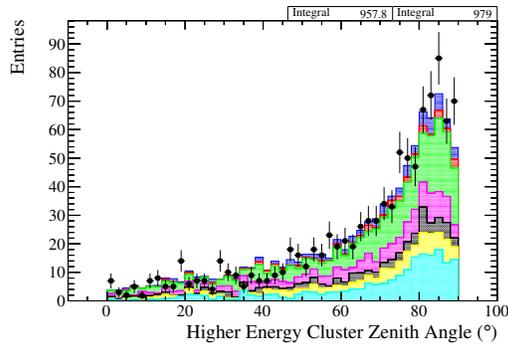
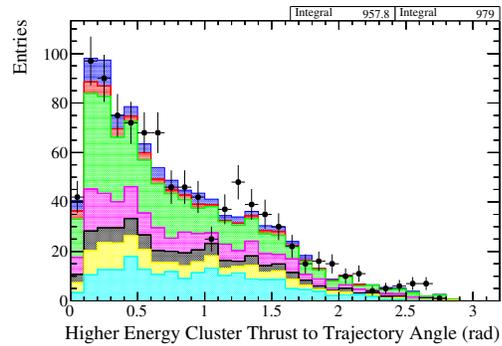
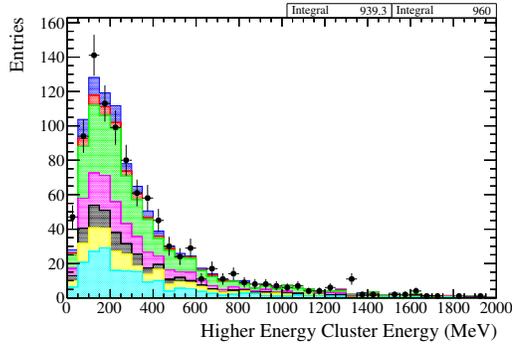
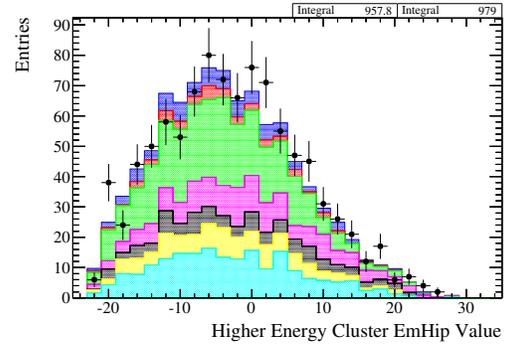
(a) E_h , first layer hit.(b) E_h , last layer hit.(c) E_h , average z position.(d) E_h , thrust.(e) E_h , zenith angle.(f) E_h , thrust to trajectory angle.

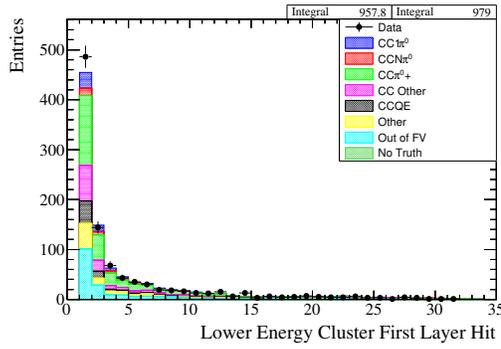
Figure C.2: Data-MC comparison plots for events that could contribute to the final FGD1DsBr event topology, after the Reconstructed Mass cut has been applied.



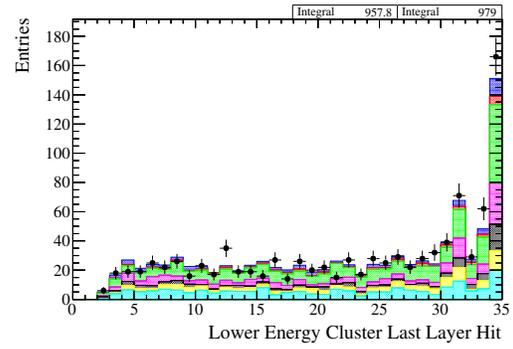
(g) E_h , EM energy.



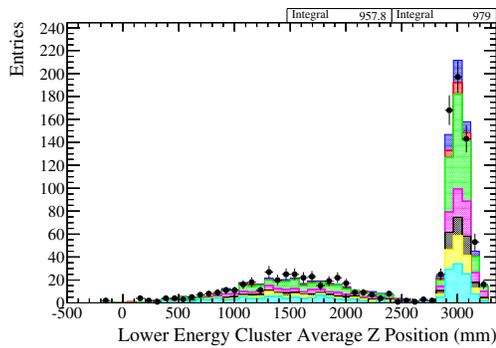
(h) E_h , EmHip.



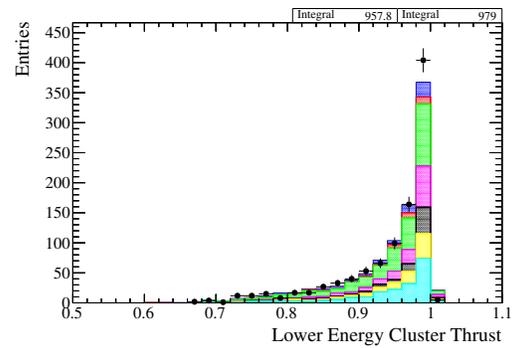
(i) E_l , first layer hit.



(j) E_l , last layer hit.

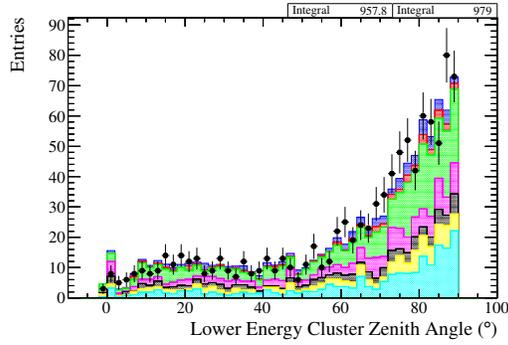
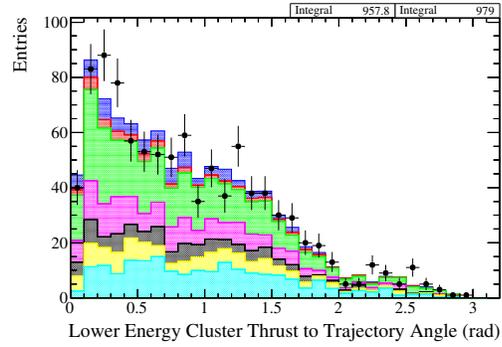
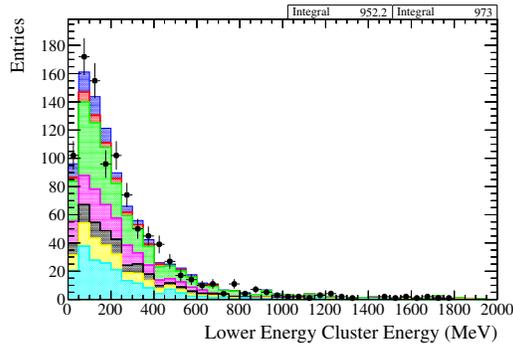
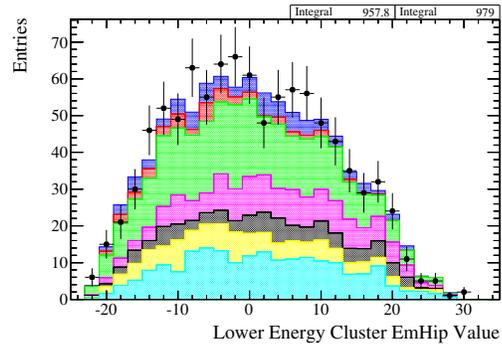
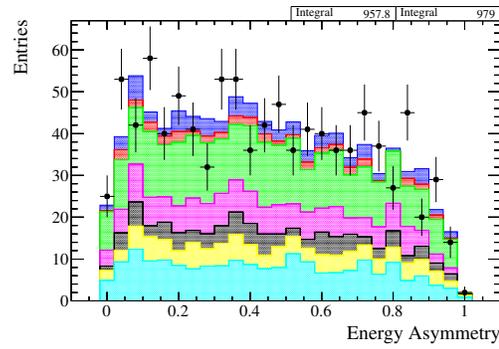


(k) E_l , average z position.



(l) E_l , thrust.

Figure C.2: Data-MC comparison plots for events that could contribute to the final FGD1DsBr event topology, after the Reconstructed Mass cut has been applied.

(m) E_l , zenith angle.(n) E_l , thrust to trajectory angle.(o) E_l , EM energy.(p) E_l , EmHip.

(q) Energy asymmetry.

Figure C.2: Data-MC comparison plots for events that could contribute to the final FGD1DsBr event topology, after the Reconstructed Mass cut has been applied.

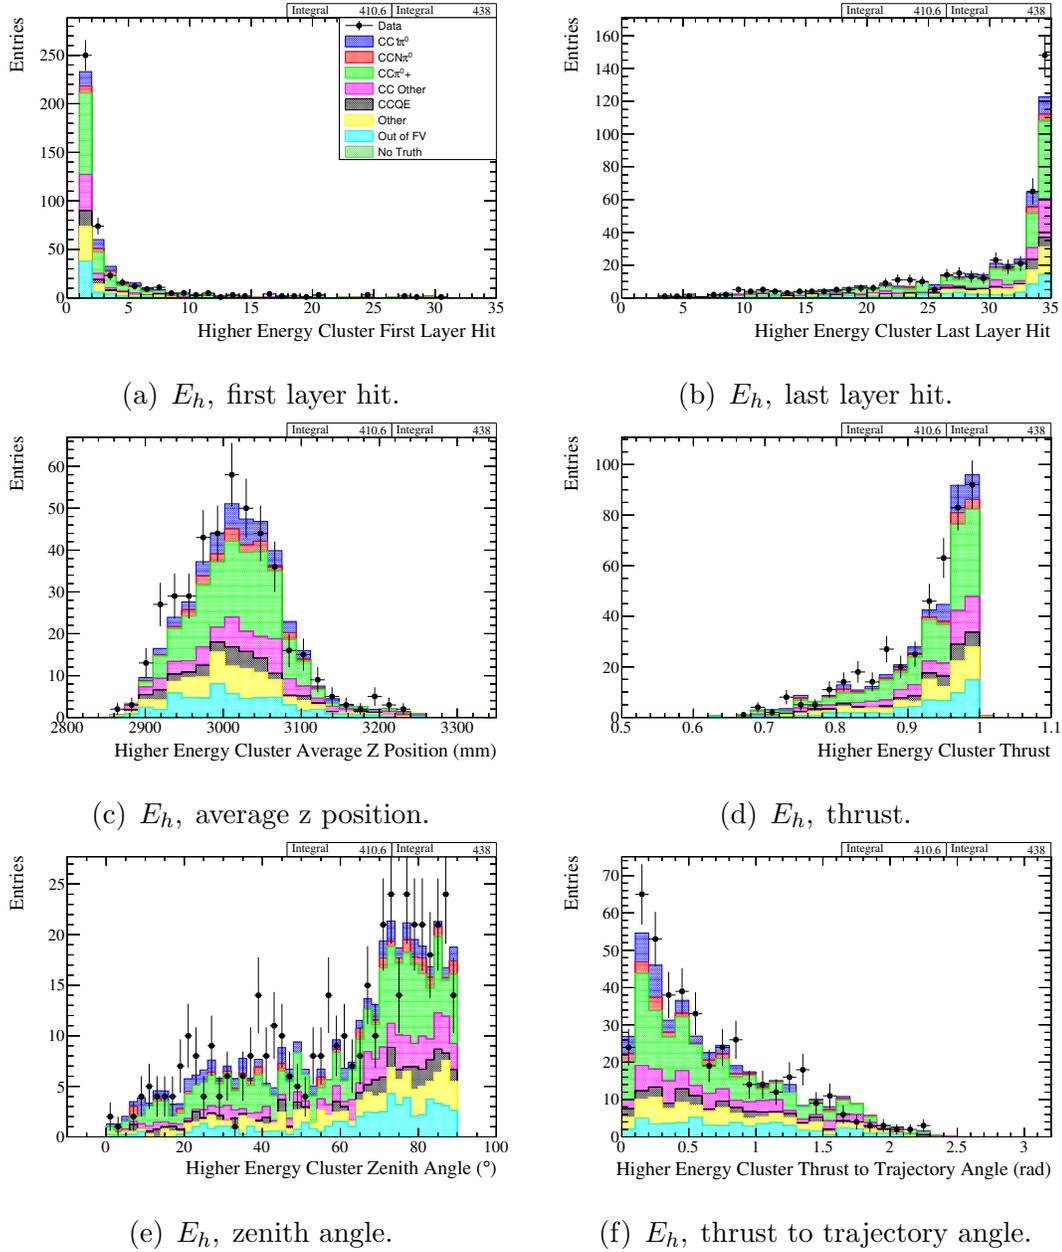


Figure C.3: Data-MC comparison plots for events that could contribute to the final FGD2DsDs event topology, after the Reconstructed Mass cut has been applied.

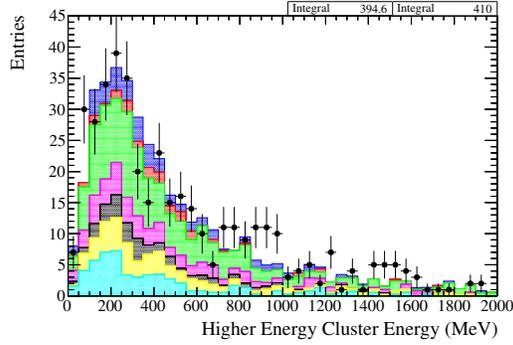
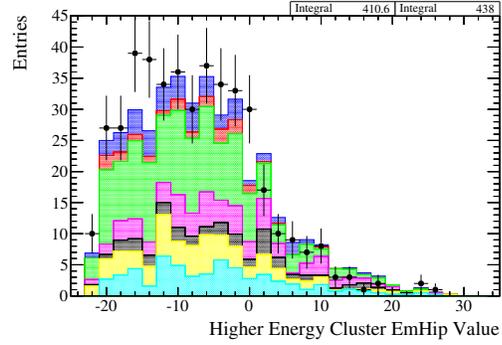
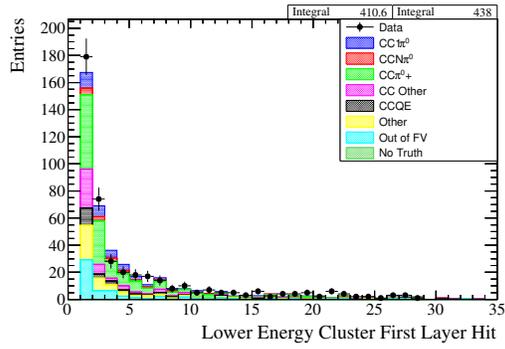
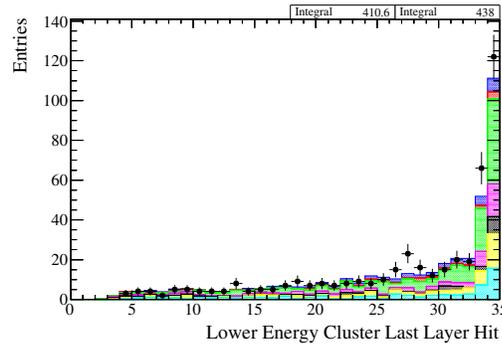
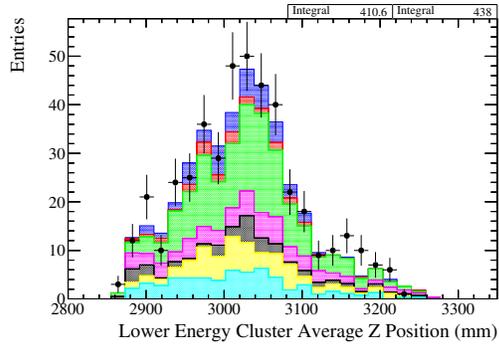
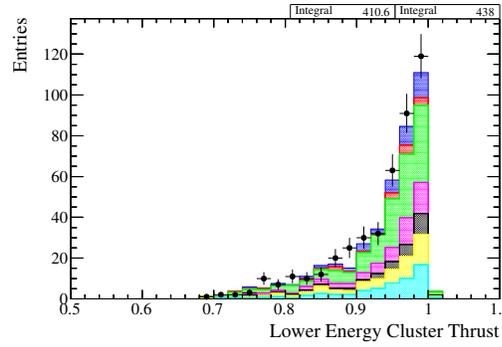
(g) E_h , EM energy.(h) E_h , EmHip.(i) E_l , first layer hit.(j) E_l , last layer hit.(k) E_l , average z position.(l) E_l , thrust.

Figure C.3: Data-MC comparison plots for events that could contribute to the final FGD2DsDs event topology, after the Reconstructed Mass cut has been applied.

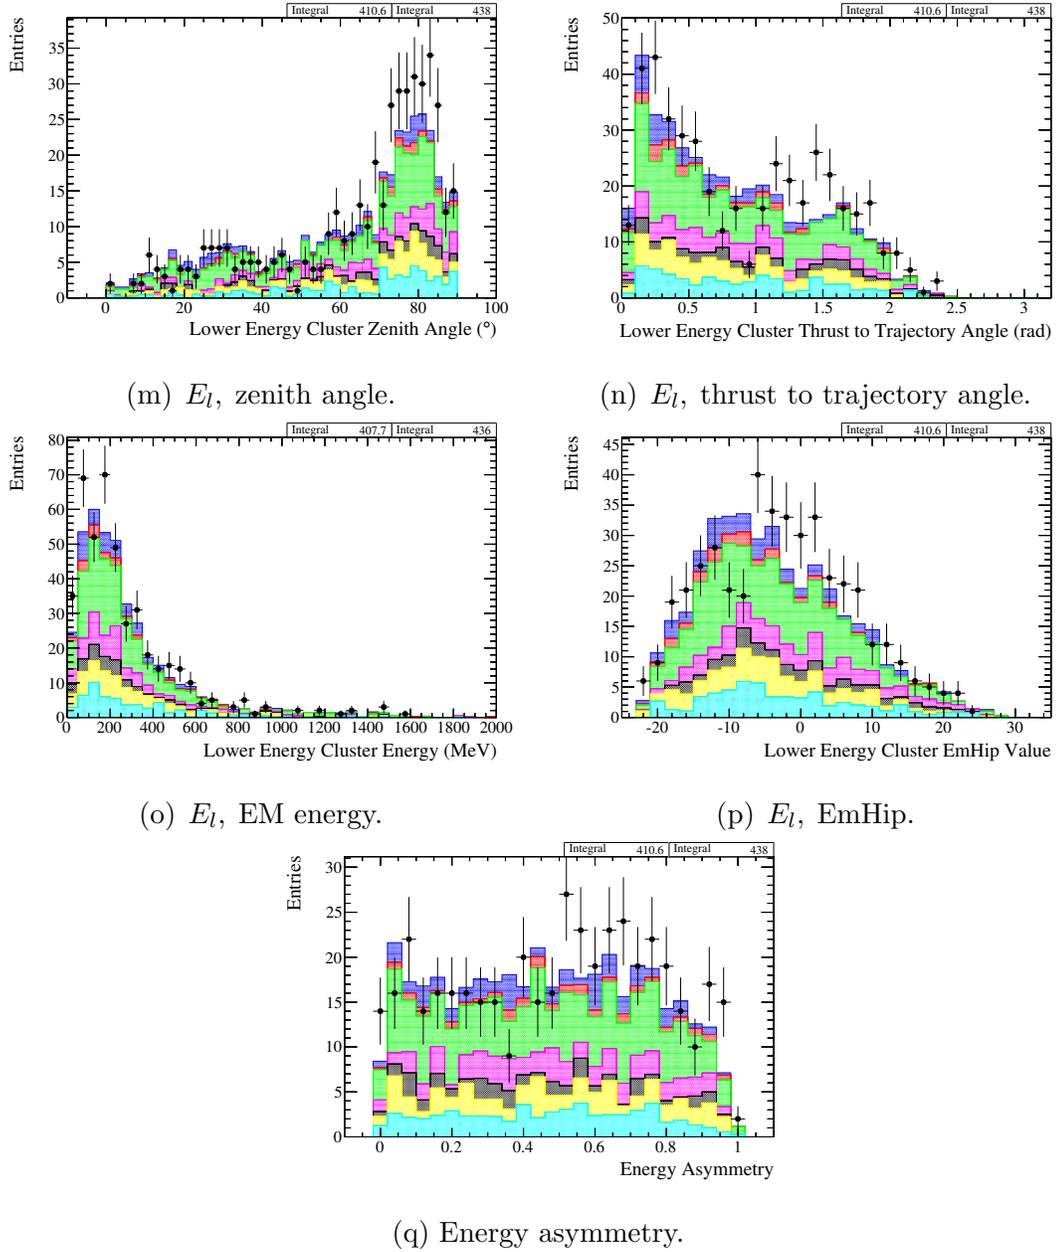


Figure C.3: Data-MC comparison plots for events that could contribute to the final FGD2DsDs event topology, after the Reconstructed Mass cut has been applied.

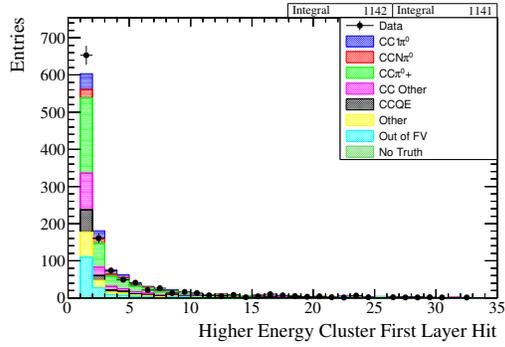
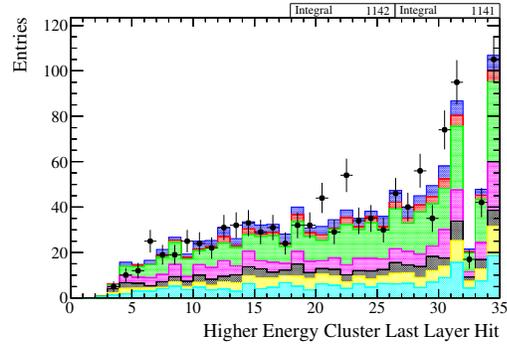
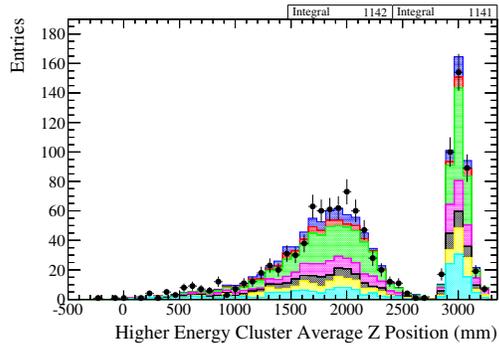
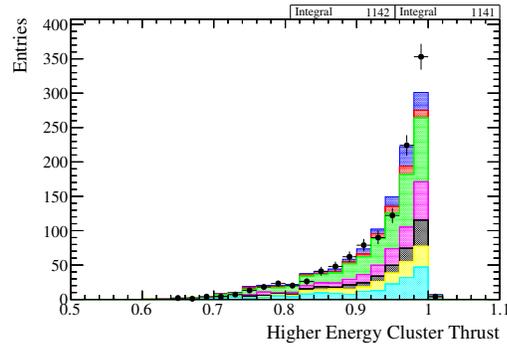
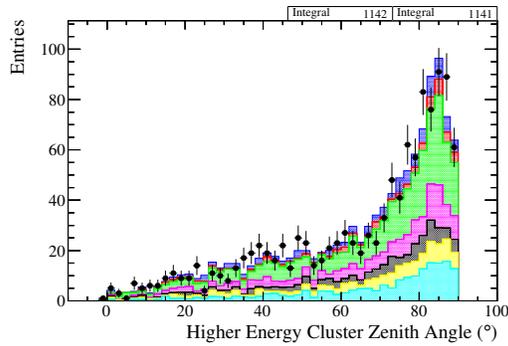
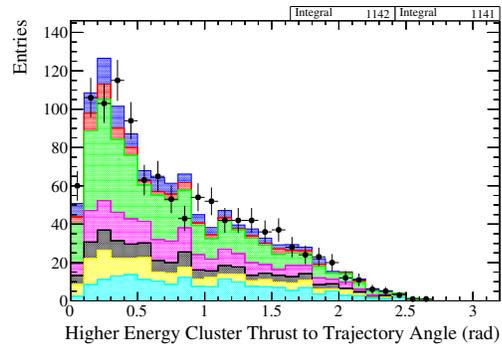
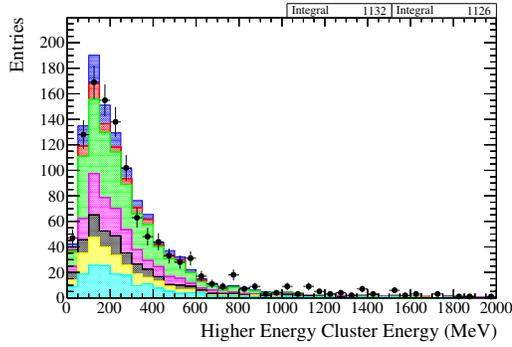
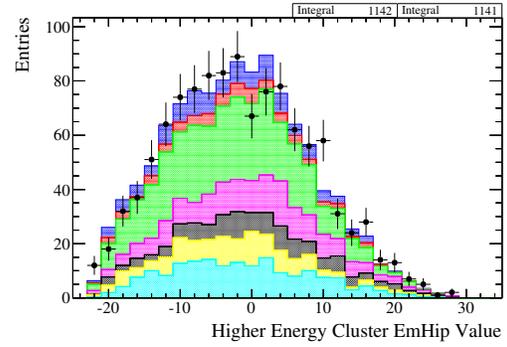
(a) E_h , first layer hit.(b) E_h , last layer hit.(c) E_h , average z position.(d) E_h , thrust.(e) E_h , zenith angle.(f) E_h , thrust to trajectory angle.

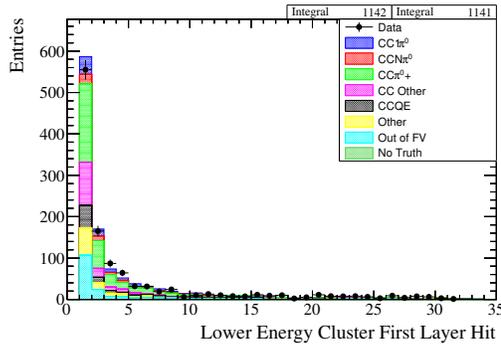
Figure C.4: Data-MC comparison plots for events that could contribute to the final FGD2DsBr event topology, after the Reconstructed Mass cut has been applied.



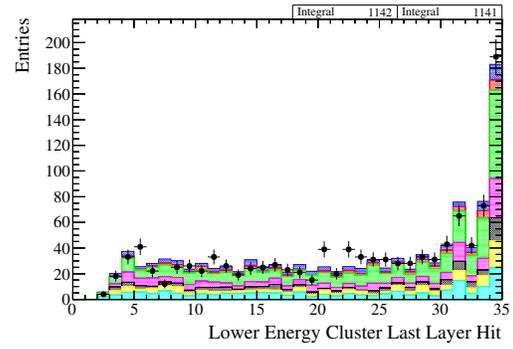
(g) E_h , EM energy.



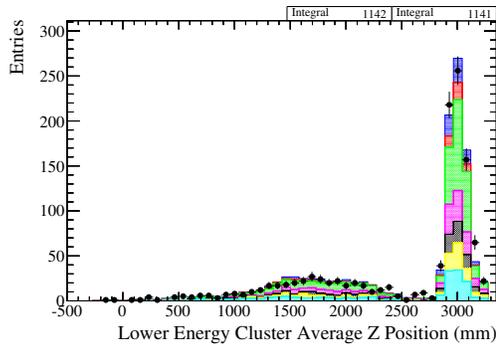
(h) E_h , EmHip.



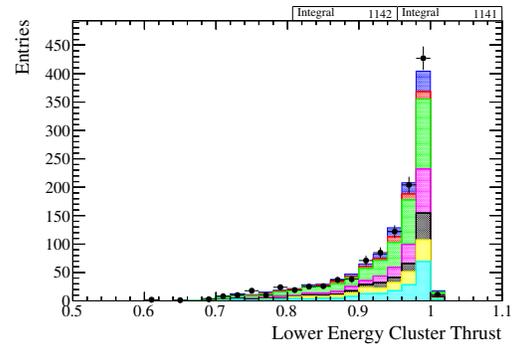
(i) E_l , first layer hit.



(j) E_l , last layer hit.



(k) E_l , average z position.



(l) E_l , thrust.

Figure C.4: Data-MC comparison plots for events that could contribute to the final FGD2DsBr event topology, after the Reconstructed Mass cut has been applied.

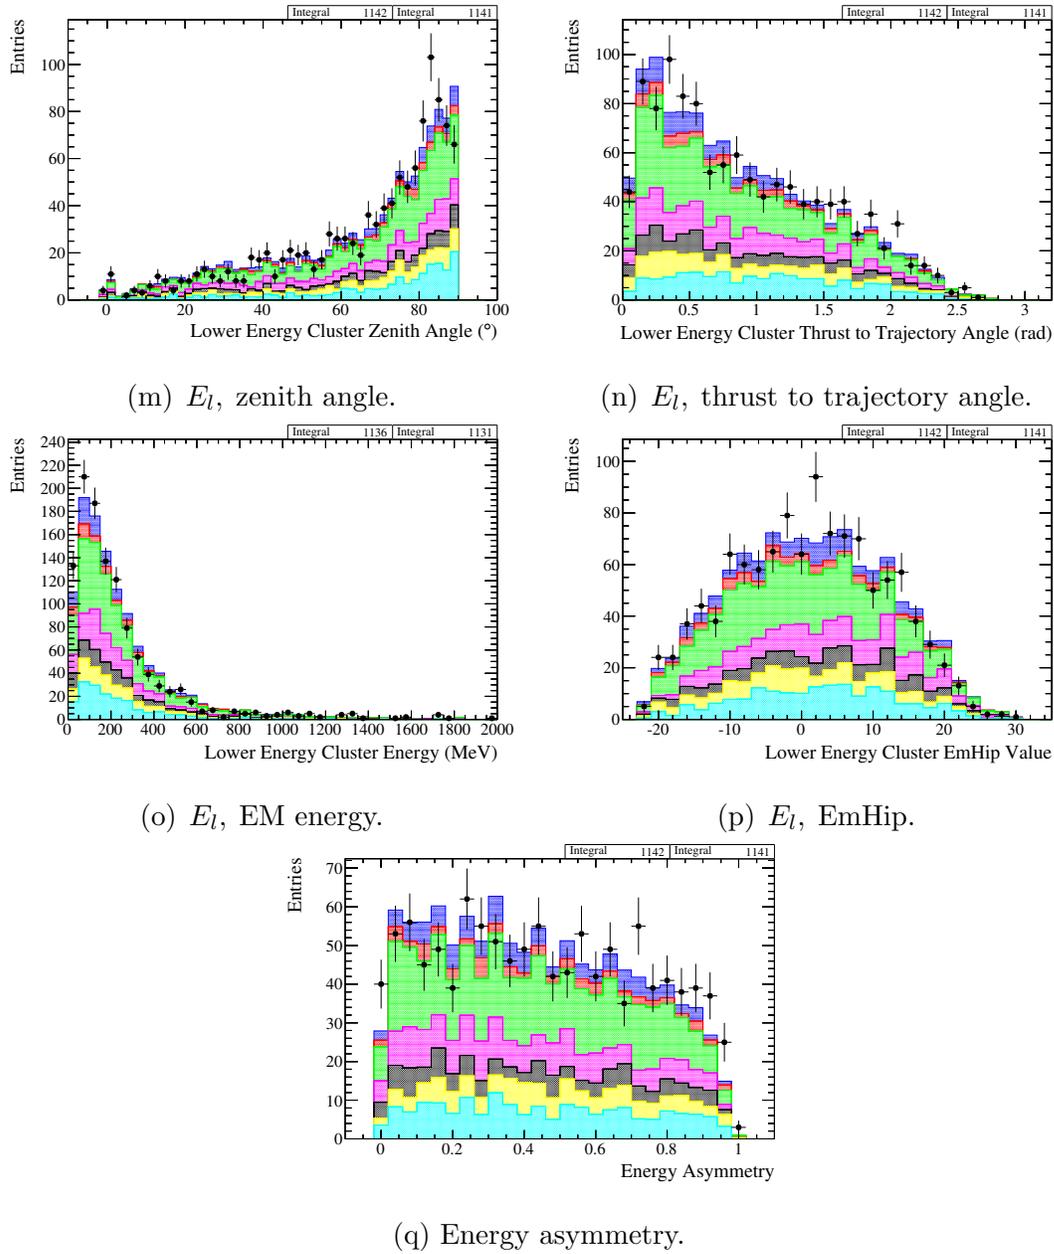
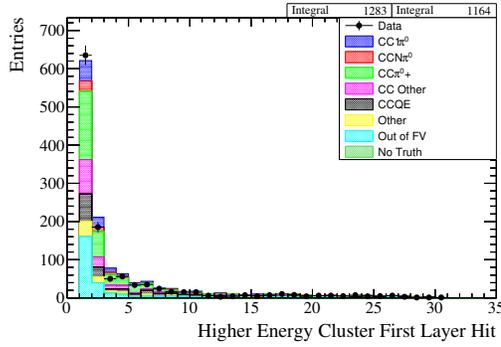
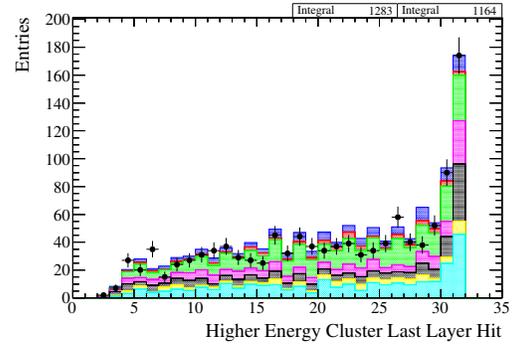


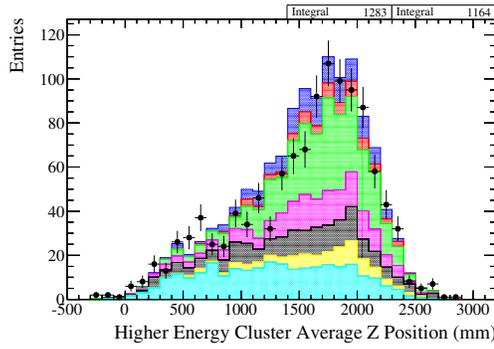
Figure C.4: Data-MC comparison plots for events that could contribute to the final FGD2DsBr event topology, after the Reconstructed Mass cut has been applied.



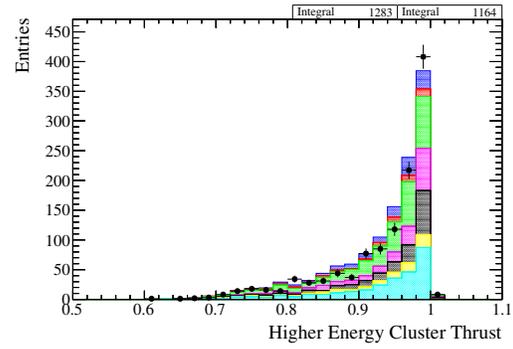
(a) E_h , first layer hit.



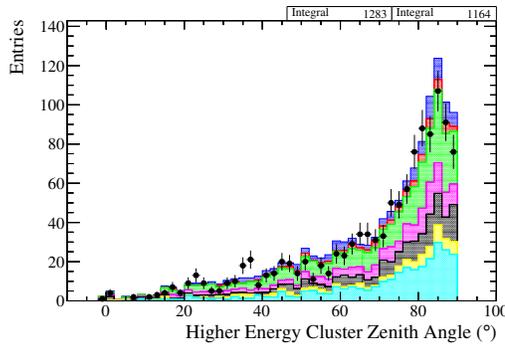
(b) E_h , last layer hit.



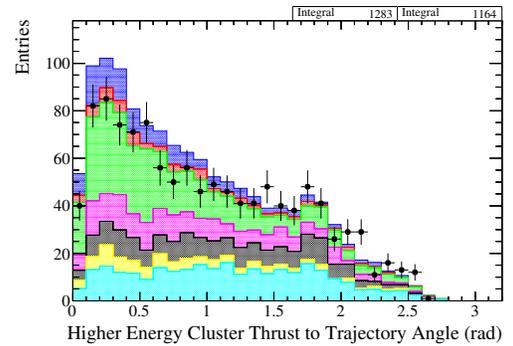
(c) E_h , average z position.



(d) E_h , thrust.



(e) E_h , zenith angle.



(f) E_h , thrust to trajectory angle.

Figure C.5: Data-MC comparison plots for events that could contribute to the final FGD2BrBr event topology, after the Reconstructed Mass cut has been applied.

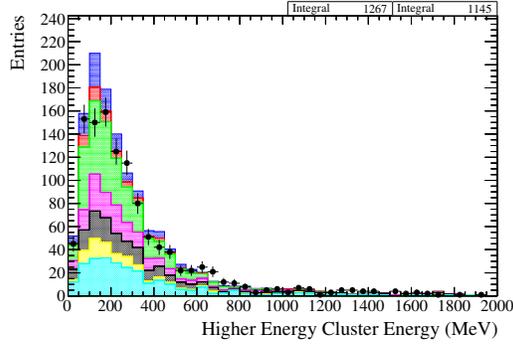
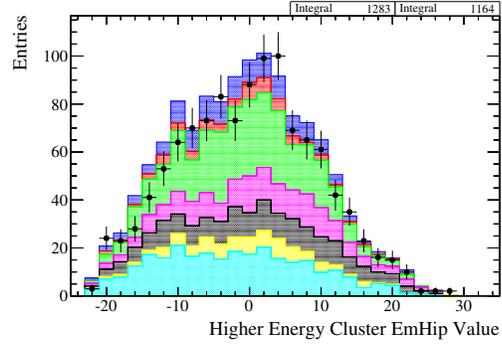
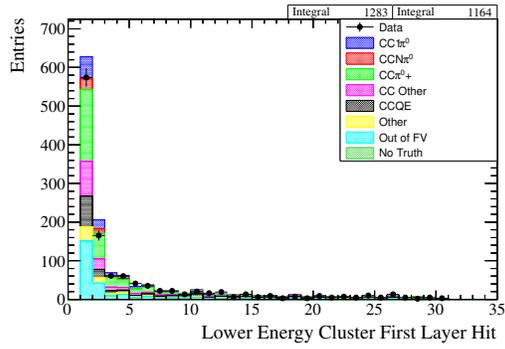
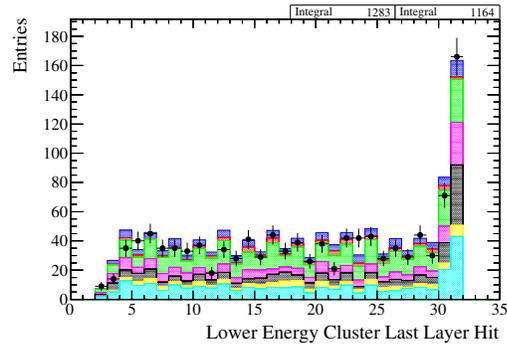
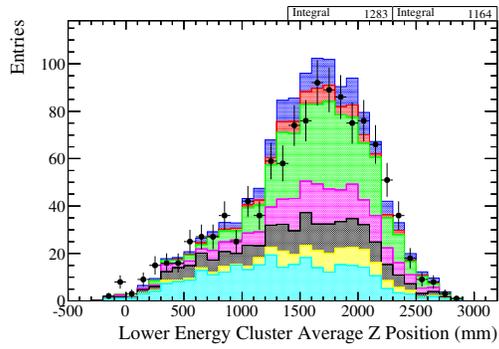
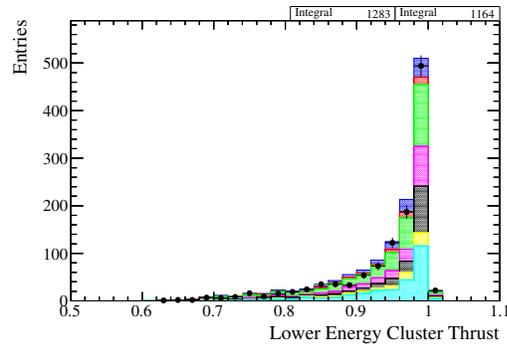
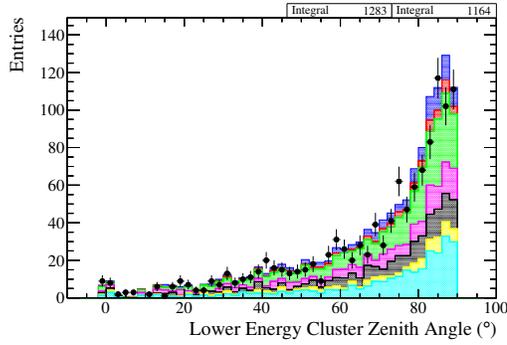
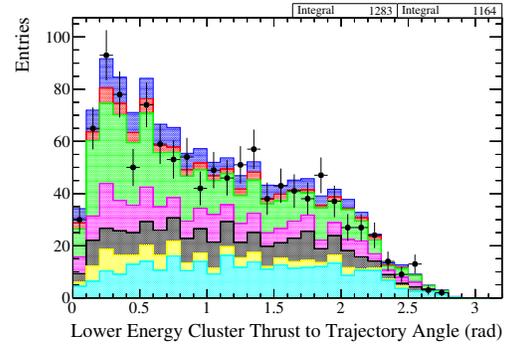
(g) E_h , EM energy.(h) E_h , EmHip.(i) E_l , first layer hit.(j) E_l , last layer hit.(k) E_l , average z position.(l) E_l , thrust.

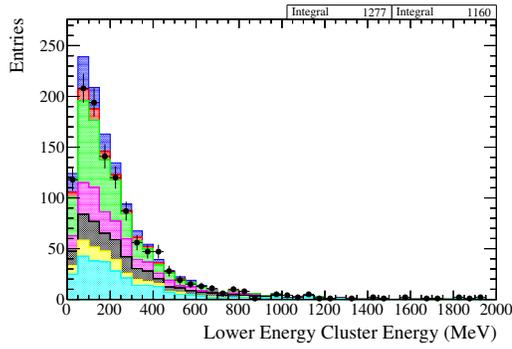
Figure C.5: Data-MC comparison plots for events that could contribute to the final FGD2BrBr event topology, after the Reconstructed Mass cut has been applied.



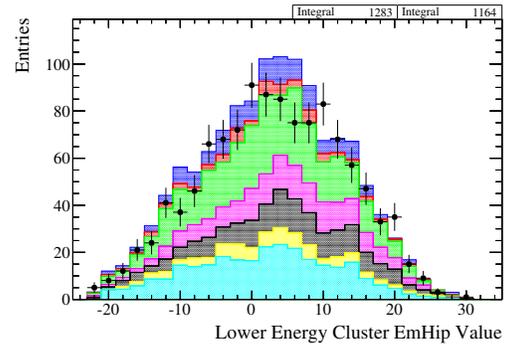
(m) E_l , zenith angle.



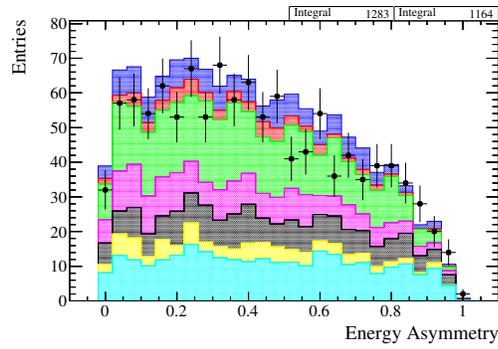
(n) E_l , thrust to trajectory angle.



(o) E_l , EM energy.



(p) E_l , EmHip.



(q) Energy asymmetry.

Figure C.5: Data-MC comparison plots for events that could contribute to the final FGD2BrBr event topology, after the Reconstructed Mass cut has been applied.

Appendix D

π^0 Invariant Mass Calculation

Figure D.1 shows the decay of a π^0 in flight to two photons, γ_1 and γ_2 . If the energy of the two photons, and the angle between them are known, it is trivial to reconstruct the invariant mass of the π^0 as follows.

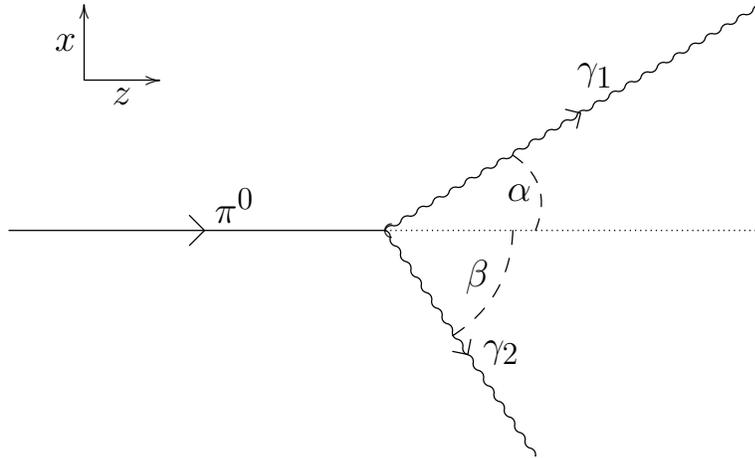


Figure D.1: The schematic decay of a π^0 in flight into a pair of photons. The π^0 is initially travelling in the positive z direction, and the two decay photons move away from the decay vertex at angles α and β from the z direction in the xz plane.

The four momentum of the initial π^0 state is

$$\mathbf{p}_{\pi^0} = (E_{\pi^0}, \vec{p}_{\pi^0}). \quad (\text{D.1})$$

The four momenta of the final γ states are

$$\mathbf{p}_{\gamma_1} = (E_{\gamma_1}, \vec{p}_{\gamma_1}), \quad (\text{D.2a})$$

$$\mathbf{p}_{\gamma_2} = (E_{\gamma_2}, \vec{p}_{\gamma_2}). \quad (\text{D.2b})$$

Through conservation of energy momentum, the following must be true

$$E_{\pi^0} = E_{\gamma_1} + E_{\gamma_2}, \quad (\text{D.3a})$$

$$\vec{p}_{\pi^0} = \vec{p}_{\gamma_1} + \vec{p}_{\gamma_2}. \quad (\text{D.3b})$$

Using the relation between energy, mass and momentum (in natural units $\hbar = c = 1$)

$$E^2 = m^2 + \vec{p}^2. \quad (\text{D.4})$$

The π^0 rest mass can be expressed as a function of its energy and momentum, and these can be further expressed in terms of the energy and momentum of the decay photons from Equation D.3

$$m_{\pi^0}^2 = E_{\pi^0}^2 - \vec{p}_{\pi^0}^2 \quad (\text{D.5a})$$

$$= (E_{\gamma_1} + E_{\gamma_2})^2 - (\vec{p}_{\gamma_1} + \vec{p}_{\gamma_2})^2. \quad (\text{D.5b})$$

Breaking down the two photon momenta into their three directional components we find that

$$\vec{p}_{\gamma_1} = (E_1 \sin \alpha, 0, E_1 \cos \alpha), \quad (\text{D.6a})$$

$$\vec{p}_{\gamma_2} = (-E_2 \sin \beta, 0, E_2 \cos \beta). \quad (\text{D.6b})$$

Substituting Equation D.6 back into Equation D.5 and expanding we find

$$m_{\pi^0}^2 = (E_{\gamma_1} + E_{\gamma_2})^2 - ((E_1 \sin \alpha, 0, E_1 \cos \alpha) + (-E_2 \sin \beta, 0, E_2 \cos \beta))^2 \quad (\text{D.7a})$$

$$= (E_{\gamma_1} + E_{\gamma_2})^2 - ((E_1 \sin \alpha - E_2 \sin \beta), 0, (E_1 \cos \alpha + E_2 \cos \beta))^2 \quad (\text{D.7b})$$

$$= E_1^2 + E_2^2 + 2E_1E_2 - (E_1 \sin \alpha - E_2 \sin \beta)^2 - (E_1 \cos \alpha + E_2 \cos \beta)^2 \quad (\text{D.7c})$$

$$= E_1^2 + E_2^2 + 2E_1E_2 - (E_1^2 \sin^2 \alpha + E_2^2 \sin^2 \beta - 2E_1E_2 \sin \alpha \sin \beta) - (E_1^2 \cos^2 \alpha + E_2^2 \cos^2 \beta + 2E_1E_2 \cos \alpha \cos \beta). \quad (\text{D.7d})$$

By grouping common terms and using the trigonometric identities $\sin^2 x + \cos^2 x = 1$ and $\cos x \cos y - \sin x \sin y = \cos(x + y)$ we can simplify Equation D.7

$$m_{\pi^0}^2 = E_1^2 (1 - (\sin^2 \alpha + \cos^2 \alpha)) + E_2^2 (1 - (\sin^2 \beta + \cos^2 \beta)) + 2E_1E_2 (1 + \sin \alpha \sin \beta - \cos \alpha \cos \beta) \quad (\text{D.8a})$$

$$= E_1^2 (1 - (1)) + E_2^2 (1 - (1)) + 2E_1E_2 (1 - \cos(\alpha + \beta)) \quad (\text{D.8b})$$

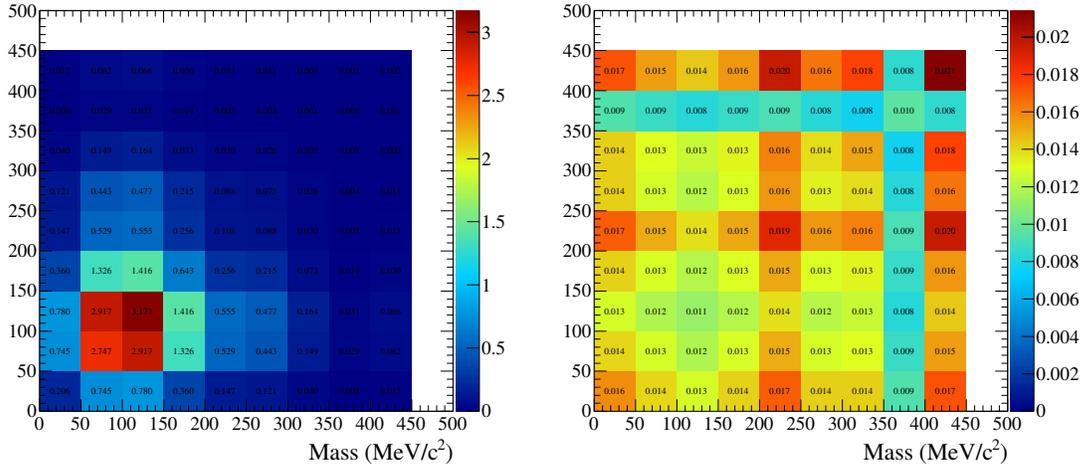
$$= 2E_1E_2 (1 - \cos \theta), \quad (\text{D.8c})$$

where $\theta = \alpha + \beta$, the angle between the clusters, and thus the invariant π^0 mass is simply

$$m_{\pi^0} = \sqrt{2E_1E_2 (1 - \cos \theta)}. \quad (\text{D.9})$$

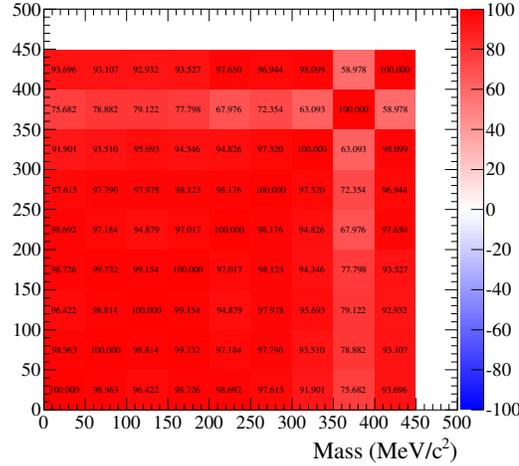
Appendix E

Correlation and Covariance Matrices for Flux Systematic Error Calculation



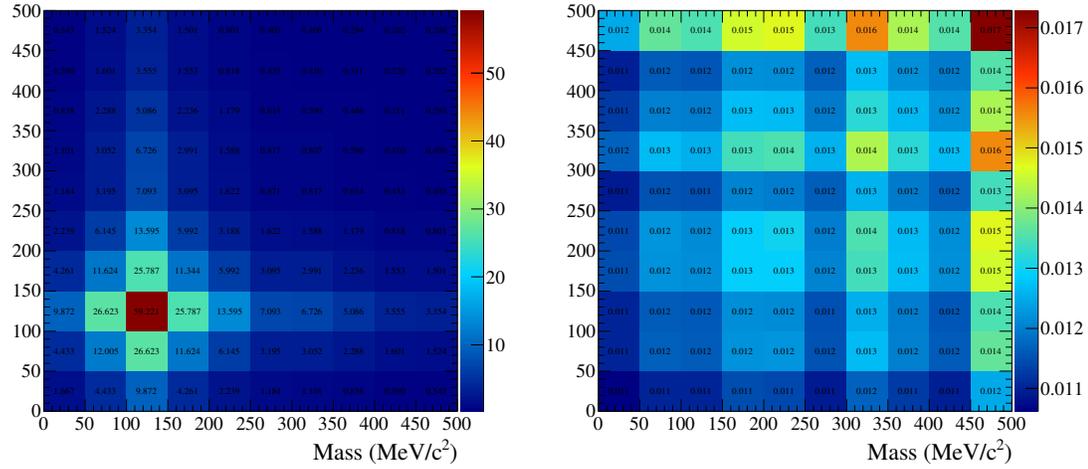
(a) Covariance matrix.

(b) Fractional covariance matrix.



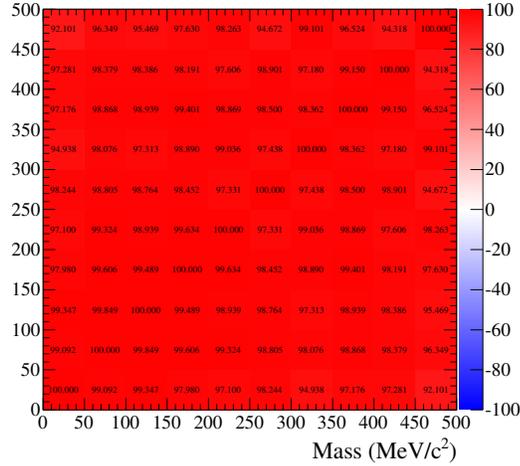
(c) Correlation matrix (%).

Figure E.1: Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD1DsBr event topology.



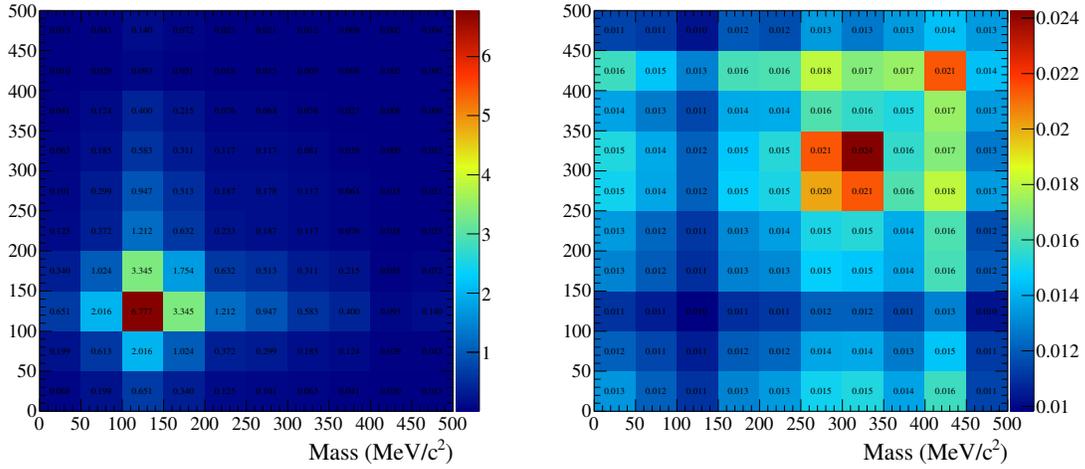
(a) Covariance matrix.

(b) Fractional covariance matrix.



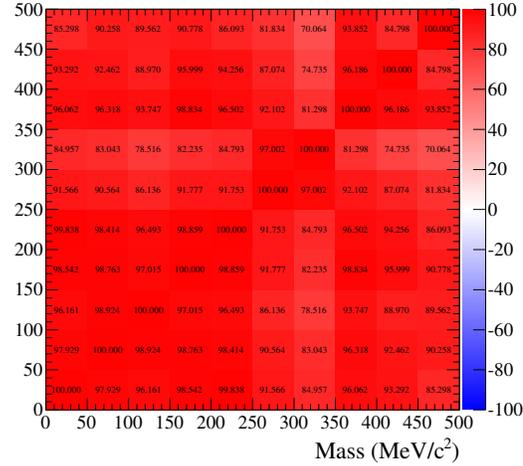
(c) Correlation matrix (%).

Figure E.2: Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD1BrBr event topology.



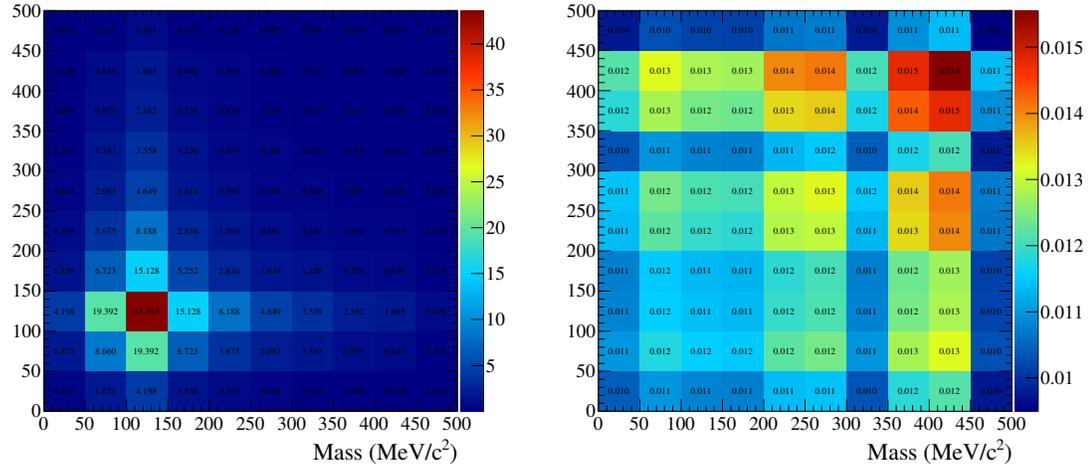
(a) Covariance matrix.

(b) Fractional covariance matrix.



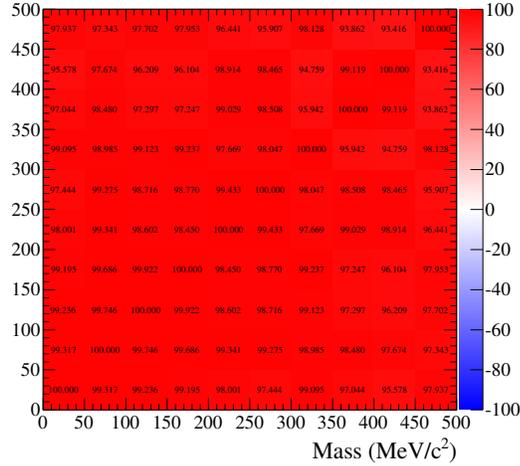
(c) Correlation matrix (%).

Figure E.3: Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD2DsDs event topology.



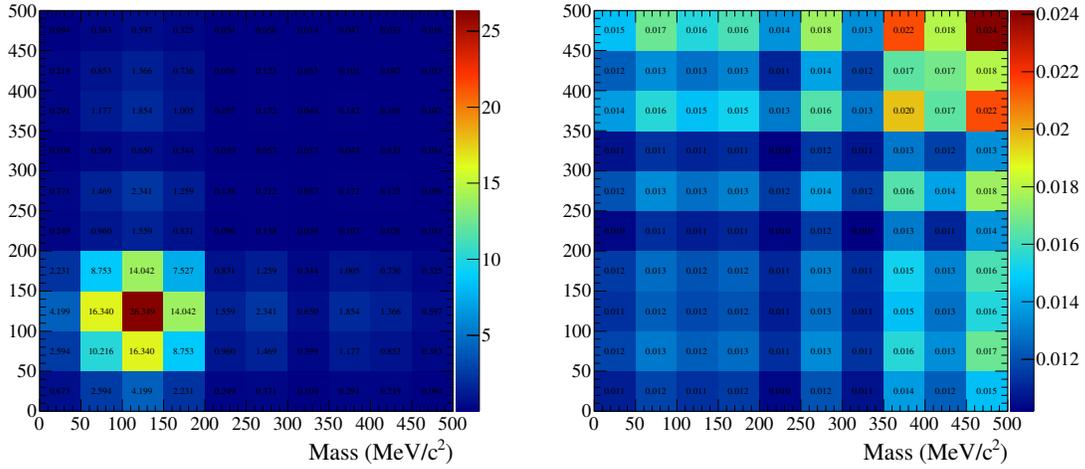
(a) Covariance matrix.

(b) Fractional covariance matrix.



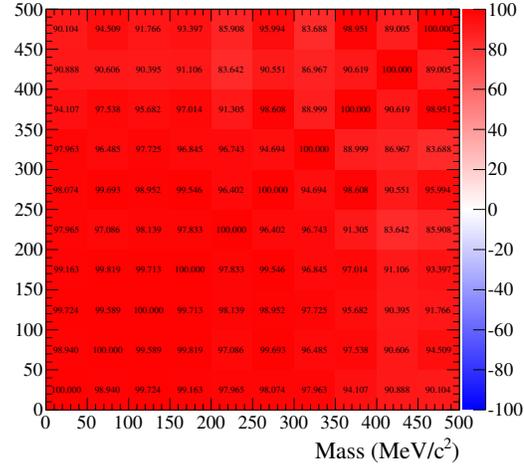
(c) Correlation matrix (%).

Figure E.4: Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD2DsBr event topology.



(a) Covariance matrix.

(b) Fractional covariance matrix.



(c) Correlation matrix (%).

Figure E.5: Covariance and Correlation matrices formed during the evaluation of the flux systematic error for the FGD2BrBr event topology.