Neutrino induced neutral current single π^0 production at the near detector of the T2K experiment

Leon James Pickard

A thesis submitted to the University of Sheffield for the Degree of
Doctor of Philosophy in Physics

Supervisor: Prof. Lee F. Thompson

July 2016

Abstract

T2K is a long baseline neutrino oscillation experiment located at the J-PARC facility in eastern Japan. The neutrino beam has a peak energy of ~0.60 GeV and is directed 2.5° off-axis, passing through the near detector (ND280), 280 m from the source, towards Super-Kamiokande, a 50 kt water Cherenkov detector, 295 km away. One of the main goals of T2K is to measure $\nu_\mu \rightarrow \nu_e$ oscillations and so determine the value of θ_{13}. The first indication of a non-zero value was found in June 2011 and the first measurement in November 2013.

The analysis presented in this thesis is a ν_μ induced neutral current single π^0 (NC1π^0) cross-section measurement which utilises the tracking detector region of ND280. This is of importance to T2K as NC1π^0 interactions represent the second largest uncertainty in the ν_e appearance measurement - behind only the intrinsic ν_e component of the beam. Furthermore, it is also of use to the wider neutrino community as few such measurements exist in the field. Thus, studies of this nature are invaluable in testing and improving cross-section models.

This analysis considers 5 distinct measurements with different topologies of interest and compares the total cross-sections, alongside the $d\sigma/dp_{\pi^0}$ and $d\sigma/dcos\theta_{\pi^0}$ distributions, to the Monte Carlo prediction. T2K Run 2, Run 3 and Run 4 data taking periods are used. The target consists of a combination of C_8H_8 and H_2O, with an average target atomic mass of 13.2. The main result of interest is the NC1π^0 inclusive cross-section which was found to be 2.145 ± 0.109 (stat) ± 0.284 (syst) $\times 10^{-40}$ cm^2/nucleon (when unfolded with respect to momentum) and 2.037 ± 0.110 (stat) ± 0.273 (syst) $\times 10^{-40}$ cm^2/nucleon (when unfolded with respect to $\cos\theta$), compared to a predicted value of 3.087 $\times 10^{-40}$ cm^2/nucleon - a discrepancy of 3.097σ and 3.567σ respectively. Furthermore, for all topologies, the data shows a deficit to that predicted, with the largest and smallest discrepancies of any topologies found to be 4.230σ and 1.412σ respectively.

The algorithm for the initial stage of the timing calibration for the Electromagnetic Calorimeters, Pi-Zero Detector, and Side Muon Range Detector is presented in this thesis. The algorithm attempts to locate, and apply offsets for, 10 ns shifts in the hit timing - induced during resynchronisation of the clock modules. An emphasis is placed on the final step of this algorithm as it was developed by the author.

Acknowledgements

First and foremost I would like to thank my supervisor Prof. Lee Thompson for his invaluable guidance in directing my work over the past four years, alongside providing excellent support and numerous opportunities. However, the area in which I offer most thanks is that it was Lee who introduced me to Daiso, without which I would not know where to buy a nose massager or indeed forearm warmers.

Further gratitude must be given to Dr. Jon Perkin, without his continuous support and computing wizardry I would not have an analysis to present. More importantly, I would not have returned home had it not been for our "Winter Roman Holiday". Furthermore, I am extremely grateful to Dr. Susan Cartwright who has provided much insight and advice.

Working in Sheffield has been a thoroughly enjoyable experience, made all the more fun by sharing an office with a great group of colleagues. Although there are too many to name, they all shared my passion for a delicious cup of tea. Travelling to and from Japan would not have been the same without the Sheffield group, Dr. Matthew Lawe, Dr. Callum "Cazza-Man-W" Wilkinson, Patrick "P-Dawg" Stowell and Mike Wallbank, and my life would have far fewer anecdotes.

Many more collaborators of T2K deserve thanks, however particular appreciation must go to Dr. Helen O'Keeffe, Prof. Clark McGrew and the π^0 group as a whole, who have helped me greatly throughout my PhD.

I am extremely grateful to all my friends and family, in particular my parents, who have encouraged me to do anything I have ever wanted to do. Special thanks must go to Hannah, as her constant support, and ability to proof read a thesis she does not understand in order to make it appear that I have the slightest grasp of the English language, has been admirable.

Contents

List of Tables

List of Figures

List of Abbreviations

BDT - Boosted decision tree.

Brl - Barrel.

BrlECal - Barrel ECal.

CC - Charged current.

CTM - Cosmic trigger module.

DIS - Deep inelastic scattering.

Ds - Downstream.

DsECal - Downstream ECal.

ECal - Electromagnetic calorimeter.

ES - Elastic scattering.

FGD - Fine grained detector.

FSI - Final state interaction.

MC - Monte Carlo.

MCM - Master clock module.

MPPC - Multi-pixel photon counter.

MVA - Multivariate analysis.

NC - Neutral current.

P0D - Pi-zero detector.

P0DECal - Pi-zero detector ECal.

PCA - Principal component analysis.

PEU - Pixel equivalent unit.

PID - Particle identification.

POCA - Point of closest approach.

POT - Protons on target.

QES - Quasi-elastic scattering.

ROC - Receiver operating characteristic.

RMM - Readout merger module.

SCM - Slave clock module.

SMRD - Side muon range detector.

SVD - Singular value decomposition.

TFB - Trip-t front end board.

TPC - Time projection chamber.

TRTT - Time with respect to trigger.

Chapter 1

Neutrino Theory

1.1 Neutrino Background

The neutrino was first postulated by Wolfgang Pauli in 1930 [1] in an attempt to understand the energy distribution of β-decay[1]. However, it was not until 1953 that it was first discovered [2], the subsequent confirmation coming in 1956 [3], by Frederick Reines and Clyde Cowan through analysing the inverse β-decay,

$$\bar{\nu}_e + p^+ \rightarrow e^+ + n, \tag{1.1}$$

driven by electron anti-neutrinos produced by nuclear reactors. In 1962, Leon Lederman, Melvin Schwartz and Jack Steinberger, using the Brookhaven Alternating Gradient Synchrotron, discovered the ν_μ by impinging protons on a beryllium target to produce pions which decay to produce neutrinos [4]:

$$\pi^\pm \rightarrow \mu^\pm + \overset{(-)}{\nu}_\mu. \tag{1.2}$$

The discovery of the τ lepton in 1974 [5], and its confirmation in 1977 [6], produced the first evidence for a third neutrino flavour - this evidence was further compounded when studies of τ decays showed missing energy and momenta. In 1989, following measurements of the Z^0 decay width at LEP and the Stanford Linear Collider, see Figure 1.1, it was found that there must indeed be three active neutrino species, each with a mass less than half of the Z^0. The elusive ν_τ was eventually observed by the DONUT collaboration in 2000 [8]. This experiment utilised a mixed flavour neutrino beam and searched for charged tracks emerging from a charged current interaction vertex, with the ν_τ being inferred by observing a track to have a kink - where the outgoing τ had decayed to another charged particle and multiple unobserved particles.

---

[1] Initially Pauli named this hypothetical particle the 'neutron', however in 1932 it was renamed the 'neutrino' by Enrico Fermi as James Chadwick used 'neutron' for his newly discovered particle.

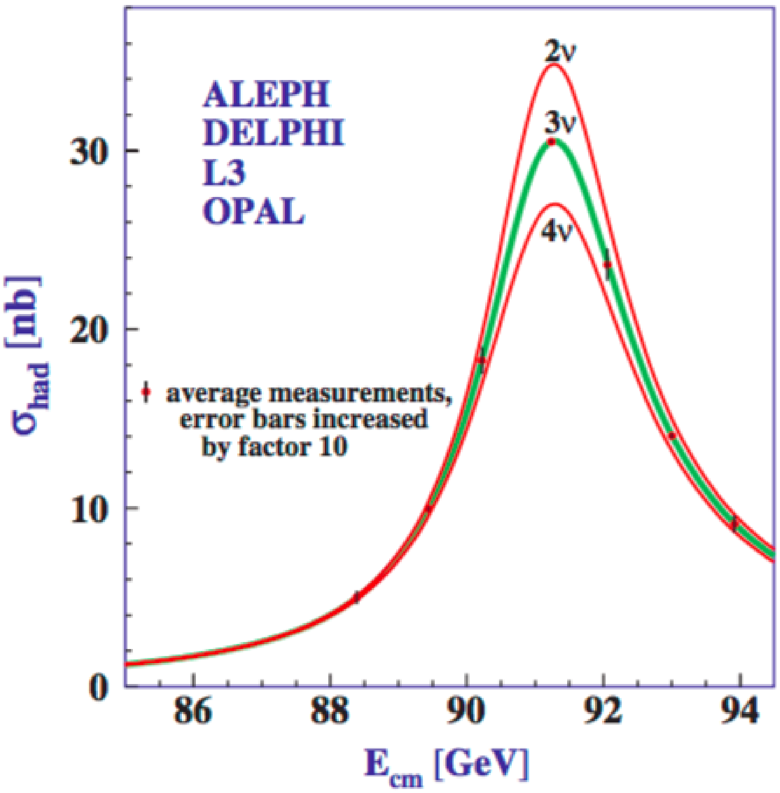

Figure 1.1: The effect of the number of neutrino species on the hadron production cross-section, in e^+e^- collisions, as a function of the centre of mass energy. The data agrees to a high level of precision with the 3-flavour scenario. Figure taken from [7].

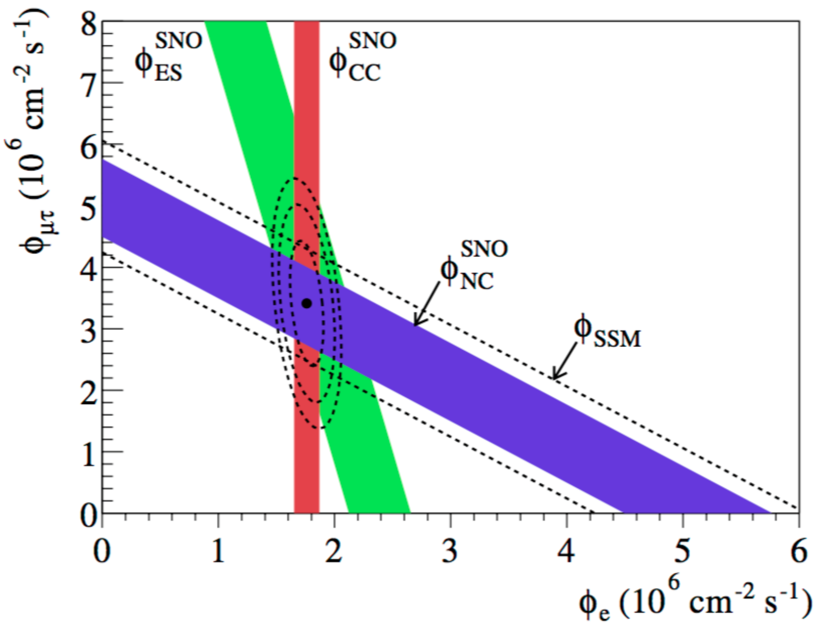

Figure 1.2: Results published by the SNO collaboration, showing the neutrino flux, from the sun, measured in the elastic scattering (ES), charged current (CC) and neutral current (NC) channels. The NC channel is sensitive to all three neutrino flavours and agrees perfectly with the SSM flux prediction. The CC channel is only sensitive to ν_e's whilst the ES channel is sensitive to ν_e's and, partially, to ν_μ's and ν_τ's. Figure taken from [18].

1.2 Neutrino Oscillations

Evidence for the phenomenon of neutrino oscillations was initially observed in 1998 by Super-Kamiokande [9], 41 years after first being proposed by Bruno Pontecorvo in 1957 [10].[2] This followed several decades of uncertainty resulting from the 'solar neutrino problem' and the 'atmospheric neutrino anomaly'. The solar neutrino problem first arose in 1967 when Ray Davis, using the Homestake experiment, found a deficit in the number of ν_e's emitted from the Sun [11] - approximately one third the number expected from the standard solar model (SSM) prediction [12]. This was later confirmed by SAGE [13], GALLEX [14], Kamiokande [15] and Super-Kamiokande [16], all showing a deficit of approximately 40-60%. It was not until 2001 when the SNO experiment, with its ability to observe flavour-independent neutrino interactions, was able to provide conclusive evidence for neutrino oscillations, and as such, that the deficit could be explained specifically as $\nu_e \rightarrow \nu_\mu/\nu_\tau$ oscillations [17], see Figure 1.2.[3]

The atmospheric neutrino anomaly arose when the IMB [19], MACRO [20], Soudan2 [21], Kamiokande [22] and Super-Kamiokande [23] experiments, studying cosmic-ray in-

---

[2] Pontecorvo presented a framework for neutrino anti-neutrino oscillation.

[3] In fact, Takaaki Kajita and Arthur B. McDonald shared the 2015 Nobel Prize in physics for the discovery of neutrino oscillations at Super-Kamiokande and SNO respectively.

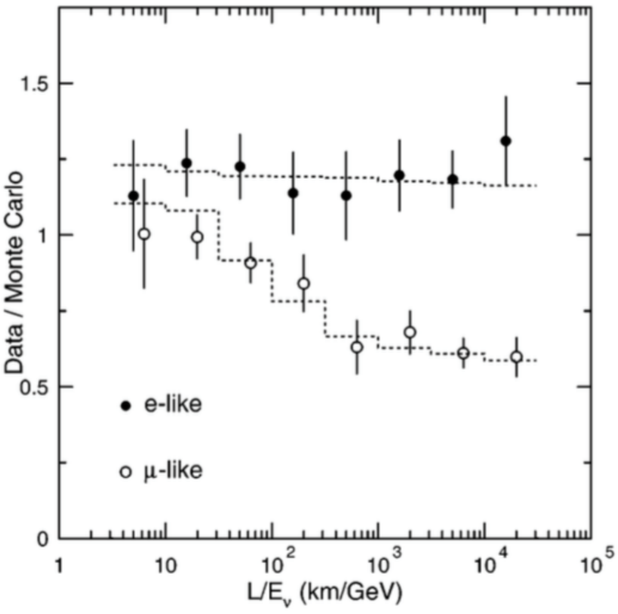

Figure 1.3: The data/MC ratio for e-like and μ-like events at Super-Kamiokande. Note, the dashed line represents the predicted shape under the assumption of ν_μ and ν_τ oscillations in which $\Delta m^2 = 2.2{\times}10^{-3}$ and $\sin^2 2\theta = 1$. Figure taken from [25].

duced neutrinos, found a deficit in the measured ratio of ν_μ/ν_e compared to the predicted value of $\sim$2. This effect was again found to be due to neutrino oscillations as Kamiokande and Super-Kamiokande (in 1994 and 1998 respectively) produced evidence for $\nu_\mu \to \nu_\tau$ oscillations [24]. This was achieved by demonstrating that μ-like events were dependent on the baseline, L, and neutrino energy, E_ν, whilst showing this not to be the case for e-like events - see Figure 1.3.

1.3 Oscillation Theory

Pontecorvo's work on the oscillation of neutrinos to anti-neutrinos was developed by Ziro Maki, Masami Nakagawa and Shoichi Sakata to propose flavour oscillations [26]. The premise is that the neutrino mass and flavour eigenstates are not related by a diagonal mixing matrix but are instead a non-diagonal matrix. Hence the flavour eigenstates are a linear superposition of the mass eigenstates:

$$|\nu_\alpha\rangle = \sum_i U^*_{\alpha i} |\nu_i\rangle \,. \tag{1.3}$$

Analogously the mass states are a linear superposition of the flavour states:

$$|\nu_i\rangle = \sum_{\alpha} U_{\alpha i} |\nu_\alpha\rangle \,. \tag{1.4}$$

This leads to the following 'PMNS' mixing matrix for the three neutrino flavours:

$$U_{\alpha i} \equiv \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & e^{-i\delta}s_{13} \\ -s_{12}c_{23} - e^{-i\delta}c_{12}s_{23}s_{13} & c_{12}c_{23} - e^{-i\delta}s_{12}s_{23}s_{13} & s_{23}c_{13} \\ s_{12}s_{23} - e^{-i\delta}c_{12}c_{23}s_{13} & -c_{12}s_{23} - e^{-i\delta}s_{12}c_{23}s_{13} & c_{23}c_{13} \end{pmatrix} \tag{1.5}$$

where $c_{ij} = \cos\theta_{ij}$, $s_{ij} = \sin\theta_{ij}$ and δ is the charge-parity (CP) violating phase. The PMNS matrix can be decomposed into its three axial rotations: the atmospheric, solar and reactor/accelerator sectors - so named due to the sources used to measure the parameters in each:

$$U_{\alpha i} \equiv \overbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}^{\text{Atmospheric Sector}} \overbrace{\begin{pmatrix} c_{13} & 0 & e^{-i\delta}s_{13} \\ 0 & 1 & 0 \\ -e^{-i\delta}s_{13} & 0 & c_{13} \end{pmatrix}}^{\text{Reactor/Accelerator Sector}} \overbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}^{\text{Solar Sector}} . \tag{1.6}$$

Due to the superposition of mass and flavour states, neutrinos can be shown to oscillate. First consider the wave function of a neutrino, the flavour state of which has been evolved to any position and time:

$$|\Psi(x,t)\rangle = \sum_{i} U^*_{\alpha i} |\nu_i\rangle \, e^{-i\mathbf{p}_\nu \cdot \mathbf{x}_\nu} \tag{1.7}$$

where $\mathbf{p}_\nu$ and $\mathbf{x}_\nu$ are the neutrino 4-momentum and 4-position respectively. By assuming the neutrino mass is negligible compared to its momentum, the following approximation can be made:

$$E_i \simeq |\mathbf{p}_\nu| + \frac{m_i^2}{2|\mathbf{p}_\nu|} \tag{1.8}$$

where E_i is the energy, m_i is the mass and $c = \hbar = 1$. This leads to the following wavefunction:

$$|\Psi(x,t)\rangle = \sum_{i} U^*_{\alpha i} |\nu_i\rangle \, e^{-i\frac{m_i^2}{2|\mathbf{p}_\nu|}} . \tag{1.9}$$

When substituting Eq. 1.4 into Eq. 1.9 the wavefunction is described by the flavour eigenstate:

$$|\Psi(x,t)\rangle = \sum_i \sum_\beta U^*_{\alpha i} e^{-i\frac{m_i^2}{2\mathbf{p}_\nu}x} U_{\beta i} \, |\nu_\beta\rangle \, . \tag{1.10}$$

The probability of a neutrino, of initial flavour α, to propagate a distance x, and be observed in flavour state β is:

$$P(\alpha \to \beta) = | \, \langle \nu_\alpha | \nu_\beta(x) \rangle \, |^2 \tag{1.11}$$

$$P(\alpha \to \beta) = \left[\sum_j U_{\alpha j} e^{-i\frac{m_j^2}{2\mathbf{p}_\nu}x} U^*_{\beta j} \right] \left[\sum_i U^*_{\alpha i} e^{i\frac{m_i^2}{2\mathbf{p}_\nu}x} U_{\beta i} \right] \tag{1.12}$$

$$P(\alpha \to \beta) = \sum_{i,j} U^*_{\alpha i} U_{\alpha j} U_{\beta i} U^*_{\beta j} e^{i\frac{m_i^2 - m_j^2}{2\mathbf{p}_\nu}x} . \tag{1.13}$$

At the relativistic limit, when using a ν_μ beam (such as at T2K), the disappearance probability is:

$$P(\nu_\mu \to \nu_\mu) \simeq 1 - \cos^4\theta_{13} \sin^2 2\theta_{23} \sin^2 \left[1.27 \frac{\Delta m^2_{32} L}{E} \right] \tag{1.14}$$

where $\Delta m^2_{32} = m^2_3 - m^2_2$ is the mass squared difference in eV^2, L is the distance travelled in km and E is the energy in GeV. Similarly, the ν_e appearance probability is:

$$P(\nu_\mu \to \nu_e) \simeq \sin^2 2\theta_{13} \sin^2\theta_{23} \sin^2 \left[1.27 \frac{\Delta m^2_{13} L}{E} \right] . \tag{1.15}$$

Hence, a neutrino, initially of one flavour, can undergo flavour change, the probability of which is dependent on the respective values in the PMNS matrix, the mass squared differences, the distance travelled and the energy of the neutrino.

1.4 Current State of Neutrino Physics

Currently, the best constraints on θ_{12} and Δm^2_{21} are provided by solar experiments and KamLAND [28]. Super-Kamiokande [29], MINOS [30] and T2K [31] provide the leading measurements for θ_{23} and $|\Delta m^2_{32}|$. Until 2012 it was not known if θ_{13} was greater than zero, however it is now the most accurately measured mixing angle, with the world leading results being provided by the reactor experiments: Daya Bay [32] and RENO [33] alongside

Table 1.1: T2K measurements and methods used [31, 35] compared to the global averages [27].

Measurement	**Value**	**Method**	**Global Average**
θ_{13}	$\sin^2 2\theta_{13} = 0.140^{+0.038}_{-0.032}$ (NH) $\sin^2 2\theta_{13} = 0.170^{+0.045}_{-0.037}$ (IH)	ν_e Appearance	$\sin^2 2\theta_{13} = (9.8\pm0.8)\times10^{-2}$
θ_{23}	$\sin^2\theta_{23} = 0.514^{+0.055}_{-0.056}$ (NH) $\sin^2\theta_{23} = 0.511\pm0.055$ (IH)	ν_μ Disappearance	$\sin^2\theta_{23} = 0.437^{+0.033}_{-0.023}$ (NH) $\sin^2\theta_{23} = 0.455^{+0.039}_{-0.031}$ (IH)
$\|\Delta m^2_{32}\|$	$(2.51\pm0.10)\times10^{-3}\text{eV}^2$ (NH) $(2.38\pm0.06)\times10^{-3}\text{eV}^2$ (IH)	ν_μ Disappearance	$(2.43\pm0.06)\times10^{-3}\text{eV}^2$ (NH) $(2.38\pm0.06)\times10^{-3}\text{eV}^2$ (IH)
$\bar{\theta}_{23}$	$\sin^2\bar{\theta}_{23} = 0.45^{+0.19}_{-0.07}$	$\bar{\nu}_\mu$ Disappearance	See Figure 1.6
$\|\Delta\bar{m}^2_{32}\|$	$\|\Delta\bar{m}^2_{32}\| = (2.51^{+0.29}_{-0.25})\times10^{-3}\text{eV}^2$	$\bar{\nu}_\mu$ Disappearance	See Figure 1.6

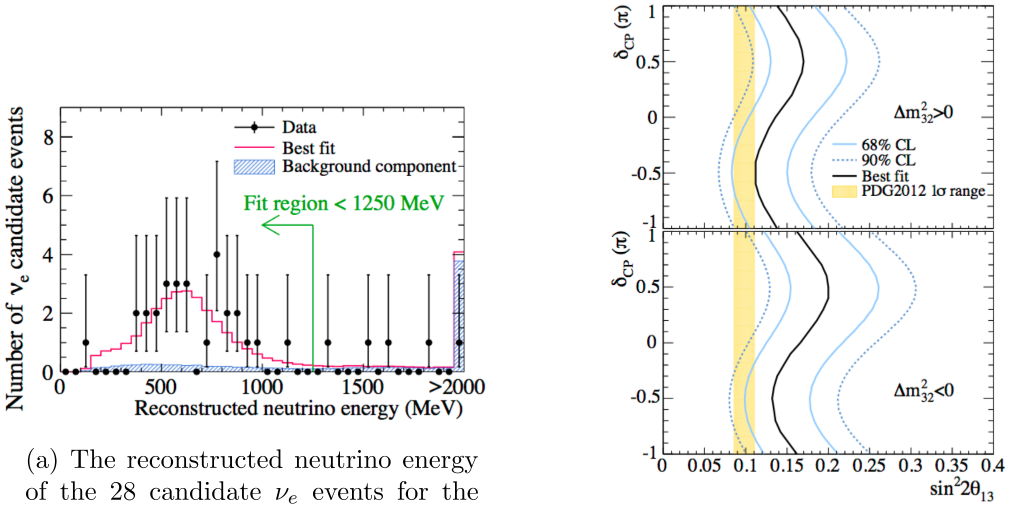

(a) The reconstructed neutrino energy of the 28 candidate ν_e events for the T2K appearance measurement, used to determine θ_{13}.

(b) The best fit value of $\sin^2\theta_{13}$ for the T2K ν_e appearance measurement in the case of normal (top) and inverted (bottom) hierarchies.

Figure 1.4: T2K appearance measurement results [36].

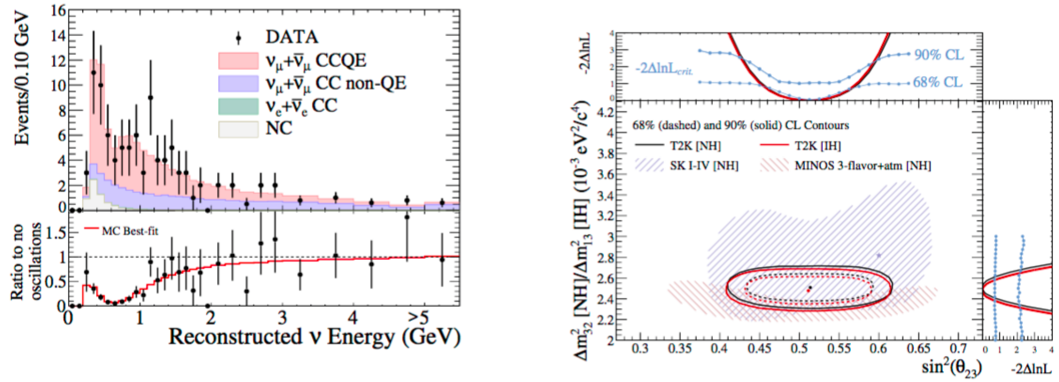

(a) The reconstructed neutrino energy of the candidate ν_μ events for the T2K disappearance measurement (top) and the MC best fit used to determine θ_{23} (bottom).

(b) The best fit values of $\sin^2\theta_{23}$ and Δm^2_{32} for normal (black) and inverted (red) hierarchies.

Figure 1.5: T2K disappearance measurement results [37].

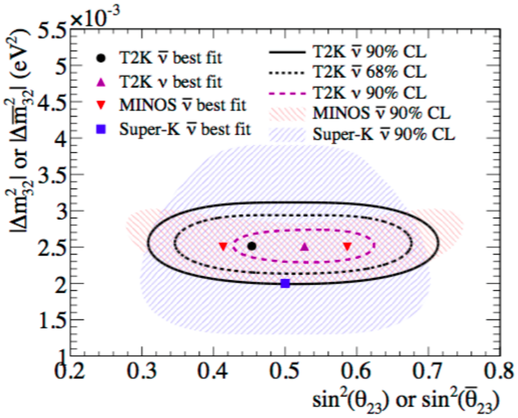

Figure 1.6: The best fit values for $\sin^2\bar{\theta}_{23}$ and $|\Delta\bar{m}^2_{32}|$ for the T2K $\bar{\nu}_\mu$ disappearance result, compared to analagous measurements by MINOS and Super-Kamiokande and the T2K ν_μ disappearance result. Note, all values are in agreement [35].

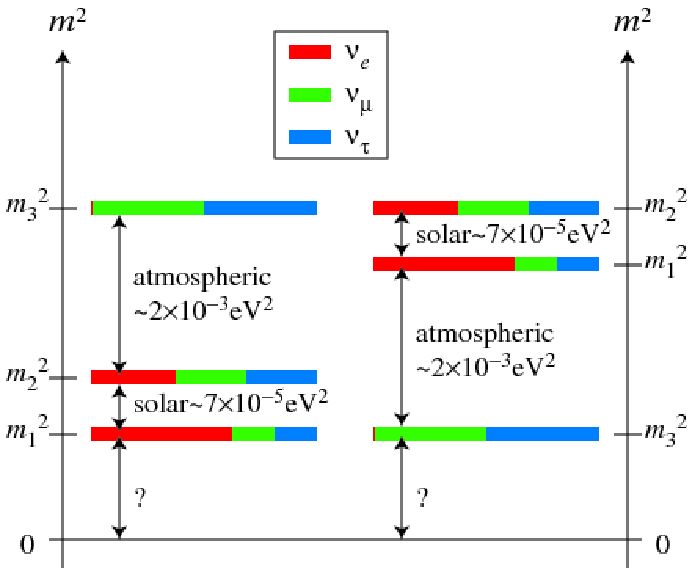

Figure 1.7: The two possible neutrino mass hierarchies. The normal hierarchy, shown on the left, exists in nature if $m_{\nu_3} > m_{\nu_2} > m_{\nu_1}$, whilst the inverted hierarchy exists if $m_{\nu_2} > m_{\nu_1} > m_{\nu_3}$. Figure taken from [38].

the long baseline experiments: T2K [31] and NOvA [34]. For the latest results produced by T2K and the global averages, see Table 1.1 and Figures 1.4, 1.5 and 1.6.

Despite extensive efforts in the field of neutrino physics, many open questions are still apparent. Most pertinent to T2K are: the octant problem, determining the neutrino mass hierarchy and the value of the CP complex phase. The octant problem exists due to the uncertainty of θ_{23}, which allows the possible value to reside in either the $>45^\circ$ or $<45^\circ$ octant.[4] The neutrino mass hierarchy structure, whether it is normal ($m_1 \approx m_2 \ll m_3$) or inverted ($m_3 \ll m_1 \approx m_2$), is not known as Δm^2_{21} has been found to be positive, whilst the sign of Δm^2_{31} has not been identified, see Figure 1.7. Numerous current (T2K, NOvA, MINOS and MINOS+ [39]) and future experiments (Hyper-Kamiokande [40], DUNE [41] and PINGU [42]) have capabilities to either shed light on or definitively solve this problem. In the case of the long baseline experiments, the MSW effect is exploited in which the propagation through matter alters the oscillation probabilities - increasing the $\nu_\mu \rightarrow \nu_e$ probability for the normal hierarchy and reducing it for the inverted.[5] Finally, the CP complex phase, which could help explain the matter-antimatter asymmetry within the universe (provided the other Sakharov conditions are met [43]), is possible to infer with the current generation of experiments [39] and should be possible to measure using the next generation [40, 41, 42]. In the case of the long baseline experiments, by comparing $P(\nu_\mu \rightarrow \nu_e)$ to $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ any discrepancy will help to elucidate the value of δ.

In addition to the open problems highlighted that T2K can begin to explore, many other questions still exist. Although accurate measurements of the mass squared differences have been produced for the three known neutrino mass states, the absolute masses of the neutrinos are still unknown. Several experiments are attempting to measure these, however experimental uncertainty is proving a big hinderance. According to cosmological models, the absolute mass of the three neutrino flavours combined must be less than 0.23 eV [44], however, the world leading measurements cannot compete with this. The most precise limits for the $\bar{\nu}_e$ mass are produced by the Troitsk and Mainz experiments [27]. By studying tritium β-decay they have set this to be <2.05 eV and <2.3 eV respectively, both with a 95% confidence level.[6] Mass measurements of the ν_μ have also been performed by analysing charged pion-decay, with a world leading limit of <190 keV at a 90% confidence level [45]. The ν_τ has the least stringent limit at <18.2 MeV with a 95% confidence level [45].

Furthermore, the neutrino may be Dirac ($\nu \neq \bar{\nu}$) or Majorana ($\nu = \bar{\nu}$) in nature. However, this cannot be inferred from long baseline techniques, nonetheless current experiments such as EXO [46], GERDA [47] and KamLAND-Zen [48] are attempting to search for neutrinoless double-β decay, a signal only possible if the neutrino is Majorana. Furthermore, other experiments, including SNO+ [49], MAJORANA [50] and SuperNemo [51], are soon to begin taking data. Unfortunately, no conclusive detection has currently been found, though a positive result has been claimed by the Heidelberg-Moscow experi-

---

[4] In fact θ_{23} may be maximal, with a value of 45°.

[5] This effect is easier to measure at larger baselines. At 295 km the baseline of T2K produces a 10% discrepancy compared to the vacuum probability, whilst at 810 km NOVA sees a 30% difference.

[6] After 5 years of running the KATRIN experiment should have a sensitivity of 0.2 eV. Additionally, if the mass is >0.3 eV a 5σ measurement is possible [45].

ment [52].

A further open question regarding neutrinos surfaced due to the 1996 LSND experiment which, when searching for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations, found a signal consistent with a mass squared difference of 1 eV^2 [53]. This, in conjunction with the LEP results suggests that one (or possibly more) sterile neutrinos should exist. To try and resolve this issue KARMEN and MiniBooNE were designed with a similar L/E to LSND, (with MiniBooNE studying both neutrino and anti-neutrino oscillations). Initially, KARMEN found no agreement with the LSND result [54], however, for the $\bar{\nu}_e$ appearance channel MiniBooNE found an excess between 475 MeV and 1250 MeV, ruling out a portion, but not all, of the LSND phase space. Conversely, when considering the ν_e appearance channel, no excess in this energy range was found. However, an excess at 200 MeV to 475 MeV emerged [53].[7] Although attempts are still being made to understand all these results within a 3+1 sterile neutrino model, such disagreement suggests that at least 2 sterile neutrinos may exist [53].

1.5 Neutrino-Nucleon Interactions

Neutrino-nucleon interactions is a heavily researched area of particle physics that has gained particular momentum in the last decade due to its importance in the high precision era of neutrino physics. They can be divided into charged current (CC) and neutral current (NC) categories. A CC process occurs when a neutrino exchanges a $W^{\pm}$ boson, in doing so a charged lepton of the same flavour as the incident neutrino is produced. Conversely, NC interactions, upon exchanging a Z^0 boson, preserve the original neutrino.

At the energy range of T2K, CC and NC categories can be further subdivided into three dominant interactions each. The first, when considering CC processes, is quasi-elastic scattering (QES):

$$\nu_l \; n \rightarrow l^- \; p \tag{1.16}$$

$$\bar{\nu}_l \; p \rightarrow l^+ \; n \tag{1.17}$$

where n represents a neutron and p a proton. The analogous process for the NC channel is elastic-scattering (ES):

$$\nu \; p \rightarrow \nu \; p \tag{1.18}$$

$$\nu \; n \rightarrow \nu \; n \tag{1.19}$$

$$\bar{\nu} \; p \rightarrow \bar{\nu} \; p \tag{1.20}$$

$$\bar{\nu} \; n \rightarrow \bar{\nu} \; n. \tag{1.21}$$

---

[7] Later the ν_μ disappearance measurement was improved by combining SciBooNE data.

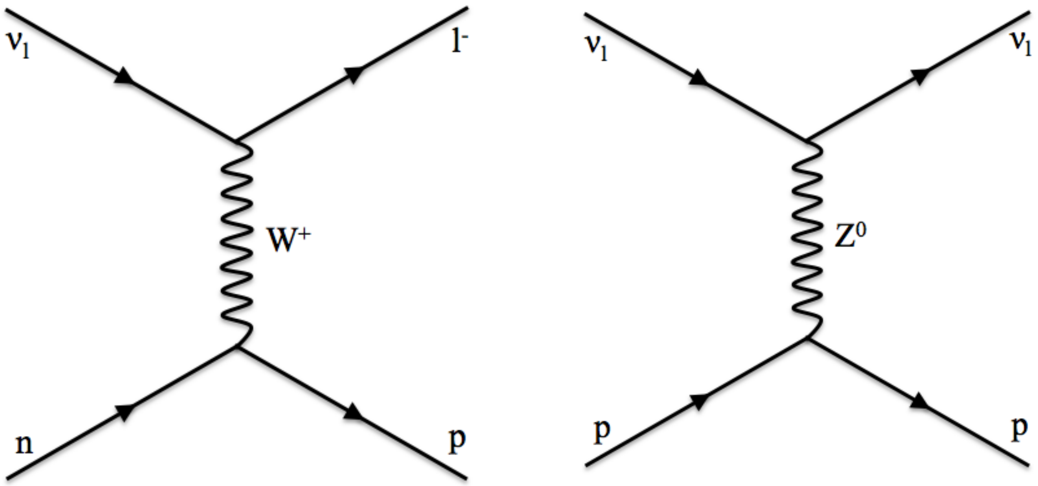

(a) A QES interaction on a neutron.

(b) An ES interaction on a proton.

Figure 1.8: Example Feynman diagrams of QES and ES interactions.

Example Feynman diagrams for QES and ES processes can be found in Figure 1.8. Below 2 GeV (Q)ES is the dominant interaction and can result in the liberation of one or multiple nucleons from the nuclear environment.

The second dominant interaction mode is resonant production, here the neutrino excites a nucleon into a resonant state - e.g. a Δ baryon or a higher energy state of the nucleon, N^*. The resonant particle can decay into many states, producing multiple possible combinations of nucleons and mesons. In the case of CC interactions, the most dominant are:

$$\nu_l \; p \rightarrow l^- \; p \; \pi^+ \tag{1.22}$$
$$\bar{\nu}_l \; p \rightarrow l^+ \; p \; \pi^- \tag{1.23}$$
$$\nu_l \; n \rightarrow l^- \; p \; \pi^0 \tag{1.24}$$
$$\bar{\nu}_l \; p \rightarrow l^+ \; n \; \pi^0 \tag{1.25}$$
$$\nu_l \; n \rightarrow l^- \; n \; \pi^+ \tag{1.26}$$
$$\bar{\nu}_l \; n \rightarrow l^+ \; n \; \pi^- \tag{1.27}$$

whilst for NC:

$$\nu_l \; p \rightarrow \nu_l \; p \; \pi^0 \tag{1.28}$$
$$\bar{\nu}_l \; p \rightarrow \bar{\nu}_l \; p \; \pi^0 \tag{1.29}$$

$$\nu_l\ p \rightarrow \nu_l\ n\ \pi^+ \tag{1.30}$$

$$\bar{\nu}_l\ p \rightarrow \bar{\nu}_l\ n\ \pi^+ \tag{1.31}$$

$$\nu_l\ n \rightarrow \nu_l\ n\ \pi^0 \tag{1.32}$$

$$\bar{\nu}_l\ n \rightarrow \bar{\nu}_l\ n\ \pi^0 \tag{1.33}$$

$$\nu_l\ n \rightarrow \nu_l\ p\ \pi^- \tag{1.34}$$

$$\bar{\nu}_l\ n \rightarrow \bar{\nu}_l\ p\ \pi^-. \tag{1.35}$$

The third dominant interaction is deep-inelastic scattering (DIS) which becomes more significant at the higher energies produced at T2K ($>\sim$GeV). At these larger energies the neutrino can resolve the individual quarks of the nucleons, producing baryonic showers.

In addition to the above interactions coherent production of pions can occur[8], though this has a much smaller cross-section. In contrast to resonant production, here the neutrino coherently scatters off the entire nucleus with negligible four-momentum transfer. Consequently, the entire, un-fragmented nucleus recoils and remains in the same state. Furthermore, a distinctly forward going pion produced:

$$\nu_l\ N \rightarrow \nu_l\ N\ \pi^0 \tag{1.36}$$

$$\bar{\nu}_l\ N \rightarrow \bar{\nu}_l\ N\ \pi^0. \tag{1.37}$$

1.6 Neutrino-Induced NC1π^0 Interactions

This analysis considers NC single π^0 production within the Fine Grain Detector (FGD) region of the ND280 detector (see Chapter 2 for a detailed description). Specifically, the topology of interest is where the π^0 decay photons convert in the Electomagnetic Calorimeters (ECals).[9] This measurement has implications for both T2K and neutrino physics as a whole. With regards to T2K, NC1π^0 interactions contribute 22% of the background for the ν_e appearance signal. This background arises since photons and electrons produce almost identical Cherenkov rings at Super-Kamiokande. As such, if either one of the π^0 decay photons is not reconstructed or if both photons are collinear - producing two highly overlapping rings - the decay could be mistaken for an electron. In this scenario a NC1π^0 event could be mistaken as a ν_e induced QES event. The main production mechanisms for NC1π^0 interactions at T2K are Δ-resonances produced from neutrons and protons with a small contribution via coherent interactions, see Figure 1.9.

In addition to reducing the systematic uncertainty associated with ν_e appearance, measuring the absolute NC1π^0 cross-section will allow improvements to be made to MC

---

[8] Although rare at T2K energy ranges, kaons can also be produced.

[9] The π^0 has an average lifetime of 8.52×10^{-17} s, and with a branching ratio of 98.9%, produces two back-to-back photons in its rest frame.

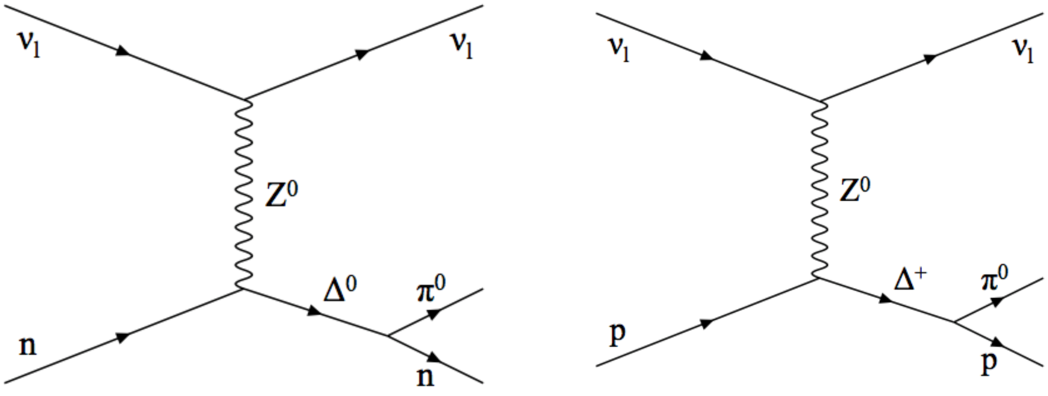

(a) Resonant π^0 production on a neutron.

(b) Resonant π^0 production on a proton.

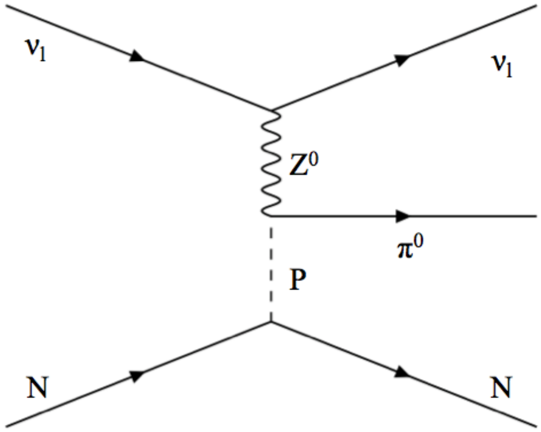

(c) Coherent π^0 production on a nucleus. Here ‘P’ represents a pomeron.

Figure 1.9: Feynman diagrams of the three neutral current single π^0 production mechanisms at T2K.

simulations alongside providing additional data to constrain T2K's flux and cross-section models, the results of which are fed into the T2K oscillation analysis.

Less pertinent to T2K, this measurement provides a useful channel to compare with existing results from a range of experiments, thus improving understanding of neutrino nucleon interactions. Most important are the NC1π^0/CC inclusive and NC1π^0/CC1π ratios as historically these have been the easiest measurements to make and so provide additional data, invaluable for testing existing models.

1.7 Current and Future Measurements

Due to the challenges in measuring NC processes in neutrino experiments, few results exist. The difficulty arises as the interaction can occur anywhere in the target medium, compounded by the fact that the NC channel does not produce a charged lepton, which would enable the reconstruction of a vertex. The NC1π^0 topology provides additional complications in that the only charged particles produced are protons. Often they do not exit the nuclear environment due to final state interaction (FSI) effects, or in the case where they do, they may have too low an energy to be reconstructed.

Between 1974 (the earliest publication) and the turn of the century, only seven experiments produced measurements - often only in the form of NC/CC ratios to constrain systematic uncertainties. The earliest NC1π^0 results not represented as a ratio were obtained in 1978 by the Gargamelle bubble chamber experiment with neutrino interactions on $C_3H_8CF_3Br$ and were followed 5 years later in 1983 at Aachen, where both neutrino and antineutrino interactions on Al were analysed, see Figure 1.10.

Since 2005, and the necessity to fully understand neutrino-nucleus cross-sections for precision neutrino oscillation measurements, six experiments have produced results. In fact MiniBooNE has been able to produce absolute cross-section measurements (of both neutrinos and antineutrinos on CH_2) and also $\frac{d\sigma}{dp_{\pi^0}}$ and $\frac{d\sigma}{d\cos\theta_{\pi^0}}$ distributions (see Figure 1.11).

A summary of all the experimental results that are not coherent-exclusive can be found in Table 1.2. In the case of the pre-K2K measurement the only like-for-like comparison has been highlighted in Figure 1.10, with the remaining measurements representing various different cross-section ratios. Furthermore, although some of the later experiments measure similar cross-sections to Gargamelle and Aachen, an equal comparison is not possible due to the enhancements in the analytical techniques, such as the consideration of FSI effects - for more information see [27, 56]. In contrast, the coherent results, found in Table 1.3, can often be compared, as shown in Figure 1.12 which presents the MINOS coherent $\nu_\mu NC1\pi^0$ cross-section to seven other experiments [72].

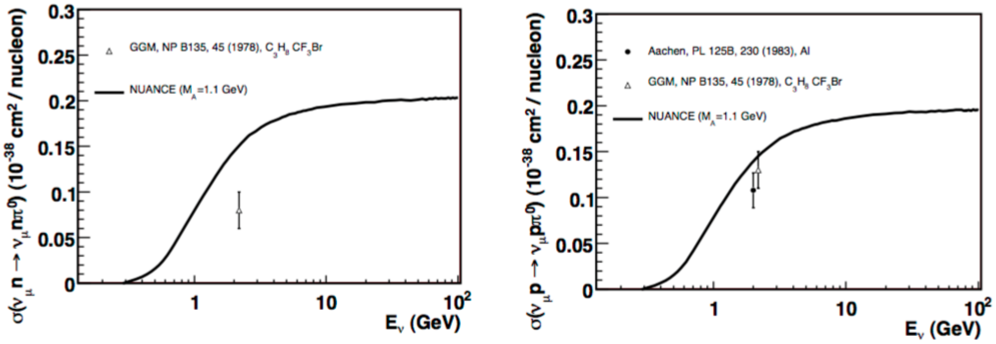

(a) The Gargamelle [57] ν_μ-induced NC1π^0 cross-section on neutrons.

(b) The Gargamelle [57] and Aachen [58] ν_μ-induced NC1π^0 cross-sections on protons.

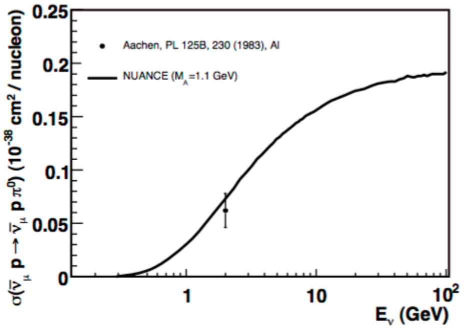

(c) The Aachen [58] $\bar{\nu}_\mu$-induced NC1π^0 cross-section on protons.

Figure 1.10: Early measurements of the absolute NC1π^0 cross-sections compared to the NUANCE prediction proposed by [55]. Figures taken from [56].

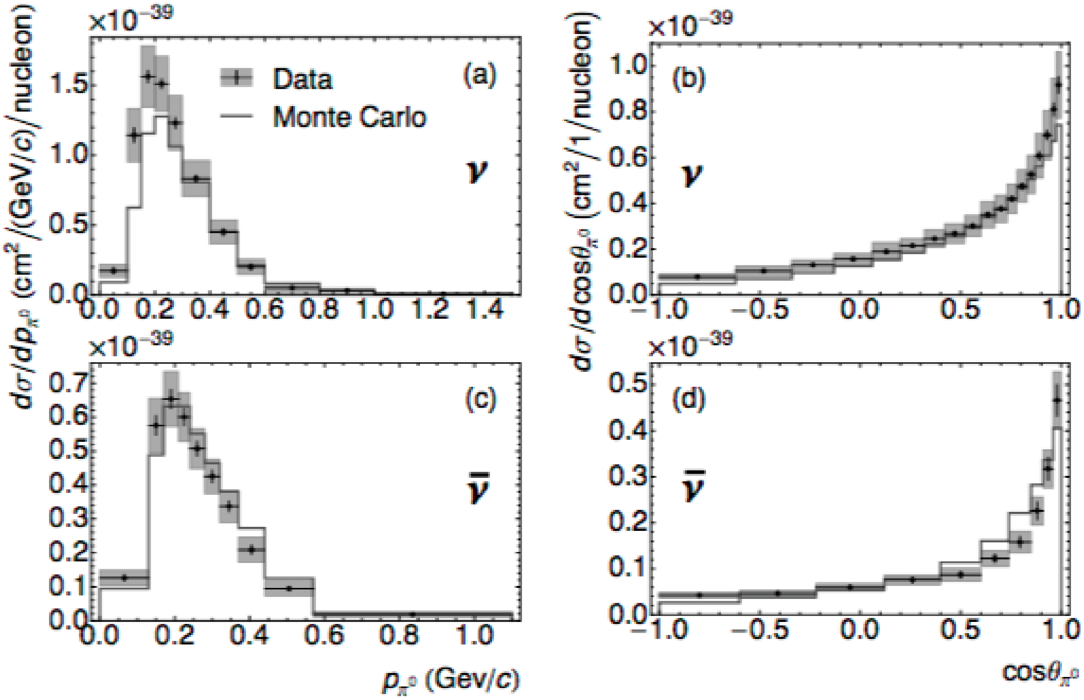

Figure 1.11: The differential NC1π^0 cross-section measurements produced by MiniBooNE. Here, (a) and (b) are the ν_μ-induced $\frac{d\sigma}{dp_{\pi^0}}$ and $\frac{d\sigma}{d\cos\theta_{\pi^0}}$ cross-sections respectively whilst (c) and (d) are the analogous $\bar{\nu}_\mu$ induced measurements. Figures taken from [63].

Table 1.2: Current NCπ^0 measurements. Note, when one energy value is given, it is the peak value of the neutrino energy spectrum.

Experiment	Target	Energy	Measurement	Result	Reference
ANL	H_2	E(ν_μ) = 0.2 - 2.5 GeV	$\frac{\sigma(\nu_\mu p \to \nu_\mu p \pi^0)}{\sigma(\nu_\mu p \to \mu^- p \pi^+)}$	0.09±0.05	[59]
Gargamelle	$C_3H_8CF_3Br$	E(ν_μ) $\simeq$ 1 - 8 GeV	$\frac{\sigma(\nu_\mu N \to \nu_\mu N \pi^0)}{\sigma(\nu_\mu n \to \mu^- p \pi^+)}$	0.45±0.08	[57]
			$\sigma(\nu_\mu n \to \nu_\mu n \pi^0)$	See Figure 1.10a	
			$\sigma(\nu_\mu p \to \nu_\mu p \pi^0)$	See Figure 1.10b	
BNL	Al	E(ν_μ) $\simeq$ 1 - 7 GeV	$\frac{\sigma(\nu_\mu N \to \nu_\mu N \pi^0)}{2\sigma(\nu_\mu N \to \mu^- p \pi^0)}$	0.17±0.04	[60]
			$\frac{\sigma(\nu_\mu N \to \nu_\mu N \pi^0)}{2\sigma(\nu_\mu N \to \mu^- p \pi^0)}$	0.25±0.09	[61]
Aachen	Al	E(ν_μ) = 2 GeV	$\sigma(\nu_\mu p \to \nu_\mu p \pi^0)$	See Figure 1.10b	[58]
		E($\bar{\nu}_\mu$) = 2 GeV	$\sigma(\bar{\nu}_\mu p \to \bar{\nu}_\mu p \pi^0)$	See Figure 1.10c	
K2K	C_8H_8	E(ν_μ) = 1.3 GeV	$\frac{\sigma(\nu_\mu NC\pi^0)}{\sigma(CC)}$	0.065±0.001±0.007	[62]
MiniBooNE	CH_2	E(ν_μ) = 808 MeV	$\sigma(\nu_\mu NC\pi^0)$	$(4.76 \pm 0.05 \pm 0.76) \times 10^{-40}\text{cm}^2/\text{nucleon}$	[63]
		E($\bar{\nu}_\mu$) = 664 MeV	$\frac{d\sigma(\nu_\mu NC\pi^0)}{dP_\pi}$	See Figure 1.11	
			$\frac{d\sigma(\nu_\mu NC\pi^0)}{d\text{cos}\theta_\pi}$	See Figure 1.11	
			$\sigma(\bar{\nu}_\mu NC\pi^0)$	$(1.48 \pm 0.05 \pm 0.23) \times 10^{-40}\text{cm}^2/\text{nucleon}$	
			$\frac{d\sigma(\bar{\nu}_\mu NC\pi^0)}{d\text{P}_\pi}$	See Figure 1.11	
			$\frac{d\sigma(\bar{\nu}_\mu NC\pi^0)}{d\text{cos}\theta_\pi}$	See Figure 1.11	
SciBooNE	C_8H_8	E(ν_μ) = 0.8 GeV	$\frac{\sigma(\nu_\mu NC\pi^0)}{\sigma(CC)}$	$(7.7 \pm 0.5^{+0.4}_{-0.5}) \times 10^{-2}$	[64]
ArgoNeut	Ar	E(ν_μ) = 9.6 GeV	$\frac{\sigma(\nu_\mu NC\pi^0)}{\sigma(CC)}$	$0.094 \pm 0.022 \pm 0.015$	[65]
		E($\bar{\nu}_\mu$) = 3.6 GeV	$\frac{\sigma(\bar{\nu}_\mu NC\pi^0)}{\sigma(CC)}$	$0.042 \pm 0.022 \pm 0.008$	
			$\sigma(\nu_\mu NC\pi^0)$	$(7.1 \pm 1.7 \pm 1.3) \times 10^{-40}\text{cm}^2$	
			$\sigma(\bar{\nu}_\mu NC\pi^0)$	$(0.5 \pm 0.2 \pm 0.1) \times 10^{-40}\text{cm}^2$	

Table 1.3: Current Coherent NC1π^0 measurements. Table adapted from [66].

Experiment	Target	Mean Energy	Measurement	Result	Reference
Aachen	Al	$E(\nu_\mu) = 2$ GeV	$\sigma(NC \text{ Coherent } \pi^0)$	$29 \pm 10 \times 10^{-40} \text{cm}^2/\text{nucleus}$	[58]
Gargamelle	C	$E(\nu_\mu) = 2$ GeV	$\sigma(NC \text{ Coherent } \pi^0)$	$31 \pm 20 \times 10^{-40} \text{cm}^2/\text{nucleus}$	[68]
CHARM	$CaCO_3$	$E(\nu_\mu) = 30$ GeV	$\sigma(NC \text{ Coherent } \pi^0)$	$96 \pm 42 \times 10^{-40} \text{cm}^2/\text{nucleus}$	[69]
SKAT	CCl_2F_2	$E(\nu_\mu) = 7$ GeV	$\sigma(NC \text{ Coherent } \pi^0)$ $\frac{\sigma(NC \text{ Coherent } \pi^0)}{\sigma \nu_\mu CC}$	$79 \pm 28 \times 10^{-40} \text{cm}^2/\text{nucleus}$ $4.3 \pm 1.5 \times 10^{-3}$	[70]
15’ BC	NeH_2	$E(\nu_\mu) = 20$ GeV	$\frac{\sigma(NC \text{ Coherent } \pi^0)}{\sigma \nu_\mu CC}$	$0.20 \pm 0.04 \times 10^{-3}$	[71]
NOMAD	C	$E(\nu_\mu) = 24.8$ GeV	$\sigma(NC \text{ Coherent } \pi^0)$ $\frac{\sigma(NC \text{ Coherent } \pi^0)}{\sigma \nu_\mu CC}$	$(72.6 \pm 8.1 \pm 6.9) \times 10^{-40} \text{cm}^2/\text{nucleus}$ $3.21 \pm 0.46 \times 10^{-3}$	[66]
MiniBooNE	CH_2	$E(\nu_\mu) = 808$ MeV	$\sigma(NC \text{ Coherent } \pi^0)$	$(7.7 \pm 1.6 \pm 3.6) \times 10^{-40} \text{cm}^2/\text{CH}_2$	[67]
SciBooNE	C_8H_8	$E(\nu_\mu) = 0.8$ GeV	$\frac{\sigma(NC \text{ Coherent } \pi^0)}{\sigma(NC\pi^0)}$ $\frac{\sigma(CC \text{ Coherent } \pi^+)}{\sigma(NC \text{ Coherent } \pi^0)}$	$(1.16 \pm 0.24) \times 10^{-2}$ $0.14^{+0.30}_{-0.28}$	[64]
MINOS	Fe	$E(\nu_\mu) = 4.9$ GeV	$\sigma(NC \text{ Coherent } \pi^0)$	$77.6 \pm 5.0 \text{ (stat) } \pm^{15.0}_{16.8} \text{ (syst) } \times 10^{-40} \text{cm}^2/\text{nucleus}$	[72]

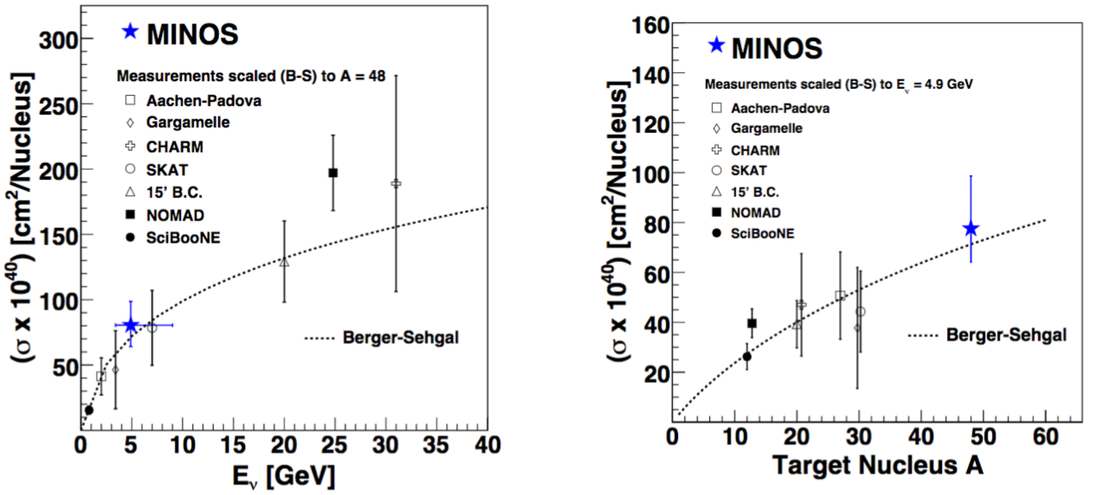

(a) Results scaled to match the MINOS target mass number.

(b) Results scaled to match the MINOS neutrino energy.

Figure 1.12: A comparison of the various coherent $NC1\pi^0$ cross-sections, compared to the Berger-Sehgal model (see Section 1.8) [75], scaled to match the MINOS target medium mass number (1.12a) and energy (1.12b). Figures taken from [72].

1.8 Cross-Section Models

The Monte Carlo (MC) generator used in the analysis presented in Chapter 5 is NEUT [73]. NEUT implements the Rein-Sehgal model [74] to simulate both resonant and coherent pion production. In the case of the former, the model assumes baryonic resonances dominate pion production. In total, 18 baryonic resonances (with an invariant mass less than 2 GeV) are included.[10] Regions where resonances overlap are taken into account, however the final lepton mass is assumed to be zero. To predict the NC1π^0 cross-section, NEUT calculates the amplitude of production from all resonances and multiplies each by the probability of the baryon to decay to produce a π^0.

In the case of coherent π^0 production (which has a predicted cross-section of ~1% of that of the resonant channel), the Rein-Sehgal model is based on the partially conserved axial-vector currents (PCAC) theorem [76] - which describes coherent production cross-sections in terms of pion-nucleus cross-sections. Rein-Sehgal developed this idea to describe pion-nucleon cross-sections in terms of these pion-nucleus cross-sections. Effects of pion absorption in the nucleus are taken into account (as this would result in a non-coherent interaction, by changing the energy state of the nucleus). However, a deficiency in this model is that the four-momentum transfer, Q^2, is assumed to be zero. Although not implemented in NEUT, the Berger-Sehgal model improves upon this. Not only does it not assume a value of zero for Q^2, it also takes advantage of the latest available pion-nucleus

---

[10] The dominant resonance being $\Delta(1232)$.

cross-section data to produce an improved neutrino-nucleon cross-section prediction.

Chapter 2

The T2K Experiment

2.1 Introduction

T2K (Tokai-to-Kamioka) is a long baseline neutrino experiment located in Japan. Its primary aim is to understand neutrino oscillations by determining the parameters of the PMNS mixing matrix and helping to determine the hierarchical nature of neutrinos. Additionally, T2K provides cross-section measurements of neutrinos on matter, further information of which can be found in [77]-[81].

The T2K experiment consists of a neutrino beam produced at the Japan Proton Accelerator Research Centre (J-PARC) at Tokai, a small town found in the Ibaraki Prefecture on the east coast of Japan, see Figure 2.1, 280 m downstream and 2.5° off-axis is the near detector complex, whose primary goal is to measure and understand the flux and composition of the neutrino beam. The beam then travels through 295 km of the earth's crust to the far detector (Super-Kamiokande) where the composition of the oscillated beam is determined.

The current status of T2K's measurements, methods used and a comparison to the

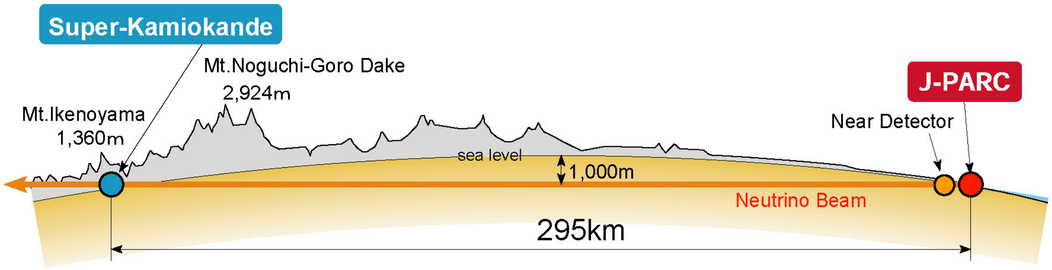

Figure 2.1: Cross-section of the path taken by the T2K neutrino beam [82].

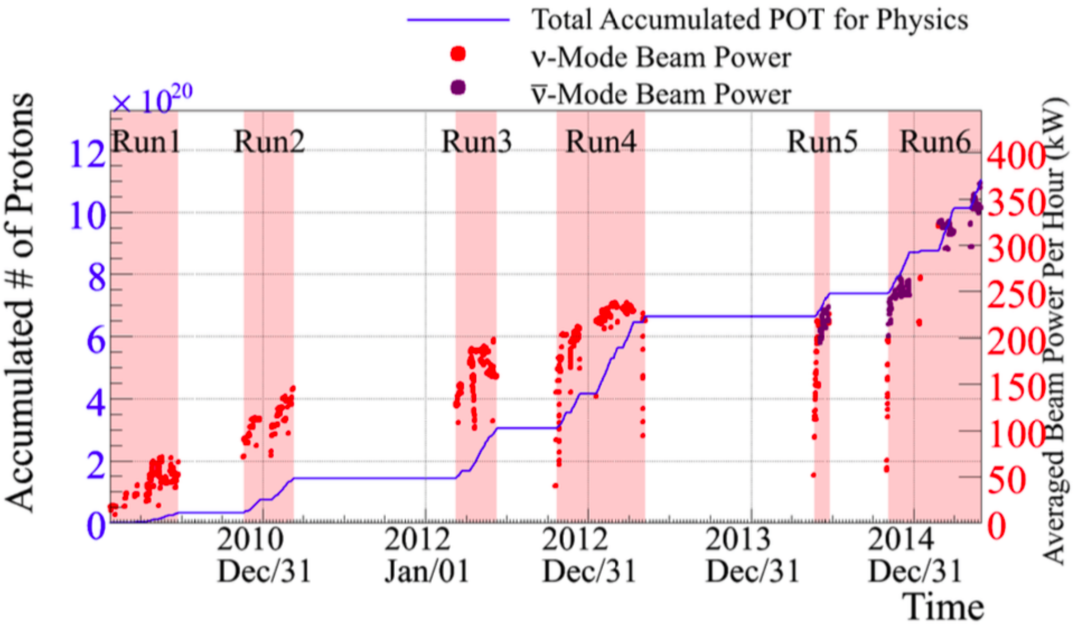

Figure 2.2: The total accumulated POT, beam power and mode of operation delineated by run period [83].

global averages, has been highlighted in Table 1.1. As of Run 6, which concluded in July 2015, T2K has achieved 11.04×10^{20} Protons On Target (POT), of which 7.00×10^{20} are in neutrino mode and 4.04×10^{20} in anti-neutrino mode. A full breakdown of the POT, beam power and mode of operation for each run can be found in Figure 2.2.

2.2 Experimental Setup

2.2.1 Accelerator

The purpose of the accelerator complex [82] is to produce a beam of 30 GeV protons (which, when impinging on a target, produce neutrinos via pion decay). This is achieved using a three stage acceleration procedure at the J-PARC facility. Initially, the Linear Accelerator (LINAC) accelerates H^- anions to 180 MeV. Using charge-stripping foils, the anions are stripped of electrons to produce protons. These enter the Rapid-Cycling Synchrotron (RCS) and are accelerated, 2 bunches at a time, to 3 GeV at an operating frequency of 25 Hz. 5% of bunches are passed to the Main Ring (MR), a synchrotron capable of accelerating protons to 50 GeV (however, it is currently running at 30 GeV). The main ring produces 8 bunches of protons (for Run 1 a six bunch structure was used) each of width 58 ns with 581 ns windows between. The narrow beam width, interspersed with large windows, in addition to accurate timing, allows for the study of cosmic events - useful for detector calibration. After reaching 30 GeV, all 8 bunches (constituting a spill) are diverted at the fast extraction point towards the neutrino beamline by 5 kicker magnets, a process taking approximately 5μs.

2.2.2 Beamline

The T2K beamline [82] has two main sections, the primary and secondary beamline, each of which are further subdivided into 3 sections, see Figure 2.3. The purpose of the primary beamline is to aim the beam towards Super-Kamiokande. To do this, the incoming protons enter the 57 m long preparation section where they are aligned by 11 conduction magnets. They then pass into the 147 m long arc section, which has a 104 m radius of curvature. This section utilises superconducting combined-function magnets to steer the protons through 80.7°, horizontally aligning them with Super-Kamiokande. To prevent dispersion of the beam within this section, horizontal and vertical superconducting steering magnets are employed. Upon entering the final focusing section, 10 conducting magnets perform the final vertical steering - a 3.637° declination.

Throughout the primary beamline, constant monitoring takes place to ensure the highest-quality, most-efficient beam possible is produced. This is achieved using the following systems:

- 5 Current Transformers (CTs) which measure, to an accuracy of 2%, the intensity of the beam
- 21 Electrostatic Monitors (ESMs) which measure, to an accuracy of 450 μm, the beam position
- 19 Segmented Secondary Emission Monitors (SSEMs) which measure, to an accuracy of 200 μm, the beam profile
- 5 Beam Loss Monitors (BLMs) which measure, to an accuracy of 16 mW, the beam loss.

At this point protons enter the first section of the secondary beamline, the target station. Here, they pass through the beam window, baffle and the Optical Transition Monitors (OTRs), before impinging on the Proton Target. The OTR consists of titanium alloy films, which produce transition radiation that can be used to provide final monitoring of the beam direction and profile (to an accuracy of < 1 mm and 0.5 mrad respectively).

The Proton Target [82], located within the first magnetic focusing horn, is a 91.4 cm long, 2.6 cm diameter graphite rod ($\rho = 1.8$ gcm^{-3}), encased in a 2 mm thick sleeve and a 0.3 mm thick titanium casing. Due to the extreme temperatures reached, 700 °C when operating with a beam current of 750 kW, He gas is pumped between the target layers to provide cooling. When the protons impinge on the target, mesons are produced[1] and are "collected" by the first focussing horn and "directed" by the second and third. The

---

[1] Predominantly pions are produced, though kaons and also other more exotic particles are also created.

focussing horns have a significant effect, producing a 16-fold increase in the neutrino flux received at Super-Kamiokande. They are manufactured from an aluminium alloy and each consists of a pair of coaxial conductors, producing a toroidal 2.1 T field (when a 320 kA current is applied) that decreases by $1/r$ when deviating from the horn centre. The beauty of this design is that by changing the sign of the current, a predominantly neutrino or anti-neutrino beam can be selected.

Downstream of the target station is the decay volume [82], providing a region for mesons to decay and produce neutrinos. The volume is water cooled (to ensure operating temperatures remain below 100 °C), 96 m long, composed of 16 mm thick steel walls and has a tapered shape with upstream and downstream cross-sections of 1.4×1.7 m^2 and 3.0×5.0 m^2 respectively. Furthermore, it is surrounded by 6 m thick concrete shielding. The main decay producing neutrinos is:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \; (BR = 99.99\%), \tag{2.1}$$

and for anti-neutrinos:

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \; (BR = 99.99\%). \tag{2.2}$$

Additionally, higher energy neutrinos are produced by kaon decay channels:

$$K^+ \rightarrow \mu^+ + \nu_\mu \; (BR = 63.55\%), \tag{2.3}$$

$$K^+ \rightarrow \pi^0 + \mu^+ + \nu_\mu \; (BR = 3.35\%). \tag{2.4}$$

Analogously, for anti-neutrinos:

$$K^- \rightarrow \mu^- + \bar{\nu}_\mu \; (BR = 63.55\%), \tag{2.5}$$

$$K^- \rightarrow \pi^0 + \mu^- + \bar{\nu}_\mu \; (BR = 3.35\%). \tag{2.6}$$

Furthermore, more complicated decay chains can occur, for example:

$$K^+ \rightarrow \pi^+ + \pi^0 \; (BR = 20.67\%), \tag{2.7}$$

followed by the decay given in Eq. 2.1.

Similarly, for anti-neutrinos:

$$K^- \to \pi^- + \pi^0 \ (BR = 20.67\%), \tag{2.8}$$

followed by the decay given in Eq. 2.2.

Unfortunately, the kaon decay channel is less pure than that of the pion; this, alongside muon decay, leads to an approximately 1% inherent ν_e contamination in the ν_μ beam:

$$K^+ \to \pi^0 + e^+ + \nu_e \ (BR = 5.07\%), \tag{2.9}$$

$$\mu^+ \to e^+ + \nu_e + \bar{\nu}_\mu \ (BR \approx 100\%). \tag{2.10}$$

Similarly, the charge conjugate of these decays leads to an approximately 1% $\bar{\nu}_e$ contamination in the $\bar{\nu}_\mu$ beam:

$$K^- \to \pi^0 + e^- + \bar{\nu}_e \ (BR = 5.07\%), \tag{2.11}$$

$$\mu^- \to e^- + \bar{\nu}_e + \nu_\mu \ (BR \approx 100\%). \tag{2.12}$$

Following the decay volume is the beam dump, the purpose of which is to increase the neutrino beam purity by removing hadrons and low energy muons ($\lesssim 5$ GeV). It has two regions of dense material: the upstream region is a water cooled, 75 t graphite mass with dimensions $3.17\times1.94\times4.67$ m^3, and the downstream region consists of 15 Fe plates of total length 2.40 m. Higher energy muons are not stopped and pass through to the Muon Monitor (MUMON) where they are monitored.

2.2.3 Off-Axis Method

To optimise T2K's effectiveness the off-axis method is employed, with the near and far detector complexes deviating from the beam axis by 2.5°. This produces a quasi-monochromatic neutrino source peaking at 0.6 GeV, see Figure 2.4. The energy distribution can be better understood by considering that the majority of muon neutrinos come from pion decay (with the charge conjugate equivalent decay during the production of anti-neutrinos):

$$\pi^+ \to \mu^+ + \nu_\mu. \tag{2.13}$$

By assuming decay along the beam axis, the following neutrino energies are produced (assuming $m_\nu = 0$):

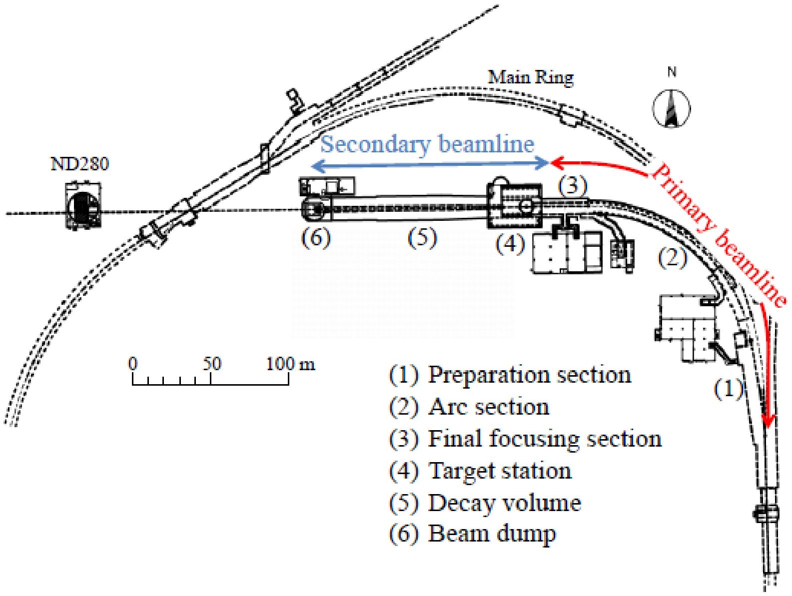

Figure 2.3: The T2K beamline separated into its 6 separate sections [82].

$$E_{\nu} = \frac{m_{\pi}^2 - m_{\mu}^2}{2(E_{\pi} - |\mathbf{p}_{\pi}|\cos\theta)} \tag{2.14}$$

where m is mass, $\mathbf{p}$ is three momentum, θ the angle from the beam axis and E the energy. Thus, when deviating from the neutrino beam axis, a strongly peaked neutrino energy spectrum is produced. This configuration maximises the oscillation probability at 295 km, increasing the sensitivity of the ν_{μ} disappearance measurement, see Figure 2.4. Other benefits of this setup include: a decrease in the proportion of higher energy neutrinos reaching the detectors and a significant reduction in the background of the $\nu_{\mu} \rightarrow \nu_{e}$ measurement (as both the intrinsic ν_{e} contamination within the neutrino beam and NC1π^0 production are reduced).

2.2.4 MUMON

The MUMON [82], located downstream of the beam bump, profiles the beam direction using two detector planes; the former consists of silicon PIN photodiodes and the latter ionisation chambers. These measure the direction of the muons, the majority of which are produced in a 2-body decay with the neutrinos, and as such can be used to produce a "position profile" of the neutrino beam. To ensure high quality data the beam direction must be known to 1 mrad, the accuracy of the MUMON surpasses this by ensuring a 0.25 mrad accuracy alongside a measurement of the beam intensity known to 2.9%.

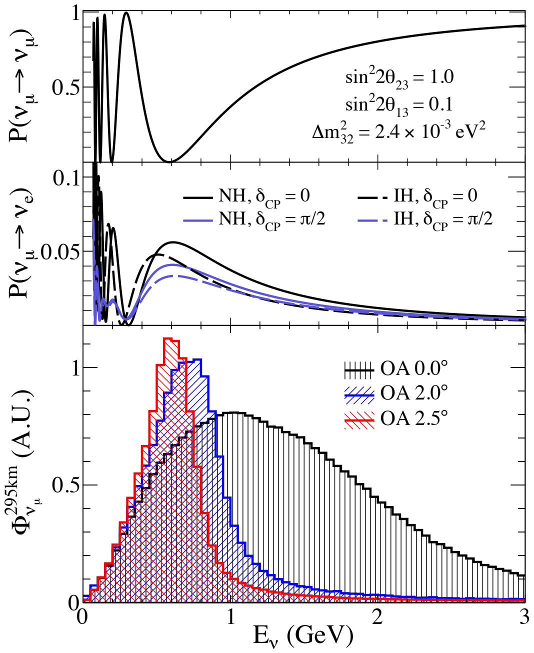

Figure 2.4: The survival probability of ν_μ as a function of energy (top), the probability of detecting a ν_e for different hierarchical structures and varying complex phases (centre) and the neutrino energy flux, in arbitrary units, for varying off-axis angles (bottom). At increasing off-axis angles the total neutrino flux reduces, however an increasingly peaked spectrum is attained. Note, each neutrino flux in this plot is not POT normalised, with the total flux for the 2.0° and 2.5° distributions greatly increased to highlight the peaked energy nature of the beam. Figure taken from [84].

2.2.5 The T2K Near Detector Suite

The T2K near detector suite is 280 m from the beam target and consists of two detectors, the Interactive Neutrino Grid (INGRID) [82, 86] and ND280 [82], which are on-axis and 2.5° off-axis respectively. Both are located in the same cylindrical pit, 17.5 m in diameter and 37 m deep, see Figure 2.5.

2.2.5.1 Scintillator bars and MPPCs

As the majority of the detectors in the T2K near detector suite utilise scintillator bars with a Multi-Pixel Photon Counter (MPPC) readout [82], this section will be dedicated to such.[2] The bars were produced at Fermilab and made of extruded polystyrene, doped with 1% 2,5 diphenyloxade (PPO) and 0.03% 1,4-di-(5-phenyl-2-oxazoyl)-benzene (POPOP). The PPO acts as the primary scintillator and the POPOP the secondary, shifting the wavelength of photons produced by the PPO. Each bar is coated with a thin reflective

---

[2] Although the designs are very similar, the scintillator bars for each detector have different cross-sections. As such, the exact dimensions are given in the relevant detector sections.

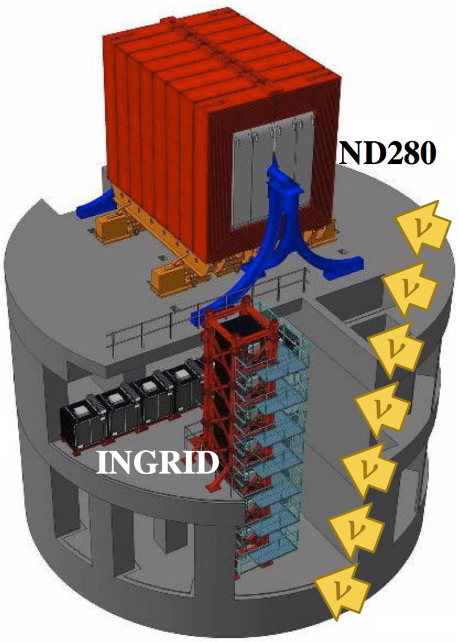

Figure 2.5: The direction of the neutrino beam with respect to the position of INGRID and ND280 [82].

layer of TiO to reduce light loss, thus increasing light collection efficiency - see [85] for a discussion of the light yield studies.

Running down the centre of the bars is a 2×3 mm^2 elliptical hole to allow a 1 mm diameter Wavelength Shifting Fibre (WLS) to be inserted. One end of the fibre has a reflective coating, whilst the other is connected to an MPPC. The absorption spectrum of the fibre peaks at 420 nm (blue light), matching that emitted by the secondary dopant, and re-emits at 476 nm (green light). This difference between the absorption and emission spectra ensures self-absorption is minimised. The MPPC detects light emitted by photons striking any of the 667 pixels, creating a Geiger avalanche and a binary response for that pixel. The final MPPC output is simply the total number of pixels hit, a value proportional to the number of photons impinging on the MPPC. MPPCs were chosen for three main properties: their detection efficiency is well matched to the WLS emission spectra, they are small (1.3×1.3 mm^2) allowing them to be attached to individual fibres, and they are one of the few photodiodes able to work in a 0.2 T magnetic field - essential for ND280.

2.2.6 INGRID

INGRID [86] is a non-magnetised detector consisting of 16 modules (7 vertical, 7 horizontal and 2 off-diagonal), see Figure 2.7. INGRID's primary purpose is to monitor the on-axis beam profile, to which it can measure the beam direction to an accuracy of 0.4 mrad (10 cm). The interesting layout of the modules serves two purposes; the cross pattern allows a wide area of the beam to be profiled whilst using as few modules as possible and

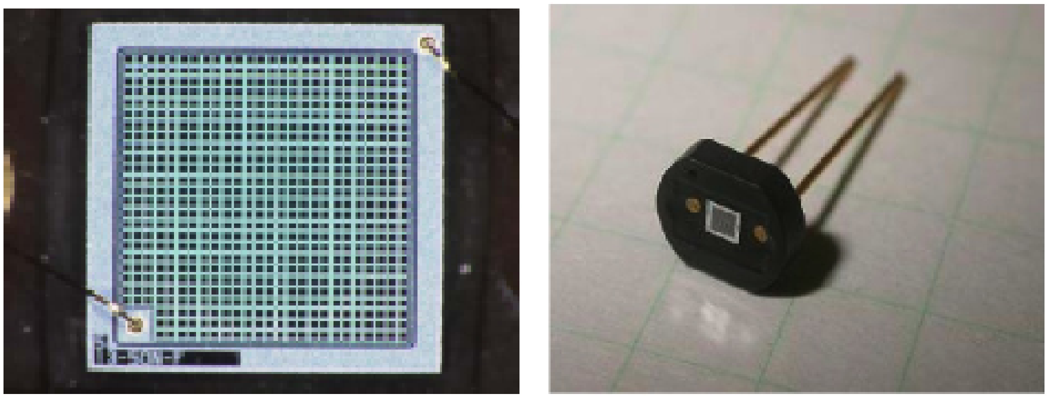

Figure 2.6: A photograph of an MPPC showing the 667 pixel, 1.3×1.3 mm^2 sensitive surface (left) and this surface fitted within its ceramic casing (right) [82].

the 2 off-diagonal modules allow the beam axial symmetry to be determined.

Each of INGRID's 16 modules are identical, and are comprised of alternating layers of plastic scintillator and iron plates, of which there are 11 and 9 respectively (due to weight restrictions there is no iron plate between the 10th and 11th scintillator layer). Each scintillator layer is composed of 48 identical bars, of dimensions 1.0×5.0×20.3 cm^3, 24 in a vertical and horizontal orientation respectively, with each read out by an MPPC. The iron plates are 124×124×6.5 cm^3 with a total per module mass of 7.2 t and provide a large interaction medium.

Above, below and on either side of each module are scintillator veto planes, enabling the rejection of interactions outside the target mass. The layout of the planes varies slightly, each consists of 22 bars with cross-sections of 1.0×5.0 cm^2, however the length of those on the bottom plane are 111.9 cm, and 129.9 cm for the remaining three. Ideally, both the beam intensity and position need to be stable; INGRID proves this to be the case by measuring the day-to-day event rate and beam position constant to 1.7% and 1 mrad (28 cm) respectively.

Situated between the two central modules is the proton module, designed to detect muons and protons produced via CCQE interactions for Monte Carlo validation. Unlike the others, this module has no iron plates and uses smaller scintillator bars to provide a higher granularity.

For an example INGRID event display of a candidate ν_μ event, see Figure 2.8.

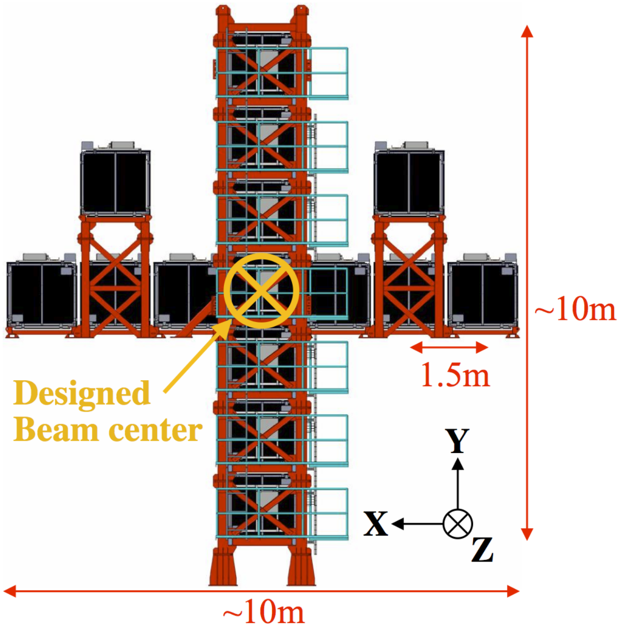

Figure 2.7: The orientation of the 16 modules of INGRID with respect to the designed beam centre [82].

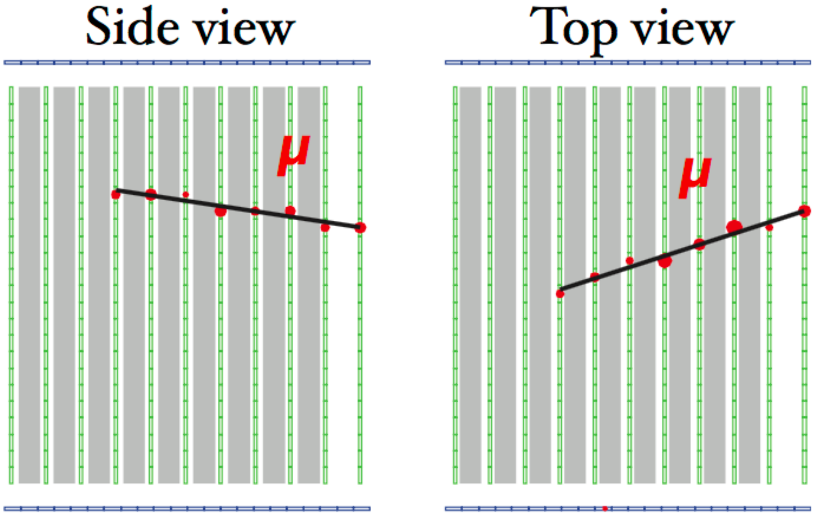

Figure 2.8: An INGRID event display for a candidate ν_μ event. Here, the hits are represented by the red circles and the reconstructed track by the black line. Figure taken from [87].

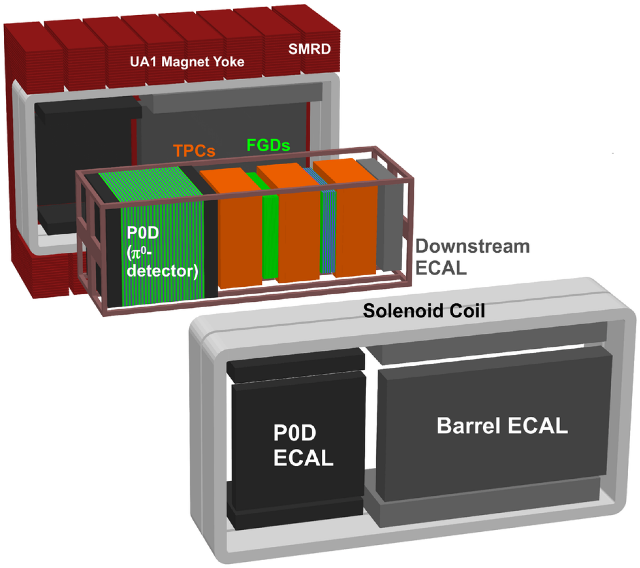

Figure 2.9: An exploded view of the ND280 detector [82]. Note, the neutrino beam travels from left to right in this image.

2.2.7 ND280

The ND280 detector [82] is angled 2.5° off-axis and is designed to measure the flux, flavour content and energy spectrum of the neutrino beam, alongside producing numerous neutrino-nucleus cross-section measurements [77]-[81]. By characterising the beam, some systematic uncertainties at Super-Kamiokande can be measured and reduced [88]. ND280 is 7.6 m long, 5.6 m wide, 6.1 m tall and consists of multiple subdetectors, see Figure 2.9. The most upstream is the Pi-Zero Detector (P0D), which is surrounded by the P0D ECal. Further downstream are three time projection chambers (TPCs) interleaved with 2 fine grain detectors (FGDs), these - alongside the P0D - constitute the tracker region. The TPCs and FGDs are surrounded above, below and on either side by the Barrel ECal (BrlECal) and capped on the downstream end by the Downstream ECal (DsECal). All are situated inside the 0.2 T UA1 magnet, which itself has Side Muon Range Detectors (SMRDs) fitted within the return yoke.

For an example ND280 event display of a simulated CCQE event see Figure 2.10.

2.2.7.1 P0D

The P0D is designed to measure π^0 production cross-sections on water, particularly NC1π^0 production, the second largest background for the ν_e appearance measurement (only behind the flux uncertainty):

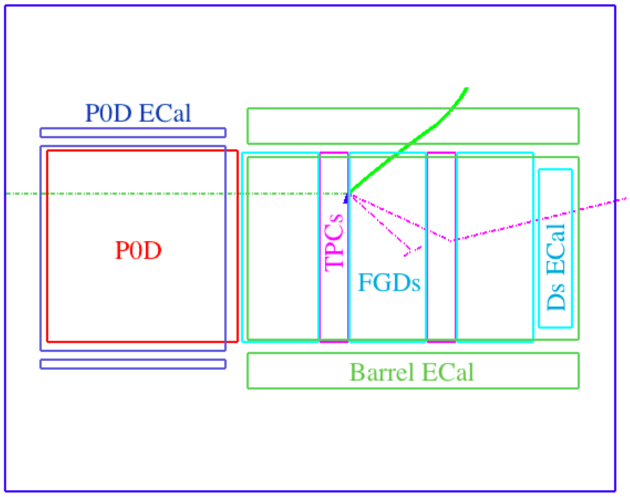

Figure 2.10: An orthographic view of an ND280 event display for a simulated CCQE interaction. The muon-neutrino (dashed green line) enters the detector from the left. Upon interacting, it produces a muon (solid green line) and a proton (solid blue line). Furthermore, two neutrons (dashed pink lines) exit the nucleus.

$$\nu_\mu + N \rightarrow \nu_\mu + N + \pi^0 + X \tag{2.15}$$

where N is a proton or neutron and X is any combination of outgoing baryons and photons.

As can be seen from Figure 2.11, the P0D is composed of layers of scintillator bars, brass, lead and water. The water is held in High Density Polyethylene (HDPE) bags which can be filled or drained to allow on-water and off-water measurements. The scintillator bars are identical to elsewhere in ND280 except all have an isosceles cross-section, $33{\times}17$ mm^2, and a single-ended readout.[3] The scintillator layers are arranged into P0Dules, i.e. two scintillator layers in alternating horizontal and vertical orientations. Each horizontal layer consists of 126 2340 mm long bars, whilst the vertical layer has 134 2200 mm long bars. This setup allows 3D track reconstruction.

The P0D is further split into four sections (from upstream to downstream): the upstream ECal, the upstream water target, the central water target and the central ECal. The upstream water target region consists of a repeating pattern of P0Dule, brass (1.5 mm thick) and HDPE water target bags (28 mm thick), for a total of 13 layers. The central water target is identical to the upstream water target, except it has 12 layers. These two regions provide the target medium for interactions. Both the upstream and central ECal are identical to one another and consist of a repeating pattern of P0Dule and lead (4 mm thick), for a total of 7 layers each. These regions provide veto capabilities for external interactions and improve containment for particles exiting the water target regions. For further information on the P0D see [89].

---

[3] Note, 33 mm is the base of the triangular cross-section and 17 mm the height.

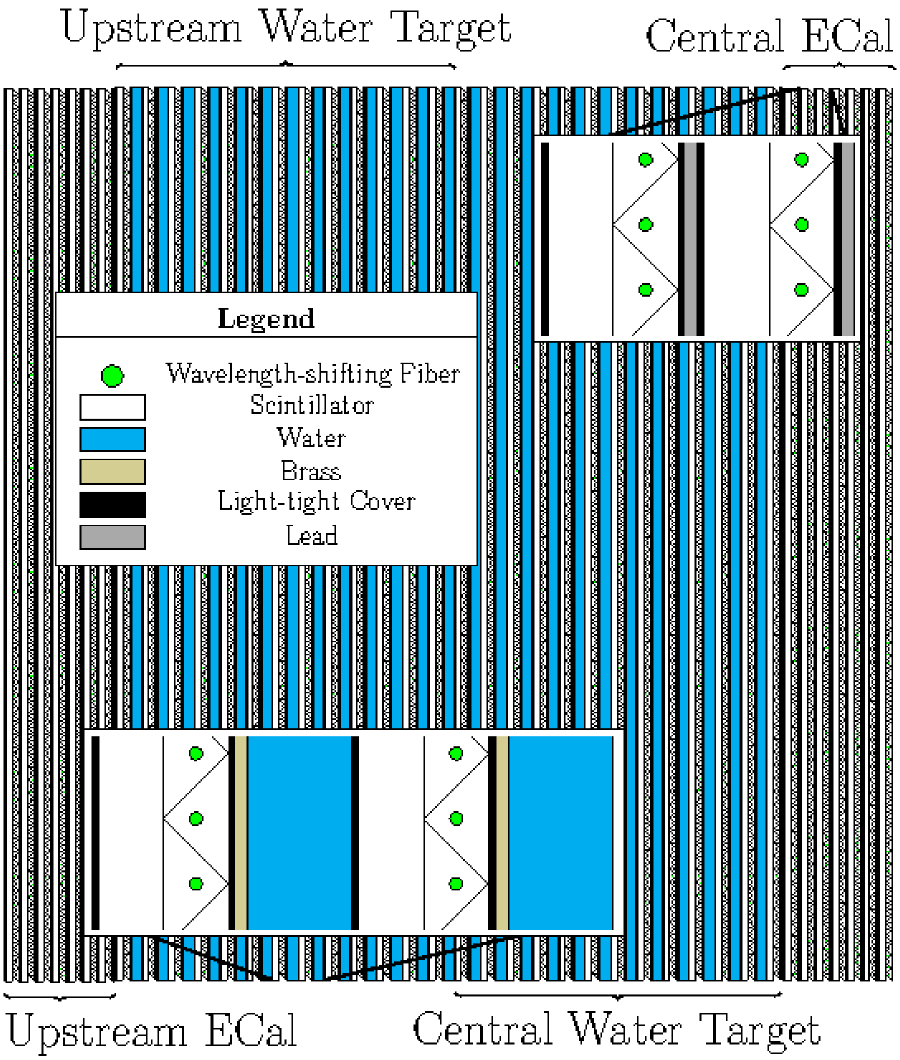

Figure 2.11: The composition of the four regions of the P0D [89].

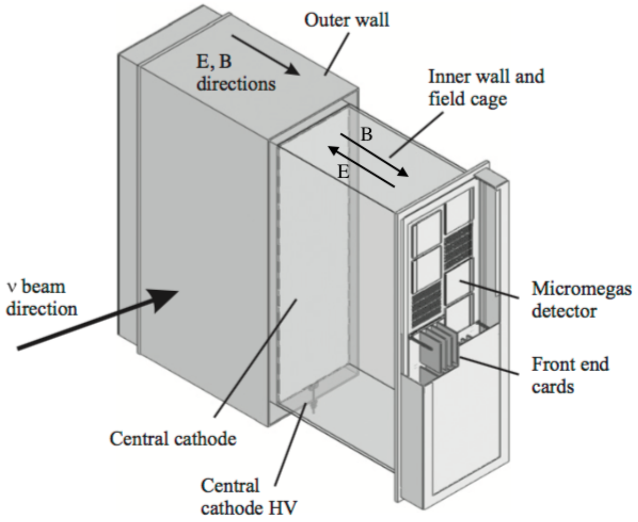

Figure 2.12: The layout of a TPC [90]. Here, the left hand portion of the diagram shows the 'complete detector' with the outer wall in place, whilst the right hand portion shows the TPC module with the outer wall removed to demonstrate the position of the central cathode. Note, either side of the cathode wall, the magnetic field remains in the same direction, yet the electric field changes. Hence, for example, an electron within the TPC will travel away from the central cathode, towards the Micromegas pads.

2.2.7.2 TPCs

The three identical TPCs (see Figure 3.5 for the detector layout) are referred to as (from upstream to downstream) TPC1, TPC2 and TPC3. Their purpose is to use high-resolution, 3D tracking of charged particles to produce a particle identification hypothesis. This is achieved by measuring the dE/dx of the particles and making use of curved tracks produced by the 0.2 T magnet to deduce the charge and measure the momentum of the particle.

Each TPC is 2302×2400×974 mm^3 and has an outer and inner volume. The outer wall surrounds the inner wall and contains a CO_2 insulator. The TPC inner volume has a central cathode splitting the TPC down the centre along the zy plane. This, in conjunction with the copper strip pattern machined onto the interior walls of the inner volume, ensures a uniform electric field that is aligned with the 0.2 T magnetic field. The inner volume is filled with an argon based drift gas, containing small amounts of tetraflurometháne (CF_4) and butane (C_4H_{10}), in the ratio 95:3:2.

To produce reconstructed tracks, charged particles ionise the drift gas and liberate electrons which flow away from the central cathode towards the anode pads of Micromegas Module (MM) detectors where the signal is amplified and measured. Each TPC has 24 MMs (measuring 342×359 mm^2), for an active surface area of 2.95 m^2 per TPC and is segmented into 1728 anode pads (7.0×9.8 mm^2). Using the hit times, hit positions and

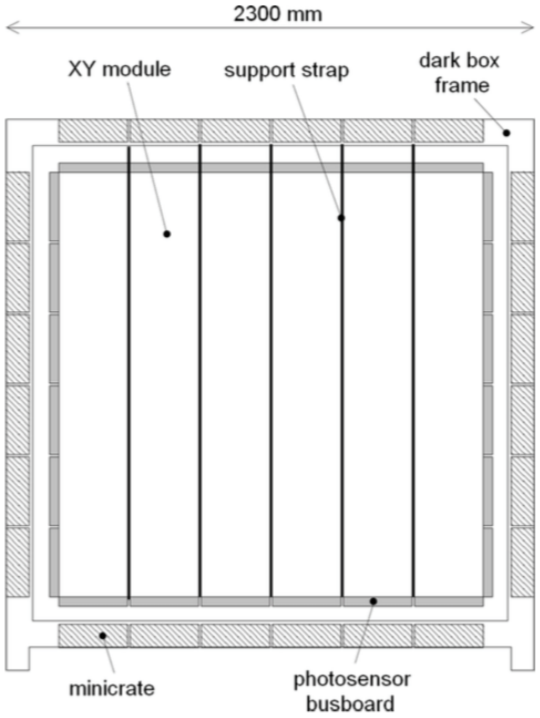

Figure 2.13: The 'XY' plane of an FGD module. The positions of the support straps, 'XY' modules, dark box, photosensor busboards and the minicrate (which stores the electronics) are shown. Figure then from [91].

known drift velocities of the liberated electrons, a full 3D reconstructed track can be deduced. For further information on the TPCs see [90].

2.2.7.3 FGDs

Interleaved between the TPCs are the two FGDs (see Figure 2.13 for the detector layout), the upstream being FGD1 and downstream FGD2, which have a total mass of 1.1 t and constitute the primary targets of ND280. FGD1 is $2300\times2400\times365$ mm^3 and consists of 30 scintillator layers arranged in 15 'XY' modules, with each module comprised of a horizontal and vertical layer to allow 3D reconstruction. Each layer is made from 192 single ended readout bars of dimension $9.61\times9.61\times1864.3$ mm^3.

FGD2 has the same dimensions as FGD1, however it only has 7 'XY' modules (14 layers of scintillator) with each module sandwiching a 2.5 cm thick corrugated polycarbonate shell (6 in total) filled with water. The shells are kept below atmospheric pressure, thus, in the event of a leak, air will enter the shell and prevent water escaping. The different designs of the two FGDs is to enable cross-section measurements on water. For further information on the FGDs see [91].

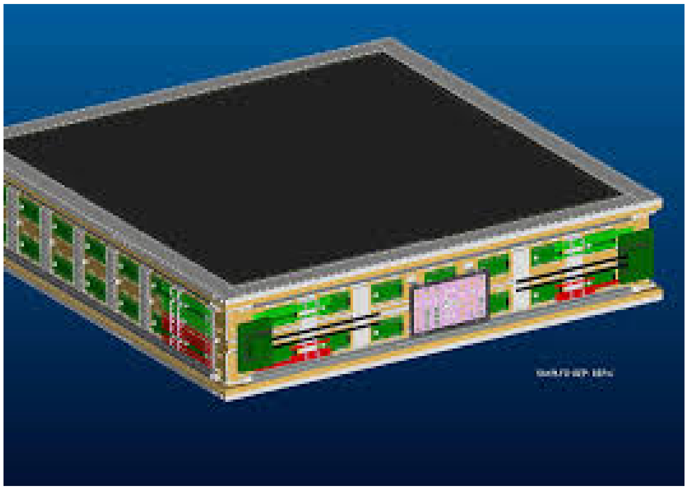

Figure 2.14: A schematic of a single ECal module. On top can be seen a black carbon fibre sheet (a similar sheet is located on the bottom), and around the sides of the module are aluminium plates. Mounted to these plates are the readout electronics (note the pink rectangle is the Readout Merger Module (RMM) and the small green rectangles are the Trip-t fronted boards (TFBs) - discussed in more detail in Section 2.2.8), the power cables and the cooling system. Figure taken from [82].

2.2.7.4 ECals

The ECals have 3 distinct sections designed to produce a near hermetic design. Most upstream is the P0D ECal (P0DECal) surrounding the P0D, further downstream is the Barrel ECal (BrlECal) surrounding the FGDs and TPCs and finally the Downstream ECal (DsECal) located downstream of TPC3. All the ECals use the same tracking sampling calorimeter design in which layers of scintillator are interleaved with lead. Note, each scintillator bar has a cross-section of 40×10 mm.

Each ECal is separated into modules: a DsECal module, six BrlECal modules and six P0DECal modules (both of which have two modules on the top and bottom and one on each side). Carbon fibre sheets encase all modules and are fitted with aluminium edges to allow electronics and cooling systems to be attached. This casing prevents light leaving the module and contamination from outside sources, see Figure 2.14.

The layout of each ECal varies slightly. The DsECal and the BrlECal (which together are referred to as the TrackerECal) are designed to provide 3D reconstruction of particle showers and tracks. Both interleave lead plates of thickness 1.75 mm, chosen as this maximises detection efficiency for π^0 decay photons, with scintillator layers orientated in alternating horizontal and vertical directions. In total the DsECal has 34 scintillator and 33 lead layers - weighing 4.80 t. The BrlECal has 31 scintillator and 30 lead layers (the interior volume of the magnet prevents more layers being added) weighing 44.9 t in total. Due to the layout of the BrlECal modules, bars of varying length are required. The longest bars spanning along the z-axis use a double ended readout and are 3840 mm

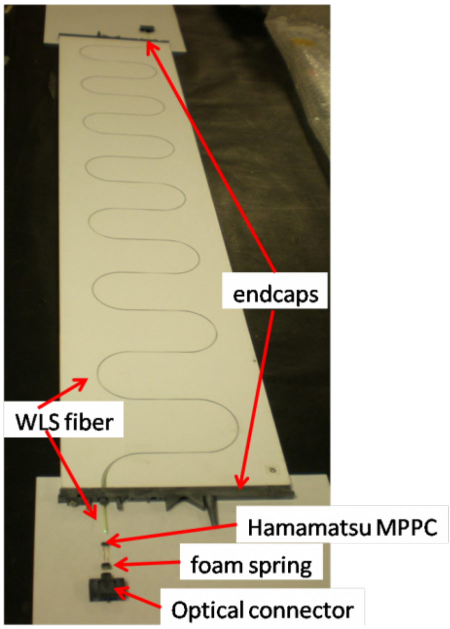

Figure 2.15: A photograph of an SMRD paddle [93].

long. In contrast, the shorter bars use a single ended readout and are 1520 mm for the top and bottom modules (orientated along the x-axis) and 2280 mm for the side modules (orientated along the y-axis).

The P0D ECal is designed to tag any particles leaving the P0D, this includes distinguishing exiting muon tracks from electron/positron/photon showers. As the P0D already provides full reconstruction, it is not necessary for the P0D ECal to do the same; thus it is much thinner than the Tracker ECals with only 6 scintillator and 5 lead layers. The width of the lead plates are also much thicker at 4 mm and were optimised using Monte Carlo simulation to compromise between the thinner layers required for photon detection, and particle discrimination and the thicker layers required for shower containment. For further information about the ECals see [92].

2.2.7.5 SMRD

The SMRD consists of 440 scintillator modules fitted in the air gaps of the magnet return yoke and has three primary uses. Firstly, it complements the other detectors within the basket by measuring the momenta of muons travelling at large angles relative to the neutrino beam. Secondly, it acts as a cosmic ray trigger, providing an essential function for detector calibration. Finally, it is a useful veto for interactions occurring outside the detector; particularly within the magnet, which produces many external interactions due to its large target mass.

The layout of the SMRDs is dependent on the position within the magnet. On the top and bottom of the magnet, three layers of scintillator paddles are inserted for all 8 yokes. On the sides of the magnet, for the first 5 yokes, 3 layers are also used. However, to improve detection of forward going particles the 6th yoke has 4 layers and the 7th and 8th yokes have 6. The scintillators in the SMRD are all single-ended readout and use the same material as bars in the other detectors, however they are much wider with dimensions of $875 \times 167 \times 7$ mm^3 for paddles orientated horizontally (these are inserted in groups of 4) and $875 \times 175 \times 7$ mm^3 for those orientated vertically (inserted in groups of 5). Additionally, to increase light capture, an S-shaped WLS fibre is used, see Figure 2.15. For further information on the SMRD see [93].

2.2.7.6 Magnet

The 850 t magnet [82] (recycled from the NOMAD experiment, and previously recycled from the UA1 experiment) produces a dipole field of strength 0.2 T, to a precision of 0.0002 T, and enables the charge and momenta of particles to be determined. The magnet can be opened along the yz centre plane, to allow maintenance work within the detector to be carried out, with each side containing 2 water cooled aluminium coils. The magnet consists of 16 C-shaped return yoke elements, 8 on each side (with the most upstream pair labelled as yoke 1, incrementing downstream until 8), each made from 16 low-carbon steel plates (thickness 4.8 cm) separated by 15 air gaps (of width 1.7 cm). The internal and external dimensions of the magnet are $7.0 \times 3.5 \times 3.6$ m^3 and $7.6 \times 5.6 \times 6.1$ m^3 respectively.

2.2.8 Electronics

As INGRID, the ECals, the P0D and the SMRD use Trip-T electronics, it is a useful exercise to understand the hierarchical structure of this readout system. The MPPCs are designed to have a large dynamic range by utilising two readout channels, the high gain which saturates at 50 photoelectrons and the low gain saturating at 500 photoelectrons. Up to 16 MPPCs are connected to a Trip-T chip which reads out each individual MPPC signal. The Trip-T chip reads out in 23, 480 ns long, integration windows (each followed by a 100 ns reset period). When the integrated charge exceeds 2.5-3.5 photoelectrons (with different sub detectors requiring different thresholds), a timestamp of accuracy 2.5 ns is produced by the high gain output. 4 Trip-T chips are connected to, and read out by, a Trip-T Frontend Board (TFB). TFBs are themselves connected to a Readout Merger Module (RMM) in groups of up to 48. Each RMM is connected to a Frontend Processing Node (FPN) which combines data from two RMMs and compresses the output. It is the RMM that controls when data is read from individual MPPCs by distributing trigger signals it receives from elsewhere in the clock hierarchy, see Figure 2.16. If ND280 is in 'global mode' (standard running), it is the Master Clock Module (MCM) that acts as the main clock for receiving and distributing triggers (e.g. beam spill triggers from the accelerator or cosmic triggers from the Cosmic Trigger Module (CTM)). The MCM is connected, via Rocket I/O transceivers, to three Slave Clock Modules (SCM), one each for the ECal, SMRD and P0D. The SCM's function is to fan out the triggers to the individual RMMs (12, 6 and 4 for the ECal, P0D and SMRD respectively) and it can also act as an

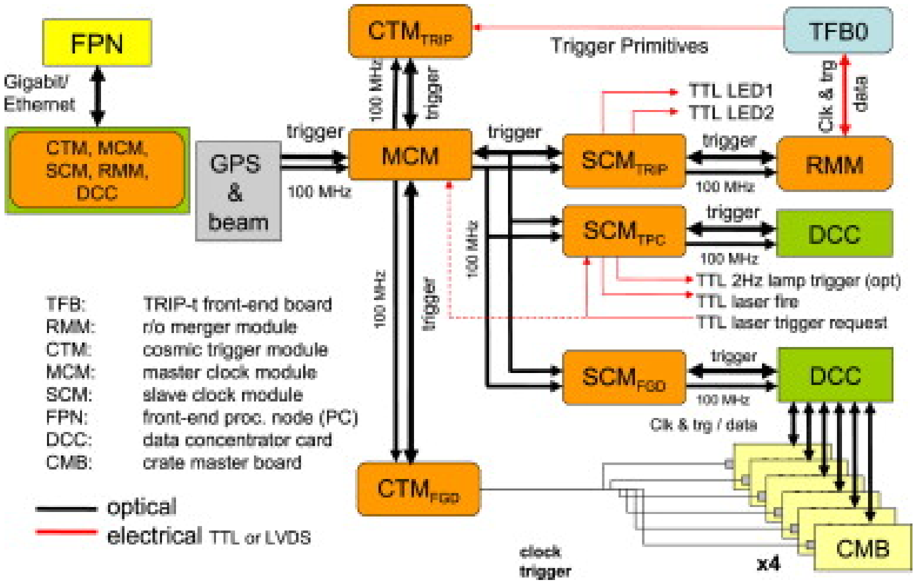

Figure 2.16: The electronics layout for ND280. Note, the hierarchy between the MCM, SCMs and RMMs [82].

MCM if one of the subdetectors is run in local mode for calibration or testing. For further information on the electronics see [82].

2.2.9 The Far Detector

Super-Kamiokande (see Figure 2.17), T2K's far detector, is located 295 km from the beam target at a depth of 1 km (a water equivalent depth of 2.7 km) below Mt. Ikenoyama, near the west coast of Japan. This cylindrical water Cherenkov detector contains 50 kt of ultra pure water, with a fiducial volume of 22.5 kt, and over 13,000 PMTs. It is split into an Inner Detector (ID) and Outer Detector (OD), in-between which is a stainless steel scaffold covered in plastic sheets to provide optical separation. The ID is 33.8 m in diameter, 36.2 m tall and has 11,129 inward facing 50 cm diameter PMTs providing a 40% surface coverage and a combined quantum and collection efficiency of 20%. The OD is 2 m wider than the ID, surrounding it and providing an almost 100% efficient cosmic ray veto. This is achieved using 1885 outward facing 20 cm PMTs, with a 7% detector coverage.

As stated, Super-Kamiokande uses Cherenkov light produced by charged particles to determine: the vertex of a neutrino interaction, the flavour of the incident neutrino and its energy (the latter of which, due to over 20 years of operation, calibration and validation is known to the percent level). A Cherenkov cone is produced when the velocity of a charged particle travelling in a medium exceeds the speed of light in that medium. To count the number of electron and muon neutrinos entering the detector during neutrino mode, CCQE events are searched for:

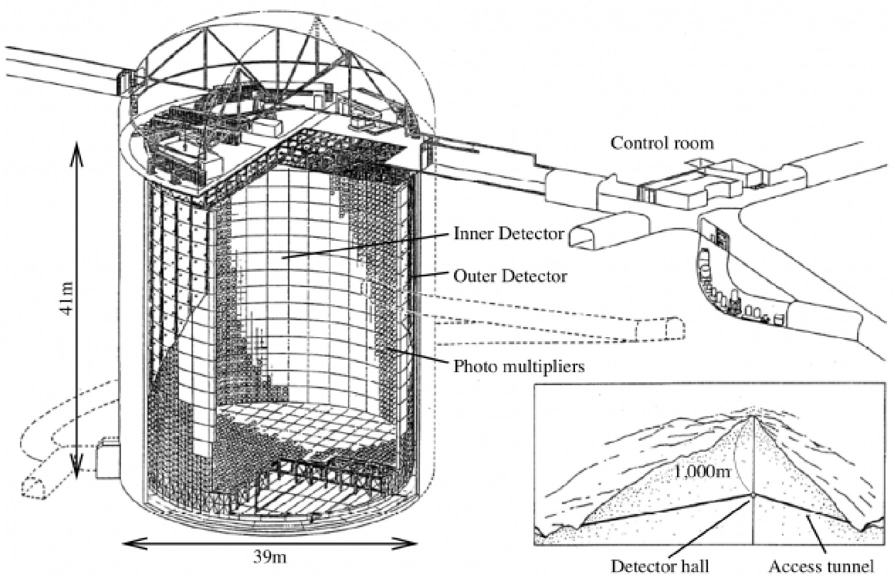

Figure 2.17: The layout of Super-Kamiokande [82].

$$\nu_e + n \rightarrow e^- + p, \tag{2.16}$$

$$\nu_\mu + n \rightarrow \mu^- + n. \tag{2.17}$$

Analogously, during anti-neutrino mode, to count the number of electron and muon anti-neutrinos, the charge conjugate CCQE interactions are searched for:

$$\bar{\nu}_e + p \rightarrow e^+ + n, \tag{2.18}$$

$$\bar{\nu}_\mu + p \rightarrow \mu^+ + n. \tag{2.19}$$

Hence, by counting the number of electrons and muons, the number of electron and muon neutrinos is deduced (with the charge conjugate equivalent holding true for anti-neutrino operation). To discriminate between electrons and muons (or positrons and anti-muons), algorithms have been developed, with an efficiency of ∼100%, to analyse the differences between the rings produced, see Figure 2.18. In the case of muons, due to their large mass, they are highly penetrating and as such do not scatter often. Consequently, the Cherenkov rings produced are well defined, see Figure 2.18a. In contrast, the much lighter electrons are able to scatter more frequently; additionally, due to the neutrino energies produced at T2K, they have a propensity to shower. The combination of these effects causes the Cherenkov ring produced to be much“fuzzier”, see Figure 2.18b.

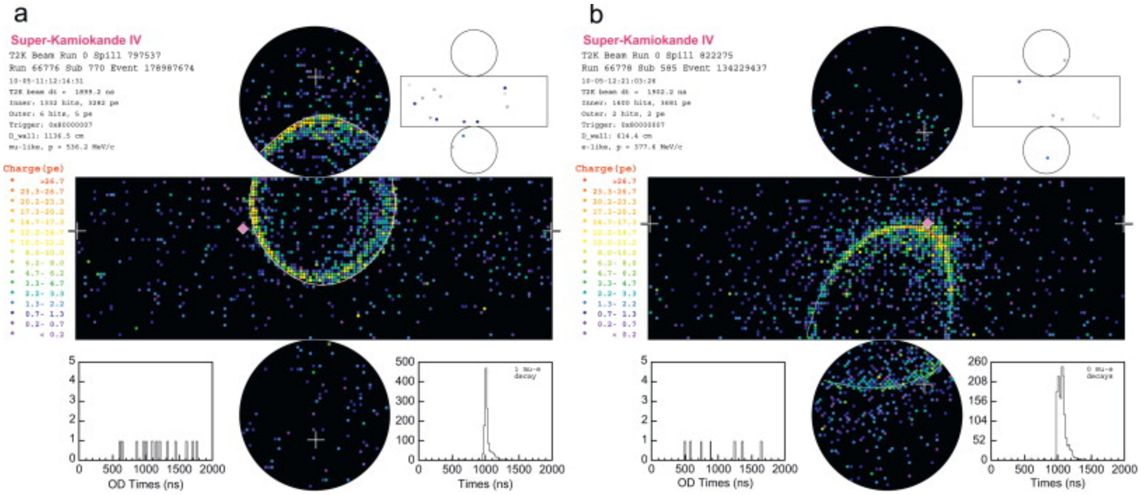

Figure 2.18: A well-defined Cherenkov ring produced by a muon (a) compared to the much 'fuzzier' electron induced ring (b) [82].

To reconstruct the neutrino energy, firstly, the number of photoelectrons collected by the PMTs is measured. This is related to the photon yield which itself is related to the lepton energy by the Frank-Tamm equation:

$$\frac{d^2N}{dxd\lambda} = \frac{2\pi\alpha z^2}{\lambda^2}(1 - \frac{c^2}{v^2n^2(\lambda)}) \tag{2.20}$$

where $\frac{d^2N}{dxd\lambda}$ is the number of photons produced per unit length per unit wavelength, z is the particle charge, v the particle velocity, n the refractive index and α the fine structure constant. Once the lepton energy has been deduced, the neutrino energy can be inferred.

Super-Kamiokande can also produce other measurements of interest. These include the detection of Michel electrons as a result of muon decay and the reconstruction of neutral pions. The latter is achieved when the Cherenkov rings of both π^0 decay photons can be resolved, allowing for the invariant mass to be deduced. However, if both cannot be resolved (either because they overlap significantly or because one of the rings is misreconstructed or missed completely) this signal can mimic a CCQE electron neutrino event and can be mistaken as a ν_e appearance signal. This is the second largest background to the electron neutrino appearance signal (only behind the contamination due to the intrinsic ν_e component of the neutrino beam).

Chapter 3

ND280 Event Reconstruction and Offline Software

3.1 Introduction

This chapter focuses initially on the framework of the ND280 software suite, specifically the production of Monte Carlo (MC) simulation and the processing of data to enable analysis work. Following this, the aspects of event reconstruction - most pertinent to the NC1π^0 analysis presented in Chapter 5 - are discussed.

3.2 The ND280 Software Suite

The general workflow of the ND280 offline software suite can be seen in Figure 3.1. This provides the production chain for two channels; the processing of raw data (stored as Maximum Integration Data Acquisition System (MIDAS) files [95]) and the production of MC. It performs numerous calibration and reconstruction algorithms to produce output that are suitable for the user to perform analysis work. The basic simulation library used by the offline software is GEANT4 [96], whilst ROOT [97] provides the underlying analysis framework.

3.2.1 Data Chain

The first step in the data chain is to "unpack" the raw data into the C++ "oaEvent" format (a format which is maintained until the final "oaAnalysis" output).[1] Following this, raw hit calibration takes place. This is a process that utilises two types of routine:

---

[1] The oaEvent format contains hit level information, whereas the oaAnalysis format contains ROOT-based objects for use during data analysis.

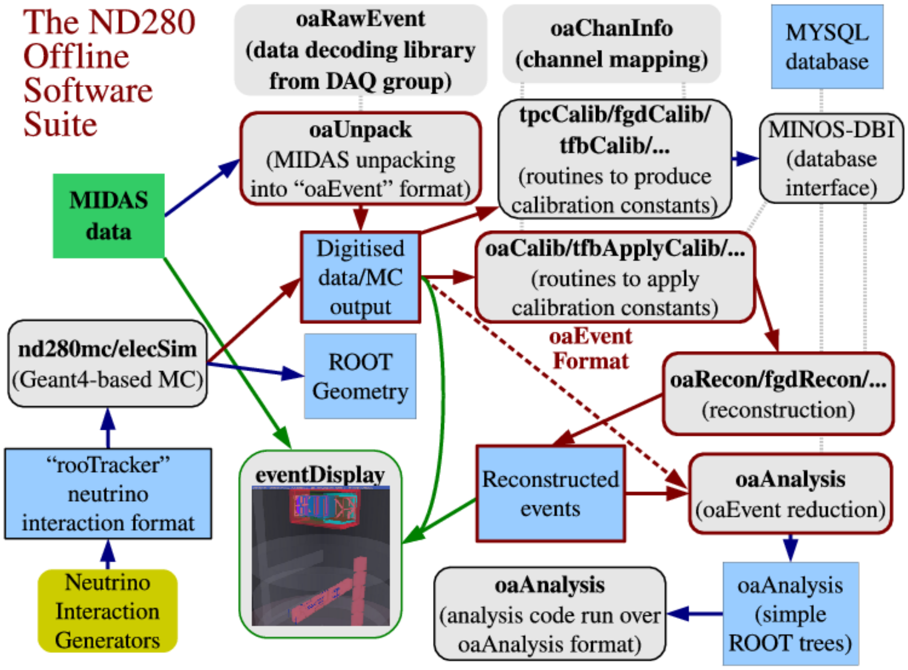

Figure 3.1: A simplified schematic of the ND280 software suite, demonstrating the flow of both data and MC to produce the final oaAnalysis format. Figure taken from [94].

time-dependent and time-independent. The former are those in which calibrations are applied that vary with time (e.g. timing offsets and charge response).[2] The relevant calibration constants for each period are stored in a MYSQL database. Conversely, time-independent routines provide calibration for non-fluctuating variables, as such constants can be hard-coded into the calibration scripts; a prime example is channel mismapping which remains the same unless the channels are physically changed.

Once calibration is complete, reconstruction of the neutrino interaction vertex and individual hits into tracks and showers is undertaken. The first stage is subdetector reconstruction, followed by global reconstruction (discussed in more detail later in the chapter). Finally, post calibration, the oaAnalysis package is used, which produces the final, data reduced, ROOT output.

3.2.2 Monte Carlo Chain

The MC chain is very similar to the data chain, however a few initial steps are required. Initially, the J-PARC Neutrino Beam package (JNUBEAM) [98] is used to simulate the neutrino beam at ND280 - it uses GEANT3 [99] to model the beamline and detectors, alongside the GCALOR package [100] to model hadronic interactions. Next, to replicate neutrino interactions within the detector and the resulting final state interactions (FSI),

---

[2] These can fluctuate throughout a run as they are affected by parameters such as detector response and environmental conditions.

the GENIE [101] and NEUT [73] event generators are employed. This MC is then passed to the nd280mc and elecSim packages. nd280mc, using GEANT4 libraries, propagates the FSI particles and produces trajectories and their associated energy deposits. Next, the elecSim package is able to model the detector response by considering effects including: light attenuation in the bars and fibres, detector noise, readout saturation and ionisation drift in the TPCs.

At this point, the format of the MC is identical to the oaEvent data, hence it can now pass through the calibration, reconstruction and oaAnalysis output stages.

3.2.3 ECal Geometries

As the analysis presented in Chapter 5 focuses heavily on the tracker ECals, it is important to understand the local and global geometries employed within the software. The global coordinate system is applied throughout ND280, it uses right-handed Cartesian coordinates in which the z-axis points from the P0D to the DsECal (i.e. downstream). Similarly, the local ECal geometry also uses right-handed Cartesian coordinates. However, for each module the z-axis is defined as pointing outwards and perpendicular to the outside face of the sub-detector, see Figure 3.2.[3] The advantage of the latter is that, when considering the ECal as a separate subdetector, each module can be compared in a homologous manner.

3.2.4 Calibration

This section focuses on the calibration of the tracker ECals, specifically the energy and timing calibration methods. These techniques are similar to the other Trip-T detectors, see [102] and [103] for more information on the P0D and SMRD respectively. For an in-depth discussion of the calibration routines for the FGD and TPC see [104] and [105] respectively.

3.2.4.1 Energy Calibration

The first ECal calibration stage is to remove the pedestal response from each channel; namely, to subtract the MPPC's ADC response when no signal is present.[4]

Next, a linearity correction is applied. This is necessary as the charge read out has to be deduced from the ADC response, however there is not a linear relationship between the two. In order to deduce the relationship, a range of known charges are injected by the TFBs into the capacitors and the response recorded (for both the high and low gain channels), see Figure 3.3. This method also provides a reference as to which channel

---

[3] This local coordinate system is also used for the P0DECals.

[4] This can be caused by multiple sources, such as electronic noise, and is monitored by the DAQ system on a three hourly timescale when there is no beam.

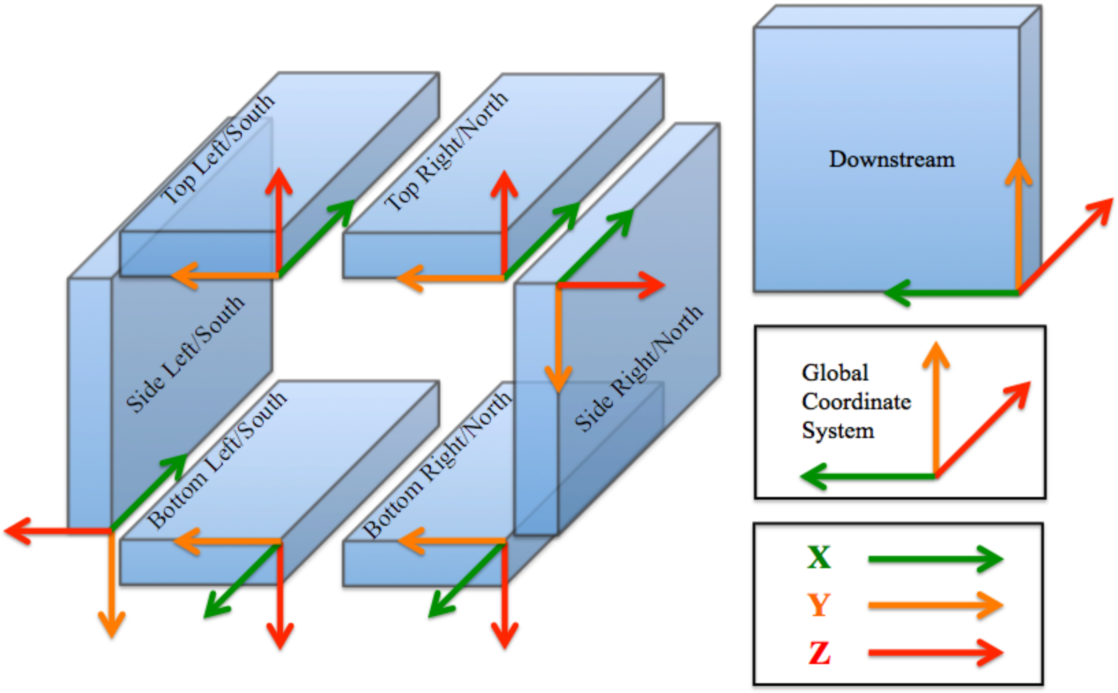

Figure 3.2: The local coordinates for each ECal module compared to the global coordinate system of ND280.

should be used - with the switchover between the high and low gain channel at ∼460 ADC (∼21 pixel equivalent units (P.E.U.)).[5]

Following this, the MPPC gain needs to be calibrated in order to allow the conversion of charge into P.E.U. To do this the MPPC gain is measured by dividing the ADC count of the pedestal peak by the ADC count of the peak produced by a single MPPC pixel being fired. Now, by simply dividing the MPPC output by this result, a measure of the number of photoelectrons produced by the MPPC can be deduced.

The MPPC response calibration is the final calibration stage for the MPPCs. At this point many other, smaller, effects are considered, for instance: light transmission down the WLS fibre, fibre coupling and pixel crosstalk. The calibration constants associated with these effects were determined by test bench measurements.

Next, bar to bar calibration is required to produce a uniform response for all scintillator bars. The method is to measure MIP-like particles, produced by cosmics, and fit the resultant energy spectrum to a Landau distribution (convolved with a Gaussian smear). The fitted peak of all bars can then be aligned to the average for all bars of that orientation within that module.

---

[5] A P.E.U. is the MPPC charge required to fire one pixel.

The final stage is the fibre scaling correction, which aligns the energy response of the different bar orientations (as calculated in the previous stage). To do this, a fibre scaling factor is applied:

$$\text{Fibre Scaling Factor} = \frac{\text{Fitted MIP peak for DsECal}}{\text{Fitted MIP peak for specific orientation}}. \tag{3.1}$$

This is done to each orientation so as to align the energy response to that of the DsECal.

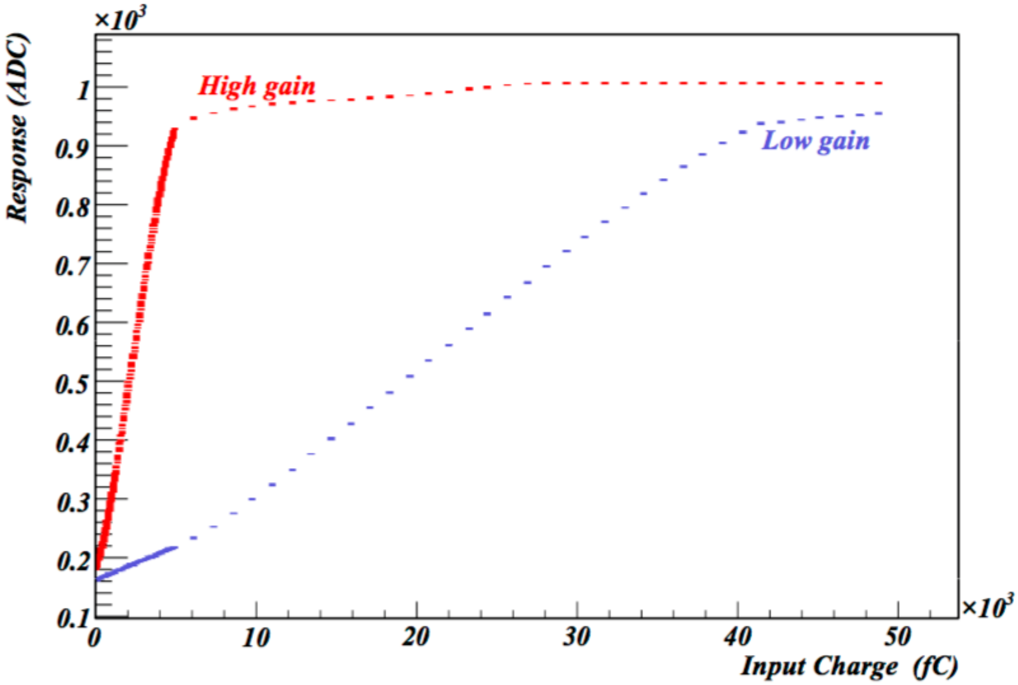

Figure 3.3: The characteristic MPPC response for a range of input charges, for both the high and low gain channels. Figure taken from [85].

3.2.4.2 Timing Calibration

Timing calibration is essential in order to allow effective hit clustering, provide directionality of tracks and showers and to efficiently remove noise (these procedures are performed by a dedicated calibration group, for a detailed discussion of the routines see [106]). The first stage in this process is to account for the electronic timing offset. This occurs as different cable lengths are used for different TFBs. By using cosmic ray events to measure the difference between the expected and observed hit times, this effect can be accounted for (see Chapter 4 for more details).

The second stage is to apply the time-walk correction to account for the effect in which higher charged events produce an earlier timestamp. This occurs as a Trip-T timestamp is produced when the integrated charge exceeds 2.5 photoelectrons (p.e.), hence interactions with larger charges will exceed this threshold sooner. To find the correlation between the time of the timestamp and the charge of the event, a range of known charges are injected into the MPPCs.

The final correction to apply is the fibre time walk. This attempts to account for large hits having a higher probability of delay in the recorded hit time. This effect occurs as reemitted photons produced by the WLS fibres follow an exponential decay. Hence, for small hit charges all the re-emitted photons will be emitted almost instantly, as such the average time of the photons will be close to the original hit time. Conversely, at higher energies, producing more scintillation photons, the wavelength shifted photons will be re-emitted over a larger period of time. Thus the average time will be shifted away from the original hit time. As this process is not constant, even for events of equal energy, it has to be analytically modelled.

3.2.5 Reconstruction

The reconstruction of objects in ND280 consists of two stages; local reconstruction where objects are formed in each subdetector separately, and global reconstruction which tries to produce global objects (i.e. those spanning multiple subdetectors). To do this a similar procedure is used across all subdetectors, in which algorithms, using hit information, produce 2D objects which are then combined to produced 3D objects. In this section the focus is on tracker[6] and tracker ECal reconstruction, for the P0D see [107] and the SMRD see [103].

3.2.5.1 Tracker and Global Reconstruction

Reconstruction in the tracker uses the RecPack [108] package. This undertakes track fitting procedures by using hit information, whilst taking into account numerous parameters such as the magnetic field, the detector geometry and particle kinematics.

To reconstruct a track, RecPack utilises a step by step procedure. Firstly, the reconstruction algorithm bunches hits together that are close in time and space - this is done within each TPC separately. Next it attempts to group these bunches together using pattern recognition algorithms. To extract particle ID information from the track (i.e. track coordinates, track angle, curvature due to the magnetic field and $\frac{dE}{dX}$), a likelihood method is employed. Once the TPC tracks have been reconstructed, a Kalman filter tries to match FGD hits to them. As this is not always possible, a separate matching algorithm for unmatched FGD hits produces FGD isolated (FGDIso) tracks. Here, XY and YZ hits are grouped separately, before a full 3D matching is attempted. Where unused FGD hits

---

[6] The TPCs and FGDs.

Figure 3.4: A flowchart of the tracker reconstruction procedure.

or unmatched 2D tracks remain, the charge weighted position is saved. Once all 3D FGD tracks have been reconstructed a Kalman filter is used once more to try and match all TPC and FGD tracks. For a flowchart of this procedure, see Figure 3.4.

To establish the start time of TPC objects, initially if it has been matched to an FGD, the FGD timing information is used. If it contains no FGD information but crosses the TPC central cathode, the maximum drift time gives timing information. If this is not the case, but it is matched to other TPCs, that timing information enables the calculation of the start time. However, if none of the above scenarios are valid a default time is used.

The first stage of global reconstruction is to match tracker tracks to ECal and P0D objects. Much like the TPC to FGD matching algorithm, this step extrapolates the tracker tracks into these subdetectors. A χ^2 fit using position and direction information is performed on nearby objects, with a match occurring if the χ^2 result is < 100 for ECal objects (< 200 for P0D and SMRD objects) and the timing between the two is ≤ 300 ns. If matching occurs, a Kalman filter recalculates the particle identification (PID) information. The global package repeats the above process until all tracks have been matched.[7]

---

[7] Note, the package also provides matching for non-tracker tracks and is able to match P0D and SMRD tracks separately.

3.2.5.2 TPC Particle Identification

By measuring $\frac{dE}{dx}$ and momenta of tracks passing through the TPC, then comparing the results with theoretical predictions for different particle species - the electron, kaon, muon, pion and proton - particle identification is possible. To calculate $\frac{dE}{dx}$, a truncated mean charge for the track is found from the charges of the lowest 70% of clusters within the track - corrections are applied depending on effects such as number of clusters and the angle of the track. This gives a measure of energy loss and allows five different "pulls" to be produced:

$$\mathrm{Pull}_\alpha = \frac{(\frac{dE}{dx})^{\mathrm{meas}} - (\frac{dE}{dx})^{\mathrm{exp}}_\alpha}{\sigma_\alpha}, \tag{3.2}$$

where α is the particle species, $(\frac{dE}{dx})^{\mathrm{meas}}$ is the measured rate of energy loss, $(\frac{dE}{dx})^{\mathrm{exp}}$ is the expected rate of energy loss and σ is the "α hypothesis" resolution.[8] This represents the number of sigma deviations between $(\frac{dE}{dx})^{\mathrm{meas}}$ and $(\frac{dE}{dx})^{\mathrm{exp}}_\alpha$. For a comparison of the theoretical $\frac{dE}{dx}$ and momentum distributions and the MC results at ND280 see Figures 3.5a and 3.5b.

3.2.5.3 FGD Particle Identification

As with the above method, pulls are also produced for FGD objects. In this case however, only three particle hypothesis are considered; the muon, pion and proton.[9] The hypotheses are given by:

$$\mathrm{Pull}_i = \frac{E - E_i(x)}{\sigma_i(x)}. \tag{3.3}$$

Here, i is the particle species, E is the measured energy of the track, $E_i(x)$ is the expected energy of a track of length x and $\sigma_i(x)$ is the expected resolution for a track of length x. For an example of reconstructed FGD1 energy deposition as a function of track length compared to the expected distributions, see Figure 3.6.

3.2.6 ECal Reconstruction

For the tracker ECals, reconstruction follows a seven step procedure performed by the ecalRecon package. With regards to searching for NC1π^0 interactions, maximising ECal reconstruction efficiency and performance is essential, not only due to the difficulty in detecting the typically low energy π^0 decay photons, but also in producing directionality of the photon clusters to enable adequate vertexing capabilities.

---

[8] For reference, MIPs have a resolution of ∼7.8%.

[9] As <1% of kaons and electrons are contained within the FGD.

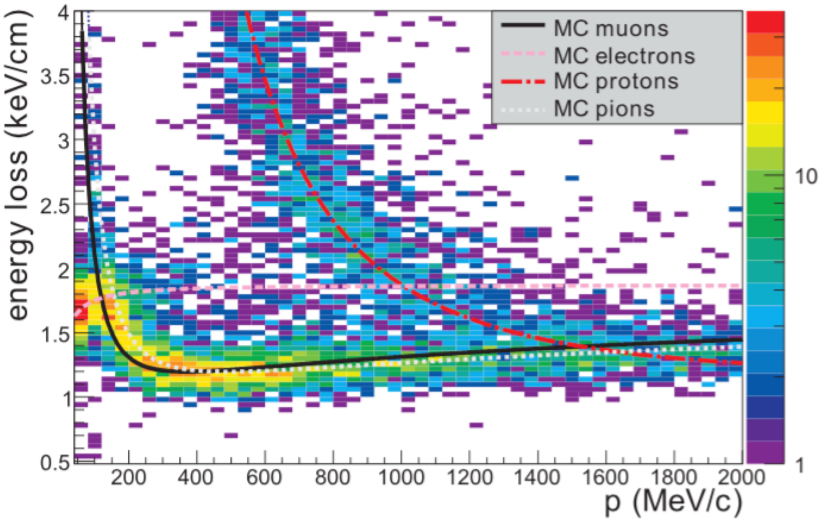

(a) Positively charged particles.

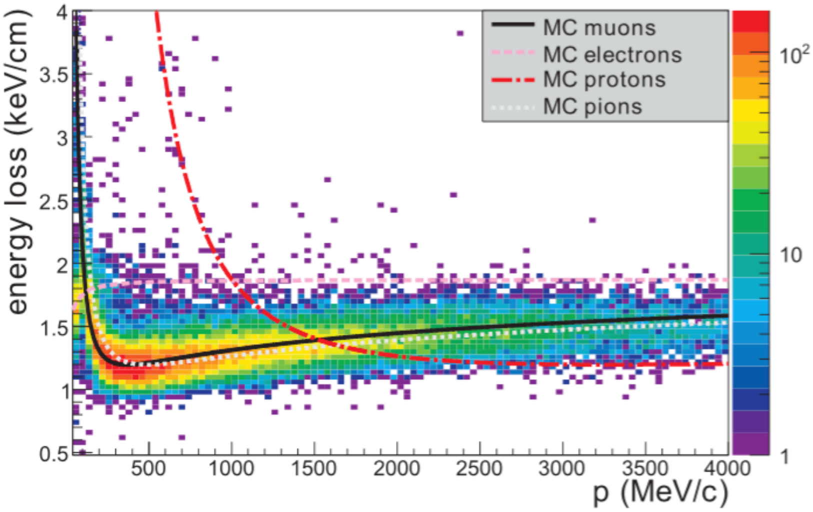

(b) Negatively charged particles.

Figure 3.5: Rate of energy loss as a function of particle momentum for positively and negatively charged particles. Here, the data is from Run1 neutrino interactions and is compared to the expected distributions (solid and dashed lines) of various particle species. Figures taken from [105].

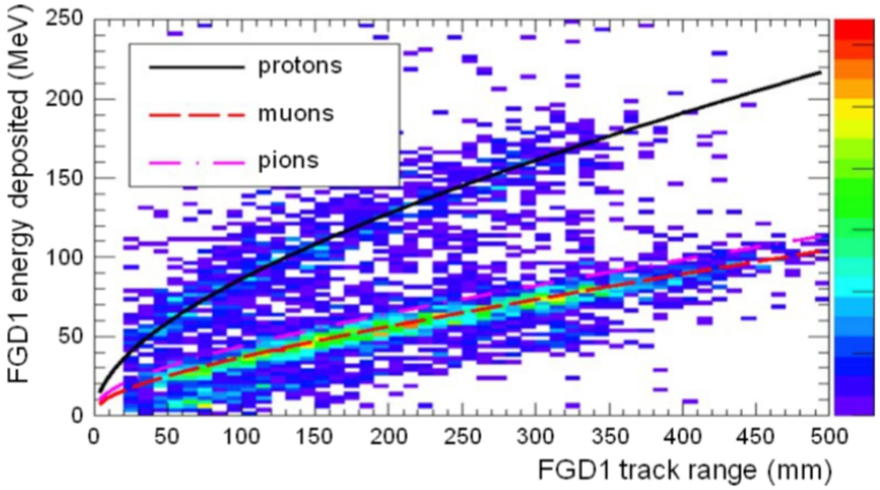

Figure 3.6: Energy deposited by tracks stopping in FGD1 that have crossed >3 layers. Note, if the tracks also have a TPC1 component and the proton hypothesis is used, the tracks must have: $|\mathrm{Pull}_p| < 2.5$, $|\mathrm{Pull}_\mu| > 2.5$ and $|\mathrm{Pull}_\pi| > 2.5$. Conversely, in the selection of muons and pions, the tracks must have: $|\mathrm{Pull}_p| > 2.5$, $|\mathrm{Pull}_\mu| < 2.5$ and $|\mathrm{Pull}_\pi| < 2.5$. Here, the data is from Run1 and 2 and is compared to the expected distributions (solid and dashed lines) of the three particle species. Figure taken from [111].

3.2.6.1 Hit Preparation

The first reconstruction stage is hit preparation. Initially, hits in both the XY and YZ views are bunched together, provided they are within a 50 ns window. If a double-ended bar detects hits, the sensor information is used to deduce whether this is one hit (if the hit time recorded is similar for both sensors, the small time difference can then be used to give the hit position on the bar) or two separate hits on either end of the bar. At this point effects such as light attenuation along the WLS fibre are considered with the hit amplitudes corrected accordingly. Finally, the response in P.E.U. is converted to MIP equivalent units (M.E.U.), a unit defined as the charge deposited by a MIP passing perpendicularly through an ECal bar at a distance of 1 m from the MPPC.[10]

3.2.6.2 Basic Clustering

The initial stage of clustering is basic clustering, which considers 2D clusters only (one per view). A nearest neighbour aggregation algorithm is used, in which the highest charged hit is set as a seed. Lower charged hits are added to the cluster if they are within 30 ns and within 2 layers and are in an adjacent bar to any hit in the cluster (or in the case of the first hit to be added, within 30 ns and within 2 layers and in the adjacent bar to the

---

[10] As a reference, 1 MEU ≈ 25 PEU.

seed). Once basic clustering is complete, all 2D clusters (defined as any with 3 or more hits), alongside any unclustered hits, are passed on to the combine clusters algorithm.

3.2.6.3 Combine Clusters

The combine clusters algorithm attempts to combine the basic 2D clusters. First, principal component analyses (PCA) [112] are performed on all clusters to find their principal axes. Next, the largest cluster is selected as the "primary" cluster and another can be added to it if the principal axes of both are less than 80 mm apart and the average hit time of each is less than 40 ns. If they are joined, this new cluster becomes the new primary.[11]

The above process is repeated until no more combinations can be found; all clusters, whether combined or not, and all unused hits are passed on to the next algorithm.

3.2.6.4 Expand Clustering

The final 2D matching algorithm is expand clustering which tries to match the remaining unused hits to the combined clusters. First PCAs are again performed on all clusters to determine their primary axes along with the 1σ width of these axes. Each hit is then considered and a matching weight for all the clusters is calculated by adding in quadrature the distance, in σ, from each of the axes. The hit is matched to the cluster with the lowest matching weight, with the stipulation that the weight is less than 80 and the hit time is less than 40 ns before, or 40 ns after, the earliest and latest hit within the cluster.

3.2.6.5 3D Clustering

3D clustering follows two approaches. The first produces a seed by propagating a tracker track into the ECal. Next, a likelihood of all 2D cluster pairs within the first four layers of the ECal is produced. This is achieved by calculating the distance, Δx and time, Δt between each cluster and tracker seed - the uncertainty on Δx, σ_x, is the distance uncertainties added in quadrature and that of Δt, σ_t taken to be 100 ns. Using this information, the following distance and time likelihoods are produced:

$$\mathcal{L}_x = \frac{G_x(\Delta x|0, \sigma_x)}{G_x(0|0, \sigma_x)} \tag{3.4}$$

and

---

[11] A PCA is a statistical procedure used to parameterise a system. It can be applied to any n-dimensional data set, producing a set of n linearly uncorrelated variables, known as principal components. A simple analogy is for a two dimensional data set, which the PCA represents as an ellipsis, with the first principle component being the major axis and the second principle component being the minor axis.

$$\mathcal{L}_t = \frac{G_t(\Delta t|0, \sigma_t)}{G_t(0|0, \sigma_t)}, \tag{3.5}$$

Where G_x and G_t are the distance and time Gaussian probability density functions respectively. Now, the final seeded likelihood can be deduced:

$$\mathcal{L}_{\text{Seeded}} = \mathcal{L}_x \times \mathcal{L}_t. \tag{3.6}$$

The second method does not consider the position of the clusters compared to the first four layers of the ECals. Instead, a likelihood is produced by calculating the Q_{Ratio} of the cluster pairs - the ratio of their charges - and comparing to a log-normal distribution produced using MC truth. A second likelihood is produced by comparing the difference between the innermost layers of each cluster, Δ_{Layer}, to a probability density function produced using MC truth. Combining the two gives the matching likelihood:

$$\mathcal{L}_{\text{Matching}} = \mathcal{L}_{Q_{\text{Ratio}}} \times \mathcal{L}_{\Delta_{\text{Layer}}}. \tag{3.7}$$

If a pair is matched using both likelihoods, the final likelihood used is the product of the two:

$$\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{Seeded}} \times \mathcal{L}_{\text{Matching}}, \tag{3.8}$$

otherwise only the matching likelihood is used. The pair that has the largest likelihood is matched together, provided it is above a minimum threshold of $\mathcal{L} > 1 \times 10^{-5}$. This process is continued until all 2D clusters have been analysed.

3.2.6.6 3D Hit Reconstruction

The reconstruction algorithm is now at a point where the unknown positions of individual hits can be deduced.[12] Initially, the charge weighted position of hits in each layer are calculated. To find the position along the bar for a particular hit, a linear fit of the average hit positions for the four surrounding layers is produced, the position of the unknown hit is the point at which the fit crosses the layer of that hit.

3.2.6.7 Low Energy Clustering

The final algorithm is low energy clustering, which has particular importance for π^0 decay photon reconstruction (these low energy photons often do not produce 3 hit clusters in

---

[12] The unknown position being the coordinate along the length of the bar. In the case of single-ended bars this is the first calculation of such a value, whilst for double-ended bars it is a recalculation.

two separate views). This process considers the remaining 2D clusters and unused hits that have not been matched by the previous algorithms. It combines the hits to the 2D clusters, provided they are within one layer and less than 10 ns before, or less than 10 ns after, the earliest and latest hit of the cluster.

3.2.7 Photon Reconstruction Efficiency

Efficient photon reconstruction is paramount when analysing π^0 decay photons. In order to understand ND280's ability to do this, Figure 3.7 and 3.8 are presented. Each figure was produced using 'cherry-picked'[13] MC, selecting NC1π^0 events originating in either FGD, and in which the π^0 decay photons convert in the tracker ECals.[14] The figures are broken down into three separate modules; where conversion has taken place in the DsECal, the top left ECal and the left ECal of the barrel. These modules were selected as they represent the different possible orientations (i.e. the left and the right barrel are identical as are the remaining barrel modules with the top left module). Additionally, each plot is separated into hard (higher energy) and soft (lower energy) decay photons to demonstrate the subtle kinematic differences between the two.

Figures 3.7a, 3.7c and 3.7e display the energy distributions. The soft photons - as expected - peak at a much lower energy and have a much narrower distribution than the higher energy counterparts. Figures 3.8a, 3.8c and 3.8e evince the reconstruction efficiency as a function of true energy. Figure 3.8a shows hard photons entering the DsECal have an almost flat distribution with a consistent reconstruction efficiency of ~80-90%. In contrast, soft photons highlight the ECals' inability to reconstruct low energy particles, with anything $\lesssim$50 MeV being reconstructed with an efficiency of ~10%. As 100 MeV is approached, the reconstruction efficiency of soft photons approaches that of the hard photons, though with increasing energies this begins to drop. This decline is indicative of a highly Lorenz boosted pair, with both photons entering the DsECal at a similar point, depositing their energy close together and as such the algorithms cannot separate the two clusters.

When considering Figures 3.7c and 3.7e, the top left and left barrel modules show similar energy distributions to each other - and indeed the DsECal - in that the soft photon peaks at a lower energy with a narrower width. However, the hard photons for the two modules peak at a much lower energy when compared to the DsECal. This reiterates the fact that highly energetic photons are Lorentz boosted and are thus more likely to only enter the most downstream portion of these modules, or indeed the DsECal. When considering Figures 3.8c and 3.8e, the reconstruction efficiency at low energies is poor. However, as the energy approaches 100 MeV the plateau is higher than in the DsECal, reaching ~85-95%.

Figures 3.7b, 3.7d and 3.7f also emphasise the reconstruction capabilities of the ECal, however this is now considered with respect to the angle at which the photon was pro-

---

[13] MC in which only NC1π^0 events are simulated. For further details, see Section 5.1.2.

[14] Note, these figures have not been cross-checked with data.

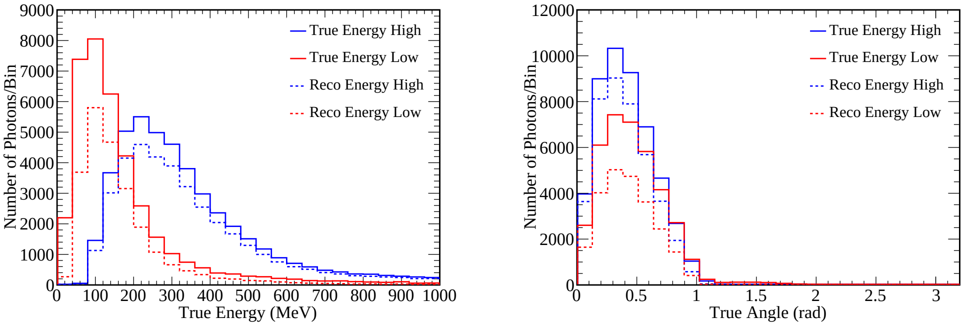

(a) DsECal energy distributions.

(b) DsECal angle distributions.

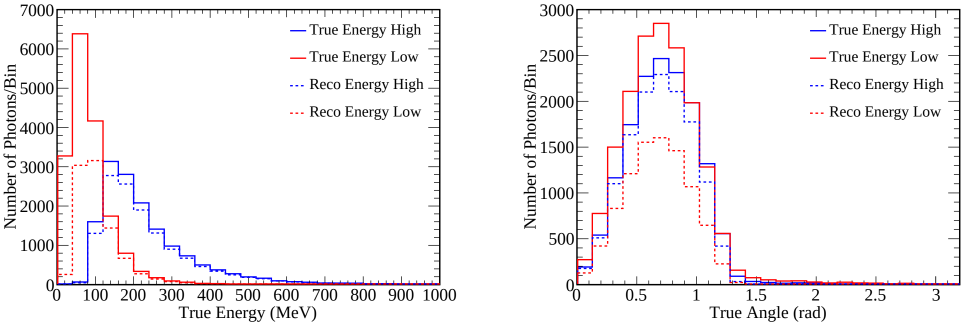

(c) Brl left energy distributions.

(d) Brl left angle distributions.

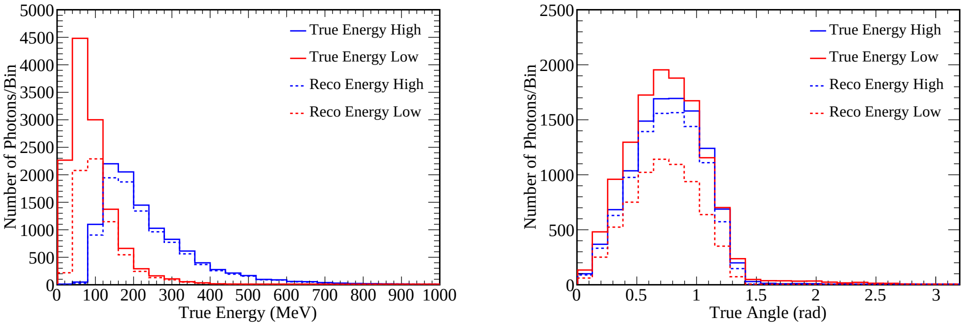

(e) Brl top left energy distributions.

(f) Brl top left angle distributions.

Figure 3.7: The true and reconstructed energy and angle distributions for both the hard and soft decay photons associated with $NC1\pi^0$ interactions. The distributions shown are for the DsECal, barrel left ECal and barrel top left ECal modules.

duced.[15] Note the small interesting tails of the distributions, where the true angle is $> \pi/2$. These occur when the photon does not travel on a straight trajectory into the module, but undergoes any combination of Compton scattering and brehmsstralung radiation. This highlights another level of difficulty when using reconstructed photons to select and locate the NC1π^0 events, in that, not only is the reconstruction efficiency often low, but also extrapolating the clusters back to an interaction vertex is sometimes impossible.

Figures 3.8b, 3.8d and 3.8f show the reconstruction efficiency as a function of true angle. Figure 3.8b shows the DsECal is able to reconstruct at an approximately constant efficiency over angles $\lesssim$0.8 rad (the hard photons being reconstructed at ~80-90% and the low energy clusters ~60-70%). On the other hand, at angles higher than this the efficiency drops sharply. Similarly Figures 3.8d and 3.8f show the reconstruction efficiencies of the top left and left barrel modules have a similar trend to the DsECal, however the plateaus are lower for the soft photons (~50-60%) and higher for the hard photons (~90-95%). Additionally, these modules do not suffer an efficiency drop-off until a larger angle is reached (~1.2 rad).

3.2.8 Energy Estimation

To estimate the cluster energy, a likelihood function is employed with three inputs: the total charge, the RMS of the charge of the hits divided by the mean charge and the skew of the charge hits. The fitting procedure was tuned with MC particle gun photons ranging from 75 MeV to 25 GeV.[16]

To analyse the effectiveness of this procedure, Figure 3.9 was produced showing a per module comparison of the true and reconstructed energy for both the hard and soft π^0 decay photons. It is apparent that the energy reconstruction algorithm is effective as, for the majority of the clusters, the energy residual is tightly centred around zero.

3.2.9 Angular Reconstruction

There are two methods of angular reconstruction. The first, which is discussed in more depth in the following section uses a PCA on the cluster ellipse with the major axis giving its direction. The second method, which has been adapted from jet physics, is to calculate the thrust-axis of the cluster. This method is used to assist reconstruction of low energy photons from π^0 decay. It assumes outward-going trajectories with the thrust being calculated by finding the charge-weighted position within the innermost layer containing hits. The thrust is then given by finding the maximum of:

---

[15] A coordinate transformation was applied to convert the direction vector from the global coordinate system to the local coordinate system of the module.

[16] A particular emphasis was placed on the region below 2 GeV.

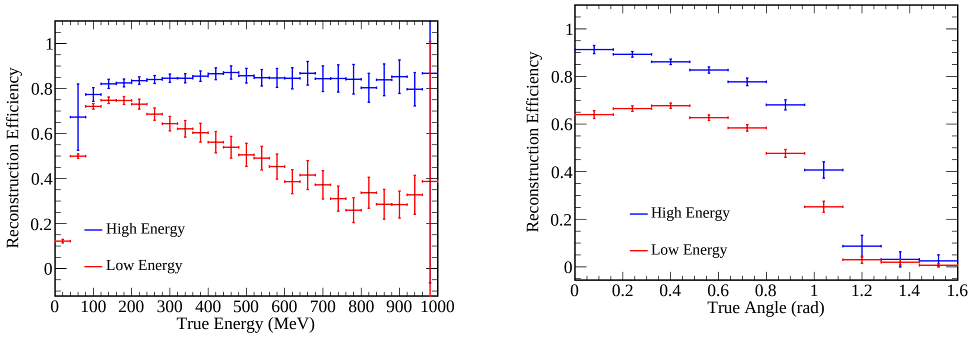

(a) DsECal energy reconstruction efficiency.

(b) DsECal angular reconstruction efficiency.

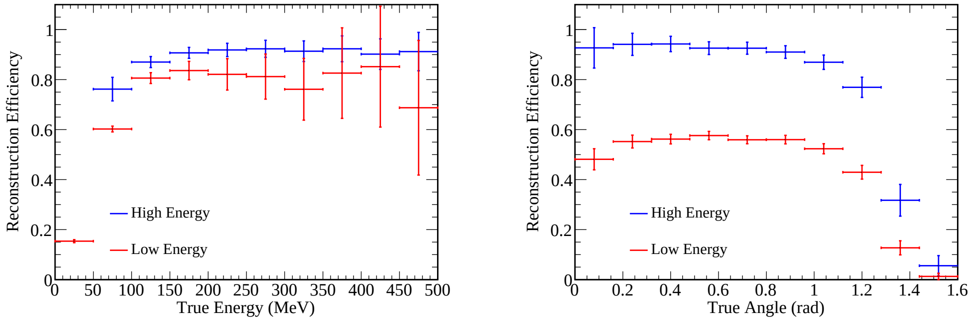

(c) Brl left energy reconstruction efficiency.

(d) Brl left angular reconstruction efficiency.

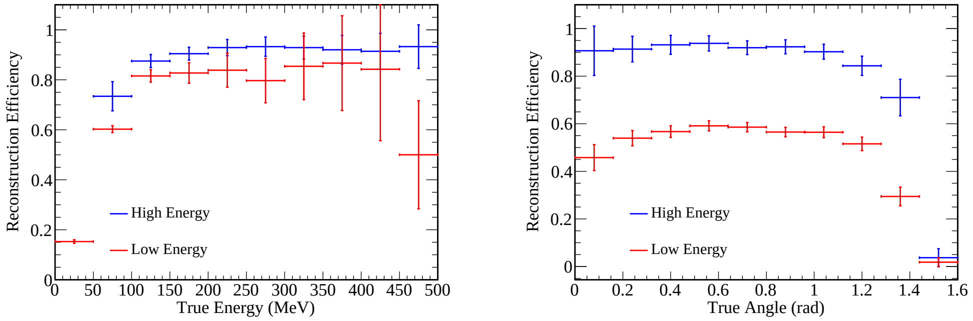

(e) Brl top left energy reconstruction efficiency.

(f) Brl top left angular reconstruction efficiency.

Figure 3.8: The energy and angular reconstruction efficiency for hard and soft decay photons associated with NC1π^0 interactions. The distributions shown are for the DsECal, barrel left ECal and barrel top left ECal modules.

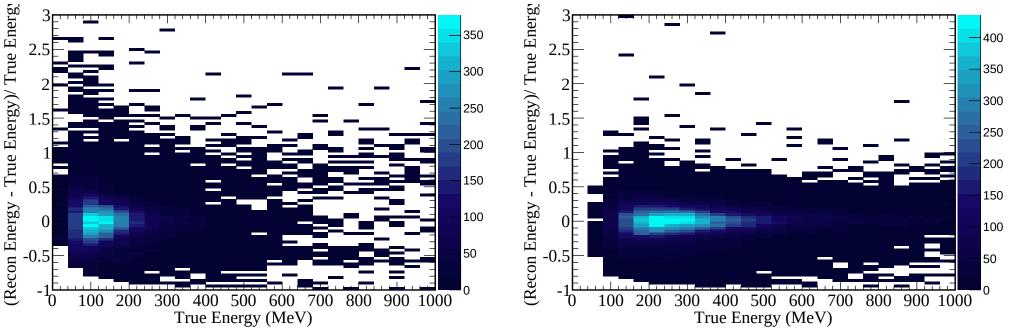

(a) DsECal soft photon energy residual. (b) DsECal hard photon energy residual.

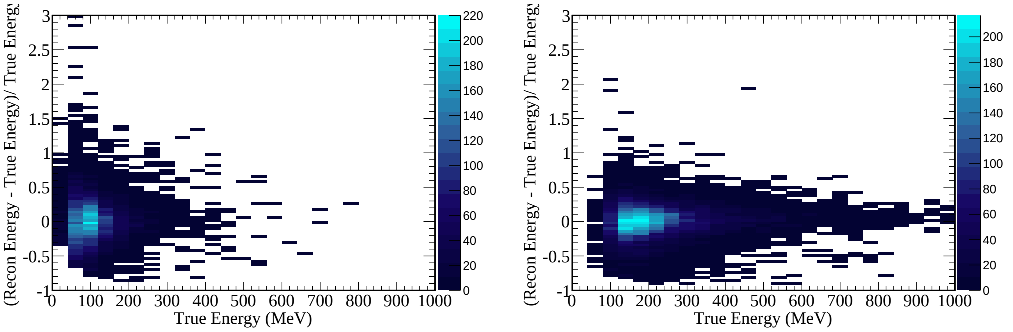

(c) Brl left soft photon energy residual. (d) Brl left hard photon energy residual.

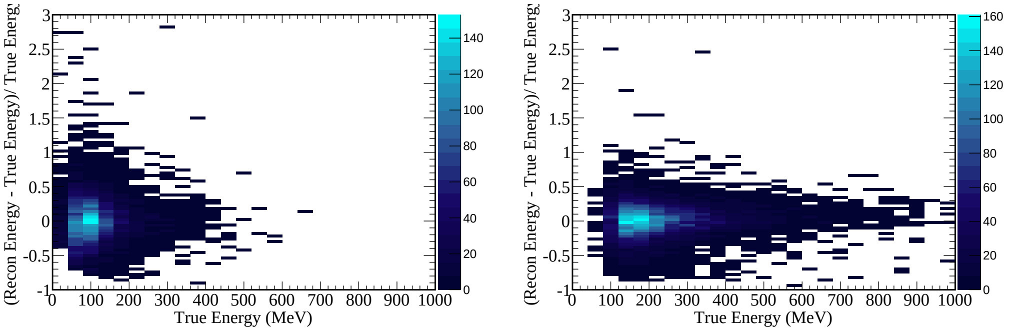

(e) Brl top left soft photon energy residual. (f) Brl top left hard photon energy residual.

Figure 3.9: ECal energy reconstruction capability (represented as energy residuals) for the hard and soft decay photons associated with NC1π^0 interactions. The distributions shown are for the DsECal, barrel left ECal and barrel top left ECal modules.

$$t(\theta,\varphi)=\frac{\sum\limits_{i} q_i|\overrightarrow{n}(\theta,\varphi)\cdot(\bar{x}_i-\bar{o})|}{\sum\limits_{i} q_i|(\bar{x}-\bar{o})|}, \tag{3.9}$$

where i is the sum over hits, q_i is the hit charge, x_i is the hit position, $\bar{o}$ is the thrust origin (the charge weighted position of hits in the innermost layer) and $\overrightarrow{n}$ is the unit vector, with polar angles θ and φ, which represents the thrust axis when $t(\theta,\varphi)$ is maximal.

To analyse the effectiveness of the thrust-axis reconstruction, Figure 3.10 was produced. From this figure, several points become apparent. Firstly, this method does provide an effective angular reconstruction method as in general the angular residual is centred on zero. Secondly, the higher energy clusters are better reconstructed as can be seen from the smaller spread in the angular residual. Thirdly, across all modules, the reconstruction is more accurate at higher angles. Finally, the two barrel modules are more effective at reconstructing the clusters than the DsECal (though this point is linked with the previous as photons in the DsECal, due to the detector geometry and also Lorentz boosting, are more likely to be forward-going).

3.2.10 Particle Identification

Numerous PID variables are available in the tracker ECals, variables which are essential for NC1π^0 analysis as there is often no information within the tracker itself. These variables have been developed by dedicated working groups, for more information see [109, 110]. Note, the variables prefixed with PID are produced using a 3D PCA. In this method, hits that are added to the PCA are charge-weighted with them being added multiple times depending on their charge. Those containing LLR are the log-likelihood-ratio variables. To calculate these, initially, the log-likelihood is found:

$$\ln|\lambda(x)|=\ln|P(\underline{\mathbf{x}}|H_0)|-\ln|P(\underline{\mathbf{x}}|H_1)|, \tag{3.10}$$

where P is the probability density function (PDF), H_0 and H_1 are the hypotheses being tested, $\underline{\mathbf{x}}$ is the set of input variables[17] and λ is the likelihood, defined as:

$$\lambda(x)=\frac{P(\underline{\mathbf{x}}|H_0)}{P(\underline{\mathbf{x}}|H_1)}. \tag{3.11}$$

To calculate the PDF, the assumption is made that the input variables are uncorrelated and so can be defined as:

$$P(\underline{\mathbf{x}}|H_0)=\prod_{i}^{N} P(x_i|H_0). \tag{3.12}$$

---

[17] These are: circularity, truncated max ratio, Q$_{RMS}$ and front back ratio, all of which are defined later in this section.

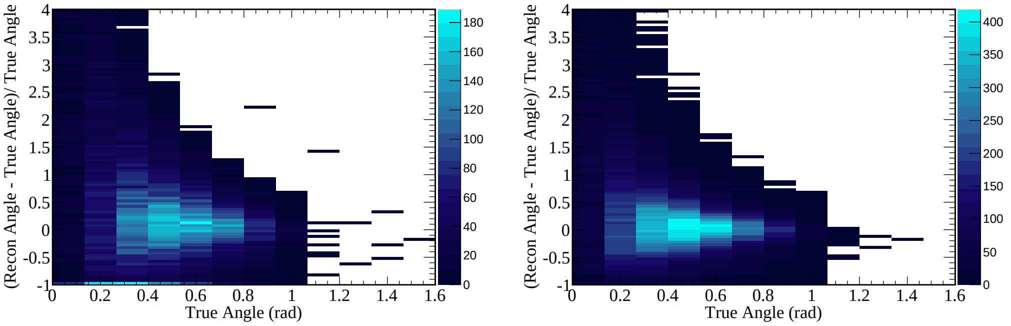

(a) DsECal soft photon angular residual. (b) DsECal hard photon angular residual.

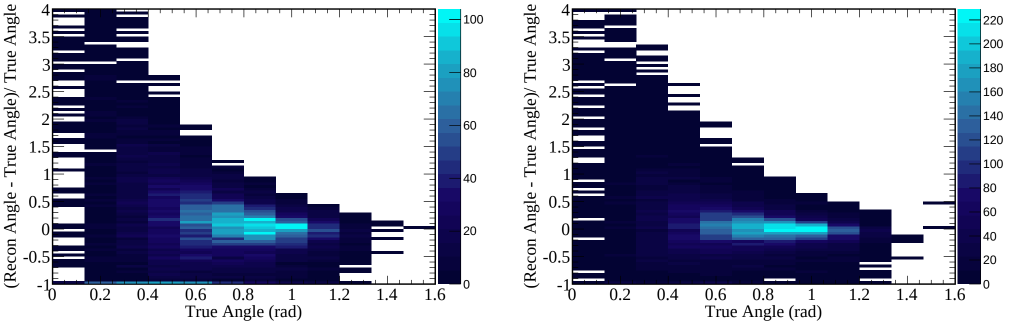

(c) Brl left soft photon angular residual. (d) Brl left hard photon angular residual.

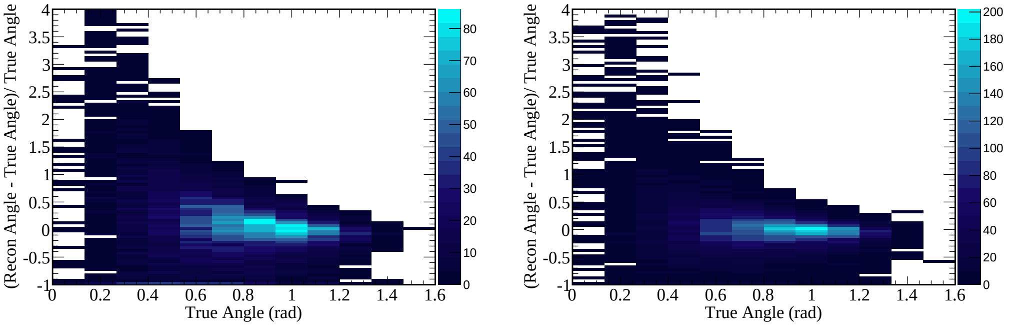

(e) Brl top left soft photon angular residual. (f) Brl top left hard photon angular residual.

Figure 3.10: The ECal angular reconstruction capability (represented as an angular residual), using the thrust axis method, for hard and soft decay photons associated with NC1π^0 interactions. The distributions shown are for the DsECal, barrel left ECal and barrel top left ECal modules. Note, the true angle is with respect to the z-axis of the ECal module.

Using the hypotheses of a minimum ionising particle (MIP), an electromagnetic shower (EM), a showering pion (Pion) and a highly ionising particle (HIP), different PDFs are generated. The MIP hypothesis was modelled using muons with momentum >300 MeV. The EM hypothesis was modelled using electrons. The Pion hypothesis used charged pions showering in the ECal (where the showering caveat was passed if the shower, in truth, was contained in the ECal). Finally, the HIP hypothesis used stopping protons.[18]

An overview of the ECal variables used in Chapter 5 are as follows:

Average Z Position
This variable is simply the non-charge weighted average Z-position of all hits within the cluster.

First Layer
The innermost layer containing cluster hits.

Last Layer
The outermost layer containing cluster hits.

NHits
The total number of hits within a cluster.

Object Length
The reconstructed length of the object (track or shower) in mm.

PID Angle
The zenith angle of the cluster with respect to each detector module (see Section 3.2.3).

PID Asymmetry
A measure of the asymmetry of the cluster, calculated by:

$$\text{Asymmetry} = \frac{3^{rd}\text{ PCA Component}}{2^{nd}\text{ PCA Component}}. \tag{3.13}$$

PID Circularity
A measure of how “circular” the cluster is. First the circularity in each view is calculated:

$$\text{Circularity}_i = (2 \times (2^{nd}\text{ Principal Component})) - 1, \tag{3.14}$$

where i is the view (x or y). The total circularity of the 3D cluster is then given by:

$$\text{Circularity} = \text{Circularity}_x \times \text{Circularity}_y. \tag{3.15}$$

---

[18] Note, for all four hypotheses particle gun MC was used for modelling.

PID ShowerAngle

The angle subtended by the width at the charge centre, from the start of the object:

$$\theta = \tan^{-1}\left(\frac{2^{nd}\text{ Principal Component}}{1^{st}\text{ Principal Component}}\right). \tag{3.16}$$

PID FrontBackRatio

A measure of the charge distribution of the track. This variable is simply the ratio of the charge at either end of the track (calculated using equal length blocks).

PID TransverseChargeRatio

A variable to determine the shower direction based on the distance and charge of hits relative to the principal axis:

$$\text{Transverse Charge Ratio} = \frac{\text{Charge of outer 50\% of hits}}{\text{Charge of inner 50\% of hits}}. \tag{3.17}$$

PID Truncated Max Ratio

A measure of the distribution of charge per layer. This variable is calculated by initially truncating 20% of the highest and lowest charge hits (the optimisation of the truncation value was found using particle gun MC). The PID Truncated Max Ratio is then:

$$\text{Truncated Max Ratio} = \frac{\text{Highest Layer Charge}}{\text{Lowest Layer Charge}}. \tag{3.18}$$

Q RMS

This variable is the root mean square of the charge in the cluster:

$$\frac{1}{\bar{q}}\sqrt{\sum_{i}^{N}\frac{(q_i-\bar{q})^2}{N}}, \tag{3.19}$$

where q_i is the individual hit charge, $\bar{q}$ is the average hit charge and N is the number of hits.

LLR MIP EM

The log likelihood ratio produced using the MIP and EM hypotheses. This variable is used to separate tracks and showers.

LLR MIP Pion

The log likelihood ratio produced using the MIP and Pion hypotheses. This variable is used to distinguish showering pions from MIP-like particles.

LLR EM HIP

The log likelihood ratio produced using the EM and HIP hypotheses. This variable is used to separate protons from electrons.

LLR MIP EM Low Momentum

This log likelihood is similar to the LLR MIP EM variable, however this variable was trained slightly differently as low momentum electrons and muons were used.

Thrust

See Section 3.2.9 and Eq. 3.9.

Thrust Axis

See Section 3.2.9 and Eq. 3.9.

Thrust Origin

See Section 3.2.9 and Eq. 3.9.

Chapter 4

RMM Timeslip Calibration

4.1 Introduction

Accurate hit timing calibration for each sub-detector is imperative for numerous reasons. Firstly, for efficient clustering (such as that explained in Chapter 3), which improves the resulting PID variables' effectiveness. Secondly, for the rejection of noise hits. Thirdly, the determination of an object's direction; this is only possible with highly accurate hit timing. Finally, for the effective matching of global objects.

In this chapter the initial stage of hit timing calibration is considered, in which RMM "timeslips" are located and the relevant offsets applied. The author was responsible for the entirety of this calibration stage from ND280 Run 4 to Run 6 and for "MCM to CTM" (MCM→CTM) slip detection (explained in Section 4.2) from Run 1 to Run 6.

This process is highly beneficial to the analysis presented in Chapter 5 due to the extensive use of ECal variables - and to a lesser extent the use of SMRD variables - in the event selection. The methods described in Sections 4.3.1 to 4.3.9 were developed by [113], whilst the method in Section 4.3.10 was developed by the author.

4.2 RMM Timeslips

RMM timeslips are defined as sudden 10 ns shifts in the time relative to trigger (TRTT) for hits read out by an RMM (see Section 2.2.8 for a description of the electronics hierarchy and the function of the RMMs). This effect arises due to the RocketIO transceivers and optical links that connect the different clock modules of all Trip-T detectors, see Figure 2.16. The transceivers have an internal 100 MHz phase-lock loop (i.e. a 10 ns clock "tick"), hence when a connection is established with a TFB clock (which itself has a clock "tick" of 2.5 ns) the phase-lock can be 1 tick out of synchronicity, inducing an instantaneous 10 ns shift in hits relative to trigger times. The known scenarios when phase-locking takes place

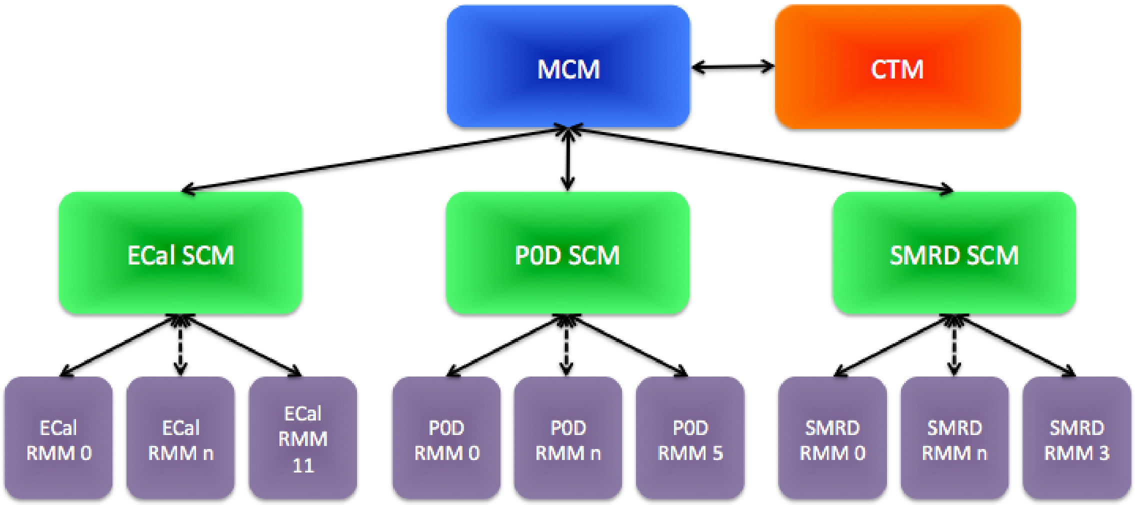

Figure 4.1: A simplified schematic of the apposite modules of the ND280 electronics layout. Here, the coloured blocks represent the different modules and the black arrows show the optical links connected to RocketIO transceivers. As a timeslip can occur at any black arrow, different possible timeslips can occur. If desynchronisation occurs between the CTM and MCM, the TRTT of all RMMs across all sub-detectors are simultaneously shifted in the same direction; if it occurs between the MCM and a specific SCM, all RMMs downstream of that SCM are simultaneously shifted, if it occurs between an SCM and an individual RMM, only that RMM is affected.

are during sub-detector power cycling and when link loss occurs.[1]

As a slip can be induced by any optical link with a RocketIO transceiver, different types of slips can arise. These are MCM→CTM, "MCM to SCM" (MCM→SCM) and "RMM" timeslips, see Figure 4.1. Additionally, it should be noted that in rare circumstances multiple slips can occur throughout the hierarchy, potentially causing 20 ns or 30 ns shifts.

Finding timeslips and applying the relevant offsets, alongside improving the calibration of the sub-detectors affected, benefits inter-detector timing (including non-Trip-T sub-detectors). The application of these offsets is now the first stage of calibration; stages further downstream in the production chain were well established prior to the incorporation of calibrating timeslips, hence this was the least invasive approach.

---

[1] Thus, during the start of run periods, where individual sub-detectors are often power cycled for calibration, timeslips occur much more frequently.

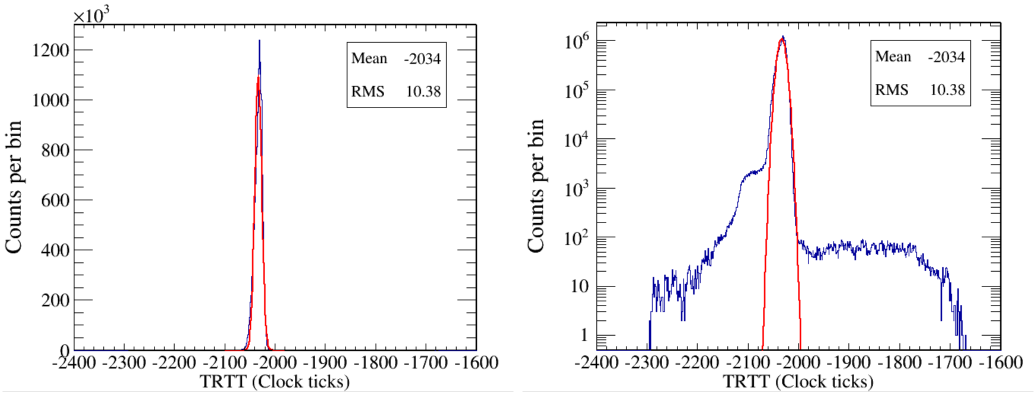

(a) Linear scale. (b) Logarithmic scale.

Figure 4.2: The Gaussian fit (red) of the TRTT (where each clock tick is of length 2.5 ns) for the individual hit times used for truncation. The data shown is of ECal RMM0 for the entirety of Run 5.

4.3 Timeslips Detection

Many stages are required for detection of timeslips: the initial collection of raw data, its subsequent preparation and the final locating of timeslips.

4.3.1 Raw Data Collection

To collect and store the raw data, a "flattened tree"[2] is produced containing cosmic-triggered Trip-T information provided the high gain read out is >900 ADC counts - this is to reduce noise hits and other pedestal level phenomena.[3]

4.3.2 Constraining Hit Times

The first stage is to constrain the hit times, i.e. find the times attributed to cosmic events and remove any unwanted noise passing the >900 ADC counts requirement. Histograms are filled according to the TRTT and Unix time of cosmic events (FGD triggered cosmic events for Run 1 and Trip-T triggered cosmics for Run 2 onwards).[4] This is performed separately on all RMMs, across all sub-detectors. The subsequent distributions are fitted to a Gaussian to calculate the mean and RMS, see Figure 4.2. Truncation is undertaken where any hits outside ±3×RMS of the mean are removed.

---

[2] A lightweight file containing only basic data types.

[3] The information stored includes: the sub-detector in question, the relevant RMM, TFB and Trip-T chip, the unix time and TRTT, along with various other information such as the high and low ADC counts.

[4] The Brl ECal was not installed for Run 1, preventing effective Trip-T triggering.

4.3.3 Calculation of Hit Time Averages

The next stage is to calculate the hit time average of every five minute period, in an attempt to reduce the effect of intrinsic variation. These intervals are used as they provide a good compromise between reducing the variation (larger periods prove more effective) and maintaining the resolution for detecting timeslips (where smaller periods are preferred).

4.3.4 TRTT Histogram Smoothing

Next, the TRTT histograms are smoothed using a boxcar moving average. This produces more accommodating distributions for future algorithms to determine timeslips times. The boxcar average algorithm works as follows:

1. Systematically select the mean TRTT of each five minute period; $\bar{t}_n$.
2. Find the mean of the three previous periods:[5]

$$\bar{t}_{\text{pre}} = \frac{\sum\limits_{i=n-3}^{n-1} \bar{t}_i}{3}. \tag{4.1}$$

3. Repeat the above process for the three subsequent periods:[6]

$$\bar{t}_{\text{post}} = \frac{\sum\limits_{i=n+1}^{n+3} \bar{t}_i}{3}. \tag{4.2}$$

4. Compare $\bar{t}_n$ to $\bar{t}_{\text{pre}}$ and $\bar{t}_{\text{post}}$. If $\bar{t}_n$ lies between the two, its value is maintained. If this is not the case, it is recalculated as the mean of the three-period averages:

$$\bar{t}_n = \frac{\bar{t}_{\text{pre}} + \bar{t}_{\text{post}}}{2}. \tag{4.3}$$

This is repeated a total of three times for each RMM. Figure 4.3 demonstrates the effectiveness of the procedure, as each iteration reduces the noise whilst maintaining, and indeed highlighting, the position and number of timeslips. Three iterations is a compromise between the smoothness of the distributions and the computation time needed to produce the ever-reducing benefit of applying the procedure.

---

[5] This is not possible for the first three periods.

[6] This is not possible for the final three periods.

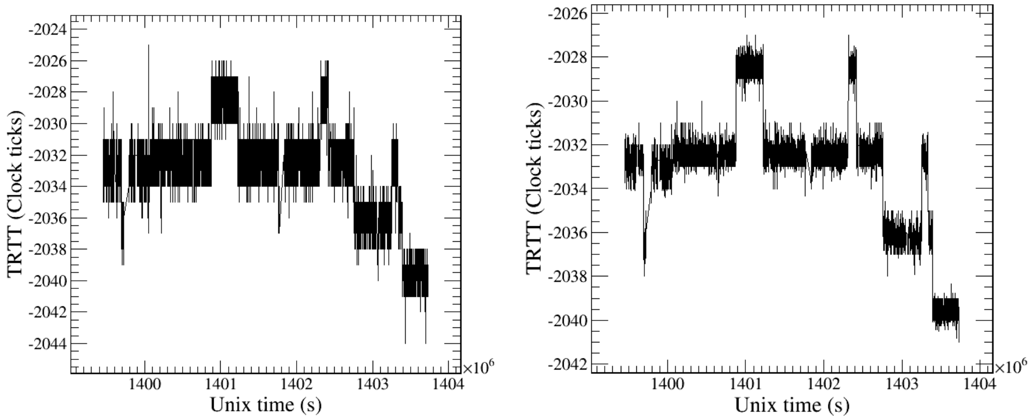

(a) The constrained, unsmoothed data.

(b) One iteration.

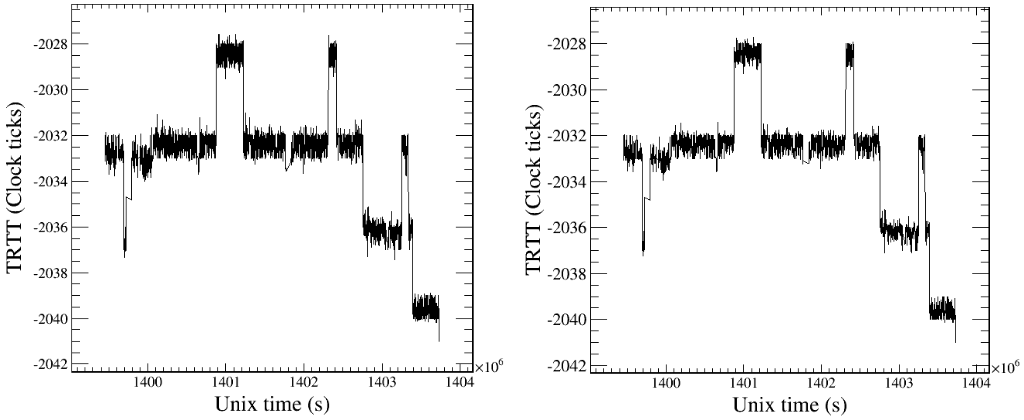

(c) Two iterations.

(d) Three iterations.

Figure 4.3: The effect of applying the boxcar moving average, using 5 minute box widths, on Run 5 data for ECal RMM0. Note, the 10 timeslips in this period become more apparent after each iteration.

4.3.5 Moving Standard Deviation

The first stage, post data preparation, is to calculate the moving standard deviation of the TRTT for each RMM - the slips correspond to discontinuities in these distributions, see Figure 4.4. The method used is as follows:

1. The mean TRTT for each five minute period is sequentially selected.
2. The mean of this, the six previous, and six subsequent periods are calculated:[7]

$$\bar{t}_{range} = \frac{\sum\limits_{i=n-6}^{n+6} \bar{t}_i}{13}. \tag{4.4}$$

3. The moving standard deviation is evaluated:

$$\sigma_n = \sqrt{\frac{\sum\limits_{i=n-6}^{n+6} (\bar{t}_i - \bar{t}_{\text{range}})^2}{13}}. \tag{4.5}$$

Note, by increasing the number of periods from 13 the remaining noise could be reduced further. However in doing this, the granularity of the detection method decreases, thus timeslips closer together will be indistinguishable. Hence, this value is a compromise between these two competing effects.

4.3.6 Peak Search and Timing

The following step is to locate the positions of the peaks and calculate the occurrence times. To do this a peak finding algorithm is used:

1. Each standard deviation is sequentially found; σ_n.
2. The time at which σ_n becomes greater than 1 is designated as the rising limb time; t_{rise}.
3. The time immediately preceding the point where σ drops below 1 is designated as the falling limb time; t_{fall}.
4. The peak is deemed to have occurred in between t_{rise} and t_{fall}, so the average time taken:

$$T = \frac{t_{\text{rise}} + t_{\text{fall}}}{2}. \tag{4.6}$$

Once all peaks have been located, the next stage determines whether they were induced by timeslips.

---

[7] This is not possible for the first or last 7 periods.

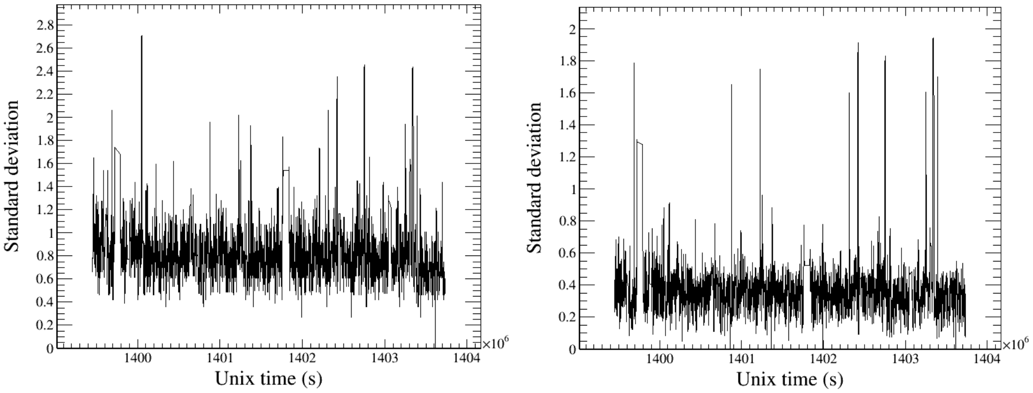

(a) Unsmoothed data.

(b) One smoothing iteration.

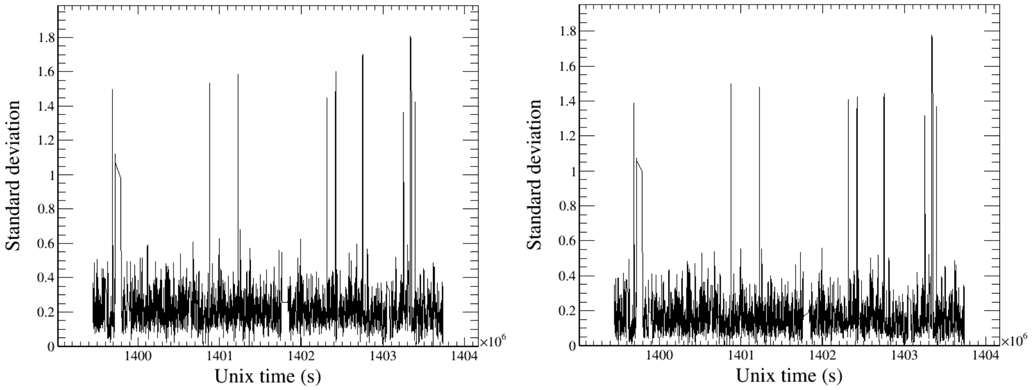

(c) Two smoothing iterations.

(d) Three smoothing iterations.

Figure 4.4: The moving standard deviation of Run 5 for ECal RMM0 after several iterations of smoothing. Note, the data after three iterations of smoothing is used.

4.3.7 Selecting Timeslips

Using the information gained from the previous stage, the following algorithm selects timeslips:

1. Scanning across the raw distributions, the TRTT associated with the point where σ exceeds 1 is found; $\bar{t}_n^{\text{rise}}$.
2. The average TRTT of the six previous periods is calculated:
$$\bar{t}_{\text{rise}} = \frac{\sum\limits_{i=n-6}^{n-1} \bar{t}_i}{6}. \tag{4.7}$$
3. The TRTT associated with the point immediately preceding σ falling below 1 is found; $\bar{t}_n^{\text{fall}}$.
4. The average TRTT for the subsequent 6 periods is calculated:
$$\bar{t}_{\text{fall}} = \frac{\sum\limits_{i=n+1}^{n+6} \bar{t}_i}{6}. \tag{4.8}$$
5. The difference between $\bar{t}_{\text{rise}}$ and $\bar{t}_{\text{fall}}$ determines the slip magnitude:
$$\Delta\bar{t}_{\text{slip}} = \bar{t}_{\text{fall}} - \bar{t}_{\text{rise}}. \tag{4.9}$$
6. If this is larger than 2.4 clock ticks, equivalent to 6 ns, the shift is deemed as a timeslip and stored, with the direction given by the sign of $\Delta\bar{t}_{\text{slip}}$.

2.4 clock ticks is a compromise to account for the noise in the distributions. If increased, a greater number of timeslips will be missed as noise may cause the shifts to appear less than 10 ns in height. Conversely, if decreased, fluctuations producing large standard deviations are more likely to be incorrectly selected. Figure 4.5 shows a variety of timeslips that have been selected along with various shortcomings of the algorithm.

4.3.8 Validation

Once this automated procedure is complete, the next step is to manually check that all timeslips have been located. This is necessary as the procedure is not perfectly accurate, nonetheless it is still extremely useful in that it correctly identifies the vast majority of timeslips - significantly reducing the manual effort required.[8] Any missed or falsely identified slips must be corrected, along with changing the magnitude of slips which are shifted incorrectly (accounting for double, or potentially triple, slips).

---

[8] The overall accuracy of the automated procedure is >95%. During detector turn-on the method is less accurate due to more frequent power cycling. However, later in the run period, the accuracy improves significantly.

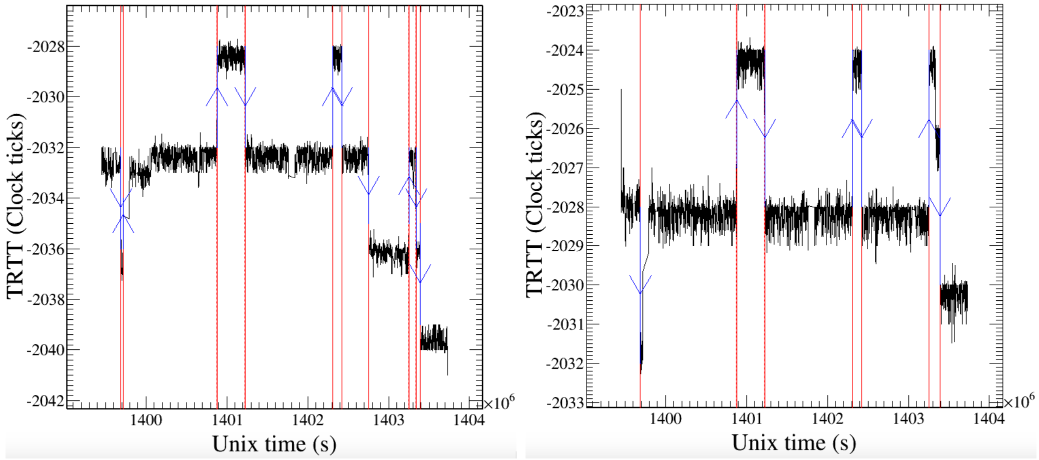

(a) ECal RMM0 (b) ECal RMM7

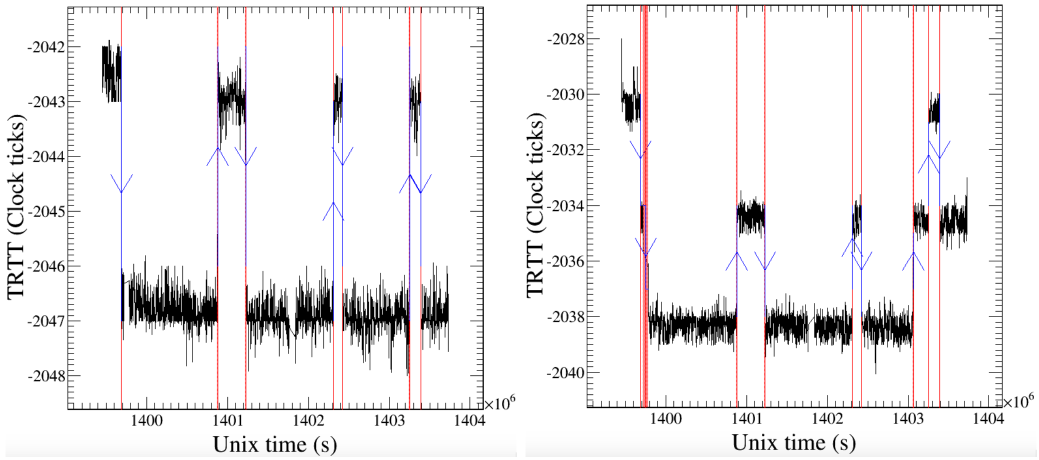

(c) P0D RMM0 (d) SMRD RMM0

Figure 4.5: The selected timeslips for four different RMMs across three different subdetectors for Run 5. Figures 4.5a, 4.5c and 4.5d show the instances in which all have been correctly identified. Figure 4.5b highlights the case where one is missed, at $\sim 1399.8\times10^6$ s, along with a 5 ns slip, at $\sim 1403.5 \times 10^6$ s. Unfortunately, the source of 5 ns slips is still not understood.

4.3.9 Flattening the Timing Distribution

The relevant corrections can now be applied to the raw data. Provided all slips have been correctly identified, this produces flattened timing distributions, see Figure 4.6.[9] An important point when considering the distributions is that sudden spikes often occur, a symptom of the 10-15 minute accuracy of the procedure. However, the BrlECal now has a fully installed light injection (LI) system, thus in future runs there will be improvement of the timing accuracy for BrlECal RMM and ECal MCM→SCM timeslips.

4.3.10 Removing the MCM→CTM Corrections

As explained previously, see Figure 4.1, MCM→CTM timeslips affect all RMMs across all sub-detectors simultaneously. However, such slips do not occur in beam triggered events, as only cosmic triggered events receive timing information from the CTM. Thus, as the procedure uses cosmic triggered events for calibration, if a correction is applied for MCM→CTM desynchronisation there will be an unwanted offset applied between the Trip-T detectors and the other sub-detectors, such as the TPC and FGD. This will affect inter-detector timing and as such affect global matching. To account for this, the following procedure is used:

1. Manually count the number of MCM→CTM timeslips[10]; this information is needed by the peak finder in the next step.[11]
2. Due to the simultaneous nature, when binning timeslips as a function of time (a histogram initially of size 10,000 bins is used),[12] peaks should exist corresponding to MCM→CTM slips. A weight is applied to each addition of the histogram to prevent MCM→SCM timeslips being misidentified as MCM→CTM slips:

$$W = \frac{1}{\text{Number of RMMs connected to the SCM where timeslip occurred}}. \tag{4.10}$$

3. A peak finder searches for the number of peaks with a width > 2 standard deviations (σ) and a peak maximum $> 0.4\times$maximum bin height. These again are compromises as a larger σ, or smaller bin height threshold, would accommodate the 10-15 minute uncertainty. Conversely, a smaller σ, or a larger bin height threshold, would reduce the chance of selecting non-MCM→CTM slips.

---

[9] If this is not the case, further manual changes are necessary.

[10] The number can vary from run to run. However, there are typically between 5 and 10.

[11] An automated counting method could be applied, however MCM→CTM timeslips, particularly during unstable periods, would be difficult to locate this way. Also, as this is a short stage in the process, it does not increase processing time and adds an additional validation of the procedure.

[12] This is a compromise between the larger number of bins needed for a higher resolution and the smaller number of bins to account for the 10-15 minute accuracy of the process which naturally increases the width of the peaks.

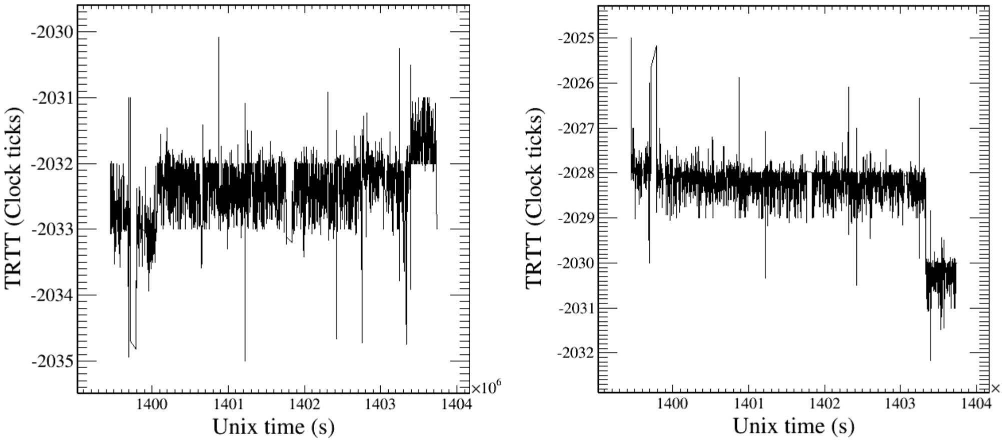

(a) ECal RMM0 (b) ECal RMM7

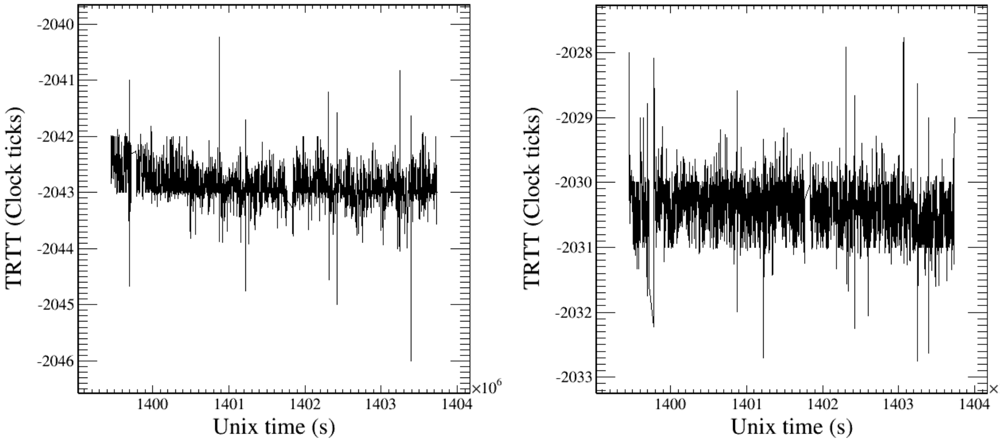

(c) P0D RMM0 (d) SMRD RMM0

Figure 4.6: The flattened distributions for different RMMs on various sub-detectors for Run 5 (note the 5 ns shift unaccounted for in 4.6b). The narrow spikes in the distributions occur due to the resolution in detecting timeslip times. For example, if a timeslip is identified as occurring slightly after the actual timeslip time, this artefact is introduced. However, the duration of such spikes as a proportion of the total time period is negligible.

4. If the correct number of peaks, as defined in the first step, pass the above criteria, all timeslips within 1500 s of each peak centre are removed.

5. If the correct amount of peaks are not found, the number of histogram bins is reduced and the process repeated.

6. Any MCM→CTM slips lying outside the 1500 s range have to be manually removed and any non-MCM→CTM slips reapplied. Once complete, all distributions should have the same flattened shape, with exceptions where MCM→CTM slips occur.

Figure 4.7 demonstrates this process, whilst Figure 4.8 shows the final distributions. At this point the appropriate calibration constants are applied to the data and the Trip-T calibrations, further downstream in the production chain, can begin.

4.3.11 Effect of Applying Timeslip Corrections to final RMM Offsets

Application of timeslip corrections offer a marked improvement in the stability of the final RMM offsets (post inter-TFB and inter-RMM timing calibration). As such they have now been applied to Run 1 through to Run 6.[13] Figure 4.9 demonstrates the effect, with Figure 4.9a showing the Run 3 Tracker ECal RMM offsets without timeslip corrections and Figure 4.9b the same but with the offsets applied.

---

[13] At the time of writing it has indeed been continued on Run 7, though with improvements due to the incorporation of the LI system.

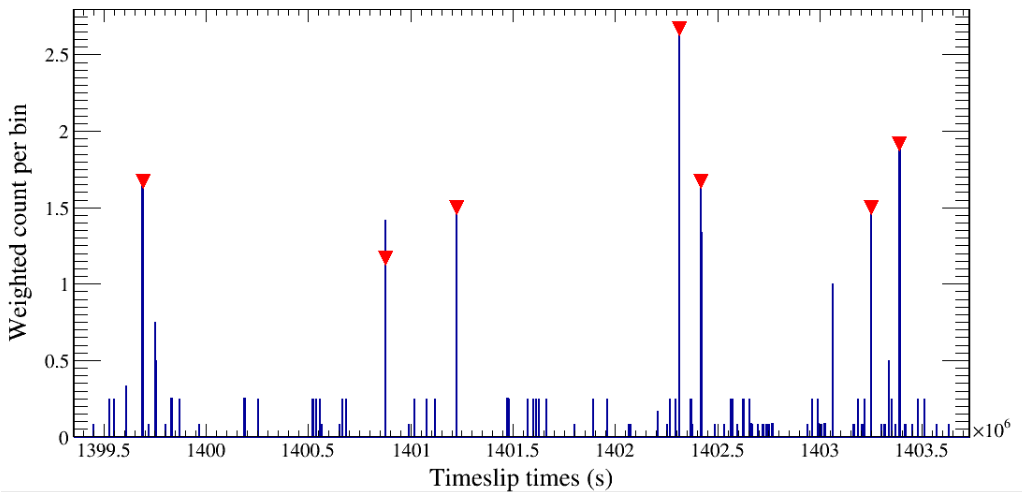

(a) The MCM→CTM peaks found (red arrows) using weighted timeslips.

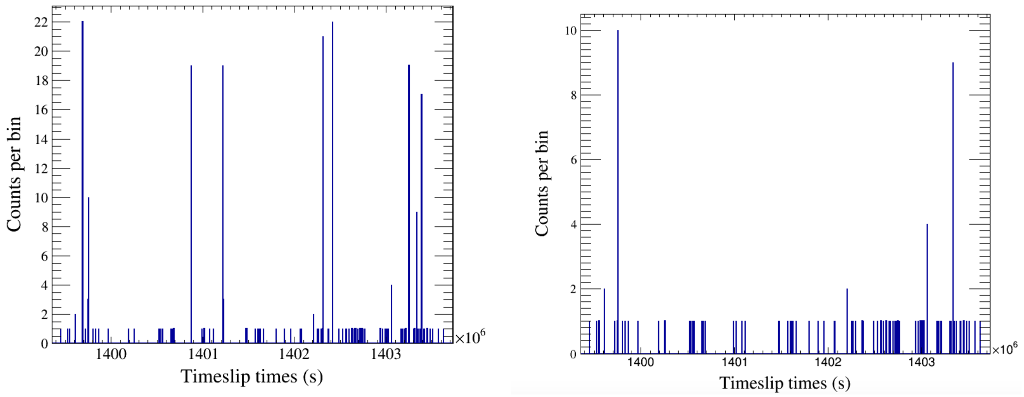

(b) All timeslips (non-weighted) prior to accounting for MCM→CTM slips.

(c) All remaining timeslips (non-weighed) after accounting for MCM→CTM slips.

Figure 4.7: The selection of MCM→CTM peaks (4.7a) and the effect that removing has on the distribution (4.7b and 4.7c). Note, the remaining peaks of 4.7c can be due to RMM and MCM→SCM timeslips.

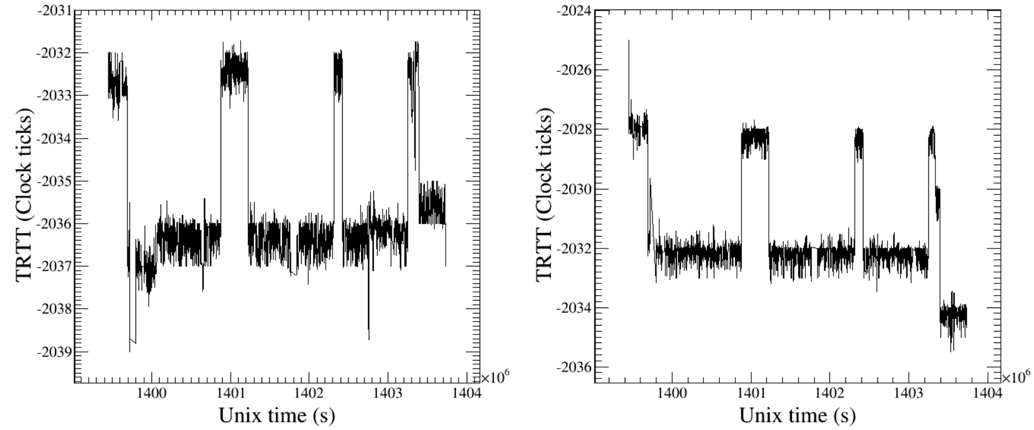

(a) ECal RMM0 (b) ECal RMM7

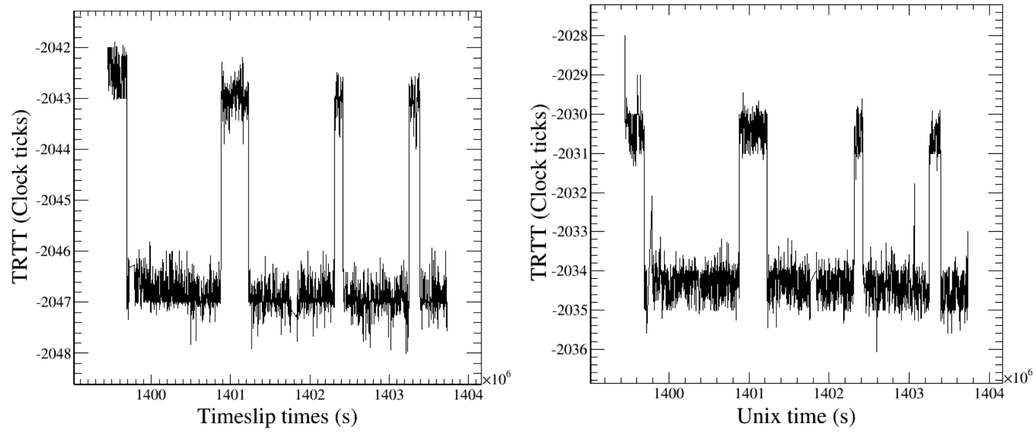

(c) P0D RMM0. (d) SMRD RMM0.

Figure 4.8: The final distributions after accounting for MCM→CTM slips during Run 5 for various RMMs on different sub-detectors. Note, these distributions are the same as Figure 4.5, but with all timeslips except for MCM→CTM slips corrected for.

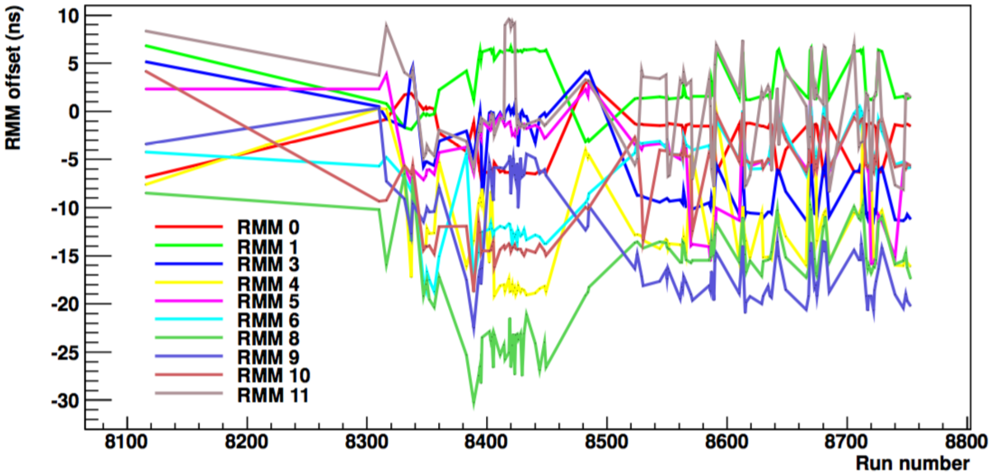

(a) RMM offsets without the calibration of timeslips.

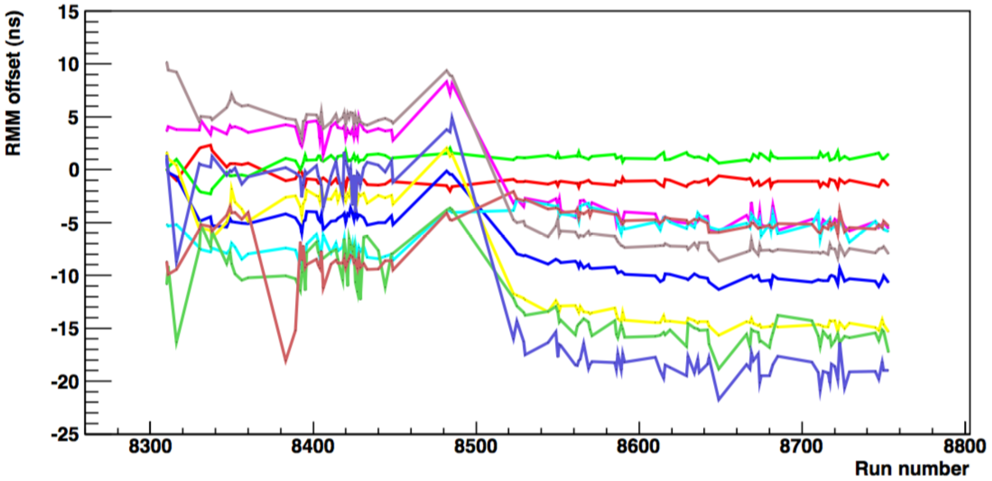

(b) RMM offsets including the calibration of timeslips.

Figure 4.9: Comparison of the distributions of Run 3 Tracker ECal RMM offsets with and without timeslip calibration. Note, 4.9a has fluctuations of the order of 10 ns, whilst for 4.9b this is reduced to 2-5 ns depending on the RMM. Figures taken from [113].

Chapter 5

NC1π^0 Analysis

5.1 Introduction

The second generation analysis presented in this thesis (for the first see [114]) is the exclusive cross-section measurement of NC1π^0 interactions within the FGD fiducial volumes (FV) of the ND280 detector (see Table 5.1 for the global coordinates of the FVs compared to the full extent of the FGDs). Of specific interest is the topology in which the π^0 decay photons convert in the Barrel and/or Downstream (tracker) ECals.[1]

The signal definition considers FSI effects and consists of any neutral current interaction in which a single π^0 exits the target nucleus along with any number of nucleons and recoil photons. As the π^0 decays to two photons with a branching ratio of ~98.8% [27][2], the analysis aims to utilise the tracker ECals to select both. This leads to three distinct topologies: where both photons convert in the BrlECal, where one converts in the BrlECal and the other in the DsECal, and where both convert in the DsECal - henceforth referred to as BrlBrl, BrlDs and DsDs topologies respectively. For an ND280 event display of each, see Figure 5.1.

Table 5.1: The minimum and maximum global coordinates (as defined in the MC simulation), in mm, of each FGD. The corresponding FV values are in brackets.

FGD	X (FV) Min	X (FV) Max	Y (FV) Min	Y (FV) Max	Z (FV) Min	Z (FV) Max
1	-932.17 (-874.51)	932.17 (874.51)	-887.17 (-829.51)	987.17 (929.51)	115.95 (136.875)	446.955 (446.955)
2	-932.17 (-874.51)	932.17 (874.51)	-887.17 (-829.51)	987.17 (929.51)	1473.95 (1473.95)	1807.05 (1799.55)

---

[1] Thus, in the future, studies of different conversion topologies can be undertaken, with the hope of combining multiple dedicate analyses.

[2] Dalitz decay ($\pi^0 \to e^+ + e^- + \gamma$) is the other main decay channel with a branching ratio of ~1.2% [27].

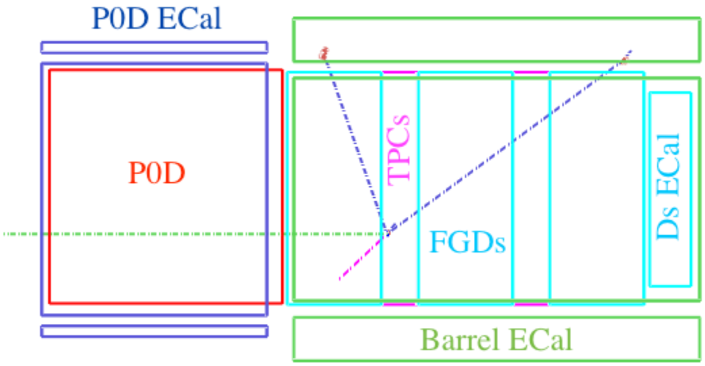

(a) The ‘golden’ BrlBrl topology.

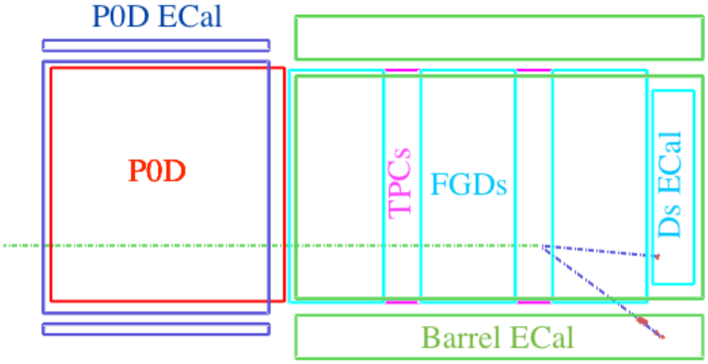

(b) The ‘golden’ BrlDs topology.

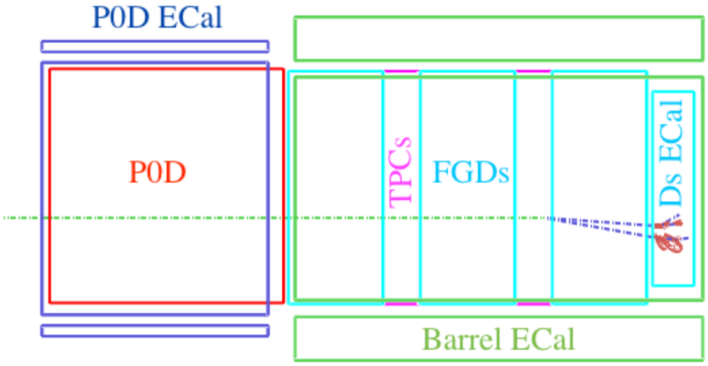

(c) The ‘golden’ DsDs topology.

Figure 5.1: The three ‘golden’ topologies of interest. Here, the ν_μ (green) enters the detector from the left and subsequently interacts in either FGD. At this point a π^0 is produced, which cannot be seen due to the short mean lifetime, along with other possible recoil photons and nucleons - note in 5.1a a neutron (pink) exits the nuclear environment. The decay photons (blue) then convert in the tracker ECals, producing showers of electrons, positrons (both mauve) and photons.

5.1.1 Data Samples

The T2K data periods used in this analysis are Runs 2, 3 and 4. Run 1, which ran from 19/03/2010 until 26/06/10 and constituted of 1.6494$\times 10^{19}$ POT that passed the quality cuts, was not used as the Brl ECal had not been installed. The first condition necessary for good quality POT is that the proton beam is running in a stable condition - this is checked on a spill-by-spill basis. The second condition is that all the individual ND280 detectors are running in a steady mode - this is monitored on a sub-run basis.

There are two configurations possible for the P0D ('water-in' and 'water-out'); as such the runs can be further subdivided. In the case of Run 2 and Run 4 both configurations were used and will be referred to as 'Run 2 Water', 'Run 2 Air', 'Run 4 Water' and 'Run 4 Air'. Conversely, Run 3 ran in the water-out configuration only and as such will be referred to as 'Run 3 Air'. For a full overview of the dates and POT for individual runs, see Table 5.2.

Table 5.2: The periods of each T2K run used in this analysis alongside the respective good quality POT.

Run	Running Period	POT
Run 2 Water	18/11/2010 - 14/02/2011	4.29×10^{19}
Run 2 Air	14/02/2011 - 11/03/2011	3.55×10^{19}
Run 3 Air	27/02/2012 - 09/06/2012	1.56×10^{20}
Run 4 Water	19/10/2012 - 06/02/2013	1.63×10^{20}
Run 4 Air	06/02/2013 - 08/05/2013	1.76×10^{20}

5.1.2 Monte Carlo Samples

For this analysis three MC samples were used. The 'magnet MC' sample was produced using the NEUT event generator and simulates all interactions that occur within the ND280 detector (including the magnet). To ensure agreement with data, it is tuned using the most up-to-date flux estimates. The 'sand MC' sample models interactions within the cavern surrounding the ND280 detector. The particles produced within this region are propagated to the detector (this particularly increases the muon flux when added to the magnet MC). The 'cherry-picked MC' sample is produced using statistically independent NEUT vector files to the magnet MC. Only NC1π^0 interactions are saved and there is only one interaction per spill. This method maintains the POT normalisation whilst providing many more signal events for a given file size (compared to magnet MC) - this makes it an invaluable tool during cut optimisation. Table 5.3 shows: the different MC samples used, the corresponding POT, the Data/MC ratios and the simulated horn current.

Table 5.3: The MC samples used in this analysis with the respective good quality POT, the corresponding Data/MC ratios and the simulated horn currents.

Run	POT	Data/MC ratio	Horn Current
Run 2 Water	1.19×10^{21}	0.036	120 kW
Run 2 Air	9.18×10^{20}	0.039	120 kW
Run 3 Air	3.06×10^{21}	0.051	178 kW
Run 4 Water	3.48×10^{21}	0.047	178 kW
Run 4 Air	3.48×10^{21}	0.051	178 kW
Sand 3 Air	1.19×10^{21}	-	178 kW
Run 4 Water NC1π^0 Cherry Picked	7.76×10^{22}	-	-

5.2 Overview of Selection

As this is a second generation analysis, the initial aim was to review the previous analysis's selection, optimise it for the current MC production, and apply additional cuts, to improve the overall efficiency × purity ($\epsilon\times\rho$), where ϵ is defined as:

$$\epsilon = \frac{\text{Number of MC signal events}}{\text{Number of MC signal events after the 2 Isolated-ECal cluster cut}}, \tag{5.1}$$

and ρ as:

$$\rho = \frac{\text{Number of MC signal events}}{\text{Total number of MC events}}. \tag{5.2}$$

The list of cuts for the first generation analysis is as follows:

1. π^0 candidate - Two isolated-ECal objects in the tracker ECals.
2. 200 ns time cut - The time difference between the two isolated-ECal objects must be $<$ 200 ns.
3. TPC1 veto - There are no reconstructed tracks in TPC1.
4. P0D veto - There are no reconstructed tracks in the P0D.
5. Proton-like tracks - Each reconstructed tracker track is required to be "proton-like".
6. Pre-TMVA - No Michel electrons in the ECal. The invariant mass of the clusters must be $<$ 500 MeV. Cuts to some Multi-Variate Analysis (MVA) inputs are applied.
7. Transformed likelihood - The transformed likelihood cut is applied; trained using 13 variables derived from the ECal objects.

8. Vertex in either FGD - Available information in the TPCs, FGDs and ECals is used to reconstruct an interaction vertex. This is then required to be within either FGD FV.

The list of cuts for this analysis is as follows:

1. π^0 candidate - Two isolated-ECal objects in the tracker ECal.
2. ECal shower cut above 140 MeV - Clusters > 140 MeV must be "shower-like".
3. 10 ns time cut - The time difference between the two isolated-ECal objects is < 10 ns.
4. SMRD activity - There are no reconstructed tracks in the SMRD.
5. TPC1 veto - There are no reconstructed tracks in TPC1.
6. P0D veto - There are no reconstructed tracks in the P0D.
7. ECal charge integral - The total ECal charge of both clusters must be < 200 MEU.
8. BDT - The Boosted Decision Tree (BDT) cut is applied; trained using 33 variables derived from the ECal objects.
9. Vertex in FGD - Available information in the TPCs, FGDs and ECals is used to reconstruct an interaction vertex. This is then required to be within either FGD FV.
10. Muon PID cut - The remaining tracker tracks with both a TPC and FGD component must not be "muon-like".

Throughout the next sections each cut of the new selection and the reasons for, maintaining, changing or selecting new cuts, will be discussed.

5.2.1 2 Isolated-ECal Clusters

The first cut of the analysis requires 2 reconstructed isolated-ECal clusters in the tracker-ECals. This is designed to select both π^0 decay photons, and remains unchanged from the original analysis.

A possible future improvement would be to consider the case of > 2 clusters, as signal events, in which more clusters are reconstructed along with the decay photons[3], would not be discarded. However, due to the combinatoric complexity associated with selecting the correct clusters to form the π^0 candidate, this was not analysed.

---

[3] A scenario which can occur when other FSI particles are reconstructed or in the occurrence of event pile-up.

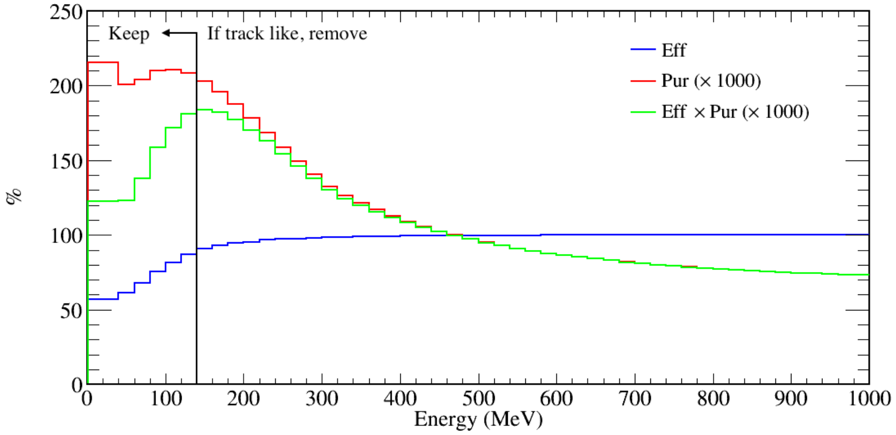

Figure 5.2: Optimisation of the ECal shower cut by maximising $\epsilon\times\rho$ as a function of the minimum energy at which clusters are required to be shower-like. Note, this is based solely on MC simulated events.

Conversely, cases in which < 2 clusters are reconstructed result in an automatic inefficiency that can only be reduced with improvements to ECal reconstruction. This is particularly difficult due to the granularity of the ECals and the low energy of the decay photons.

5.2.2 ECal shower cut

The next step is the ECal shower cut, the first new cut of the analysis. The rationale behind its implementation is that clusters are expected to be shower-like - provided the decay photons are being reconstructed. Nonetheless, due to the difficulty of reconstructing low energy photons, by requiring all clusters to be shower-like may remove events in which NC1π^0 decay photons have been mis-identified as track-like. As such the signal $\epsilon\times\rho$, as a function of the minimum energy at which clusters are required to be shower-like, was maximised and found to occur at 140 MeV - see Figure 5.2. Hence, if one, or both, of the clusters was found to be track-like with an energy above 140 MeV that event would fail to pass the cut. On the other hand, below 140 MeV there is no shower-like requirement, as the reconstruction is not deemed reliable enough at such low energies. This cut resulted in an improvement to the signal-to-background separation of $\sim$250%. Interestingly, by reducing the energy threshold further, the purity can be increased, however this results in a much reduced efficiency.

5.2.3 10 ns time cut

Re-optimisation of the ECal time cut was investigated. Originally, it required the reconstructed time difference between the two isolated ECal objects to be less than 200 ns. This was designed to reduce the chance of selecting clusters from separate bunches. However,

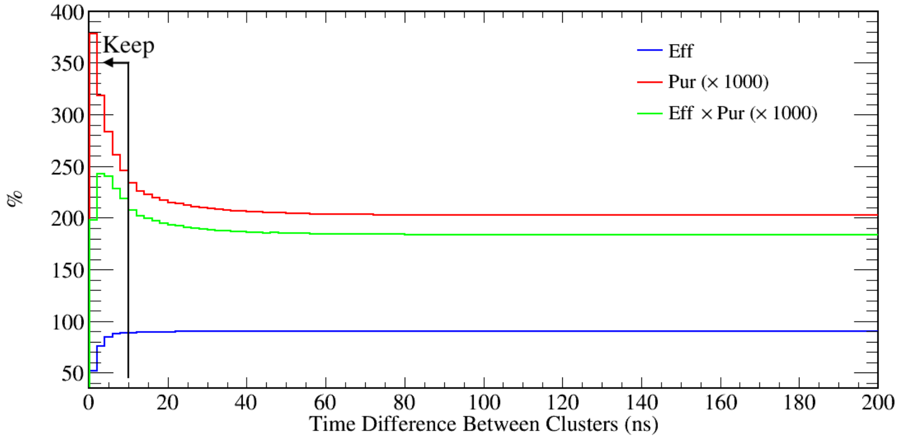

Figure 5.3: Optimisation of the ECal time cut using MC simulated events. Note, a significant reduction from the original time difference is required in order to have an appreciable effect on the signal-to-background separation. Additionally, by reducing the cut there is a possibility of increasing the $\epsilon \times \rho$, however this is dependant on the data/MC agreement.

as the two π^0 decay photons should convert almost simultaneously, it should be possible to reduce the 200 ns requirement, and, in doing so, reduce backgrounds from pile-up and events in which the clusters are created by particles with different parents.

To find the optimal cut, the ($\epsilon \times \rho$) metric was again used - see Figure 5.3. What becomes clear is that the MC requires the tightest possible cut. Unfortunately, this is not possible since at such small differences it is imperative that the data and MC timing agree to a high precision. Although Figure 5.4 shows strong agreement between data and MC, it is not perfect, particularly at small time differences. As a compromise 10 ns was implemented, representing a significant background reduction without introducing large data/MC discrepancies.

5.2.4 SMRD Activity

To attempt to reduce CC background produced inside the basket[4] (in which the exiting muons enter the magnet) and interactions occurring in the magnet itself, an SMRD activity cut was introduced. Two ideas were studied using MC simulated events. Initially the maximum number of SMRD tracks allowed was analysed - see Figure 5.5. This demonstrates that the maximum signal-to-background separation is achieved when no tracks are present. The second idea was to loosen this, and allow a single track, but have a maximum energy deposit - see Figure 5.6. However, the maximum $\epsilon \times \rho$ occurred at the lowest possible energy. Hence, the optimal cut was deemed to be when no SMRD activity is present.

---

[4] The central portion of the detector that is surrounded by the magnet.

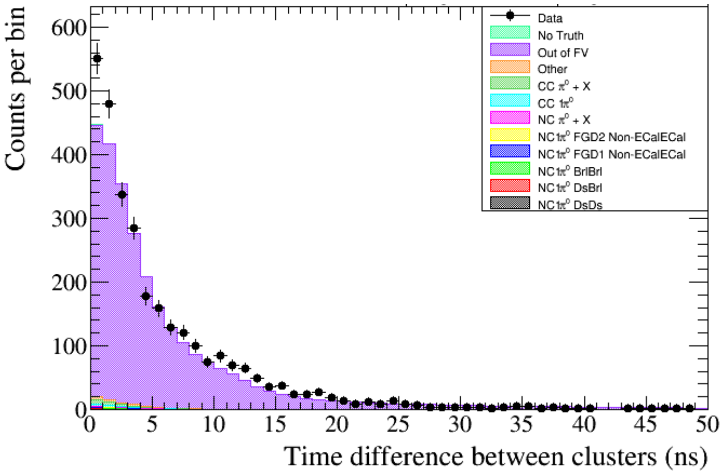

(a) Linear scale.

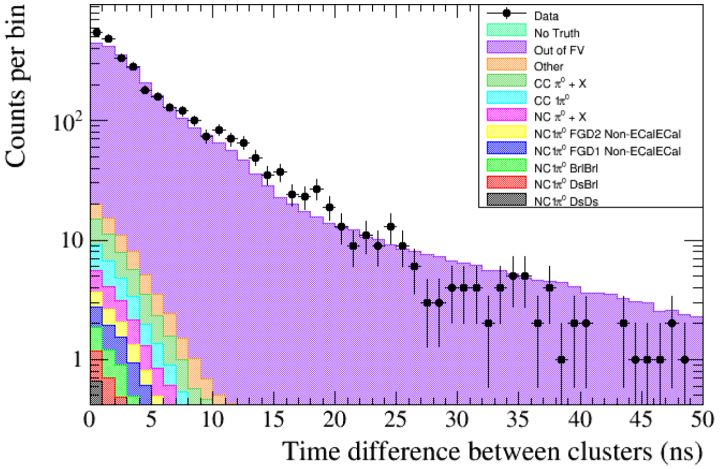

(b) Logarithmic scale.

Figure 5.4: The POT normalised Data/MC comparison of the time difference between the ECal clusters, for a combination of Run 4 Air and Run 4 Water. These show good agreement between the two distributions except at small time differences. Note the data used was separated from the final data used for the analysis. Here, “Out of FV” are any interactions not occurring within the FGD FV. The 5 interaction types preceded with “NC1π^0” give the description of where the two π^0 decay photons convert. “NC π^0 + X” and “CC π^0 + X” are interactions in which at least one π^0 and any number of mesons, exit the nucleus. “Other” represents other interactions not specified, but occur within the FGD FV.

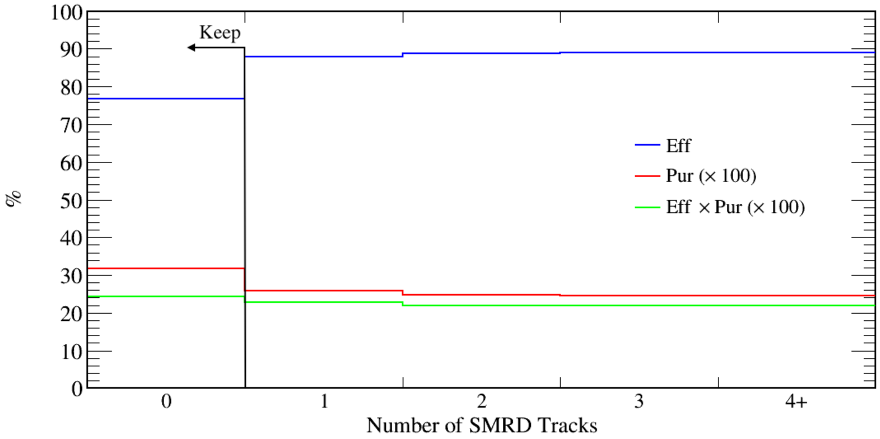

Figure 5.5: $\epsilon{\times}\rho$ as a function of the number of SMRD tracks allowed. Although little improvement can be made, the optimal cut occurs when no activity is present.

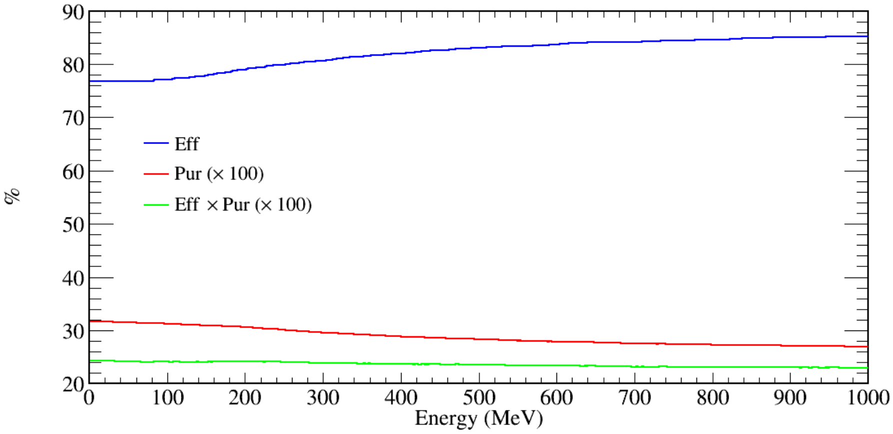

Figure 5.6: The $\epsilon{\times}\rho$, when cutting as a function of maximum allowed energy deposit, when a single SMRD track was permitted. The energy dependence has little effect on the overall signal-to-background separation. However, it reaches a maximum as the SMRD energy reaches zero. As such the optimal cut of no SMRD activity was implemented.

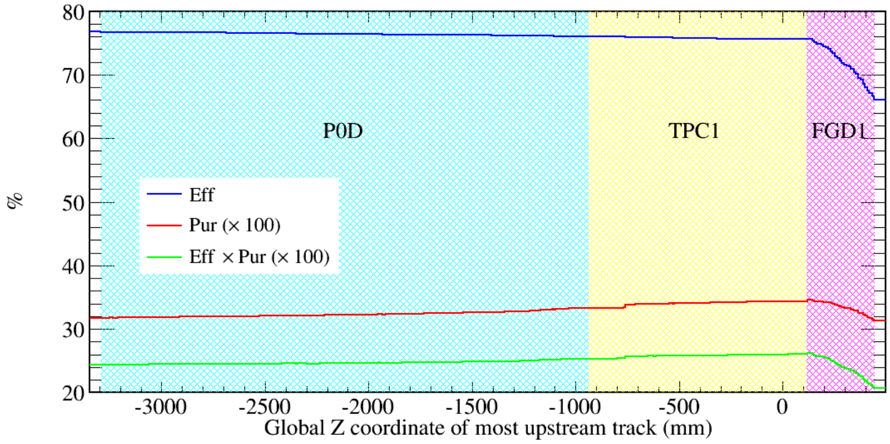

Figure 5.7: $\epsilon \times \rho$ when cutting on the most upstream Z coordinate of the most upstream track in the event. Note, the maximum occurs on the TPC1/FGD1 boundary, suggesting the TPC1 veto is optimal.

5.2.5 TPC1 Veto

To ensure the TPC1 veto is fully optimised, changing the cut position to either increase or decrease the veto region (i.e. either having a small region within the FGD1 included in, or removing some of the TPC1 volume from, the veto region) was analysed using MC simulation. To consider this, the most upstream position of the most upstream track was found. By moving the Z coordinate of the veto and removing events where the most upstream portion of the track occurs further upstream, the corresponding $\epsilon \times \rho$ was calculated - see Figures 5.7 and 5.8. Unfortunately, no improvement can be found. With this method, $\epsilon \times \rho$ is lower than the original combination of the TPC1 and P0D veto. Thus, the TPC1 veto was maintained.

5.2.6 P0D Veto

Due to the results of the studies above, the P0D veto, as with the TPC1 veto, was maintained. Keeping the TPC1 veto yet removing or changing the P0D veto would be counterproductive in increasing the $\epsilon \times \rho$ of the sample.

5.2.7 ECal Charge Integral

The two π^0 decay photons are expected to deposit little charge in the ECals - as much of the neutrino energy can be transferred to nucleons exiting the nucleus during the interaction. On the other hand, background ECal clusters are expected to have a wider range of charge deposited as many different interaction mechanisms can yield background events. This, compounded with a neutrino energy spectrum, which although peaked at 0.6 GeV has a

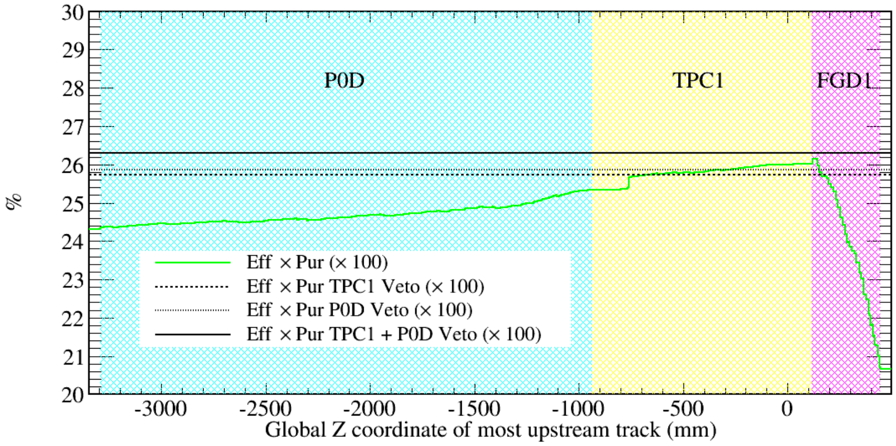

Figure 5.8: The zoomed $\epsilon \times \rho$ when cutting on the most upstream position of the most upstream track in the event, compared to the $\epsilon \times \rho$ of; the TPC1 Veto alone, the P0D Veto alone and the combination of the two. Note, the maximum does not reach the combined value of the TPC1 and P0D Vetos. As such, these vetos were maintained.

long tail (see Figure 2.4), means such background events are able to produce higher energy ECal clusters. As such, a cut, dependent on the charge integral of the two isolated ECal clusters, was analysed using MC simulation. However, there is a compromise between an increasing purity and a decreasing efficiency when reducing the charge threshold. The maximal $\epsilon \times \rho$ occurs at 80 MEU - see Figure 5.9. Yet at this point, particularly for the DsDs topology (which is expected to have more energetic decay photons due to Lorentz boosting), this results in a much decreased signal efficiency. This has consequences for the BDT cut, as a sizeable signal sample is required for training. Thus the threshold was loosened to 200 MEU as this has neglegible impact on the signal efficiency, whilst reducing the background by $\sim 1/3$. Unlike previous cuts - such as the SMRD activity - loosening this is more intuitive as the BDT is not given information relating to the SMRD, although many ECal variables (such as NHits) will be correlated with the ECal charge integral. Consequently, the BDT will make this cut by proxy.

As the invariant mass of the π^0 is reconstructed from the energies of, and angle between, the decay photons:

$$M_{\text{inv}} = \sqrt{2E_h E_s(1 - \cos\theta_{hs})} \tag{5.3}$$

where E is the photon energy, h is the hard photon, s is the soft photon and θ is the angle between the photons, the danger of this cut is that it is an ‘invariant mass cut in disguise’ - i.e. the background is being shaped to look signal-like. Figure 5.10 shows this not to be the case as there is little difference between the invariant mass distribution pre- and post-cut.

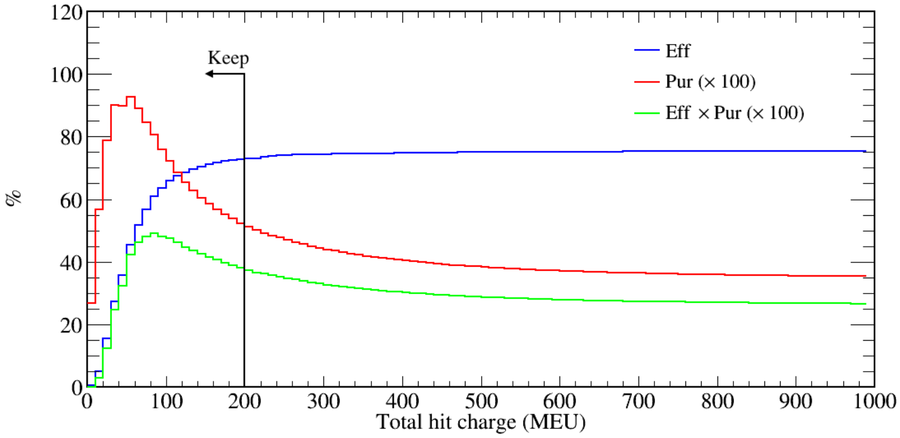

Figure 5.9: The $\epsilon\times\rho$ as a function of the total hit charge of the two ECal clusters. Note, the maximal value of 80 MEU reduced the efficiency too much, resulting in the BDT being overtrained. As such a compromise of 200 MEU was implemented as this increases the purity with very little effect on the efficiency.

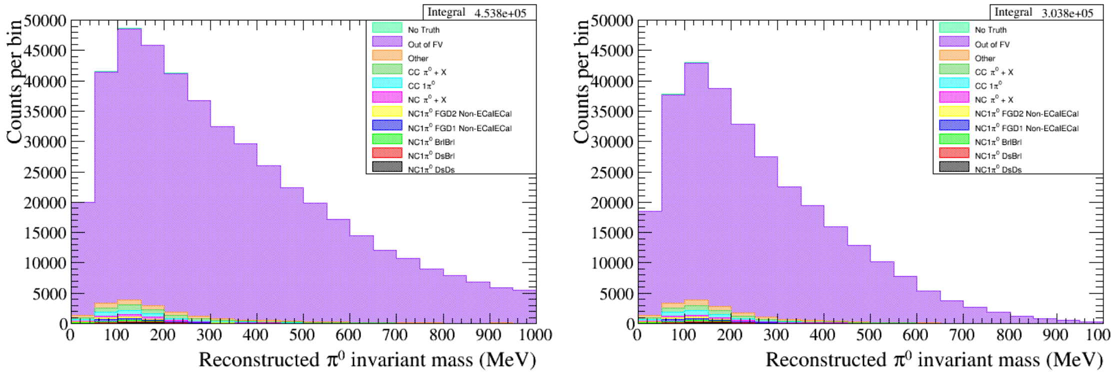

(a) Pre-ECal charge integral cut.

(b) Post-ECal charge integral cut.

Figure 5.10: The reconstructed π^0 invariant mass pre- and post-ECal charge integral cut. Note, with the exception of the first bin, the invariant mass has been reduced uniformly by $\sim$ 5000 events per bin. As such this cut cannot be deemed to be shaping the invariant mass to look signal-like. Here, "Out of FV" are any interactions not occurring within the FGD FV. The 5 interaction types preceded with "NC$1\pi^0$" give the description of where the two π^0 decay photons convert. "NC π^0 + X" and "CC π^0 + X" are interactions in which at least one π^0 and any number of mesons, exit the nucleus. "Other" represents other interactions not specified, but occur within the FGD FV.

5.2.8 BDT

A BDT was selected for each topology as a range of MVA's were tested, using MC simulation, to find which worked best 'out of the box', the logic being that this gives the best starting point. The receiver operating characteristic (ROC) was used as the metric, with Figure 5.11 showing the BDT outperforming all other methods for each topology. The ROC curve is a method to determine the performance of classifiers which make judgements on binary outcomes. Figure 5.12 demonstrates the principles involved. The green curve is an 'ideally performing' classifier, as for any signal efficiency it produces complete background rejection. Furthermore, the orange curve shows a response which is not ideal as there is a trade off between signal efficiency and background rejection - namely for a more efficient signal, the background rejection decreases. Finally, the red curve represents a classifier performing the worst of the three. Again, there is a trade off between signal efficiency and background rejection, however, now a lower efficiency results in an even worse background rejection. Hence, when selecting an MVA to use, the curve most representative of the 'ideally performing' curve is preferred.

A BDT is a binary tree structured classifier in which a range of discriminating variables are used to select 'signal-like' and 'background-like' hypercubes in a multidimensional phase space. This is achieved in numerous steps; initially, all the variables are scanned over, with repeated yes/no decisions posed in an attempt to find the best possible signal-to-background separation - see Figure 5.13. Once the best signal and background separation is achieved for that 'node', the process is repeated for each node produced until a 'stop criterion' is reached. At this point a decision tree has been produced. The stop criteria are user defined values such as the maximum depth of the tree, the minimum size of the nodes, or if the node comprises of 100% signal or background.

A 'boosting' procedure is then used in which misclassified events are weighted. Repeated iterations of scanning and boosting creates many trees - or a 'forest'. Once achieved, by using all trees of the forest, a single classifier is produced to give a response to determine if an event is signal-like or background-like.

5.2.8.1 Separation Criteria

To find which variable and which value provides the greatest signal-to-background separation, a Gini index is used:

$$\text{Gini Index} = P \times (1 - P), \tag{5.4}$$

where P is the signal purity:

$$P = \frac{S}{S + B}. \tag{5.5}$$

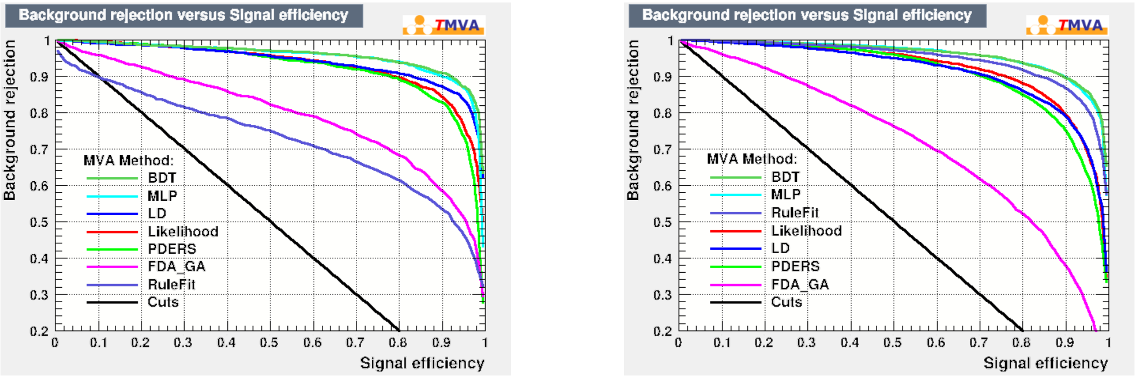

(a) DsDs ROC curves. (b) DsBrl ROC curves.

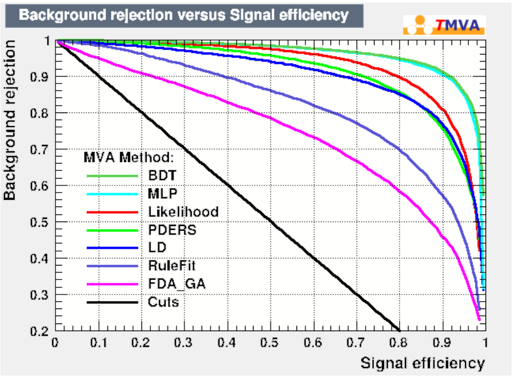

(c) BrlBrl ROC curves.

Figure 5.11: The ROC curves for 8 different MVA techniques for each of the three topologies. In each case the BDT outperforms the MLP, closely followed by the Linear Discriminant (LD), Likelihood, and Probability Density Estimator Range Search (PDERS). The RuleFit and Functional Discriminant Analysis with a Generic Algorithm fitter (FDA_GA) perform slightly worse (except for the DsBrl in which the RuleFit is the third best). Finally, the basic linear Cuts algorithm performs by far the worst in all three topologies. As the BDT performs best, it was selected as the MVA to be used in the analysis. Furthermore, as the others were not implemented, a detailed description is not provided in this thesis. Thus, for a comprehensive overview of each, see [115].

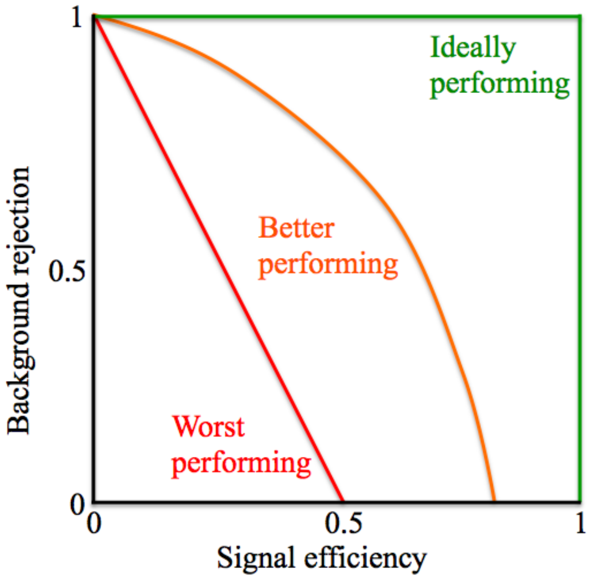

Figure 5.12: An example of a 'worst performing' (red), 'better performing' (orange) and 'ideally performing' (green) classifier.

Here S is the number of signal events and B the number of background events in the node. Note Equation 5.4 is symmetric, this is used as it is equally important to find signal nodes as it is background nodes. Hence the BDT considers each variable in turn and scans over the variable range (in increments set by the user). It then selects the optimal cut by searching for the biggest increase in the difference between the parent node Gini index and the Gini index of the two daughter nodes, summed together after weighting by the fraction of events in each node. Using this method, a variable ranking is produced to find the most important variables in the BDT; the value given is calculated by counting the number of times that variable is used in a node split, weighting each split by the square of the separation gain and by the size of the node that has been split.

5.2.8.2 Boosting

Many boosting algorithms are available, however Adaptive Boosting 'AdaBoost' was used. The first stage of 'AdaBoost' is to calculate the previous tree's misclassification weight. This is defined as the sum of the misclassified event weights, divided by the sum of all event weights:[5]

---

[5] Thus, for the boosting procedure after the first tree is produced, each weight would be equal to one.

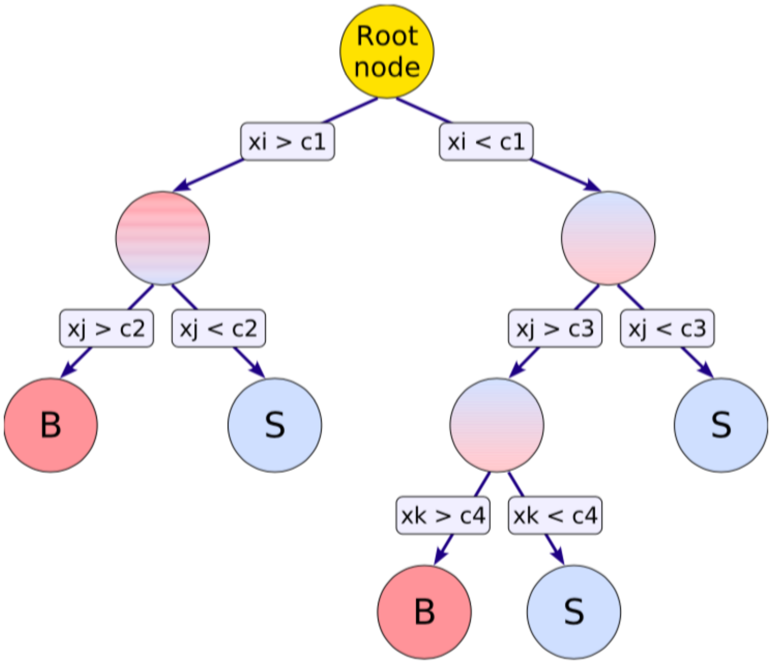

Figure 5.13: A schematic of a single 'tree', in which numerous yes/no decisions are made. This produces multiple nodes at an ever-increasing depth, with an ever-increasing signal-to-background separation. In this example, the root node represents the MC sample with a mixture of signal and background. Each of the discriminating variables are selected in turn and scanned over until the variable (x_i in this case) with the optimal cut ($c1$) is found. Two new nodes are produced at a depth of one, for the samples with $x_i > c1$ and $x_i < c1$. The process is repeated for each node until a 'stop criterion' is reached and the final nodes are allocated as signal-like (S) or background-like (B).

$$\mathrm{err} = \frac{\sum_{i=1}^{\text{Total Misclassified}} \omega_i}{\sum_{j=1}^{\text{Total}} \omega_j}. \tag{5.6}$$

Once the misclassification weight has been deduced, a 'common boost weight' is calculated:

$$\alpha = \frac{1 - \mathrm{err}}{\mathrm{err}}, \tag{5.7}$$

using this, each misclassified event is given a new weight:

$$\omega_i^{t+1} = \omega_i^t \alpha^{\beta} \tag{5.8}$$

where t is the tree and β is the 'learning factor', a value that can be set in order to increase or decrease the rate at which 'learning' occurs - i.e. a higher value will increase the weight of the event. Note, the weights are renormalised in order to preserve the total weight across all trees. The new weights give an adapted purity compared to Equation 5.5,

$$P = \frac{\sum_{i=1}^{S} \omega_i}{\sum_{j=1}^{S+B}} \tag{5.9}$$

which in turn is passed to the Gini index.

Once the forest has been generated, a 'response function' is produced which creates an output for all events, giving a larger value for signal-like, and a lower value for background-like, events:

$$R(\boldsymbol{x}) = \frac{1}{N_{\text{trees}}} \sum_{t=1}^{N_{\text{trees}}} \ln(\alpha_t) R_t(\boldsymbol{x}). \tag{5.10}$$

Here, $(\boldsymbol{x})$ is the ordered list of values for the discriminating variables and $R_t(\boldsymbol{x})$ is a classification by the decision tree that has a value of +1 if it is classed as signal and -1 if classed as background.

5.2.8.3 BDT Training

In this analysis 33 ECal variables were used in training; 15 from the high energy cluster, the same 15 from the low energy cluster and a further 3 that use information from both clusters - see Table 5.4.

Table 5.4: The variables used for BDT training, the numbered label by which they are referred to and a brief description of each.

Variable	Label	Description
PID_Angle(_Low)	1 (16)	The PID angle of the high (low) energy cluster.
AverageZPosition(_Low)	2 (17)	The average z global coordinate of the high (low) energy cluster.
PID_Circularity(_Low)	3 (18)	The circularity of the high (low) energy cluster.
FirstLayer(_Low)	4 (19)	The first layer hit of the high (low) energy cluster.
PID_FrontBackRatio(_Low)	5 (20)	The front back ratio of the high (low) energy cluster.
LastLayer(_Low)	6 (21)	The last layer of the high (low) energy cluster.
PID_LLR_MIP_EM(_Low)	7 (22)	The log likelihood response for the minimum ionising particle hypothesis of the high (low) energy cluster.
PID_LLR_MIP_EM_LowMomentum(_Low)	8 (23)	The log likelihood response for the minimum ionising particle and low momentum hypothesis of the high (low) energy cluster.
NHits(_Low)	9 (24)	The number of hits in the high (low) energy cluster.
ObjectLength(_Low)	10 (25)	The length of the high (low) energy cluster.
PID_ShowerAngle(_Low)	11 (26)	The shower angle of the high (low) energy cluster.
Thrust(_Low)	12 (27)	The thrust of the high (low) energy cluster.
ThrustAxis(_Low)	13 (28)	The z component of the thrust axis of the high (low) energy cluster.
PID_TransverseChargeRatio(_Low)	14 (29)	The transverse charge ratio of the high (low) energy cluster.
PID_TruncatedMaxRatio(_Low)	15 (30)	The truncated max ratio of the high (low) energy cluster.
Energy_Asymmetry_Between_Clusters	31	The energy asymmetry between the two clusters $E_A = \frac{E_h - E_l}{E_h + E_l}$.
Angle_Between_Clusters	32	The angle between the two clusters.
Distance_Between_Clusters	33	The distance between the two clusters.

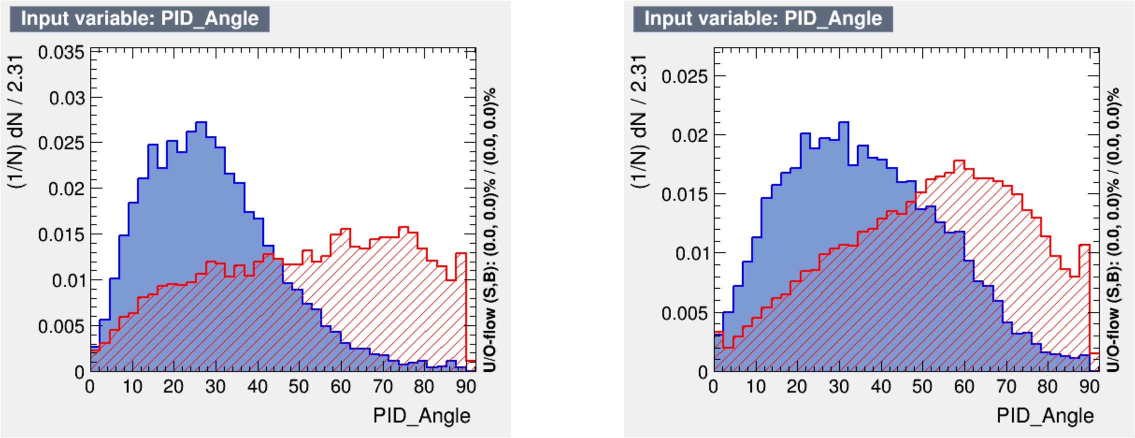

(a) DsDs topology.

(b) DsBrl topology.

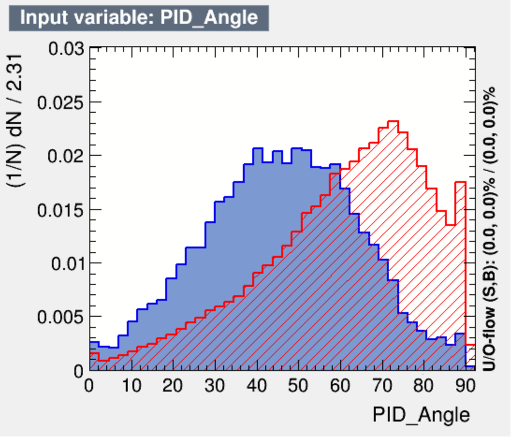

(c) BrlBrl topology.

Figure 5.14: The PID angle of the high energy cluster for signal (blue) and background (red) events. Note, each topology was trained separately as the distributions often vary greatly between the topologies; the signal and background clusters tending to higher angles when moving from the DsDs to DsBrl to BrlBrl distributions.

Due to the different kinematics of the topologies (see Figures 5.14 - 5.16)[6] the training for each was done separately in order to maximise the effectiveness of the cut. Furthermore, as the BDT was trained using MC simulation, for it to work optimally, it is important that there is strong data/MC agreement - Figures 5.17 - 5.19 show this to be the case.[7]

During training, to try to encapsulate the ideal topology, the signal definition passed to the BDT was events in which the true topology position and the reconstructed position of the clusters matched - e.g. the true photons both convert in the DsECal and the two clusters are both reconstructed in the DsECal.[8] If this was not the case, but the event

---

[6] For all distributions for each topology, see Figures A1 - A9.

[7] For all data/MC comparisons for each topology, see Figures A10 - A18.

[8] Note, this can be further improved by stipulating that the ECal clusters are truth matched to the decay photons to provide an even 'cleaner' signal.

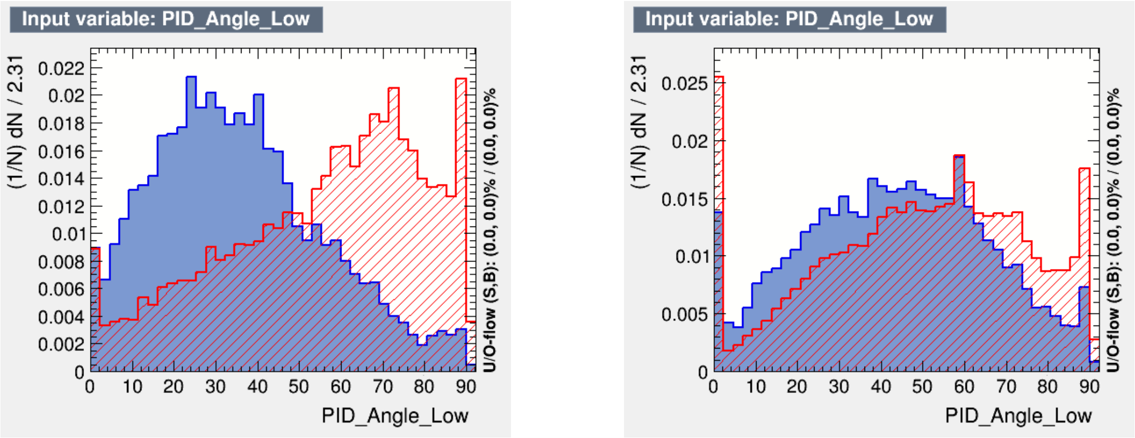

(a) DsDs topology.

(b) DsBrl topology.

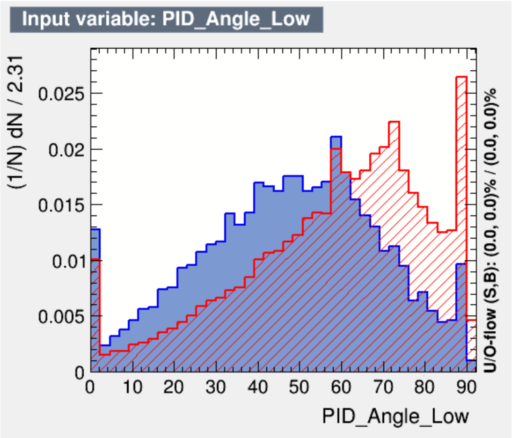

(c) BrlBrl topology.

Figure 5.15: The PID angle of the low energy cluster for signal (blue) and background (red) events. Note, each topology was trained separately as the distributions often vary greatly between the topologies; the signal and background clusters tending to higher angles when moving from the DsDs to DsBrl to BrlBrl distributions.

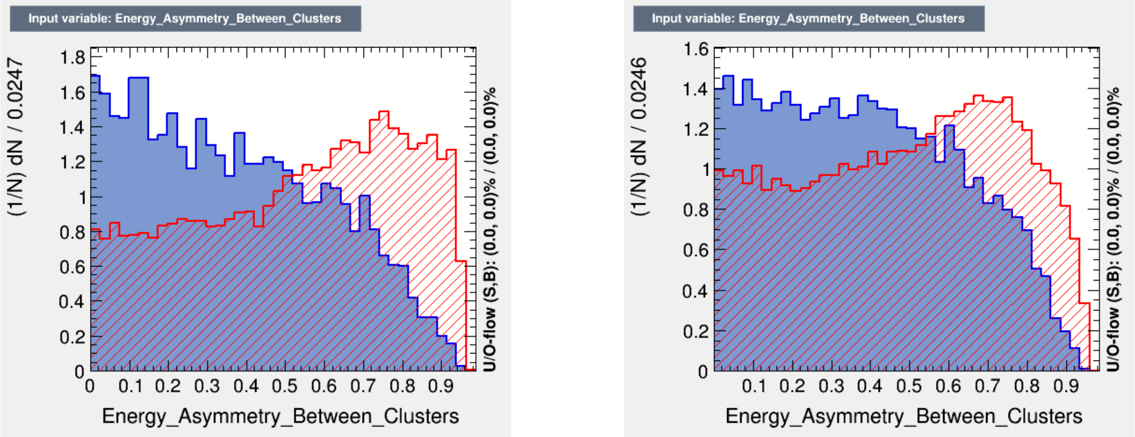

(a) DsDs topology.

(b) DsBrl topology.

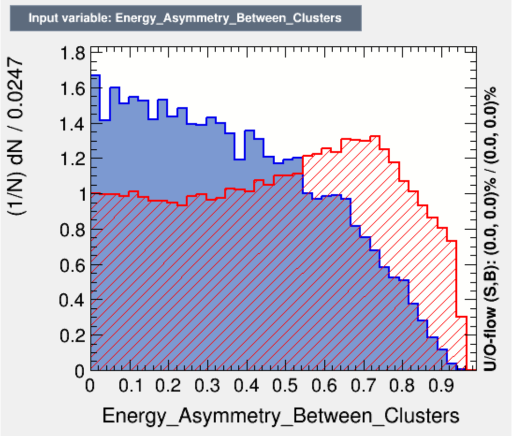

(c) BrlBrl topology.

Figure 5.16: The energy asymmetry between clusters for signal (blue) and background (red) events. Note, each topology was trained separately as the distributions often vary greatly between the topologies. However, for this variable, this is not the case as all three have similar signal and background shapes. Hence, if all distributions showed this similarity between topologies it would be possible to use a single BDT.

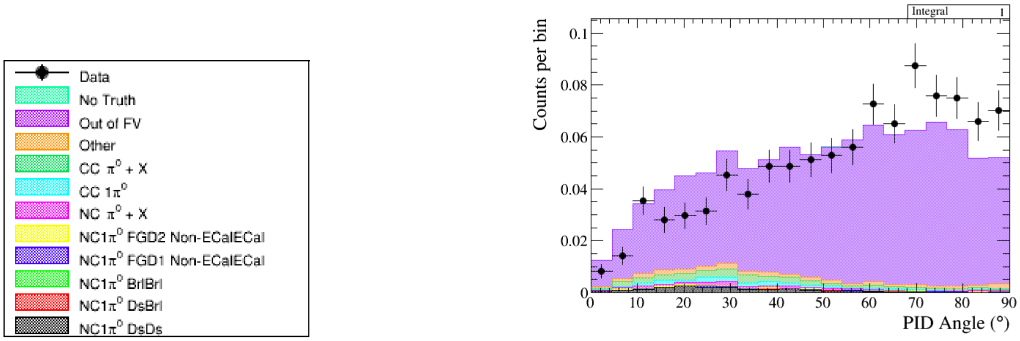

(a) DsDs topology.

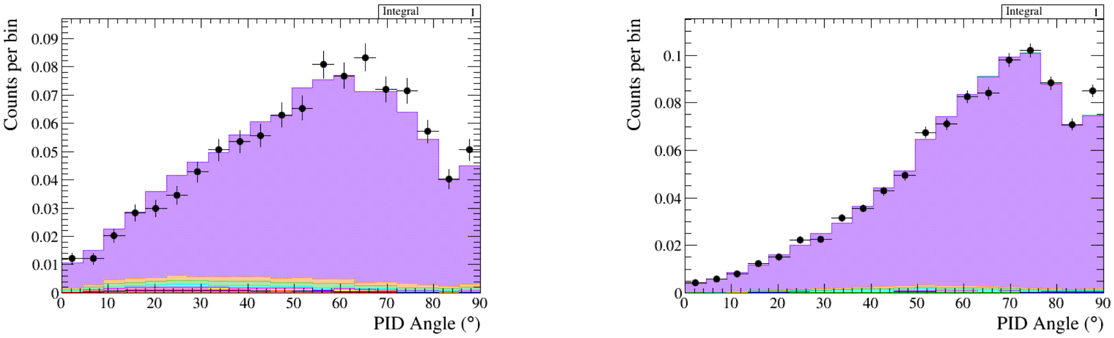

(b) DsBrl topology.

(c) BrlBrl topology.

Figure 5.17: Data/MC comparisons of the PID Angle for the high energy clusters. Here, “Out of FV” are any interactions not occurring within the FGD FV. The 5 interaction types preceded with “NC1π^0” give the description of where the two π^0 decay photons convert. “NC π^0 + X” and “CC π^0 + X” are interactions in which at least one π^0 and any number of mesons, exit the nucleus. “Other” represents other interactions not specified, but occur within the FGD FV.

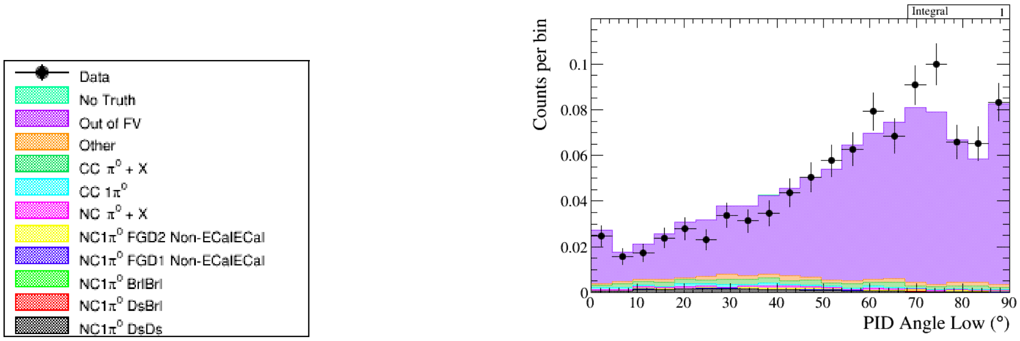

(a) DsDs topology.

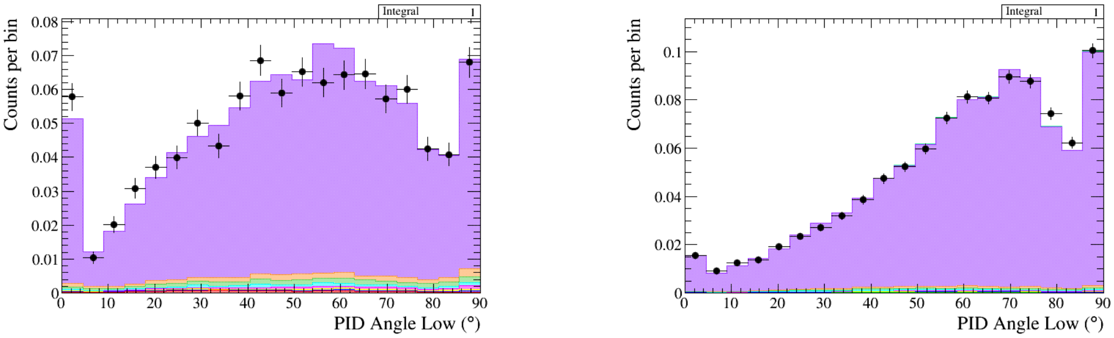

(b) DsBrl topology.

(c) BrlBrl topology.

Figure 5.18: Data/MC comparisons of the PID Angle for the low energy clusters.

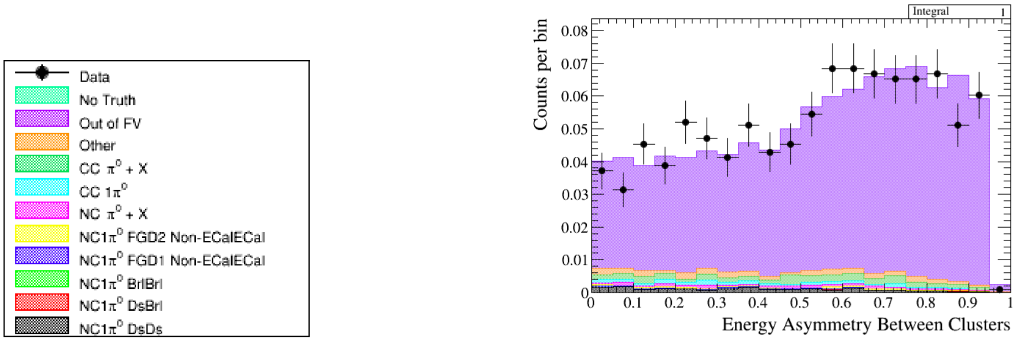

(a) DsDs topology.

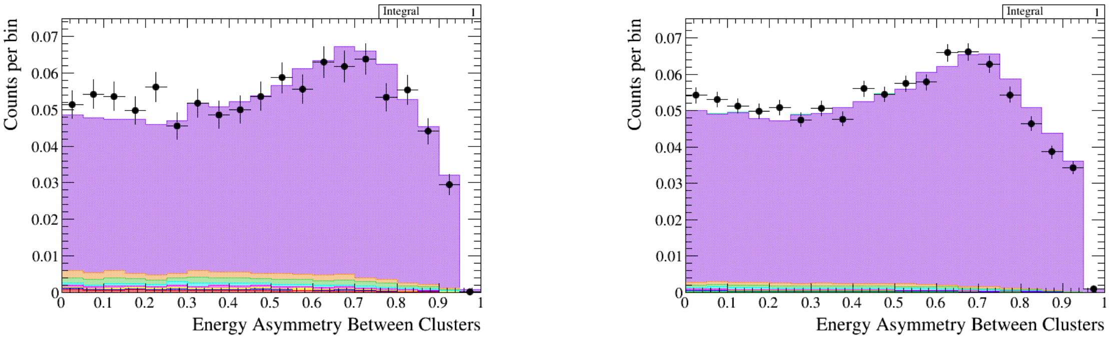

(b) DsBrl topology.

(c) BrlBrl topology.

Figure 5.19: Data/MC comparisons of the energy asymmetry between the clusters.

Table 5.5: The parameters used during BDT training. Note, the tree depth was reduced for the DsDs sample in order to prevent overtraining.

Parameter	Value
NTrees	1200
β	1
Max Tree Depth	3 for DsBrl and BrlBrl, 2 for DsDs
Cut granularity	20
Training/Testing ratio	80:20

contained a true signal, the event was not passed to the BDT. Finally, the background sample was passed to the topology corresponding to the reconstructed positions - e.g. if there were two DsECal clusters, this would be passed as a background DsDs topology.

Once trained, the response for each of the three topologies - see Figure 5.20 - show that none are overtrained. Namely, the test[9] and training sample distributions for both the signal and background match. This is most clear for the BrlBrl sample, which contains the most events for training, and less apparent for the DsDs topology, with the fewest. On the other hand, the DsDs topology BDT performs strongest and the BrlBrl the weakest - with the overlap between the DsDs signal and background distributions less. For the parameters used during BDT training, see Table 5.5.

Post-training, the rankings and signal and background correlation matrices demonstrate that a reduction in the number of variables for a similar effectiveness is possible by removing the least important variables within highly correlated groups - see Table 5.6 and Figures 5.21-5.23. For example, in the DsDs topology (Figure 5.21) variables 7 and 8 are highly correlated. Hence by removing variable 7, as it has a lower variable ranking of 19 compared to 4, it may be possible to maintain a similar signal-to-background separation of the BDT, whilst using a small set of variables. However, the maximum $\epsilon \times \rho$ is expected to increase with an increasing number of variables, as the BDT has more 'choice' at every node in every tree. Thus, providing the systematic uncertainty increase with an increasing number of variables does not compromise the increase in signal-to-background separation; a reduction in the number of variables is not necessary - see Section 5.5.5.

Variable Number	DsDs Position	DsBrl Position	BrlBrl Position
1	5	11	7
2	2	20	10
3	15	29	28
4	1	7	3
5	32	24	23
6	26	17	11
7	19	23	26
8	4	5	9

---

[9] The MC used by the BDT is split so that the response of a small separate 'test sample' can be compared to that of the 'training sample'.

9	7	2	4
10	30	18	30
11	24	22	19
12	14	12	25
13	12	10	2
14	9	13	20
15	23	9	17
16	16	21	14
17	6	14	16
18	29	33	32
19	8	3	5
20	33	30	29
21	27	19	12
22	18	27	27
23	11	6	24
24	10	4	13
25	31	32	33
26	20	28	13
27	25	31	21
28	22	16	18
29	28	26	22
30	25	25	15
31	17	15	6
32	21	8	8
33	3	1	1

Table 5.6: The parameter importance rankings in each topology.

Once trained, a weights file is produced providing the linear cut sequence that each event will undertake. However, the BDT uses cherry-picked MC, makes the assumption that an equal number of signal and background events are passed (1000 of each, see Figure 5.24), and optimises $\frac{S}{\sqrt{S+B}}$, when finding the ideal response cut. As such a re-optimisation procedure was used. Using the newly produced weights files, the entirety of the magnet MC training sample (as this simulates all interactions that occur within ND280 - including the magnet) was re-optimised with the response cut selected corresponding to the maximum $\epsilon \times \rho$, see Figures 5.25. This yielded values of: 0.06, 0.03 and 0.02 for the DsDs, DsBrl and BrlBrl topologies respectively. This represents a much purer sample (with a slightly reduced efficiency) compared to the BDT optimised values.

Post training, the BDT was implemented as the final cut in the selection. This was due to concerns about the tracker track parameter space affected by the cut and the unknown effect this would have on the validity of the TPC systematic calculations.

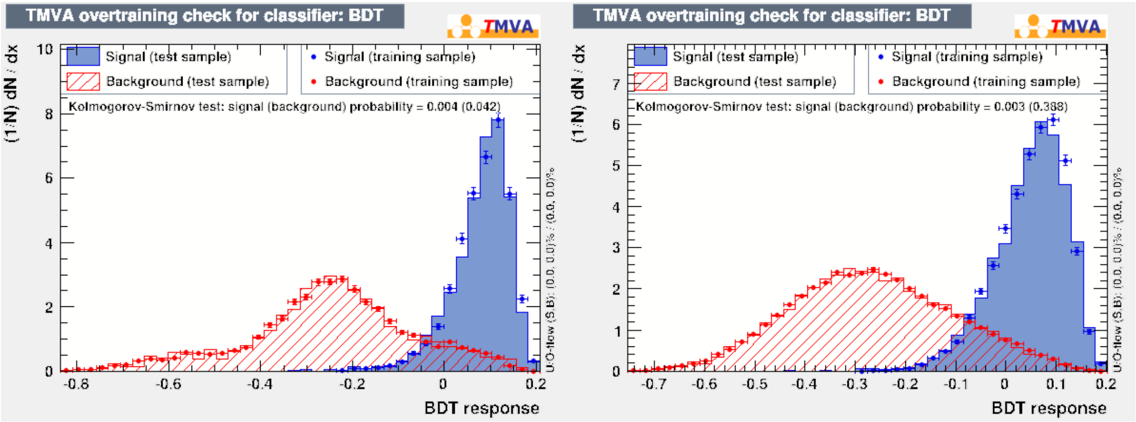

(a) DsDs topology. (b) DsBrl topology.

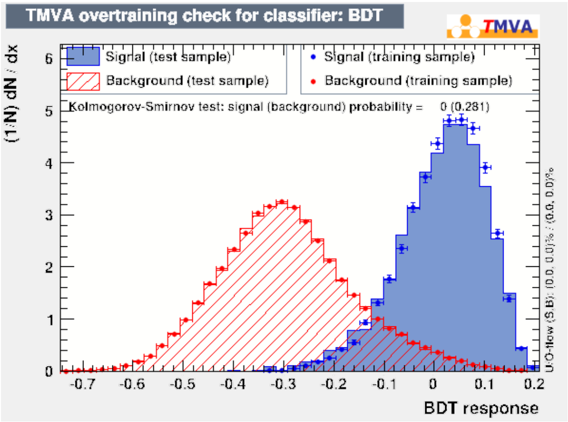

(c) BrlBrl topology.

Figure 5.20: The response distributions after training for each of the three topologies.

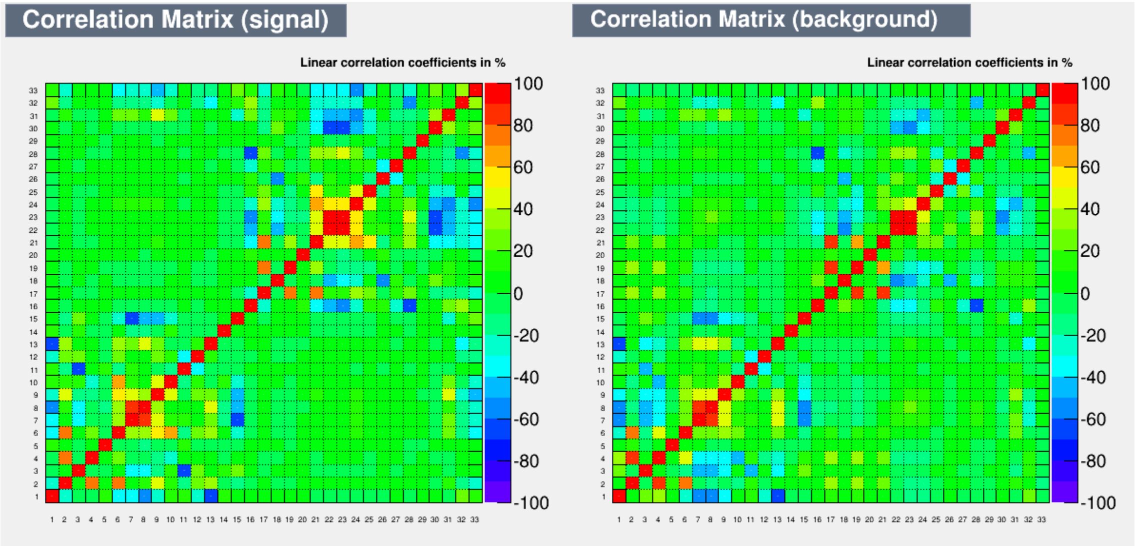

(a) Signal.

(b) Background.

Figure 5.21: The signal and background correlation matrices of the ECal variables used to train the DsDs topology.

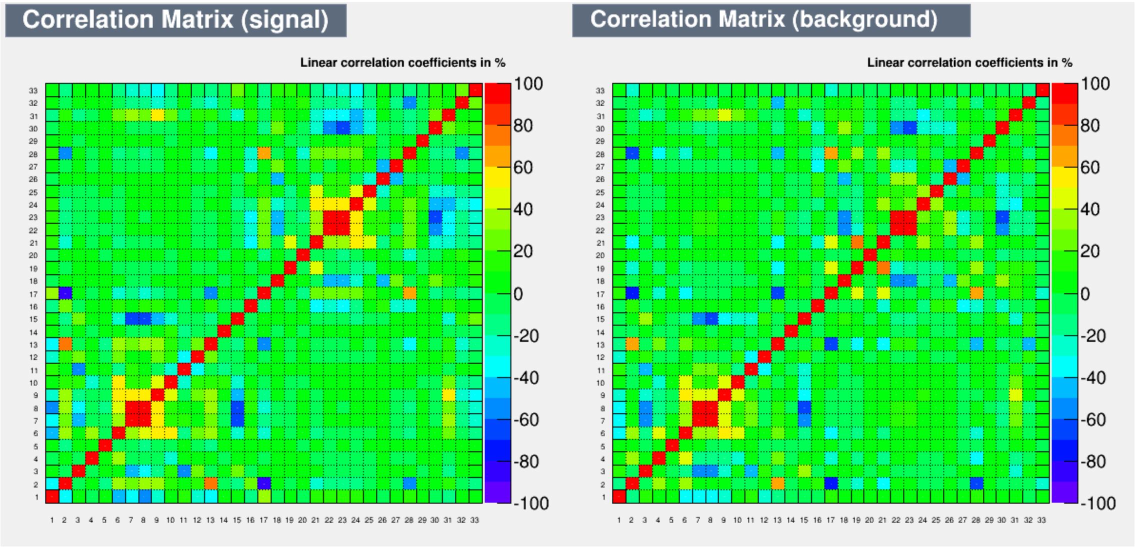

(a) Signal.

(b) Background.

Figure 5.22: The signal and background correlation matrices of the ECal variables used to train the DsBrl topology.

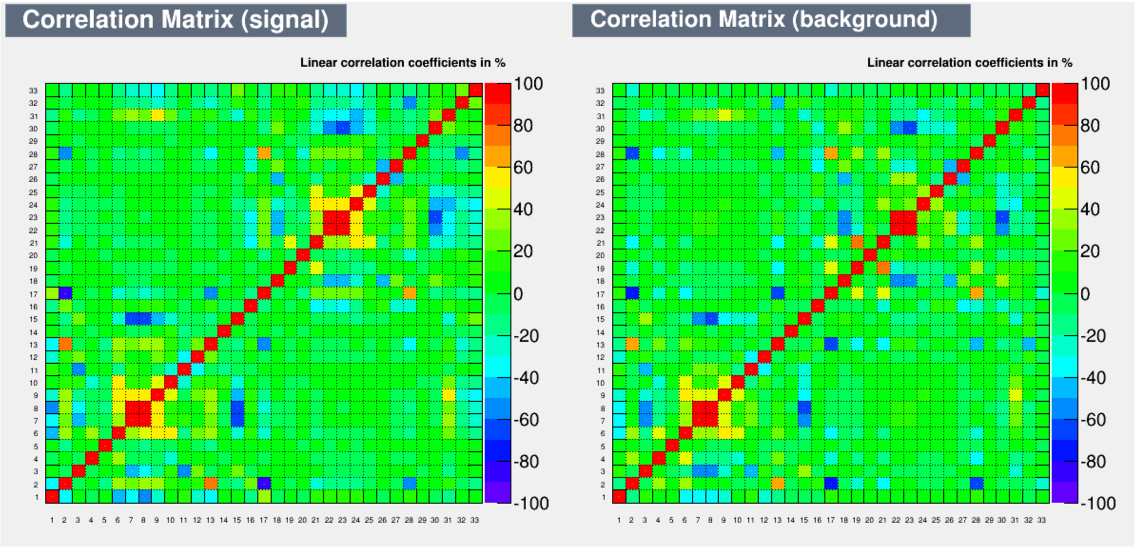

(a) Signal.

(b) Background.

Figure 5.23: The signal and background correlation matrices of the ECal variables used to train the BrlBrl topology.

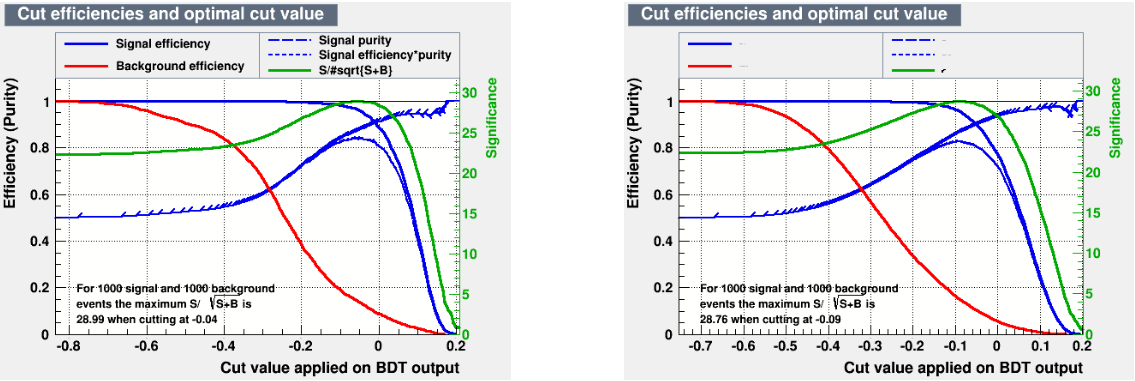

(a) DsDs

(b) DsBrl

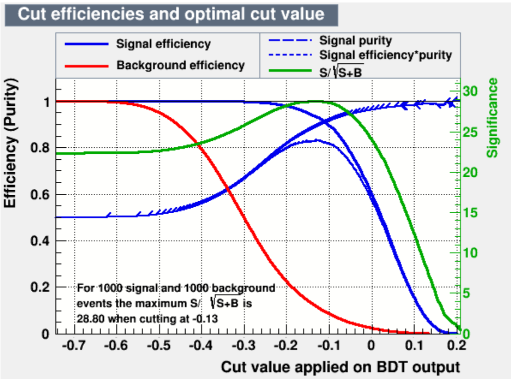

(c) BrlBrl

Figure 5.24: The optimal response cut values suggested by the BDT package for each topology. These were not used and a re-optimisation procedure was undertaken.

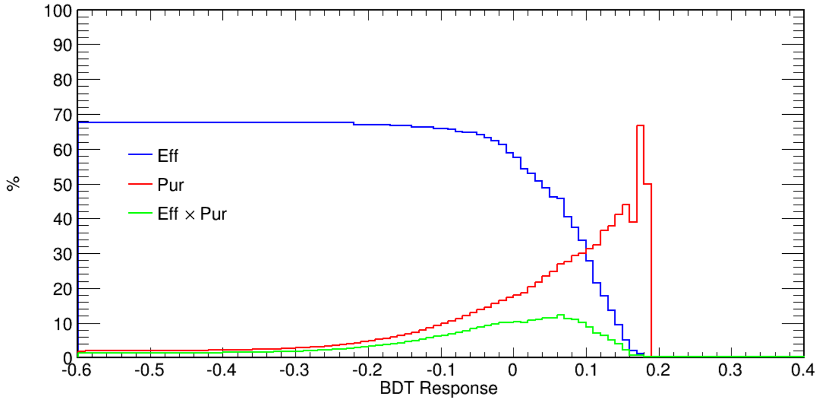

(a) DsDs

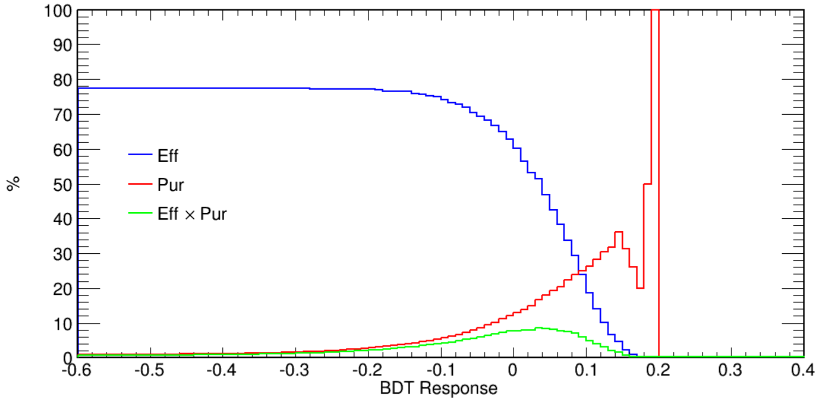

(b) DsBrl

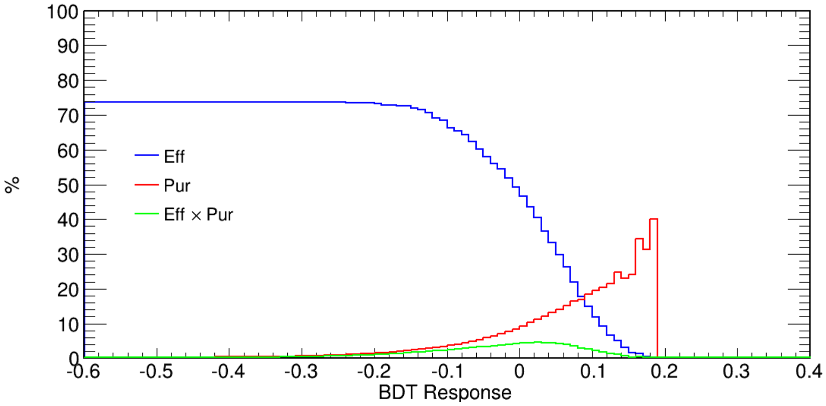

(c) BrlBrl

Figure 5.25: The optimised response values for each topology, produced by passing the magnet MC through the weights files and optimising $\epsilon \times \rho$.

5.2.9 FGD Vertexing

One of the major challenges associated with this analysis is the difficulty in accurate vertex reconstruction. Due to the low energy of the recoiling nucleons, often there is no tracker track - the start point of which can be used as a reconstructed vertex. Furthermore, due to the low energy nature, there is often no unused hit information in the tracker to indicate the interaction point. As a result, it is often necessary to use the direction of the ECal clusters to project back to the FGD to reconstruct a vertex; the inaccuracy of the method was demonstrated in Figure 3.10.

The algorithm for selecting events based upon vertex reconstruction, taken from the first generation analysis, is as follows:

1. If TPC tracks are present, the start position of the highest momentum candidate is used as the vertex.[10] If this is within the FGD FV the event passes to the next stage of the selection.
2. If no TPC tracks are present, the start position of the highest momentum FGD isolated (FGDIso) track is used.[11] Again, if this is within the FGD FV, the selection is passed.
3. If there are no reconstructed tracks within the tracker, unused FGD hits are used. The FGD with the highest total unused hit charge is selected as the candidate FGD and the charge weighted position of the unused hits within this FGD is calculated. Provided this is within the FGD FV, the selection is passed.
4. When no tracker information is present the thrust axis of the highest energy ECal cluster is extrapolated back towards the xy-plane with the plane's z coordinate corresponding to the centre of an FGD. The selection is passed if the x and y position of the plane intercept is within the x and y FV coordinates. To select an FGD to project towards, the position of the cluster is considered. If its z coordinate is further upstream than the central z coordinate of FGD1, FGD1 is the candidate FGD. If it is further downstream than the centre of FGD2, FGD2 is the candidate. If it is between these two limits, the thrust direction is used to choose the FGD; if the z coordinate of the thrust is positive, FGD1 is the selected FGD whilst if it is negative, FGD2 is selected. If the high energy cluster does not extrapolate to an FGD, the low energy cluster is used. If neither extrapolate to an FGD centre plane, the selection is not passed.

The justification for this procedure is apparent when considering Figures 5.26 - 5.29, with the most accurate method used when possible.

Several attempts were made to improve this algorithm by optimising, with MC simulated events, both the unused hit information and ECal cluster vertexing methods. Due

---

[10] If there is only one TPC track, its start position is used.

[11] If there is only one isolated FGD track, its start position is used.

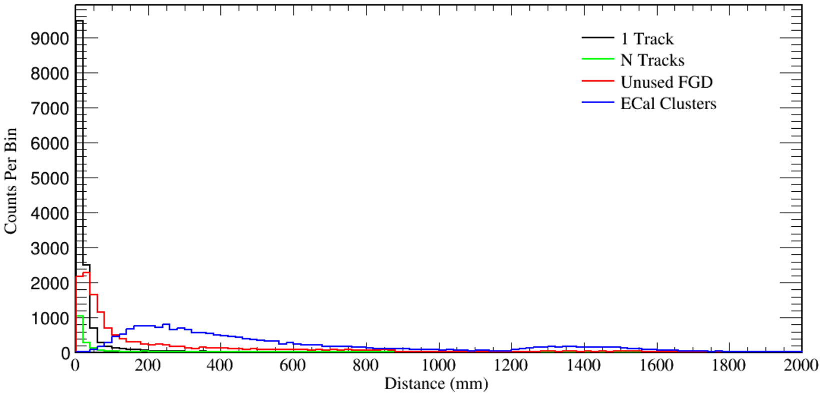

(a) Linear scale.

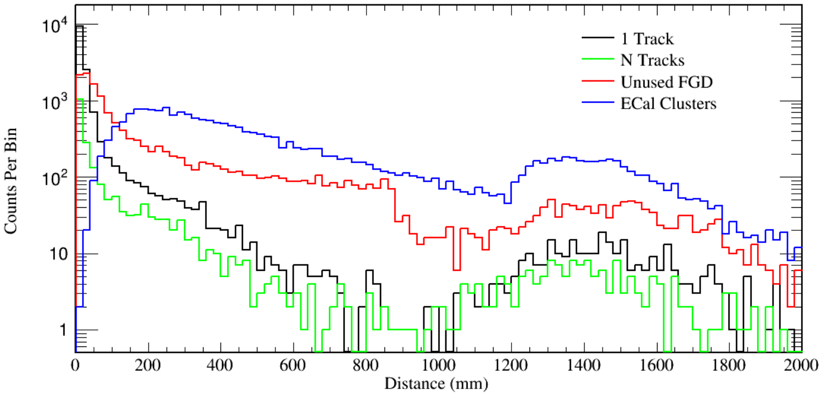

(b) Logarithmic scale.

Figure 5.26: The total distance between the reconstructed and true MC vertex position for the different techniques in the vertexing algorithm. Here, “1 Track” is where only one tracker track is present (TPC or FGDIso), “N Tracks” is where multiple tracker tracks are present (TPC and/or FGDIso), “Unused FGD” is where unused FGD hit information is used and “ECal Clusters” is where the vertex is found using the projection of the ECal clusters.

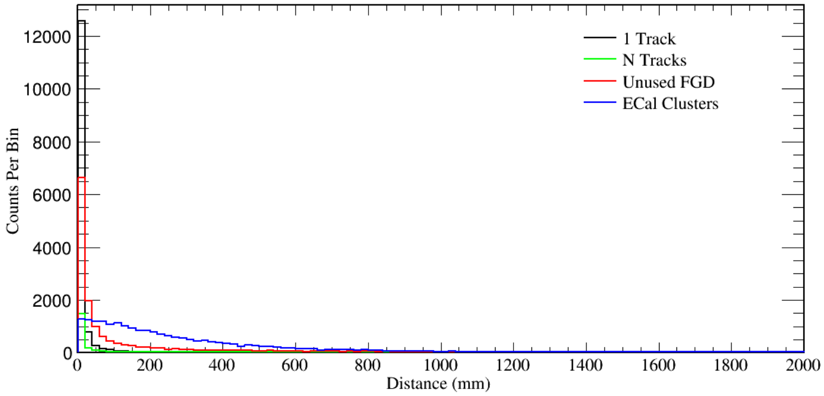

(a) Linear scale.

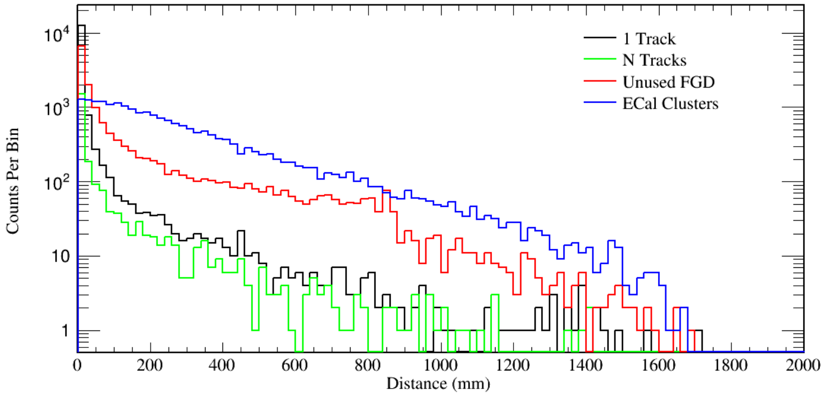

(b) Logarithmic scale.

Figure 5.27: The x-coordinate distance between the reconstructed and true MC vertex position for the different techniques in the vertexing algorithm. Here, “1 Track” is where only one tracker track is present (TPC or FGDIso), “N Tracks” is where multiple tracker tracks are present (TPC and/or FGDIso), “Unused FGD” is where unused FGD hit information is used and “ECal Clusters” is where the vertex is found using the projection of the ECal clusters.

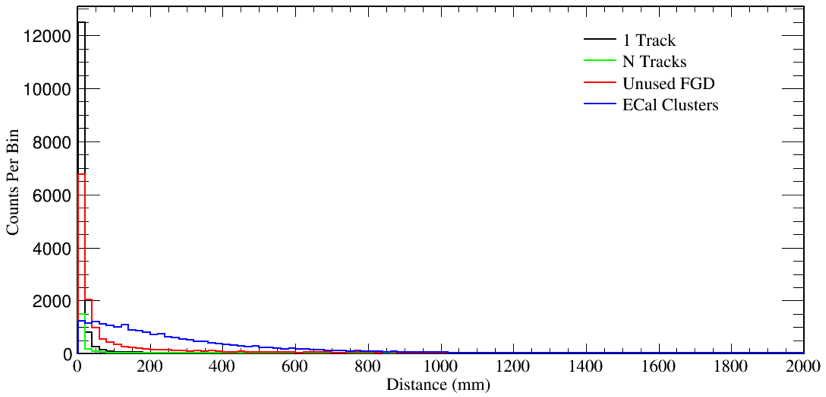

(a) Linear scale.

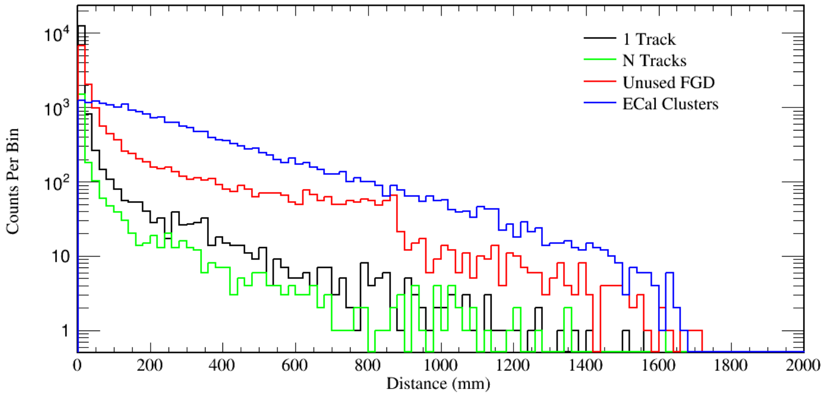

(b) Logarithmic scale.

Figure 5.28: The y-coordinate distance between the reconstructed and true MC vertex position for the different techniques in the vertexing algorithm. Here, “1 Track” is where only one tracker track is present (TPC or FGDIso), “N Tracks” is where multiple tracker tracks are present (TPC and/or FGDIso), “Unused FGD” is where unused FGD hit information is used and “ECal Clusters” is where the vertex is found using the projection of the ECal clusters.

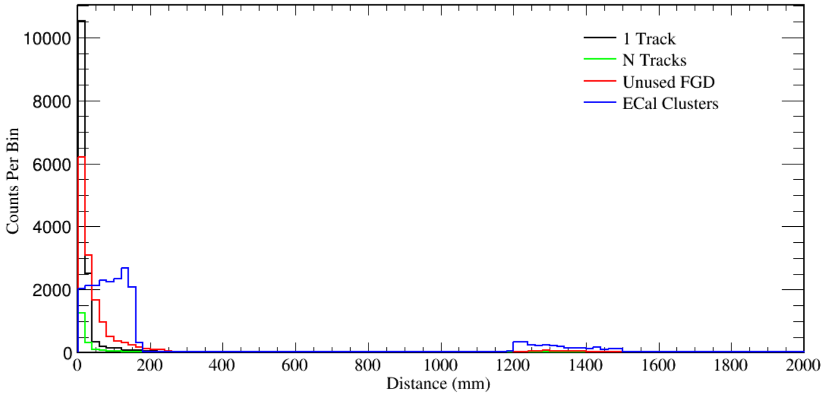

(a) Linear scale.

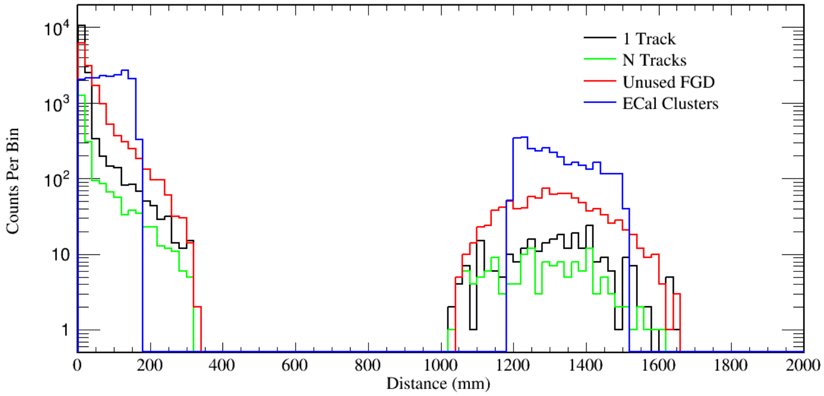

(b) Logarithmic scale.

Figure 5.29: The z-coordinate distance between the reconstructed and true MC vertex position for the different techniques in the vertexing algorithm. Here, “1 Track” is where only one tracker track is present (TPC or FGDIso), “N Tracks” is where multiple tracker tracks are present (TPC and/or FGDIso), “Unused FGD” is where unused FGD hit information is used and “ECal Clusters” is where the vertex is found using the projection of the ECal clusters.

to the accuracy of the starting position of the tracker tracks there is little scope for improvement, as such no optimisation techniques were undertaken.

By analysing the unused FGD hit information, it was found that vertexing is least accurate when lower charges are deposited - see Figure 5.30. Hence, the hypothesis that, in this low energy regime, the ECal clustering method is more accurate was pursued. However, comparing Figure 5.30 to Figure 5.26 shows this not to be the case - even when the unused hit information is least accurate, in the 0 $<$ to $\leq$ 200 MEU regime, the distance between the true and reconstructed hit positions peaks much closer to zero than for the ECal cluster method.

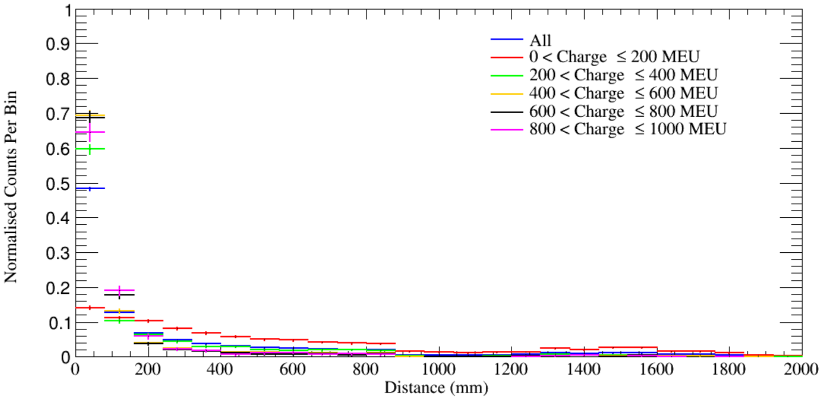

(a) Linear scale.

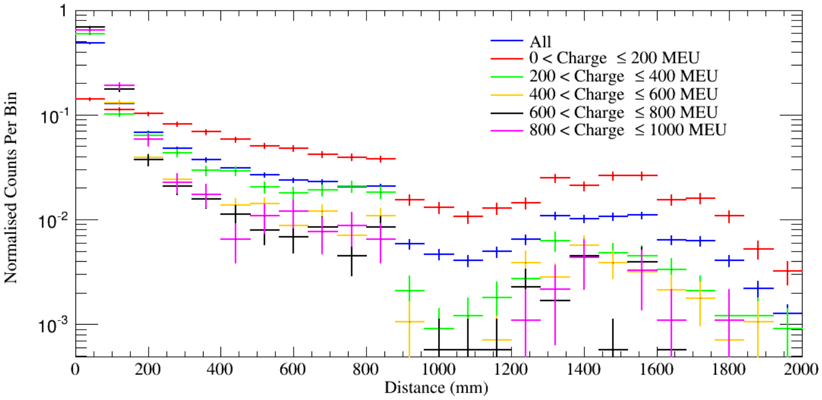

(b) Logarithmic scale.

Figure 5.30: The distance between the charge weighted unused FGD hit position and the true MC vertex position for a range of charges.

Other potential improvements were pursued by reanalysing the ECal extrapolation technique. Firstly, a point of closest approach method (POCA) was attempted. This entailed extrapolating both clusters back along the thrust axes and using the point of closest approach between the two trajectories as the reconstructed interaction vertex. Secondly, for the scenario in which the original method would result in the cut being passed, the vertex was then simply taken as the central coordinate of the associated FGD.

Figure 5.31 shows that although the POCA method produces a slightly more accurate reconstructed vertex (the mean distance between the true and reconstructed vertex being 577.7 mm compared to 587.6 mm), it reconstructs a vertex within the FGD FVs only ∼15% as often. This, in conjunction with the fact that adding an extra step to the algorithm greatly increases its complexity, was justification to not use the approach in this analysis. Furthermore, the central FGD method shows no improvement to the original method, with the mean distance increasing to 760.5 mm. As such it was also not incorporated in the second generation algorithm.

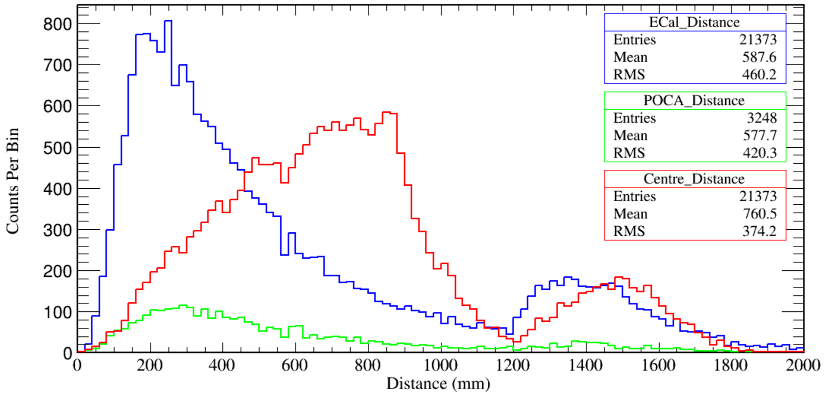

(a) Total distance.

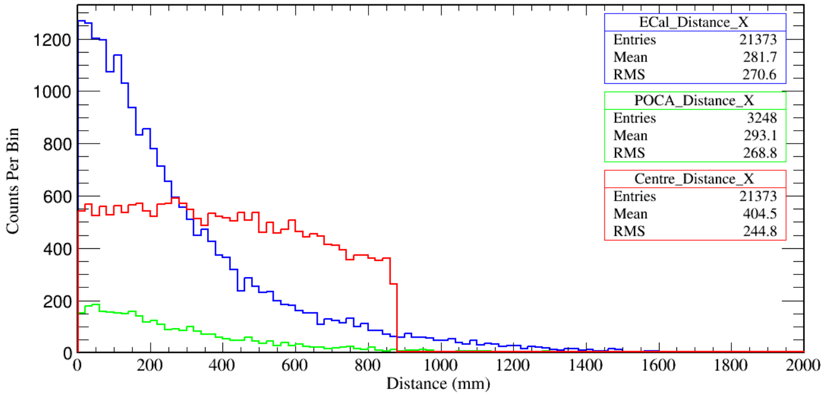

(b) x coordinate distance.

Figure 5.31: The distance between the reconstructed and true MC vertex for the three different ECal projection techniques.

Although small improvements to the vertexing are possible, the ever-increasing complexity of the selection suggests that making changes is not worthwhile. As such, until further improvements in reconstruction are available, this algorithm is a good compromise between complexity and accuracy.

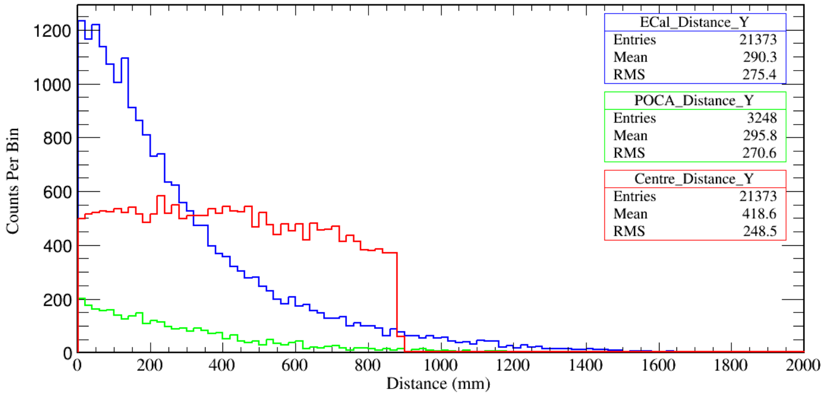

(c) y coordinate distance.

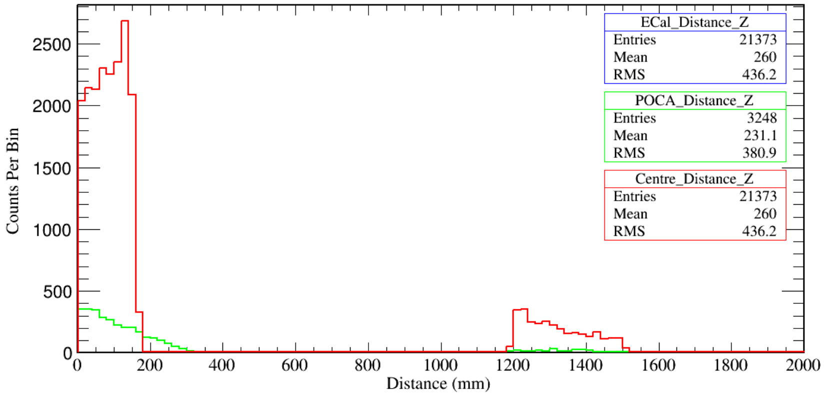

(d) z coordinate distance.

Figure 5.31: The distance between the reconstructed and true MC vertex for the three different ECal projection techniques.

5.2.10 Muon PID Cut

In the first generation analysis, to select NC1π^0 events, proton-like tracks were searched for [114]. To do this, when a good quality TPC track was present with a kaon pull $<$ 3.95, the event was cut. If this was not present, though an FGD isolated tracks was, an event was cut if that track had a muon pull of < 3.65. Finally, if neither of these track types were available, the ECal charge per unit length was used in the case where there were ECal-FGD objects, that were deemed to be track-like, available. Thus an event was cut if the charge per unit length of such objects was < 0.18 MEU mm^{-1}.[12] However, this resulted in a large efficiency loss ($\sim$ 14%) and a small loss in purity ($\sim$ 0.002%). As such for this analysis, attempts were made to improve this cut by optimising (using MC simulation) the kaon pull to select signal events - in contrast to selecting proton-like tracks. Nonetheless this method did not improve the resulting $\epsilon \times \rho$ and was removed.

However, after the BDT cut, CC interactions within the FGD constitute a large proportion of background events - see Figure 5.32a. To reduce this, an inverse muon PID cut is applied, in which any event with a muon-like track is removed. Such a track is defined as any negatively charged TPC track with an FGD component that has a muon PID likelihood > 0.05 and/or momentum > 500 MeV and/or:

$$\frac{L(\mu|\{p, \mathrm{Pull}_{\mu}\}_{t=1}^{3}) + L(\pi|\{p, \mathrm{Pull}_{\pi}\}_{t=1}^{3})}{1 - L(P|\{p, \mathrm{Pull}_{P}\}_{t=1}^{3})} > 0.8. \tag{5.11}$$

Here, the likelihood is defined as the particle probability (in each TPC) to be of type α when its momentum is p and pull is Pull$_\alpha$:

$$L(\alpha|\{p, \mathrm{Pull}_{\alpha}\}_{t=1}^{3}) = \frac{\prod_t P_t(\mathrm{Pull}_\alpha|p, \alpha)}{\sum_\beta \prod_t (\mathrm{Pull}_\beta|p, \beta)}, \tag{5.12}$$

where t is the TPC, β represents the 4 particles with TPC pulls[13] and P_t is the normalised probability in TPC t to measure Pull$_\alpha$ for a particle with momentum p and of type α:

$$P_t(\mathrm{Pull}_\alpha|p, \alpha) = exp(-0.5(\frac{\mathrm{Pull}_\alpha - \mu_{\mathrm{Pull}_\alpha}}{\sigma_{\mathrm{Pull}_\alpha}})). \tag{5.13}$$

Here, μ_{Pull_α} and σ_{Pull_α} are the mean pull and the uncertainty for that TPC and beam run.

The result of applying this cut to the MC sample can be seen in Figure 5.32b.

---

[12] Although the previous analysis searched for proton-like tracks, the kaon pull and muon pull for the TPC and FGD isolated tracks respectively were found to be more effective than using the proton pull.

[13] The e^-, μ^-, π and p.

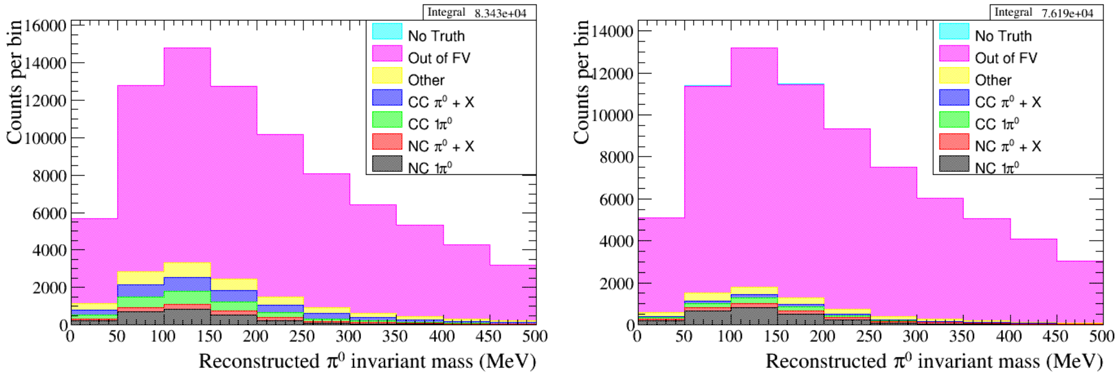

(a) Before the muon PID cut.

(b) After the muon PID cut.

Figure 5.32: The reconstructed π^0 invariant mass distribution pre- and post- muon PID cut. Note, prior to the cut being implemented, a large CC component is present. However, after the cut, the CC component is greatly reduced with very little impact on the NC sample.

5.3 Sand MC

As the magnet MC does not model the sand muon contamination, predicting the effect of including sand MC was undertaken. If the contamination constitutes a significant proportion of the total events passing all cuts, then the corresponding POT of sand MC must be added to the total MC sample. Conversely, if it is negligible, it can be discounted, thus saving computation time. Table 5.7 shows that after the final cut, the sand MC comprises only $\sim$2% of the total sample, with the 2 isolated ECal cluster cut removing much of the contamination.[14] Similar trends occur for the separate topologies, with the sand contamination for the DsDs (Table 5.8) and BrlDs (Table 5.9) topologies at $\sim$2%, whilst for the BrlBrl (Table 5.10) topology, it is slightly higher at $\sim$3%. As the combined statistical and systematic uncertainties are much larger than this, see Section 5.6, the sand MC was discounted from the analysis.

5.4 Selection Results

Using the above optimised selection, the resulting ϵ, ρ and $\epsilon \times \rho$ (and results of the first generation analysis[15]) can be found in Table 5.11. In the case of the NC1π^0 inclusive sample, for which the signal is defined as any NC1π^0 interaction regardless of the true photon conversion position, the lowest ϵ is achieved at 18.64 %, however this results in a ρ of 29.41 % and a combined $\epsilon \times \rho$ of 5.48 %. Conversely, for the NC1π^0 ECal-ECal signal, in which the two true π^0 decay photons convert in the ECals, the ϵ compared to the inclusive

---

[14] The sand events were scaled so the sand MC POT matched that of the Run 3 Air magnet MC used for this analysis.

[15] The first generation analysis included sand MC.

Table 5.7: The sand contamination when combining all topologies.

Cut	Sand Events (1.11 $\times 10^{21}$ POT)	Sand Events (4.98 $\times 10^{20}$ POT)	Magnet Events Scaled (4.98 $\times 10^{20}$ POT)	Sand Component (%)
2 Isolated ECal	145269.34	65143.20	540819.44	12.05
ECal Shower	10455.25	4688.45	167268.75	2.80
Time	8367.91	3752.43	135778.25	2.76
SMRD Activity	7513.19	3369.14	89641.45	3.76
TPC1 Veto	6445.79	2890.49	83105.00	3.48
P0D Veto	5253.73	2355.93	79813.08	2.95
ECal Charge	2710.01	1215.25	48131.20	2.52
FGD Vertex	653.99	293.27	14707.27	1.99
Muon PID	647.07	290.17	13485.98	2.15
MVA Cut	27.70	12.42	530.60	2.34

Table 5.8: The sand contamination for the DsDs topology.

Cut	Sand Events (1.11 $\times 10^{21}$ POT)	Sand Events Scaled (4.98 $\times 10^{20}$ POT)	Magnet Events (4.98 $\times 10^{20}$ POT)	Sand Component (%)
2 Isolated ECal	457.10	204.98	9583.28	2.14
ECal Shower	244.38	109.59	5640.57	1.94
Time	236.47	106.04	5254.12	2.02
SMRD Activity	227.56	102.04	3965.93	2.57
TPC1 Veto	139.51	62.56	3601.42	1.74
P0D Veto	114.77	51.47	3511.48	1.47
ECal Charge	114.77	51.47	2469.31	2.08
FGD Vertex	49.47	22.18	1139.89	1.95
Muon PID	49.47	22.18	1006.88	2.20
MVA Cut	1.98	0.89	55.18	1.61

Table 5.9: The sand contamination for the DsBrl topology.

Cut	Sand Events (1.11 $\times 10^{21}$ POT)	Sand Events Scaled (4.98 $\times 10^{20}$ POT)	Magnet Events (4.98 $\times 10^{20}$ POT)	Sand Component (%)
2 Isolated ECal	16027.74	7187.33	92517.05	7.77
ECal Shower	1677.01	752.02	31692.32	2.37
Time	1002.25	449.44	23864.26	1.88
SMRD Activity	927.06	415.72	17436.47	2.38
TPC1 Veto	611.45	274.19	16245.21	1.69
P0D Veto	479.86	215.18	15844.51	1.36
ECal Charge	353.22	158.39	10341.82	1.53
FGD Vertex	93.99	42.15	3160.77	1.33
Muon PID	91.02	40.82	2714.60	1.50
MVA Cut	8.90	3.99	190.30	2.10

Table 5.10: The sand contimantion for the BrlBrl topology.

Cut	Sand Events (1.11 $\times 10^{21}$ POT)	Sand Events Scaled (4.98 $\times 10^{20}$ POT)	Magnet Events (4.98 $\times 10^{20}$ POT)	Sand Component (%)
2 Isolated ECal	128849.20	57779.91	438714.56	13.17
ECal Shower	8534.11	3826.96	129935.64	2.95
Time	7129.36	3197.02	106660.02	3.00
SMRD Activity	6358.73	2851.49	68238.93	4.18
TPC1 Veto	5694.94	2553.78	63258.30	4.04
P0D Veto	4659.19	2089.32	60457.11	3.46
ECal Charge	2241.97	1005.37	35320.04	2.85
FGD Vertex	510.53	228.94	10406.61	2.20
Muon PID	506.57	227.16	9764.50	2.33
MVA Cut	16.82	7.54	285.12	2.64

sample almost doubles to 34.12 %, for a slight reduction in ρ, to give an increased $\epsilon \times \rho$ of 7.92 %. When breaking this down into the separate ECal-ECal topologies (in which the cluster topology matches the true photon conversion topology[16]) the DsDs topology performs by far the best with a final $\epsilon \times \rho$ of 15.02 %, compared to 10.38 % for the DsBrl topology and 5.48 % for the BrlBrl topology.

A possible technique to improve the final results would be to perform the cut optimisations on each topology separately (for this analysis optimisation was performed on the entire NC1π^0 ECal-ECal sample). This should particularly improve the DsDs topology; as it has the smallest sample size the optimisation is weighted towards the DsBrl and BrlBrl samples. Another method that should improve the DsDs response relates to the BDT's propensity to overtraining. By providing a larger MC training sample and increasing the tree depth to 3, akin to the other two topologies, an increased $\epsilon \times \rho$ should result.

For the ϵ, ρ and $\epsilon \times \rho$ as a function of cut for each topology, see Tables 5.12 - 5.16; whilst for a comparison of the total number of events passing each cut for both data and MC, see Tables 5.17 and 5.18. Interestingly, for each of the three topologies, and the combined topologies, the cut efficiencies differ between data and MC. For example, the total data cut efficiency is $\sim 5.6 \times 10^{-4}$ (420/746511), whilst for MC the analogous value is $\sim 9.7 \times 10^{-4}$ (3339.89/3451888.25). This difference could be due to a multitude of reasons, potentially due to cross-section model differences or discrepancies between the reconstruction of data and reconstruction in MC. For example, the cut sequence with the largest discrepancy is the FGD vertexing cut. One potential test to see if this is due to model dependency would be to compare both data and MC to see how often each method is used for vertexing (TPC track start position, FGD track start position, unused FGD hit information or ECal cluster extrapolation). Thus, if any is preferred in data but not MC, or if one has a larger cut efficiency in data or MC, this would suggest a non-model dependent discrepancy. However, to fully understand which cuts introduce model dependent and non-model dependent discrepancies, each would have to be analysed in depth before publication. For the systematic studies utilised for this analysis, see Section 5.5.

The reconstructed π^0 invariant mass for all three topologies combined and for each of the three topologies are shown in Figures 5.33 - 5.36. All show the same restriction regarding further improvement of the selection, namely that the vast majority of the background comes from events outside the FGD FVs. This is unsurprising considering the difficulty in reconstructing the interaction vertex. As stated earlier, often there is not tracker track activity to reconstruct the vertex. Hence the limit for the maximum $\epsilon \times \rho$ is capped by the directional accuracy of the ECal clusters. Thus, any future improvement is dependent on the ECal reconstruction capability, which is limited by the granularity of the detector and the low energy nature of the π^0 decay photons. Furthermore, for each of the ECal topologies, the MC overestimates the number of data events passing all cuts. How this translates into the final cross-section results is discussed in Section 5.6.

---

[16] Note, truth-matching between the clusters and true photons is not performed.

Table 5.11: The final ϵ, ρ and $\epsilon \times \rho$ for the different signal topologies.

Topology	ϵ (%)	ρ (%)	$\epsilon \times \rho$ (%)
Generation 1 Total	22.27	15.86	3.53
Generation 1 DsDs	32.88	43.20	14.20
Generation 1 DsBrl	29.33	15.88	4.66
Generation 1 BrlBrl	12.24	9.59	1.17
Generation 2 NC1π^0 Inclusive	18.64	29.41	5.48
Generation 2 NC1π^0 ECal-ECal	34.12	23.20	7.92
Generation 2 NC1π^0 DsDs	38.56	38.94	15.02
Generation 2 NC1π^0 DsBrl	42.17	24.61	10.38
Generation 2 NC1π^0 BrlBrl	30.21	18.14	5.48

Table 5.12: The ϵ, ρ and $\epsilon \times \rho$ as a function of cut for the NC1π^0 inclusive sample.

Cut	ϵ (%)	ρ (%)	$\epsilon \times \rho$ (%)
2 Isolated ECal	1.00	0.15	0.15
ECal Shower	87.11	0.43	0.37
Time	83.93	0.51	0.43
SMRD Activity	72.86	0.67	0.49
TPC1 Veto	71.62	0.71	0.51
P0D Veto	70.36	0.73	0.51
ECal Charge	67.95	1.15	0.78
FGD Vertex	52.58	2.93	1.54
Muon PID	52.24	3.17	1.66
MVA Cut	18.64	29.41	5.48

Table 5.13: The ϵ, ρ and $\epsilon \times \rho$ as a function of cut for the NC1π^0 ECal-ECal sample.

Cut	ϵ (%)	ρ (%)	$\epsilon \times \rho$ (%)
2 Isolated ECal	100.00	0.07	0.07
ECal Shower	88.36	0.19	0.17
Time	86.96	0.23	0.20
SMRD Activity	74.33	0.29	0.22
TPC1 Veto	73.75	0.31	0.23
P0D Veto	72.30	0.32	0.23
ECal Charge	69.19	0.51	0.35
FGD Vertex	56.18	1.35	0.76
Muon PID	56.10	1.47	0.82
MVA Cut	34.12	23.20	7.92

Table 5.14: The ϵ, ρ and $\epsilon \times \rho$ as a function of cut for the NC1π^0 DsDs sample.

Cut	ϵ (%)	ρ (%)	$\epsilon \times \rho$ (%)
2 Isolated ECal	100.00	0.07	0.07
ECal Shower	92.43	1.04	0.96
Time	91.72	1.11	1.02
SMRD Activity	77.38	1.23	0.95
TPC1 Veto	77.11	1.34	1.04
P0D Veto	75.42	1.35	1.02
ECal Charge	66.61	1.70	1.14
FGD Vertex	56.99	3.20	1.82
Muon PID	56.99	3.61	2.06
MVA Cut	38.56	38.94	15.02

Table 5.15: The ϵ, ρ and $\epsilon \times \rho$ as a function of cut for the NC1π^0 DsBrl sample.

Cut	ϵ (%)	ρ (%)	$\epsilon \times \rho$ (%)
2 Isolated ECal	100.00	0.12	0.12
ECal Shower	92.10	0.33	0.30
Time	91.41	0.43	0.40
SMRD Activity	79.77	0.52	0.41
TPC1 Veto	79.63	0.56	0.44
P0D Veto	78.35	0.56	0.44
ECal Charge	76.58	0.83	0.64
FGD Vertex	62.94	2.22	1.40
Muon PID	62.89	2.55	1.60
MVA Cut	42.17	24.61	10.38

Table 5.16: The ϵ, ρ and $\epsilon \times \rho$ as a function of cut for the NC1π^0 BrlBrl sample.

Cut	ϵ (%)	ρ (%)	$\epsilon \times \rho$ (%)
2 Isolated ECal	100.00	0.04	0.04
ECal Shower	86.49	0.11	0.09
Time	85.45	0.43	0.11
SMRD Activity	72.45	0.17	0.12
TPC1 Veto	71.69	0.18	0.13
P0D Veto	69.99	0.18	0.13
ECal Charge	68.64	0.31	0.21
FGD Vertex	54.16	0.83	0.45
Muon PID	54.02	0.88	0.48
MVA Cut	30.21	18.14	5.48

Table 5.17: The total number of data events passing each cut for all branches combined and each ECal cluster topology.

Cut	All Data Branches	DsDs Branch	DsBrl Branch	BrlBrl Branch
2 Isolated ECal	746511	15254	118770	612487
ECal Shower	211404	9235	38677	163492
Time	174191	8497	29656	136038
SMRD Activity	110754	6416	20780	83558
TPC1 Veto	101413	5880	19091	76442
P0D Veto	95231	5643	18210	71378
ECal Charge	60227	4186	12215	43826
FGD Vertex	11916	1492	2532	7892
Muon PID	10596	1298	2052	7246
MVA Cut	420	56	174	190

Table 5.18: The total number of MC events passing each cut for all branches combined and each ECal cluster topology. Note, the MC is normalised to account for the generated POT of the sample and the effect of flux tuning.

Cut	All Branches	DsDs Branch	DsBrl Branch	BrlBrl Branch
2 Isolated ECal	3451888.25	60595.91	591929.25	2799335.75
ECal Shower	1063444.50	35380.87	2011125.27	826937.50
Time	864446.81	32847.40	152220.56	679379.12
SMRD Activity	575173.38	25136.36	111319.48	438717.59
TPC1 Veto	532906.31	22809.77	103183.72	406912.66
P0D Veto	511647.09	15561.46	100456.12	388961.75
ECal Charge	311183.34	15561.46	66639.94	228981.25
FGD Vertex	94403.71	7096.74	20432.54	66874.48
Muon PID	86570.15	6270.58	17734.42	62565.17
MVA Cut	3339.89	382.80	1250.17	1706.92

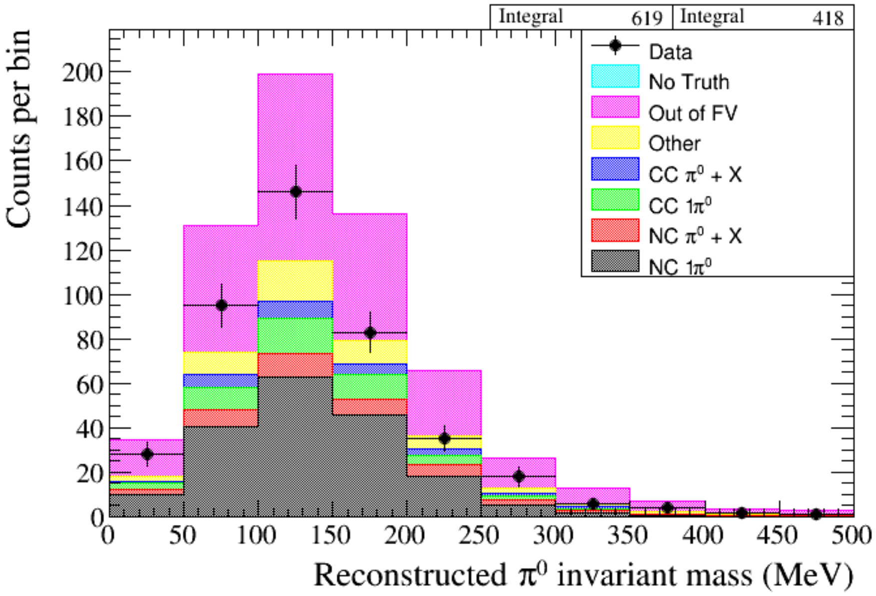

(a) Inclusive topologies.

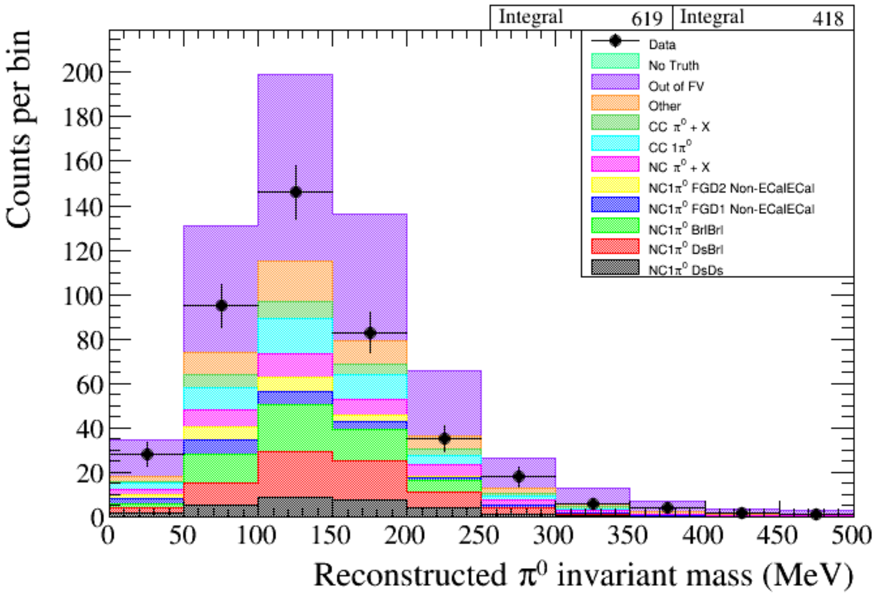

(b) Exclusive topologies.

Figure 5.33: Reconstructed π^0 invariant mass data/MC comparisons for all ECal cluster branches combined. Here, "Out of FV" are any interactions not occurring within the FGD FV. In the case of the Inclusive topologies "NC1π^0" does not consider the detectors in which the π^0 decay photons convert, however in the case of the Exclusive topologies the 5 interaction types preceded with "NC1π^0" give a description of where the two π^0 decay photons convert. "NC π^0 + X" and "CC π^0 + X" are interactions in which at least one π^0 and any number of mesons, exit the nucleus. "Other" represents other interactions not specified, but that occur within the FGD FV.

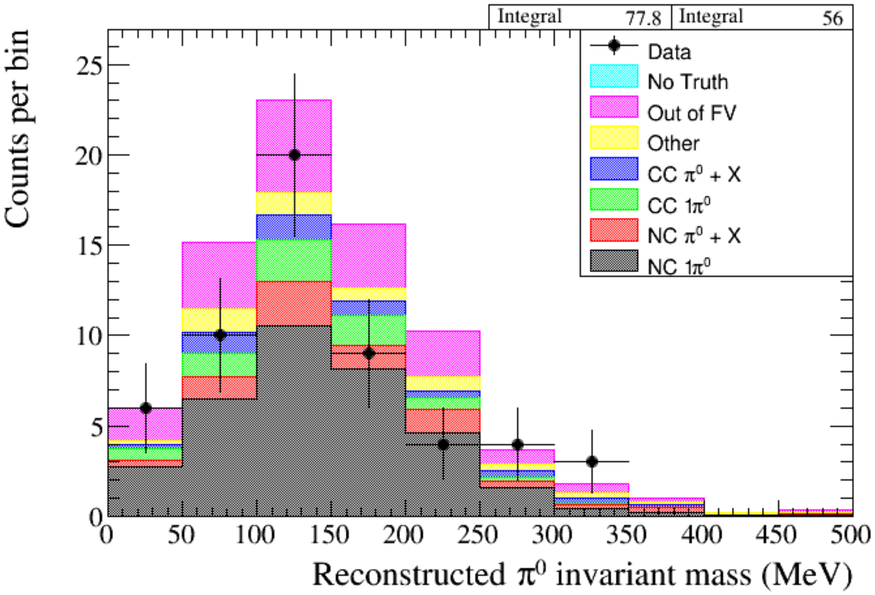

(a) Inclusive topologies.

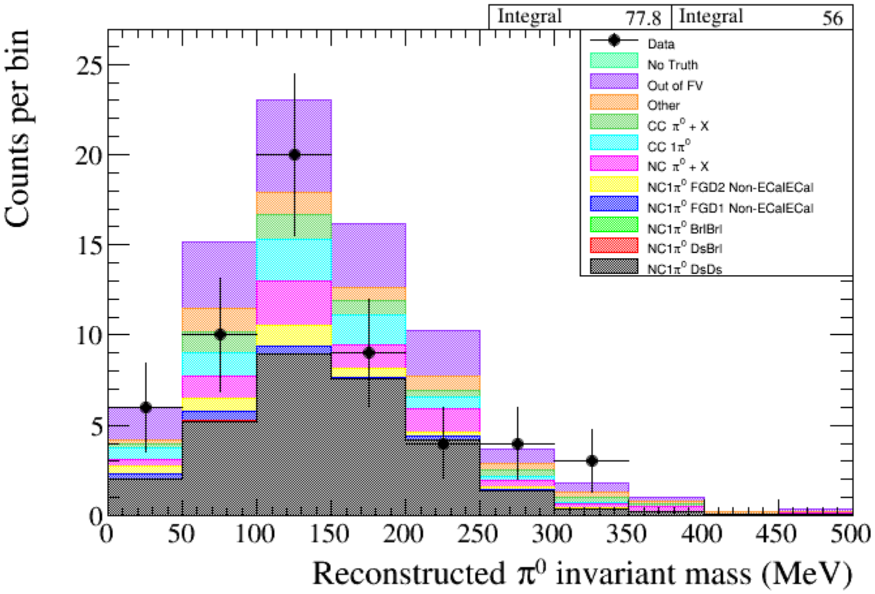

(b) Exclusive topologies.

Figure 5.34: Reconstructed π^0 invariant mass data/MC comparisons for the DsDs ECal cluster branch.

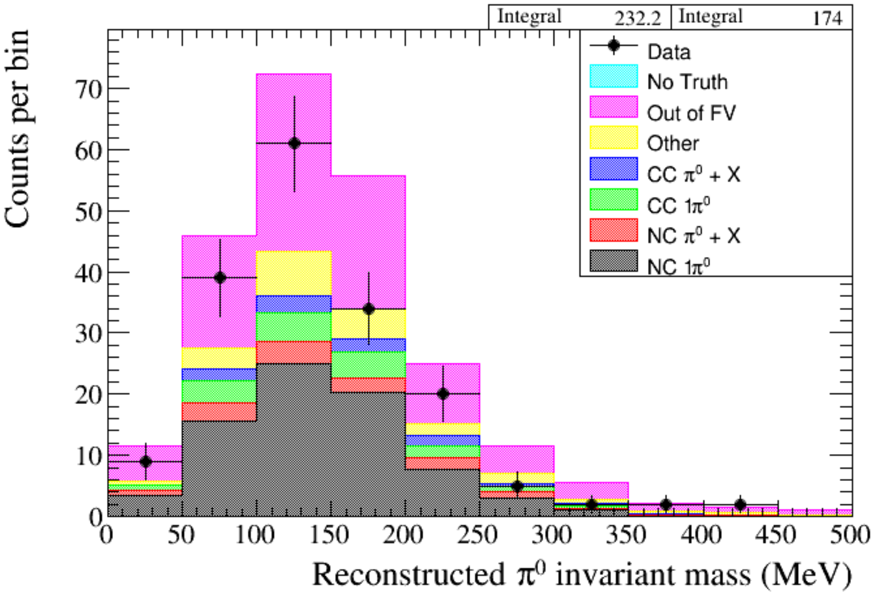

(a) Inclusive topologies.

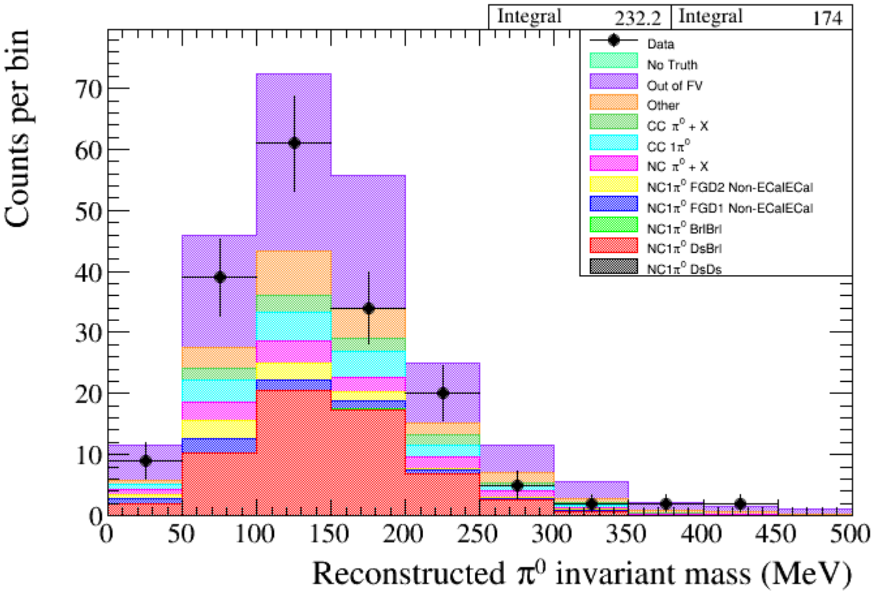

(b) Exclusive topologies.

Figure 5.35: Reconstructed π^0 invariant mass data/MC comparisons for the DsBrl ECal cluster branch.

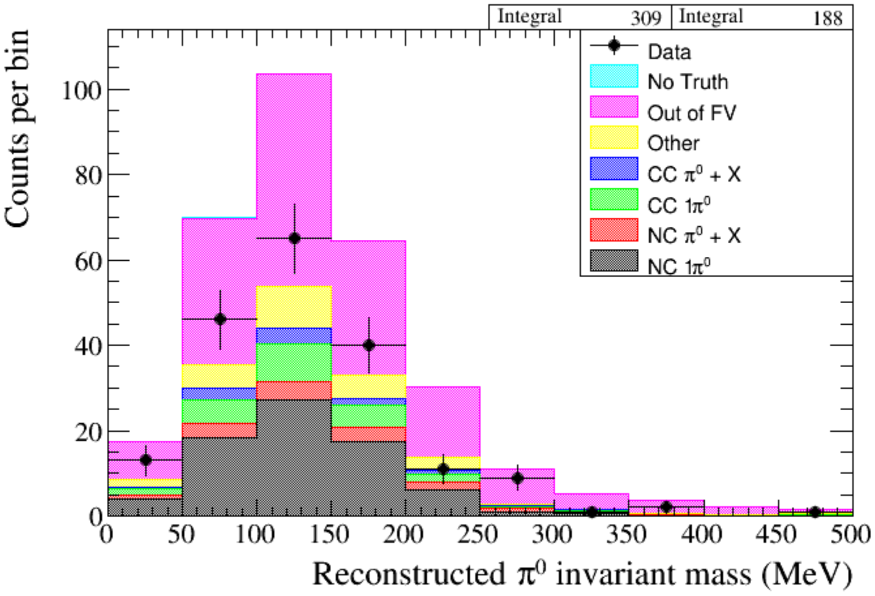

(a) Inclusive topologies.

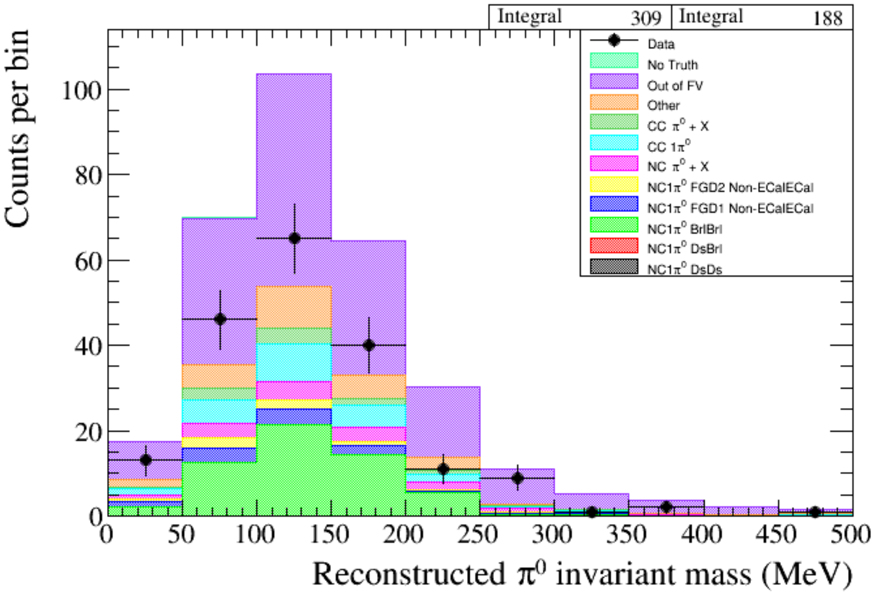

(b) Exclusive topologies.

Figure 5.36: Reconstructed π^0 invariant mass data/MC comparisons for the BrlBrl ECal cluster branch.

5.5 Systematic uncertainty

The systematic uncertainties in this analysis can be broken down into 4 main types: flux, cross-section, FSI and detector. Of the detector systematics, 6 were considered: the TPC PID uncertainty, the TPC cluster efficiency, the TPC charge ID confusion, the thrust vertexing uncertainty, the isolated ECal reconstruction efficiency and the BDT uncertainty.

Two main techniques were used when propagating the systematic uncertainties. Firstly, weight systematics in which the final event weight is varied. This can be applied to binary variables as an efficiency systematic, or in the case of event normalisation as a normalisation systematic. Secondly, variation systematics in which individual parameters in the analysis are varied, the effect of which is propagated through the analysis to change the number of events being selected for each individual throw. This method can be applied to any continuous or binary parameter.

Sections 5.5.1 - 5.5.4 explain the underlying principles behind these methods and the utilisation of covariance matrices to deduce the systematic effect. Additionally, Sections 5.5.5 - 5.5.13 discuss each of the systematics calculated in this analysis in detail. Next, Section 5.5.14 highlights potential other systematic effects that were deemed to be negligible in this analysis. Finally, Section 5.5.15 considers the principles behind cross-section extraction, Bayesian unfolding and obtaining the final uncertainties.

5.5.1 Normalisation Systematics

In the case of a systematic requiring an event normalisation, a normalisation weight can be applied to the event:

$$W = W_0(1 + \delta \times \sigma_W). \tag{5.14}$$

Here, W_0 is the nominal weight (which can differ from 1 if a correction is applied), δ is the variation to be applied in number of standard deviations and σ_W is the normalisation systematic error (the value of which must be deduced from separate studies).

5.5.2 Efficiency Systematics

Efficiency systematics apply a weight to the event to replicate the uncertainty of binary variables. To calculate this, the efficiencies of data and MC must be deduced - often achieved using a control sample. Once the analysis sample MC efficiency has been found, using the implicit assumption that the data/MC efficiency for the control and analysis sample are identical, the data efficiency is calculated as:

$$\epsilon_{\mathrm{data}} = \frac{\epsilon^{\mathrm{CS}}_{\mathrm{data}}}{\epsilon^{\mathrm{CS}}_{\mathrm{MC}}}\epsilon_{\mathrm{MC}}. \tag{5.15}$$

Here, $\epsilon^{\mathrm{CS}}_{\mathrm{data}}$ is the control sample data efficiency, $\epsilon^{\mathrm{CS}}_{\mathrm{MC}}$ the control sample MC efficiency and ϵ_{MC} the analysis sample MC efficiency. Equation 5.15 does not consider the statistical variation affecting the data efficiency, hence the following adaptation is used:

$$\epsilon^{'}_{\mathrm{data}} = \frac{\epsilon^{\mathrm{CS}}_{\mathrm{data}} + (\delta_{\mathrm{data}} \times \sigma_{\epsilon^{\mathrm{CS}}_{\mathrm{data}}})}{\epsilon^{\mathrm{CS}}_{\mathrm{MC}} + (\delta_{\mathrm{MC}} \times \sigma_{\epsilon^{\mathrm{CS}}_{\mathrm{MC}}})} \tag{5.16}$$

where $\sigma_{\epsilon^{\mathrm{CS}}_{\mathrm{data}}}$ is the statistical error of the data control sample, $\sigma_{\epsilon^{\mathrm{CS}}_{\mathrm{MC}}}$ is the statistical error on the MC control sample, δ_{data} is the variation in number of standard deviations for the data sample and δ_{MC} is the variation in number of standard deviations, for the MC sample.

The event weight in the case where the binary variable is one (e.g. a track is found) is:

$$W_{\mathrm{eff}} = \frac{\epsilon^{'}_{\mathrm{data}}}{\epsilon_{\mathrm{MC}}}. \tag{5.17}$$

Conversely, the event weight in the case where the binary variable is zero (e.g. a track not reconstructed) is:

$$W_{\mathrm{eff}} = \frac{1 - \epsilon^{'}_{\mathrm{data}}}{1 - \epsilon_{\mathrm{MC}}}. \tag{5.18}$$

5.5.3 Variation Systematic

When using a variation systematic, for each throw, the variable is given a value:

$$x^{'} = x + \Delta x + (\delta \times \sigma_{\Delta x}), \tag{5.19}$$

where x is the original value, Δx is a correction (for when the mean of the data and MC mean disagree), δ is the variation in number of standard deviations and $\sigma_{\Delta x}$ is the statistical uncertainty on Δx.

5.5.4 Calculating errors using the covariance matrix

Once either the weight or systematic errors are deduced, the effect they have on the analysis must be calculated. To find the magnitude of the systematic error, a covariance matrix is produced and its integral calculated, the square root of which gives the total uncertainty of the number of events passing all cuts. The covariance matrix is defined as:

$$C_{ij} = \sum_{t=1}^{N_{\text{Toys}}} [(N_i^t)^W - N_i^{\text{avg}}][(N_j^t)^W - N_j^{\text{avg}}].w_t. \tag{5.20}$$

Here, $(N_i^t)^W$ is the number of events (after weight systematic reweighting) for toy experiment t in bin i:

$$(N_i^t)^W = \sum_{e=1}^{N_{\text{events}}} (W^t)_e.(\delta_i^t)_e, \tag{5.21}$$

where $(W^t)_e$ is the weight for toy experiment t and event e and $(\delta_i^t)_e$ equals 1 if the event is passed, and equals 0 otherwise. Furthermore, N_i^{avg} is the average number of events in bin i, without weighting:

$$N_i^{\text{avg}} = \sum_{t=1}^{N_{\text{Toys}}} w^t.N_i^t = \sum_{t=1}^{N_{\text{Toys}}} w^t.\sum_{e=1}^{N_{\text{events}}} (\delta_i^t)_e. \tag{5.22}$$

Here, w_t is the toy experiment weight, given by:

$$w^t = \frac{1}{N_{\text{Toys}}}, \tag{5.23}$$

where N_{Toys} is the number of toy experiments.

5.5.5 BDT Systematic

To understand the effect of the discrepancies between data and MC for the BDT variables, and to take the correlations between these variables into account, an ECal variable covariance matrix was produced. A control sample of photons was used, with the low-level ECal variables thrown at the calibration stage. This effect was propagated through the MC production chain to replicate the effect on the final variables passed to the BDT. The low-level variables considered were:

Hit Efficiency
The ECal bar hit efficiency, see Figure 5.37, was calculated in [85]. Using cosmic muons, a

missed hit was found when a hit occurred in layer $n-1$ and also $n+1$ without one being registered in layer n.

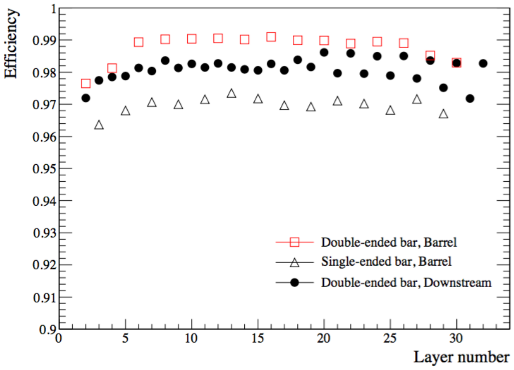

Figure 5.37: The ECal layer hit efficiency broken down by bar type and layer. Figure taken from [85].

To replicate this effect in the magnet MC, and to prevent underestimation of the resulting uncertainties, the lowest efficiency in Figure 5.37, 96.2%, was used for all bar and layer configurations. As removing hits at random is not representative of the processes involved - low energy hits are much less likely to be reconstructed - hits with energy corresponding to the lowest 3.8% were removed at the calibration stage.[17] Of course, this method could be improved by applying the exact efficiency of each individual bar and layer configuration and understanding the efficiency as a function of deposited hit charge.

Charge Resolution

To represent the charge resolution uncertainty, the hit charges were multiplied by a random factor selected from a thrown Gaussian distribution of width 0.08 centred on 1. This width was deduced using Run 3 cosmic data and is the width of the most probable value of the MIP peak [116].

Hit Timing

To represent the hit timing resolution, the hit times were multiplied by a random factor selected from a thrown Gaussian distribution of width 3.141 ns centred on 1. This corresponds to the width in the timing distribution for hits in all Tracker RMMs for Run 5.

It should be noted that each of the three variables discussed above were treated as uncorrelated. This is not an issue for the hit timing and charge resolution since, due to the time walk correction, these variables should be uncorrelated. Conversely, the hit efficiency may have unknown correlations with the other two variables. However, due to the nature of calculating the inefficiencies, finding the correlations is difficult. One possible method would be to predict the missing charge and time from the hits in layers $n-1$ and $n+1$ -

---

[17] Constituting all hits less than 7.55 PEU.

the disadvantage of this is that it will introduce model dependencies. Furthermore, these three variables are not a complete set as others will contribute to higher level data/MC discrepancies. The next generation covariance matrix must also consider: the effect of dark noise, the ECal mass uncertainty and the magnetic field uncertainty.

Once the low-level variables were thrown, the entire MC reconstruction chain was re-run[18] and a covariance matrix produced, see Figure 5.38a, using the high-level ECal variables (to improve the accuracy of the matrix, each variable was split into 4 bins). The equation used to calculate the matrix was:

$$V_{ab} = \frac{1}{N_{\text{throws}}} \sum_{i=1}^{N_{\text{throws}}} \frac{(A_a^i - A_a^{\text{nom}})(A_b - A_b^{\text{nom}})}{A_a^{\text{nom}} A_b^{\text{nom}}} \tag{5.24}$$

where A_a^i is the average value for that parameter range for the thrown MC and A_a^{nom} the average value for that parameter range for the nominal MC. For the binning information, see Table 5.19.

Variable	Limit 1	Limit 2	Limit 3
PID_Angle	60°	70°	80°
PID_Asymmetery	0.1	1	2.5
Average Z Position	1000 mm	2000 mm	2500 mm
PID_Circularity	0.6	0.85	0.95
EM Energy Fit Result	200 MeV	400 MeV	600 MeV
First Layer	3	10	25
PID_Front Back Ratio	0.2	1	2
Last Layer	10	20	30
PID_LLR_EM_HIP	-2	20	25
PID_LLR_MIP_EM	-20	0	20
PID_LLR_MIP_EM_LowMomentum	-5	20	25
PID_LLR_MIP_Pion	-5	20	30
NHits	10	20	40
Object Length	500 mm	1500 mm	2500 mm
Thrust Origin X	-750 mm	0 mm	750 mm
Thrust Origin Y	-750 mm	0 mm	750 mm
Thrust Origin Z	500 mm	1500 mm	2500 mm
PID_ShowerAngle	0.25 rad	0.3 rad	0.35 rad
Thrust	0.90	0.95	0.98
Thrust Axis Z	-0.4	-0.1	0.6
PID_TransverseChargeRatio	1	1.5	2
PID_TruncatedMaxRatio	0.2	0.4	0.6
PID_Angle_Low	60°	70°	80°
PID_Asymmetery_Low	0.1	1	2.5

---

[18] Only 1 throw was used due to time constraints.

Average Z Position_Low	1000 mm	2000 mm	2500 mm
PID_Circularity_Low	0.6	0.85	0.95
EM Energy Fit Result_Low	200 MeV	400 MeV	600 MeV
First Layer_Low	3	10	25
PID_Front Back Ratio_Low	0.2	1	2
Last Layer_Low	10	20	30
PID_LLR_EM_HIP_Low	-2	20	25
PID_LLR_MIP_EM_Low	-20	0	20
PID_LLR_MIP_EM_LowMomentum_Low	-5	20	25
PID_LLR_MIP_Pion_Low	-5	20	30
NHits_Low	10	20	40
Object Length_Low	500 mm	1500 mm	2500 mm
Thrust Origin X_Low	-750 mm	0 mm	750 mm
Thrust Origin Y_Low	-750 mm	0 mm	750 mm
Thrust Origin Z_Low	500 mm	1500 mm	2500 mm
PID_ShowerAngle_Low	0.25 rad	0.3 rad	0.35 rad
Thrust_Low	0.90	0.95	0.98
Thrust Axis Z_Low	-0.4	-0.1	0.6
PID_TransverseChargeRatio_Low	1	1.5	2
PID_TruncatedMaxRatio_Low	0.2	0.4	0.6

Table 5.19: The high-level ECal variables and the corresponding bin limits used to produce the ECal covariance matrix. Thus, to populate the first, second, third or fourth bin, the corresponding ECal variable must have a value less than Limit 1, between Limit 1 and Limit 2, between Limit 2 and Limit 3 or greater than Limit 3, respectively.

Unfortunately it is not possible to Cholesky decompose[19] this matrix as it is not positive definite[20], hence two options are available. Reconditioning using diagonal loading could be used, however, as this changes the values within the matrix this method was not preferred. The other possibility is to decompose the matrix using a different technique. Singular value decomposition (SVD) was pursued. See the Appendix for the equivalency between Cholesky decomposition and SVD and Figure 5.38b for the SVD matrix.

Once decomposed, the fractional change to apply to each variable is given by:

$$\Delta P = JR \tag{5.25}$$

where J is the decomposed covariance matrix and R is the vector of random Gaussian throws of width 1 centred on 0. For this systematic 500 throws were performed and the resulting covariance matrices for the reconstructed π^0 invariant mass and reconstructed π^0 momentum can be seen in Figure 5.39. By integrating the matrix and comparing to the

---

[19] A decomposition method which produces a lower triangular matrix and its conjugate transpose. The decomposed matrix can then be used to propagate random throws of correlated variables.

[20] Likely due to the lack of thrown samples.

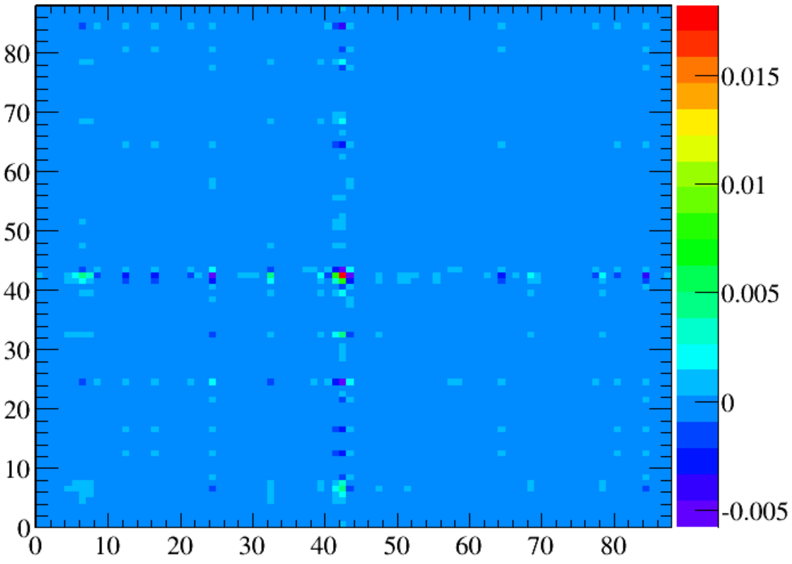

(a) The ECal covariance matrix.

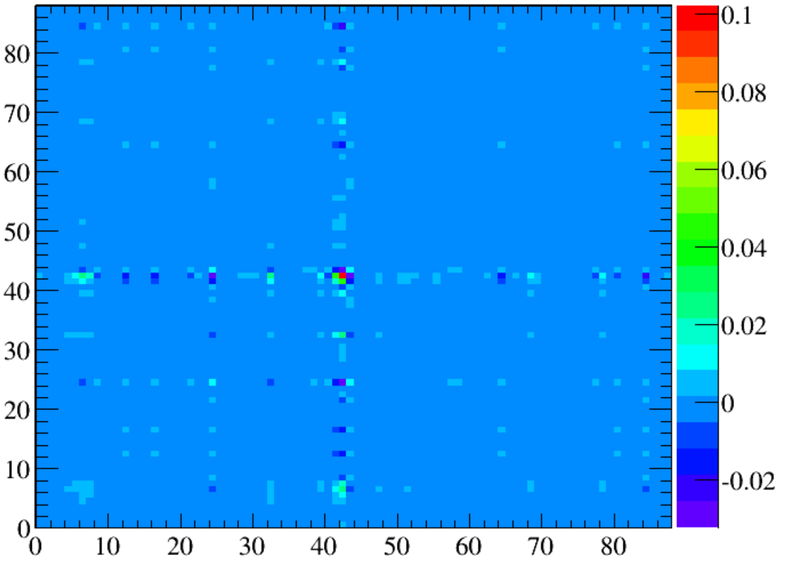

(b) The singular value decomposed ECal covariance matrix.

Figure 5.38: The ECal covariance matrix and the resulting singular value decomposition. Note, in Figure 5.38a the largest covariances occur in bins 41 to 43, corresponding to the 'PID_LLR_MIP_EM_LowMomentum' variable. This suggests that during a throw, this variable has the largest bin-to-bin variation, with the largest covariance reaching ∼0.018 (indicating a positive correlation) and the lowest reaching ∼-0.005 (indicative of a negative correlation). In general however, as expected due to the small impact that throwing the low-level variables has on the higher level ECal variables, most show covariances extremely close to 0 - suggesting little correlation.

number of events passing all cuts, the scale of the systematic error was found to be 0.62%. This is a small effect, however, when considering the pre- and post-thrown distributions, see Figure 5.40, it is unsurprising as the low-level variations have a small effect on the overall distributions.

5.5.6 Thrust Systematic

The systematic effect associated with reconstructing the interaction vertex was analysed. The uncertainty associated with the start position of the TPC track was deemed to be negligible. The reasoning for this is that the difference between the vertex and the start position of the TPC tracks is low. Furthermore, the systematic effects attributed to using FGDIso tracks was also deemed negligible. This was because the FGDIso track must be reconstructed in the FGD (unlike the TPC and thrust methods) as such it is less likely to reconstruct outside the FGD FVs and thus change the total number of events. Thirdly, as with the FGDIso tracks, the FGD unused hits must also reconstruct in the FGD, this again means it is less likely to reconstruct out of the FGD FVs and change the total number of events.

On the other hand, the ECal thrust variables are the least accurate vertexing tool. This is due to the difficulty in reconstructing such low energy clusters. To model the systematic effect, photons from the TPCs and FGDs entering the ECals were analysed. The angular difference between the true trajectory and the thrust directions was fitted using a 7th order polynomial[21], see Figure 5.41 and Table 5.20. To replicate the discrepancy, for each throw an angle was randomly selected from the fit distribution. The thrust was then rotated around an orthogonal vector to the original thrust vector. As this only rotates in one axis, a second random angle in the range of 0 to 2π was used to rotate the new thrust axis around the original. This process randomly throws the original thrust to elsewhere on an "uncertainty cone", the opening angle of which is twice the random angle selected from the 7th order polynomial fit.

Table 5.20: The 7th order polynomial fit information. Note, the functional form is: $f = p0 + p1x + p2x^2 + p3x^3 + p4x^4 + p5x^5 + p6x^6 + p7x^7$.

Variable	Value	Error
p0	11.99	22.09
p1	4411.71	274.95
p2	-11490.90	946.47
p3	13704.20	1418.90
p4	-8998.10	1092.05
p5	3318.80	451.30
p6	-642.84	95.26
p7	50.83	8.06

---

[21] Lower order polynomials were attempted, though these gave a larger χ^2/Number of degrees of freedom.

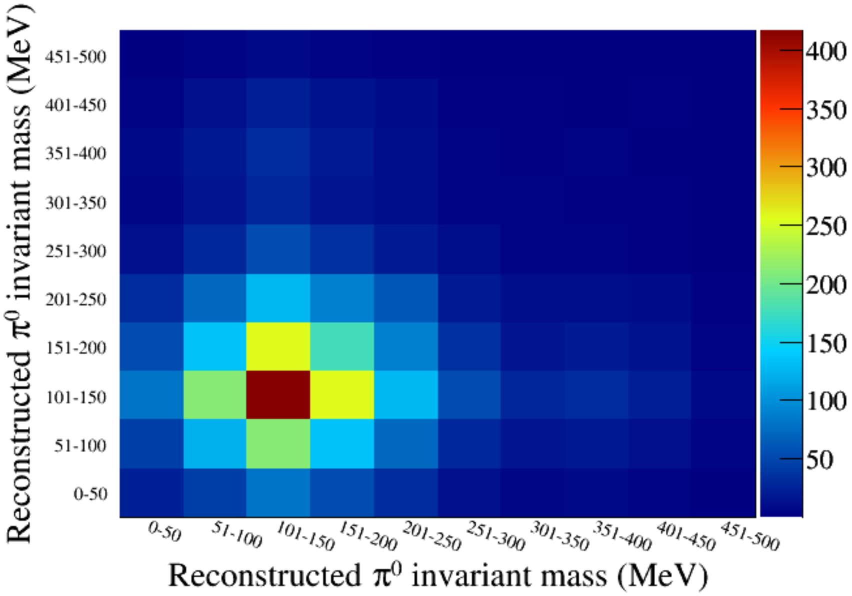

(a) Reconstructed π^0 invariant mass covariance matrix.

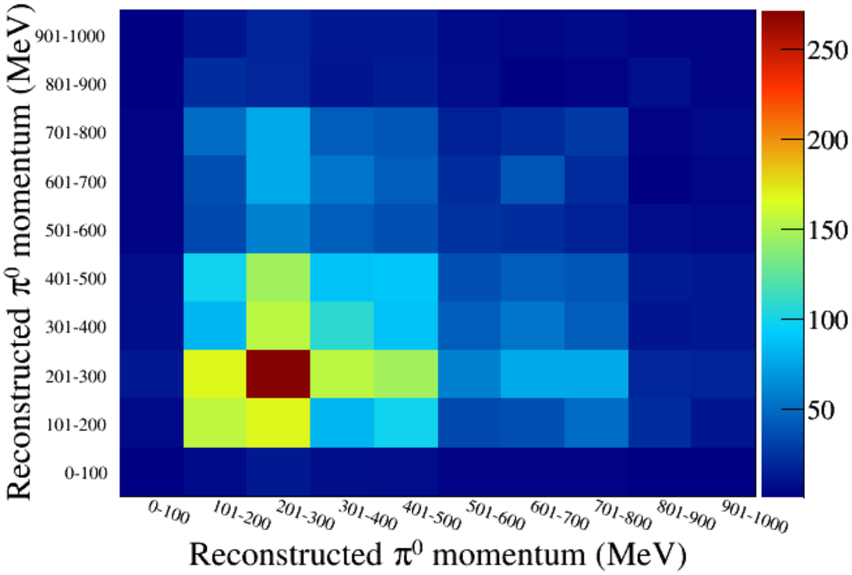

(b) Reconstructed π^0 momentum covariance matrix.

Figure 5.39: The covariance matrices used to calculate the size of the BDT systematic. In the case of Figure 5.39a, the largest covariances (reaching as high as ∼400) are found around the π^0 mass of 135 MeV. This is because of the higher statistics in the 50-200 MeV region, allowing for greater bin-to-bin variation. In contrast, due to the lower statistics above ∼250 MeV, the covariances are ∼0, suggesting almost no bin-to-bin variation and correlation. Figure 5.39b is similar, with the highest covariances occurring in the higher statistics regions, whilst for extremely high or extremely low momenta there is little to no covariance.

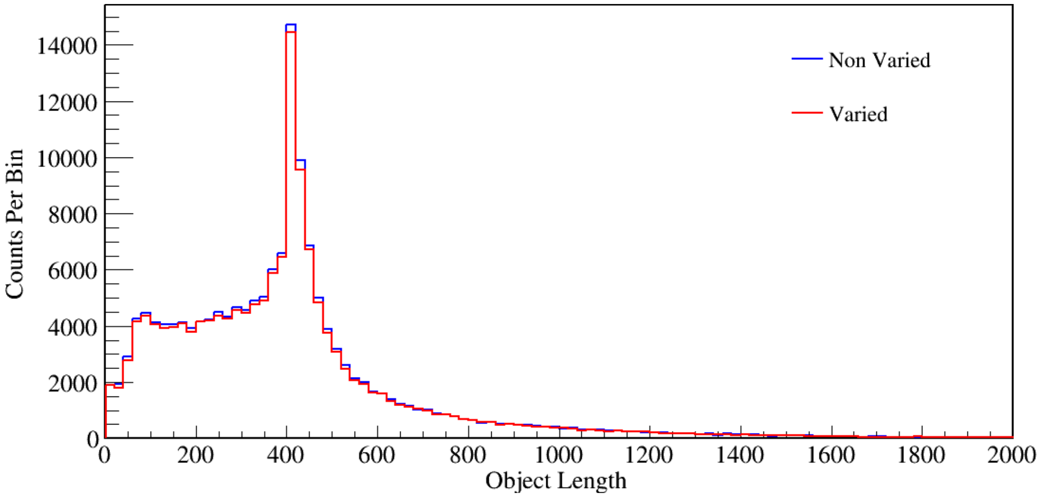

Figure 5.40: The ECal cluster object length comparison when using non-thrown and thrown low-level ECal variables.

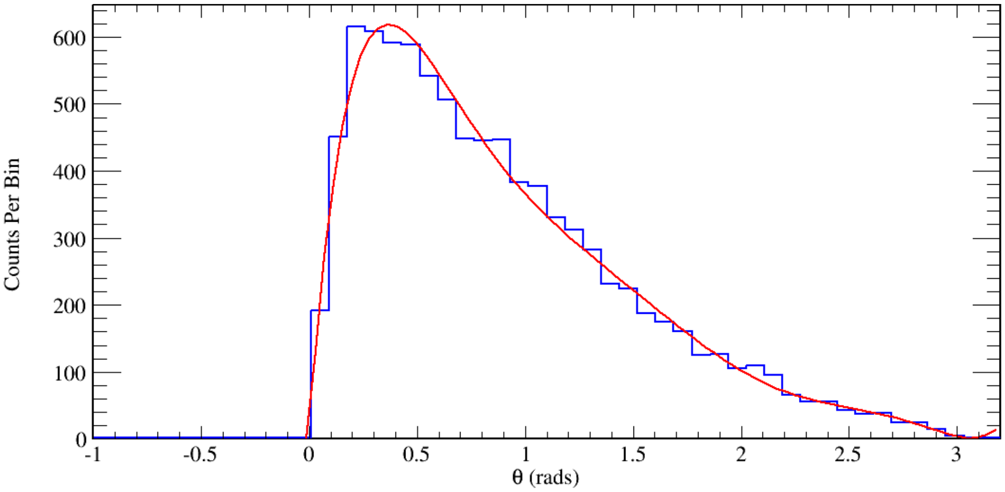

Figure 5.41: The angular difference between the thrust and the true photon direction (blue) and the corresponding 7th order polynomial fit (red). The fit gave a χ^2/Number of degrees of freedom = 39.44/29 = 1.36.

Using this method, 500 throws were used[22], with the overall size of the systematic found to be 6.9% - see Figure 5.42 for the covariance matrices used to calculate this value.

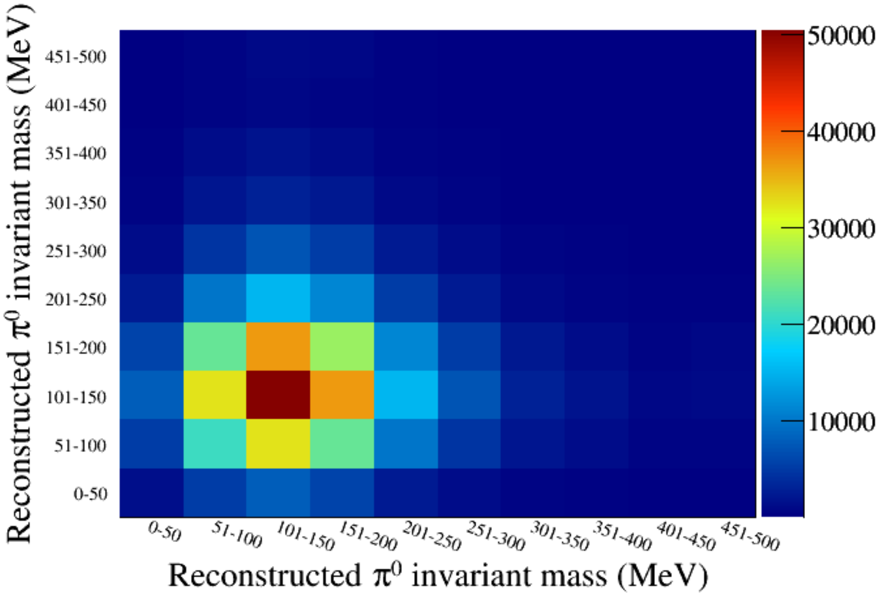

(a) Reconstructed π^0 invariant mass covariance matrix.

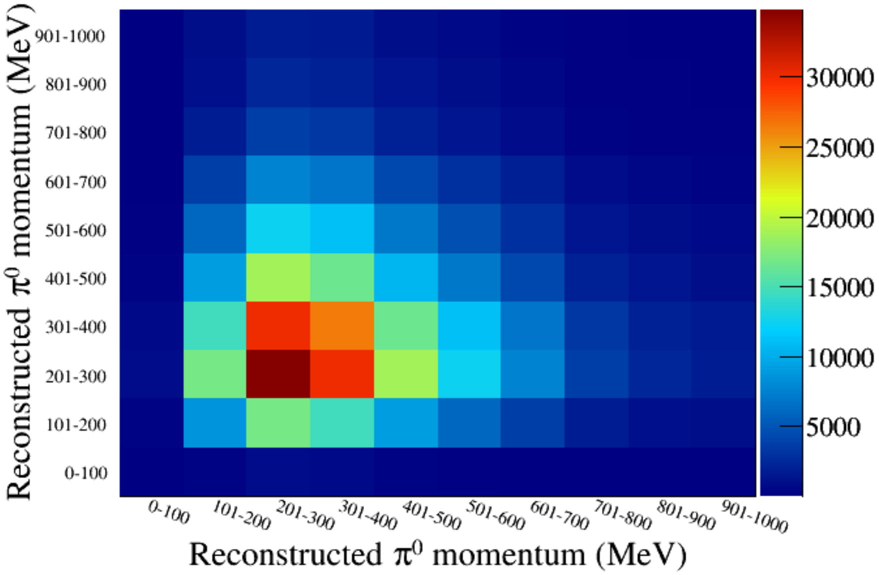

(b) Reconstructed π^0 momentum covariance matrix.

Figure 5.42: The covariance matrices used to calculate the size of the ECal thrust systematic.

---

[22] Note, the thrust variables were treated in an uncorrelated manner to the thrust variables used in the BDT systematic.

5.5.7 Isolated ECal reconstruction efficiency

The method used to understand the uncertainty when selecting isolated ECal clusters is almost identical to the efficiency systematic developed in [117]. This was formulated to understand the error associated with selecting any ECal cluster (including non-isolated clusters). This original method used good quality TPC tracks starting in the FGD and entering the ECals. The efficiency of track-like and shower-like clusters was then deduced using:

$$\text{Eff}_{\text{Shower}} = \frac{(\text{Number of shower-like TPC tracks entering ECal}) \bigcap (\text{ECal shower found})}{(\text{Number of shower-like TPC tracks entering ECal})} \tag{5.26}$$

and

$$\text{Eff}_{\text{Track}} = \frac{(\text{Number of track-like TPC tracks entering ECal}) \bigcap (\text{ECal shower found})}{(\text{Number of track-like TPC tracks entering ECal})}, \tag{5.27}$$

where $\bigcap$ is the intersection.[23] The results of the study are shown in Tables 5.21 and 5.22.

Table 5.21: Ds and Brl efficiencies and statistical errors for shower-like data and MC samples. Table taken from [117].

Type	Efficiency (%)	Statistical Error (%)
Ds MC	66.73	0.19
Ds Data	64.15	0.63
Brl MC	28.23	0.16
Brl Data	28.49	0.49

Table 5.22: Ds and Brl efficiencies and statistical errors for track-like data and MC samples. Table taken from [117].

Type	Efficiency (%)	Statistical Error (%)
Ds MC	85.39	0.09
Ds Data	84.85	0.28
Brl MC	46.03	0.18
Brl Data	51.43	0.57

When implementing the original systematic, all true tracks crossing the Ds and Brl ECals are analysed. All that have a matched reconstructed ECal object have an efficiency weight assigned; all that do not, have an inefficiency weight assigned. To tailor this systematic to this analysis, the same method was used, though any non-isolated ECal

---

[23] Hence when considering two sets, ‘A’ and ‘B’, A$\bigcap$B represents objects that belong to set A and set B.

cluster was given an inefficiency weight. This is not a perfect representation of the data and MC discrepancies. However, to fully replicate these, a new study to update Equations 5.26 and 5.27 along with Tables 5.21 and 5.22, to account only for isolated ECal clusters, must be undertaken. It should be noted, at the time of writing, work is ongoing to understand Data/MC differences as a function of TPC track angle and momentum.

For the covariance matrices binned in reconstructed π^0 invariant mass and reconstructed π^0 momentum, used to calculate the magnitude of this systematic, see Figure 5.43. The resulting uncertainty was found to be 2.7%.

5.5.8 TPC Cluster Efficiency

The TPC cluster efficiency systematic assesses each TPC segment of a TPC track and assigns a weight to account for the data/MC discrepancy of the probability of finding a group of adjacent pad hits. In this method, horizontal and vertical tracks are treated separately, for full details see [118]. The results are shown in Table 5.23.

Table 5.23: Data and MC efficiency discrepancies for vertical and horizontal clusters. Table taken from [119].

Cluster Type	$\frac{\epsilon^{\text{MC}}-\epsilon^{\text{Data}}}{\epsilon^{\text{MC}}}$	$\frac{\epsilon^{\text{Data}}}{\epsilon^{\text{MC}}}$
Vertical	0.0011±0.0002	0.9989±0.0002
Horizontal	0.0007±0.0001	0.9993±0.0001

For the reconstructed π^0 invariant mass and reconstructed π^0 momentum covariance matrices, which yield a systematic error of 0.34%, see Figure 5.44.

5.5.9 TPC PID Systematic

The TPC PID systematic is a variation systematic which changes the charge and time of each TPC segment of each track - resulting in a modification to the reconstructed dE/dx. These variations are range, momentum and particle dependent. For full information on this systematic, see [120].

For the reconstructed π^0 invariant mass and reconstructed π^0 momentum covariance matrix, which yield a systematic error for the analysis of 0.08%, see Figure 5.45.

5.5.10 TPC Charge ID

The TPC charge ID is an efficiency systematic to parameterise the charge misidentification of TPC tracks. To do this, both the probability that multi-segmented TPC tracks have different local charges and the probability that the globally reconstructed charge differs from the local charge, is parameterised. For full details see [121].

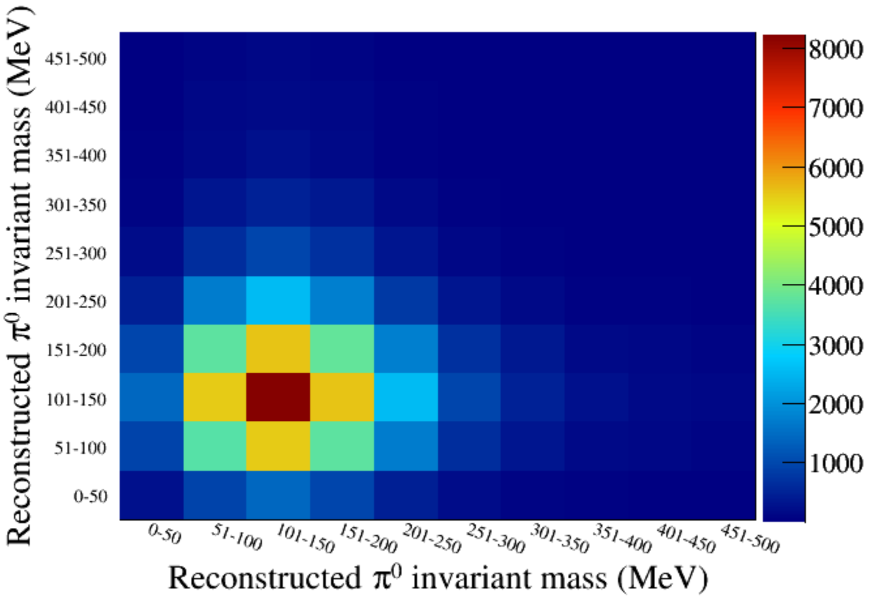

(a) Reconstructed π^0 invariant mass covariance matrix.

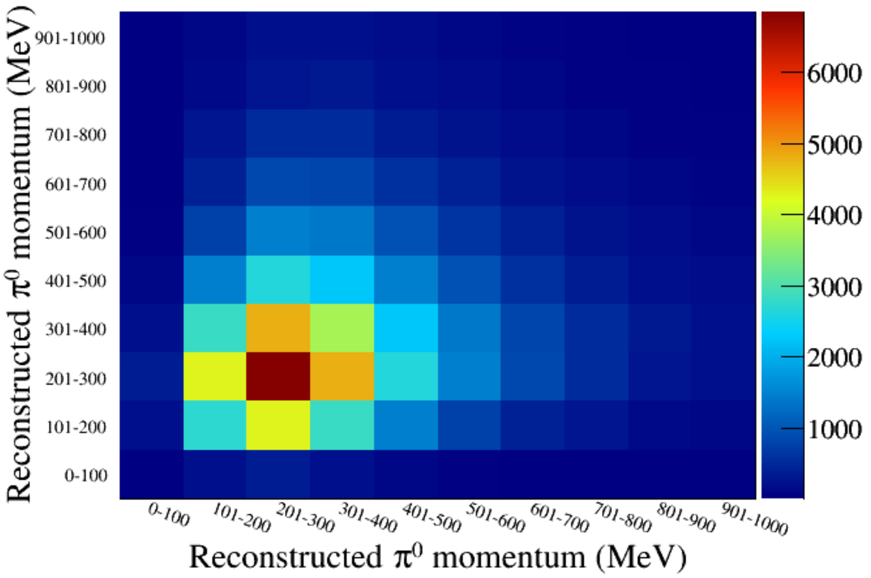

(b) Reconstructed π^0 momentum covariance matrix.

Figure 5.43: The covariance matrices used to calculate the size of the ECal efficiency systematic.

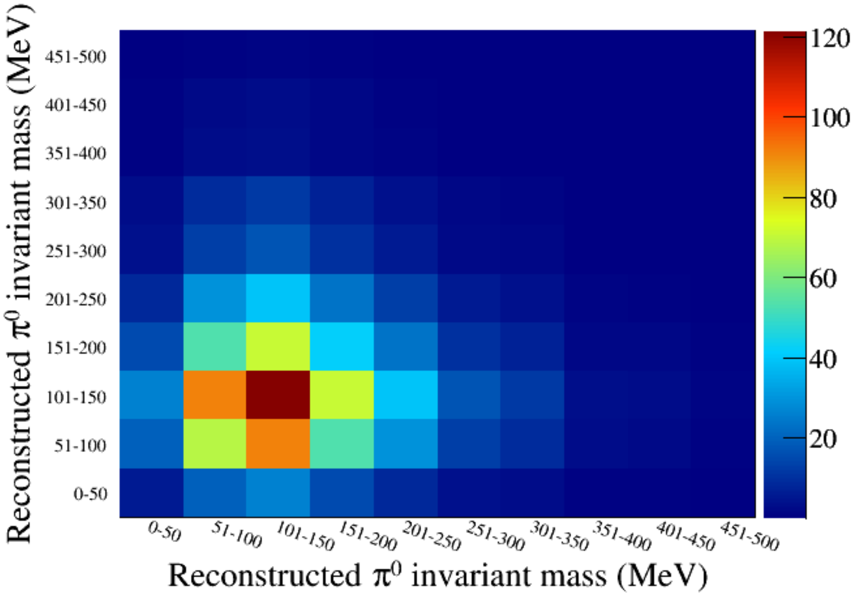

(a) Reconstructed π^0 invariant mass covariance matrix.

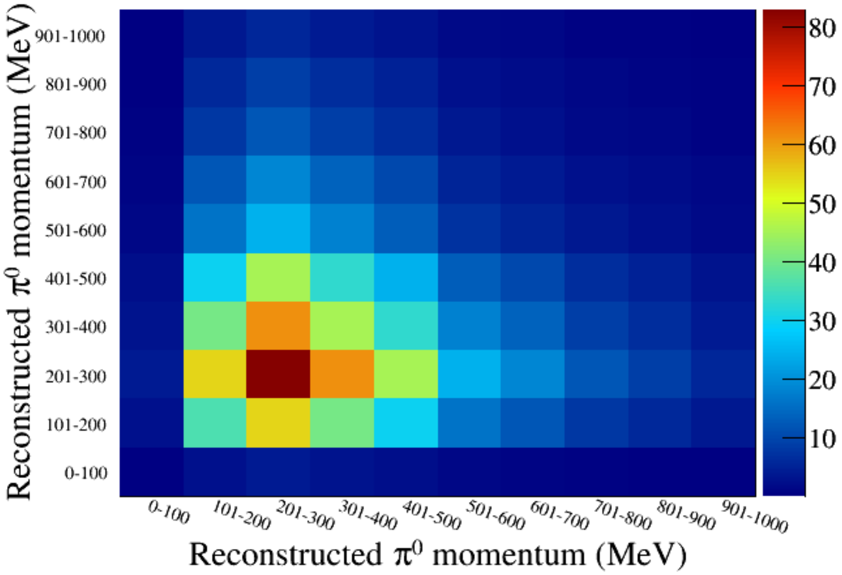

(b) Reconstructed π^0 momentum covariance matrix.

Figure 5.44: The covariance matrices used to calculate the size of the TPC cluster efficiency systematic.

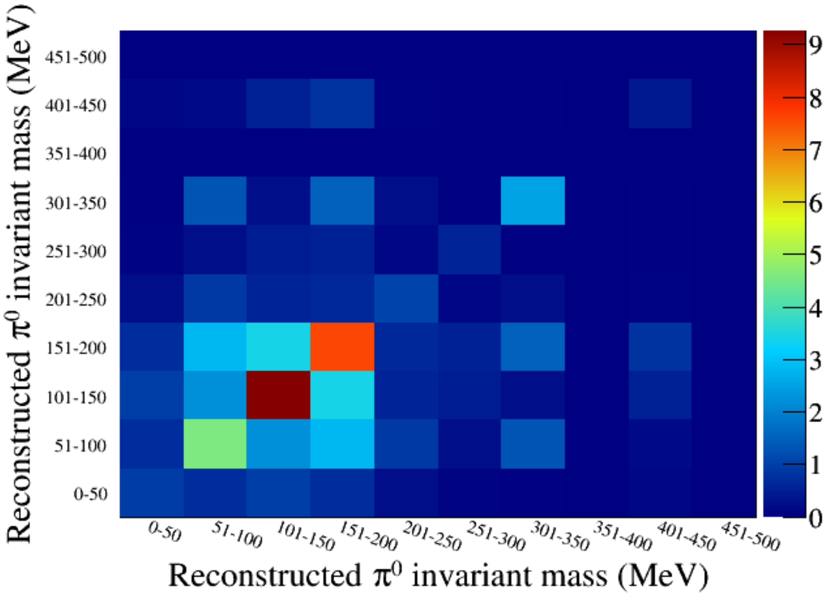

(a) Reconstructed π^0 invariant mass covariance matrix.

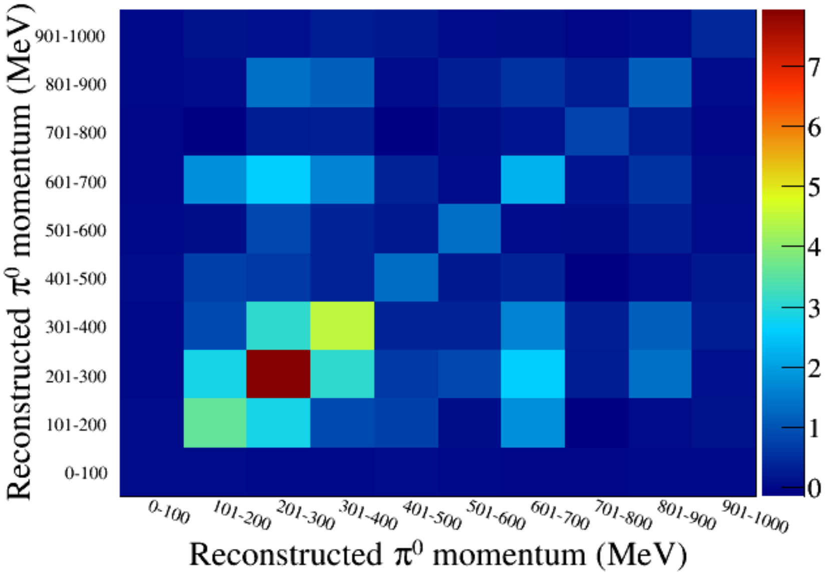

(b) Reconstructed π^0 momentum covariance matrix.

Figure 5.45: The covariance matrices used to calculate the size of the TPC PID systematic.

For the reconstructed π^0 invariant mass and reconstructed π^0 momentum covariance matrices, which yield a systematic error of 0.05%, see Figure 5.46.

5.5.11 Cross-Section Systematic

The cross-section systematic aims to assess the inherent uncertainties in the individual cross-section model parameters. Reweighting is utilised, offering two main benefits. Firstly, without the need to re-run the MC hundreds, or thousands, of times, the effect of varying the parameters can be found - this saves computational time. Secondly, it prevents conflation of MC statistical errors with systematic errors when using multiple data sets. For full details of the T2KReweight package that performs reweighting, see [122] and [123]. To calculate an event weight the following equation is used:

$$w_i = \frac{\sigma(\bar{x}^{'})}{\sigma(\bar{x})}. \tag{5.28}$$

Here, i is the event number, $\sigma(\bar{x}^{'})$ is the modified cross-section after parameter variation and $\sigma(\bar{x})$ is the nominal cross-section. For the parameters varied, see Table 5.24 and the explanations below.

Table 5.24: The cross-section parameters and the apposite errors varied in the cross-section systematic.

Parameter	E_ν Range (GeV)	Best Fit 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
S_F	all	0 (off)	1 (on)
CC Other shape	all	0.00	0.40
Pion-less Δ decay	all	0.00	0.20
CCQE E1	0.0 - 1.5	1.00	0.11
CCQE E2	1.5 - 3.5	1.00	0.30
CCQE E3	> 3.5	1.00	0.30
CC1π^+ E1	0.0 - 2.5	1.15	0.43
CC1π^+ E2	> 2.5	1.00	0.40
CC Coh	all	1.00	1.00
NC1π^0	all	0.96	0.43
NC1$\pi^\pm$	all	1.00	0.30
NC Coh	all	1.00	0.30
NC Other	all	1.00	0.30

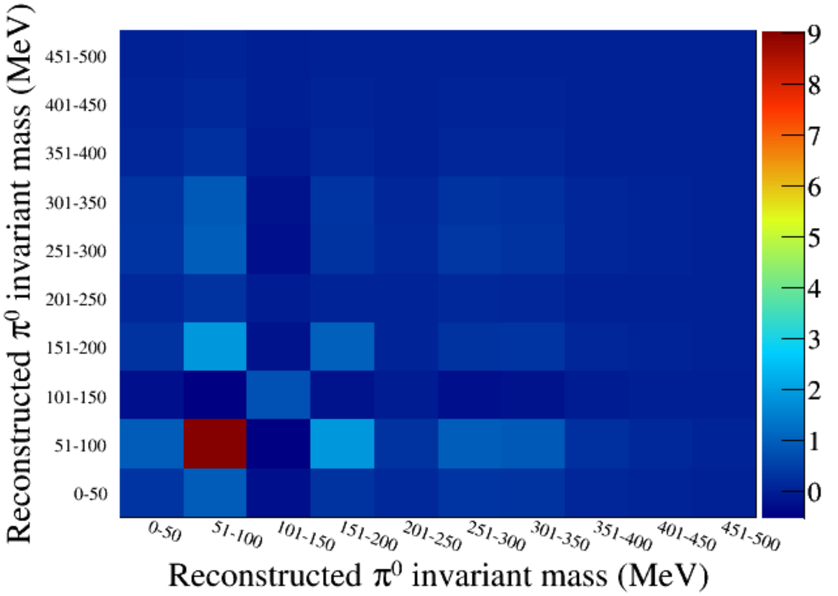

(a) Reconstructed π^0 invariant mass covariance matrix.

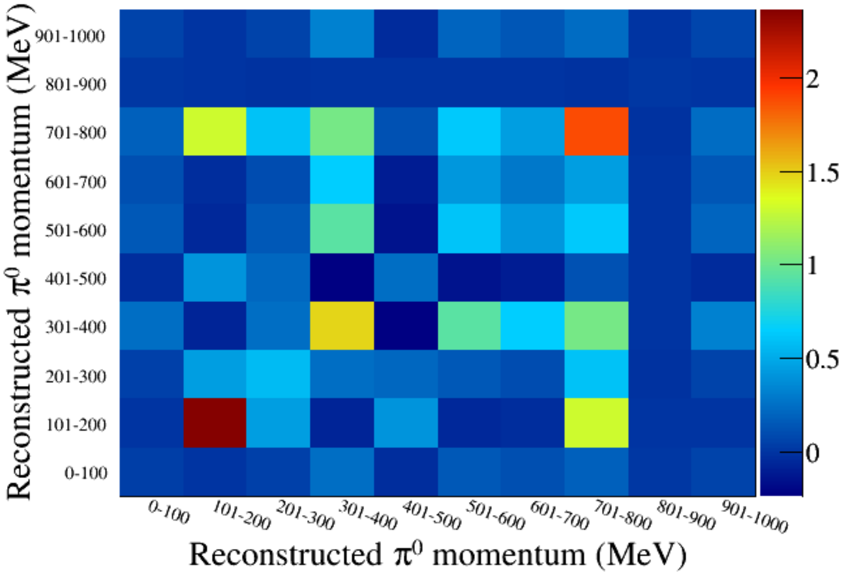

(b) Reconstructed π^0 momentum covariance matrix.

Figure 5.46: The covariance matrices used to calculate the size of the TPC charge ID systematic.

M_A^{QE} and Low Energy CCQE Normalisation (CCQE E1)

M_A^{QE} is the CCQE cross-section axial mass form factor. It is used by both the Llewellyn-Smith model for nucleon CCQE scattering [124], and also the Smith-Moniz Relative Fermi Gas (RFG) model to describe nucleons bound in a nuclear potential [125], [126]. To extract the error for M_A^{QE}, the MiniBooNE error could not be used as no covariance matrix was available. Thus, the NEUT CCQE MC was fit to the MiniBooNe CCQE corrected double-differential cross-section data, using a minimised χ^2 with M_A^{QE} and the normalisation factor as free parameters. The difference between the best fit, 1.64±0.03 GeV/c^2 (with the addition of the fitted error), and the NEUT nominal, 1.21 GeV/c^2, yields an error of 0.45. Furthermore, the CCQE E1 uncertainty was taken as 11% - the uncertainty in the MiniBooNE neutrino flux.

M_A^{RES}, CC1π^+ E1, NC1π^0, CC Coh, NC1$\pi^\pm$, NC Coh and NC Other

M_A^{RES} is the axial mass form factor for resonant single pion production from NC and CC interactions [74, 127]. MiniBooNe CC1π^0, CC1π^+ and NC1π^0 samples are fit using a χ^2 minimisation with M_A^{RES}, CC1π^+ E1 and NC1π^0 fit simultaneously.[24] As NC1π, CC other and NC other contribute such a small proportion of events, they cannot be used in the fit but instead contribute penalty terms. Additionally, NC Coh cannot be constrained as the spectrum of the π^0 momentum is almost identical to that of NC1π^0, hence contributes a further penalty term. Penalty terms are used as, although constraining is not possible, they do have an effect on the prediction of the overall MC. The error of each of these samples is found by simply scaling to the apposite data samples.

P_f and E_B

P_f is the Fermi momentum whilst E_B is the nuclear binding energy - assumed to be constant through the Fermi sphere - of the RFG mode, used by NEUT to model CCQE interactions. These values are 274 MeV/c and 25 MeV respectively, with the apposite errors deduced from electron scattering data [128].

S_F

S_F is the spectral function (SF) [129], a more realistic model of nuclear potential than the RFG model as it better replicates the electron scattering data. Instead of a constant binding energy, the SF defines a probability distribution. This distribution models the momenta and energy needed to liberate a nucleon from the nuclear environment. To calculate the systematic uncertainty due to the model dependence of the RFG and SF models, both cross-sections are renormalised, so they agree at 1 GeV, and the fractional difference taken as the error. When producing weights for the overall cross-section systematic, a positive throw sets the S_F to 1 whilst a negative throw 0.

CCQE E2, CCQE E3 and CC1π^+ E2

Due to low fluxes at these energies, the CCQE normalisations and errors are calculated by comparing MiniBooNE data [130] to that of the higher energy NOMAD

---

[24] Leading to correlations between the samples.

experiment [131]. Additionally, the CC1π^+ normalisation and error is deduced by comparing the NEUT nominal MC to the MiniBooNE data at 2 GeV.

CC Other shape

CC other shape encompasses CC multi-pion, CC DIS and CC resonant production where the outgoing particle is an η, K or γ. MINOS [132] found the error at an energy of 4 GeV to be 10%. From this, to model the energy dependent nature of this uncertainty, the following equation is used:

$$\sigma_{\text{CC Other}} = \frac{0.4}{E_\nu(\text{GeV})}. \tag{5.29}$$

Note, at decreasing energies this error will tend to infinity, as such a 0.6 GeV energy threshold is stipulated.

Pion-less Δ decay

Approximately 20% of NEUT Δ-decays are pion-less. This introduces a new source of error as CC1π^+ events can look CCQE-like - an irreducible background in CCQE fits using single-pion data. To parameterise this, pion-less Δ-decay in NEUT was reduced to 0% and the effect this had on the CC1π^+ E1 and NC1π^0 sample was found to be 20%.

5.5.12 FSI Systematic

The FSI systematic addresses the uncertainties due to strong re-interactions of hadrons within the nucleus. To understand such effects on pion FSI, NEUT uses a microscopic cascade model [133] in which particles are propagated in 'steps' through the nuclear environment - interacting according to the specific particle types' mean free path. The first of the four processes simulated is charge exchange, in which the pion undergoes a single (e.g. π^+ to π^0), or double (e.g. π^+ to π^-), charge change. The second is absorption, in which the pion produced by the neutrino interaction is absorbed by the nucleus and does not leave the nuclear environment. Thirdly, scattering of pions (inelastic only) is modelled. Finally, $\sim$30% of the time particle production occurs where multiple pions exit the nuclear environment.

To simulate such effects and produce the FSI systematic, six parameters are varied: FSIQEL, which models low energy quasi-elastic scattering; FSIQEH, which models high energy quasi-elastic scattering; FSIINEL, which models pion production for high momentum pions; FSIABS, which models absorption; FSICXL, which models low energy single charge exchange and FSICXH, which models high energy single charge exchange. However, as these are highly correlated, each cannot be varied separately. As such, 16 'parameter sets' are calculated by a separate working group, see [133] - [135], with each set of the six parameters corresponding to a throw that can be used to produce a covariance matrix.

5.5.13 Flux Systematic

The final systematic in this analysis is the flux systematic, due to the inherent uncertainty in the neutrino flux. To calculate this, a dedicated working group provides a covariance matrix binned according to neutrino flavour, detector and neutrino energy.

The flux covariance matrix is produced by summing the separate covariance matrices of each underlying parameter affecting the flux uncertainty. In the case where such parameters are correlated, a flux reweighting procedure is used to produce a covariance matrix. However, for single uncorrelated parameters, two throws are undertaken representing a throw of $\pm 1\sigma$.

In total there are 5 factors analysed to understand the flux uncertainty. The first is the production and cross-section uncertainties of pions and kaons. These are obtained predominantly from NA61/SHINE data [136], [137] along with earlier studies [138], [139]. This uncertainty undergoes a reweighting procedure as it is correlated with the proton beam uncertainty. Secondly, the uncertainty regarding the off-axis angle is monitored by MUMON and INGRID - this is treated as a single variable. Thirdly, the horn and target alignment uncertainty, also treated as a single variable, is measured with a dedicated survey. Furthermore, the proton beam uncertainty, as stated above, is correlated with the production uncertainties of pions and kaons. Variables affecting this include the magnetic field uncertainty and beam position - measured by beam monitors. The final variable is the magnetic horn current, again treated as a single variable, with the uncertainty monitored using Rogowski coils [82].

5.5.14 Further Systematics

The systematic analyses presented above are expected to be the dominant contributions to the uncertainty of the final result. Detector systematics associated with other cuts were not considered as these cuts have a negligible effect on the total number of events. For example, in the case of the SMRD activity cut, suppose the track selection efficiency has a large uncertainty, due to the small proportion of background events cut, the total error will likely be minimal. Nonetheless, for the next generation analysis, a complete systematic analysis will need to asses each cut; possible avenues are:

Shower cut As the shower cut has two parameters to consider, the energy of the cluster and whether it is reconstructed as shower-like or not, each must be assessed. A possible methodology would be to compare the width in the data and MC energy distributions and utilise a variation systematic to replicate this effect. Additionally, a weight systematic could be used to simulate the data/MC discrepancy of the shower-like ECal hypothesis. However, studies of the correlations between these two must be addressed. If the correlation is minimal two separate systematics can be used, if highly correlated a covariance matrix method will be needed.

Time uncertainty The 10 ns timing difference was used as this width results in a minimal data/MC discrepancy. Nonetheless a variation systematic would be needed to fully understand the extent of its negligibility.

SMRD activity As with the isolated ECal efficiency, there will be an efficiency systematic associated with reconstructing SMRD tracks.

TPC1 Veto Although this may be covered by the TPC cluster efficiency, a full systematic similar to the ECal efficiency may be more representative. However, double counting when using other TPC systematics would need to be scrutinised.

P0D Veto Again, much like the ECal efficiency systematic, an analogous systematic is needed for the P0D veto. However, P0D systematics have not been addressed in the current software framework.

Charge integral When assessing the BDT systematic it was found that the charge variation width was minimal, thus having little effect on higher level variables. As such, it is likely that the charge integral is also robust against such uncertainties. Nonetheless, this needs further analysis to measure the level of negligibility.

FGD Mass As the target is the FGD FV, the uncertainty in the FGD mass needs to be accounted for. However, as the component parts are known to a high degree of accuracy, this is expected to have an insignificant effect.

ECal Mass As photon conversion in the ECal is the cornerstone of this analysis, the ECal mass uncertainty needs to be analysed. However, like the FGD, each component is measured to a high accuracy and as such should have a small effect.

5.5.15 Cross-section extraction

The flux integrated cross-sections, per bin, as a function of π^0 momentum and angle are:

$$\langle\frac{d\sigma_i}{dp_{\pi^0}}\rangle = \frac{N(p_{\pi^0})_i^{\text{unfolded}}}{\Phi N_{\text{targets}}\Delta_i\epsilon_i} \tag{5.30}$$

and

$$\langle\frac{d\sigma_i}{d\cos\theta_{\pi^0}}\rangle = \frac{N(\cos\theta_{\pi^0})_i^{\text{unfolded}}}{\Phi N_{\text{targets}}\Delta_i\epsilon_i}. \tag{5.31}$$

Here, i is the bin number, $N(p_{\pi^0})_i^{\text{unfolded}}$ and $N(\cos\theta_{\pi^0})_i^{\text{unfolded}}$ are the estimated number of true events, N_{targets} is the number of target nucleons, Δ_i is the bin width and ϵ_i is the signal selection efficiency. To deduce N_i^{unfolded}, a Bayesian unfolding method is used, in which bin-to-bin migration and background subtraction is taken into account.

5.5.15.1 Bayesian unfolding

Bayesian unfolding [140] is a method designed to predict, from the reconstructed spectrum, the true spectrum. The cornerstone of this method is the smearing matrix:

$$R_{ij} = P(r_i|t_j). \tag{5.32}$$

This matrix is a representation of the probability of reconstructing an event in bin i, whilst the true event is in bin j. To predict $P(t_j)$, the number of true events in each bin is needed:

$$P(t_j) = \frac{N_{t_j}}{\sum\limits_{\alpha}^{\text{true bins}} N_{t_\alpha}}, \tag{5.33}$$

where N_{t_j} is the number of true events in bin j. Furthermore, to predict $P(r_i)$, the following equation is used:

$$P(r_i) = \sum_{\alpha}^{\text{true bins}} P(r_i|t_\alpha)P(t_\alpha) = \frac{N_{r_i}}{\sum\limits_{\alpha}^{\text{true bins}} N_{t_\alpha}} \tag{5.34}$$

where N_{r_i} is the number of reconstructed events in bin i. By utilising Bayes' theorem, the unsmearing matrix can be calculated:

$$P(t_j|r_i) = \frac{P(r_i|t_j)P(t_j)}{P_{r_i}}. \tag{5.35}$$

Now, the true distribution can be deduced from the reconstructed data distribution:

$$N_{t_j}^{\text{unfolded}} = \frac{1}{\epsilon_j}\sum_i P(t_j|r_i)(N_{r_i} - B_{r_i}). \tag{5.36}$$

Here, ϵ_j is the efficiency in true bin j, N_{r_i} is the number of reconstructed data events in bin i and B_{r_i} is the MC predicted number of background events in bin i. At this point, one iteration has been undertaken. To produce a more data driven estimate, this process can be repeated by inserting Equation 5.36 into Equation 5.33, and reproducing a new unsmearing matrix and value for $N_{t_j}^{\text{unfolded}}$.

5.5.15.2 Final uncertainties

To obtain the final errors on the measured cross-section results for each of the uncertainties described in the previous section, a covariance matrix is produced:

$$V_{ij}^{s} = \frac{1}{N}\sum_{s_n=1}^{N}(\sigma_i^{s_n} - \sigma_i^{\mathrm{nom}})(\sigma_j^{s_n} - \sigma_j^{\mathrm{nom}}). \tag{5.37}$$

Here, s is the error source, N is the number of toy throws, $\sigma_i^{s_n}$ is the thrown cross-section for bin i and σ_i^{nom} is the nominal cross-section for bin i.

5.6 Results

The final cross-section measurements on the combination of a C_8H_8 and H_2O target (Bayesian unfolded using bins of both; π^0 momentum, p_{π^0}, and of the cosine of the angle between the π^0 direction and the global z coordinate, $\cos\theta_{\pi^0}$), for each signal definition, the associated signal/MC ratio and the standard deviation discrepancy from the NEUT model prediction are presented in Table 5.25. These were produced using the T2K ν_μ beam with its peak energy of ~0.6 GeV.

For each of the topologies the same method of cross-section extraction was implemented, as explained in Section 5.5.15. Additionally, the same detector configurations were used and the same event selections implemented. However, as the event selection did not consist of one set of linear cuts, but split into three distinct 'branches' depending on the position of the reconstructed isolated ECal clusters (with the only major difference between each branch being the different BDT weights files used), it is obvious that some topologies will 'prefer' different branches. For example, it is expected that, assuming the π^0 decay photons are reconstructed correctly, the BrlBrl topology would be more likely to use the branch requiring two isolated ECal clusters reconstructed in the Brl than the DsDs topology.

In each topology the measured cross-section is lower than the NEUT model. The largest discrepancy occurs in the BrlBrl exclusive topology when unfolded with respect to $\cos\theta_{\pi^0}$, in which the predicted cross-section is 0.271 $\times 10^{-40}$ cm^2/nucleon compared to a measured value of 0.165 ± 0.012 (stat) ± 0.022 (syst) $\times 10^{-40}$ cm^2/nucleon - a discrepancy of 4.230σ. Conversely, the smallest discrepancy is in the DsDs exclusive topology, when unfolded with respect to $\cos\theta_{\pi^0}$, in which the predicted cross-section is 0.101 $\times 10^{-40}$ cm^2/nucleon compared to a measured value of 0.070 ± 0.019 (stat) ± 0.011 (syst) $\times 10^{-40}$ cm^2/nucleon - a discrepancy of 1.412σ.

For each differential cross-section distribution and the corresponding breakdown of the error sources see Figures 5.47 - 5.56 and Tables 5.26 - 5.35 respectively.[25] For the apposite total error covariance matrices and the Highland error covariance matrices, see Appendix C. Interestingly, although for each topology there is a cross-section deficit, there is remarkable agreement between the shape of the distributions, suggesting that the source of the discrepancy is a normalisation shift.

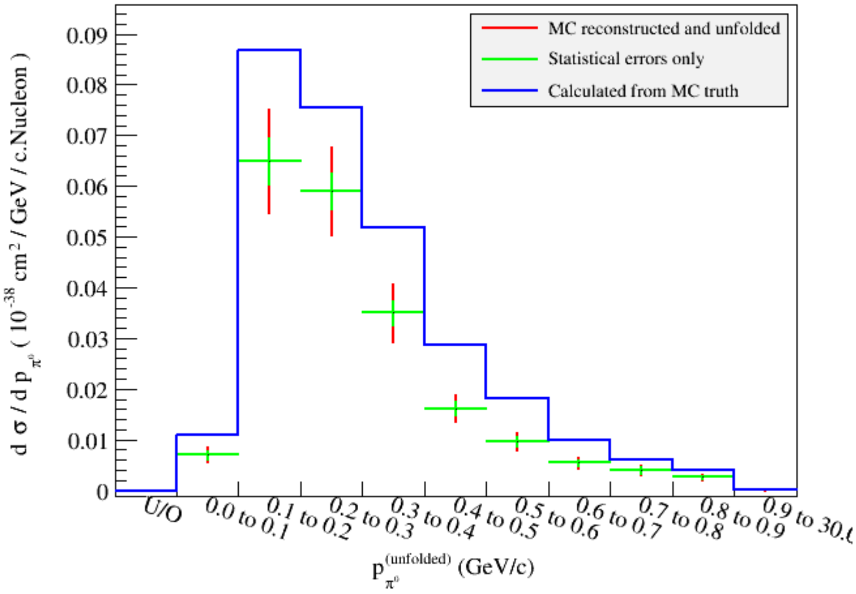

Figure 5.47: The measured differential cross-section with respect to momentum for the main, NC1π^0 inclusive, topology.

---

[25] Note, in the Figures, U/O represents the underflow and overflow bins.

Table 5.25: The predicted and measured cross-section for each topology, the apposite data/MC ratio and the discrepancy between the NEUT model and measured values. The main, inclusive, result shows a measured cross-section deficiency compared to the predicted cross-section of 3.097 and 3.567 σ (when unfolded with respect to momentum and cosθ respectively). Furthermore, each topology shows a data deficit compared to the NEUT MC prediction, with the largest discrepancy of 4.230 σ for the BrlBrl exclusive topology (when unfolded with respect to cosθ) and the smallest discrepancy of 1.412 σ for the DsDs exclusive topology (when unfolding with respect to momentum). Note, there is a small difference when unfolding with respect to momentum and cosθ, however, within each topology these values are within error.

Topology	Predicted cross-section ($\times 10^{-40}$ cm^2/Nucleon)	Measured cross-section ($\times 10^{-40}$ cm^2/Nucleon)	Data/MC Ratio	Discrepancy (σ)
Inclusive (p unfolded)	3.087	2.145 ± 0.109 (stat) ± 0.284 (syst)	0.695 ± 0.099	3.097
Inclusive ($\cos\theta$ unfolded)	3.087	2.037 ± 0.110 (stat) ± 0.273 (syst)	0.660 ± 0.095	3.567
Exclusive ECalECal (p unfolded)	0.564	0.383 ± 0.019 (stat) ± 0.051 (syst)	0.679 ± 0.096	3.326
Exclusive ECalECal ($\cos\theta$ unfolded)	0.564	0.380 ± 0.020 (stat) ± 0.052 (syst)	0.674 ± 0.099	3.303
Exclusive DsDs (p unfolded)	0.101	0.077 ± 0.012 (stat) ± 0.011 (syst)	0.762 ± 0.161	1.474
Exclusive DsDs ($\cos\theta$ unfolded)	0.101	0.070 ± 0.019 (stat) ± 0.011 (syst)	0.693 ± 0.217	1.412
Exclusive DsBrl (p unfolded)	0.192	0.144 ± 0.011 (stat) ± 0.020 (syst)	0.750 ± 0.119	2.103
Exclusive DsBrl ($\cos\theta$ unfolded)	0.192	0.145 ± 0.011 (stat) ± 0.021 (syst)	0.755 ± 0.123	1.983
Exclusive BrlBrl (p unfolded)	0.271	0.167 ± 0.013 (stat) ± 0.022 (syst)	0.616 ± 0.094	4.070
Exclusive BrlBrl ($\cos\theta$ unfolded)	0.271	0.165 ± 0.012 (stat) ± 0.022 (syst)	0.609 ± 0.092	4.230

Table 5.26: The errors associated with the NC1π^0 inclusive measurement when unfolded using π^0 momentum.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.107	4.972
MC statistics	0.024	1.118
Flux	0.181	8.444
FSI	0.047	2.195
Cross-section	0.142	6.619
Detector	0.160	7.453

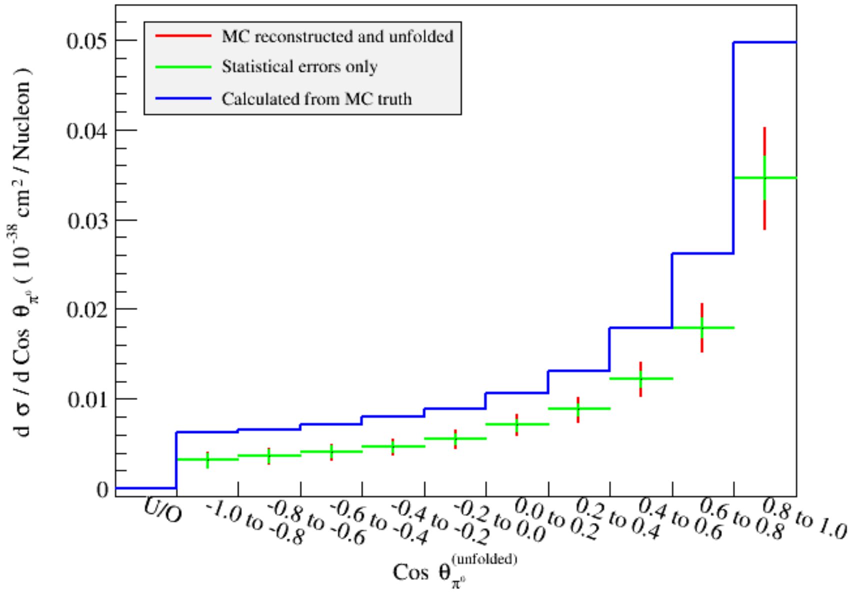

Figure 5.48: The measured differential cross-section with respect to cosθ for the main, NC1π^0 inclusive, topology.

Table 5.27: The different errors associated with the NC1π^0 inclusive measurement when unfolded using π^0 cosθ.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.107	5.245
MC statistics	0.025	1.229
Flux	0.172	8.432
FSI	0.048	2.365
Cross-section	0.138	6.765
Detector	0.154	7.565

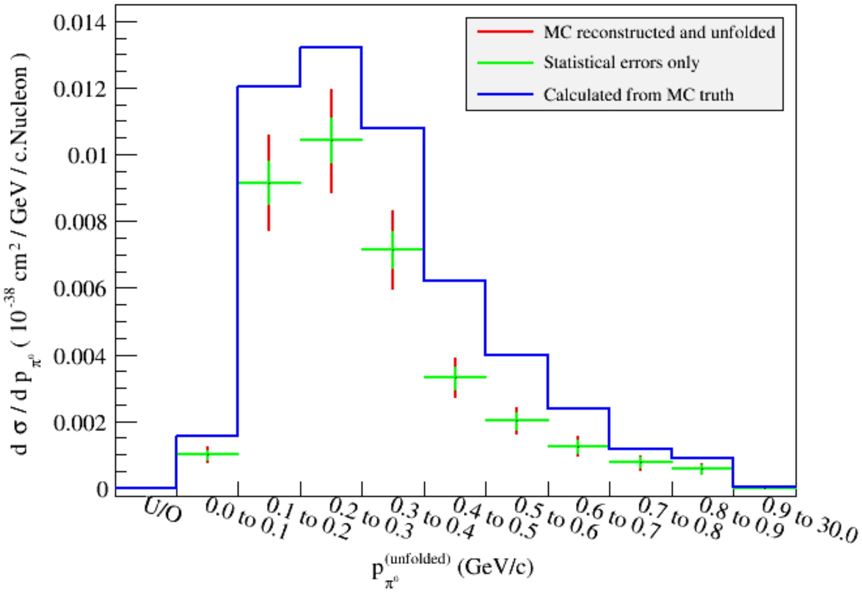

Figure 5.49: The measured differential cross-section with respect to momentum for the NC1π^0 ECalECal topology.

Table 5.28: The different errors associated with the NC1π^0 ECalECal measurement when unfolded using π^0 momentum.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.019	4.829
MC statistics	0.005	1.388
Flux	0.032	8.410
FSI	0.008	2.007
Cross-section	0.025	6.554
Detector	0.029	7.571

Table 5.29: The different errors associated with the NC1π^0 ECalECal measurement when unfolded using π^0 cosθ.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.019	4.976
MC statistics	0.005	1.295
Flux	0.032	8.452
FSI	0.008	2.233
Cross-section	0.025	6.642
Detector	0.030	7.965

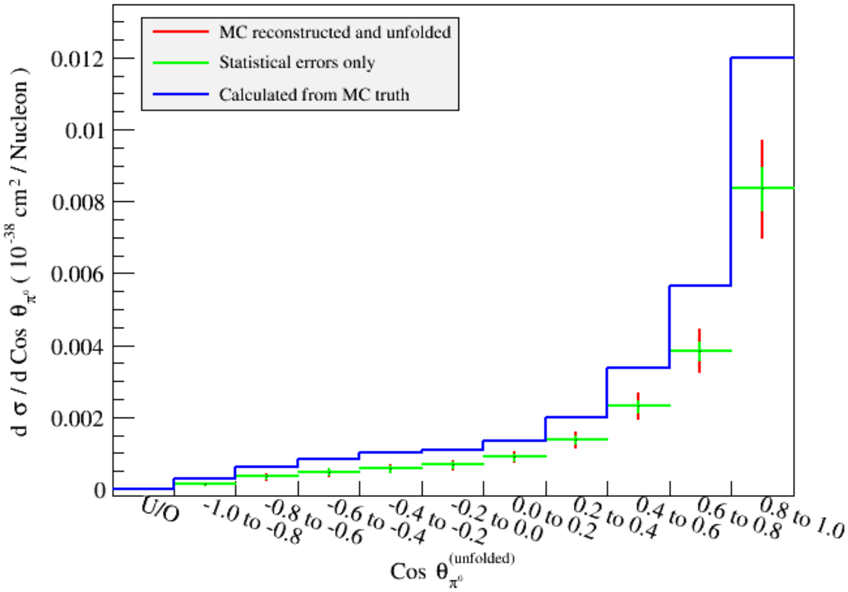

Figure 5.50: The measured differential cross-section with respect to cosθ for the NC1π^0 ECalECal topology.

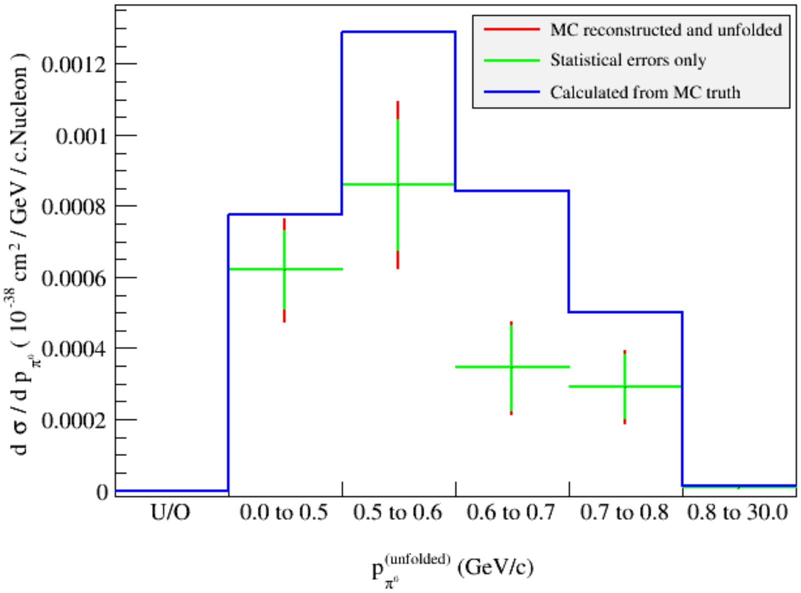

Figure 5.51: The measured differential cross-section with respect to momentum for the NC1π^0 DsDs topology.

Table 5.30: The different errors associated with the NC1π^0 DsDs measurement when unfolded using π^0 momentum.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.012	15.579
MC statistics	0.003	4.073
Flux	0.007	8.601
FSI	0.001	1.857
Cross-section	0.004	5.623
Detector	0.008	9.767

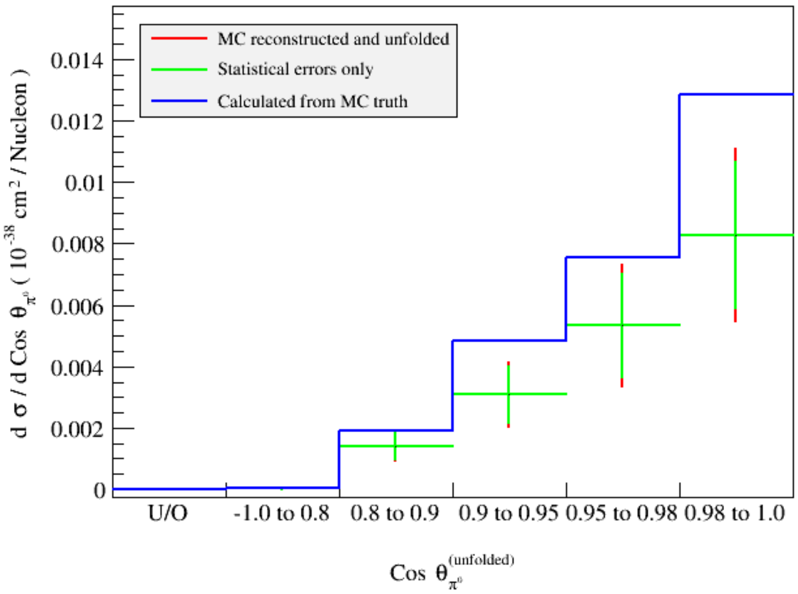

Figure 5.52: The measured differential cross-section with respect to cosθ for the NC1π^0 DsDs topology.

Table 5.31: The different errors associated with the NC1π^0 DsDs measurement when unfolded using π^0 cosθ.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.018	25.091
MC statistics	0.006	8.050
Flux	0.006	8.873
FSI	0.002	2.672
Cross-section	0.003	4.724
Detector	0.008	11.756

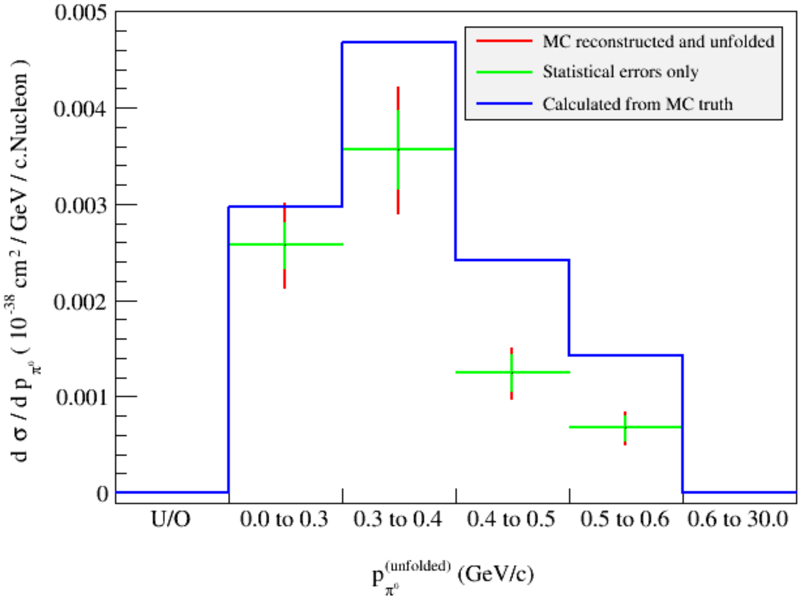

Figure 5.53: The measured differential cross-section with respect to momentum for the NC1π^0 DsBrl topology.

Table 5.32: The different errors associated with the NC1π^0 DsBrl measurement when unfolded using π^0 momentum.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.011	7.608
MC statistics	0.003	2.211
Flux	0.012	8.346
FSI	0.003	1.807
Cross-section	0.010	6.745
Detector	0.013	9.006

Table 5.33: The different errors associated with the NC1π^0 DsBrl measurement when unfolded using π^0 cosθ.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.011	7.540
MC statistics	0.003	2.170
Flux	0.012	8.337
FSI	0.003	1.948
Cross-section	0.010	6.922
Detector	0.013	8.956

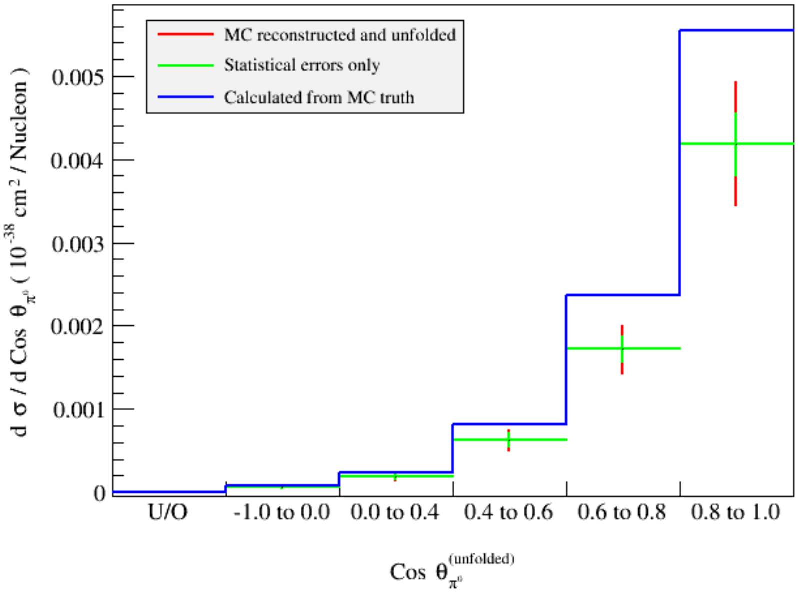

Figure 5.54: The measured differential cross-section with respect to cosθ for the NC1π^0 DsBrl topology.

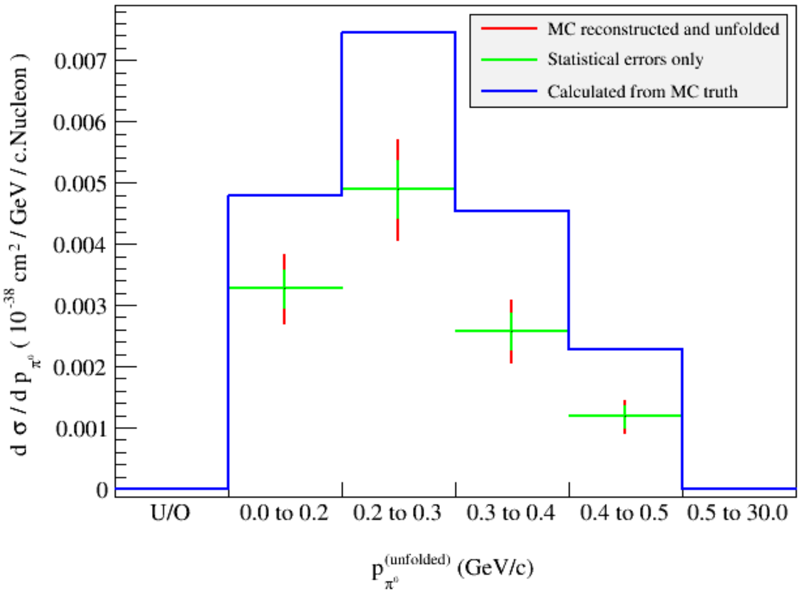

Figure 5.55: The measured differential cross-section with respect to momentum for the NC1π^0 BrlBrl topology.

Table 5.34: The different errors associated with the NC1π^0 BrlBrl measurement when unfolded using π^0 momentum.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.012	7.441
MC statistics	0.003	1.904
Flux	0.014	8.398
FSI	0.004	2.646
Cross-section	0.012	7.039
Detector	0.012	6.928

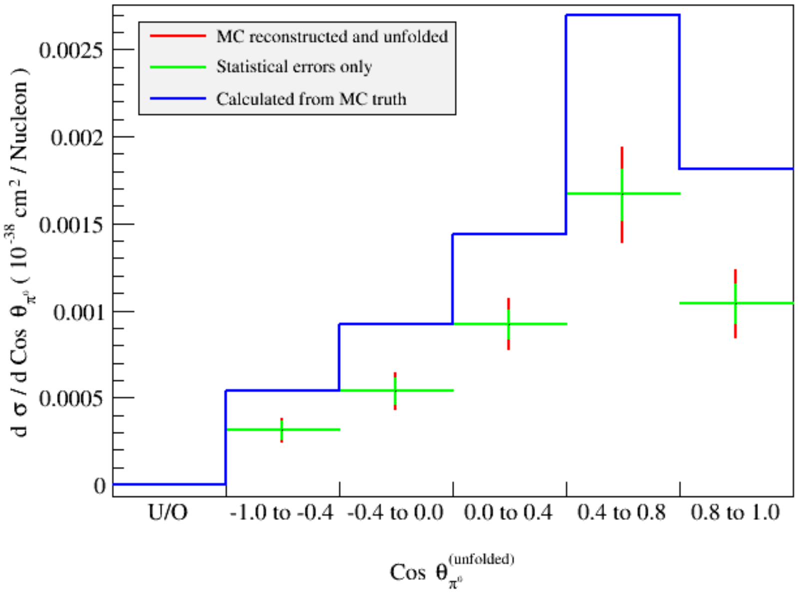

Figure 5.56: The measured differential cross-section with respect to cosθ for the NC1π^0 BrlBrl topology.

Table 5.35: The different errors associated with the NC1π^0 BrlBrl measurement when unfolded using π^0 cosθ.

Error	Absolute Error ($\times 10^{-40}$ cm^2/Nucleon)	Relative Error (%)
Data statistics	0.012	7.303
MC statistics	0.003	1.860
Flux	0.014	8.401
FSI	0.005	2.814
Cross-section	0.012	7.327
Detector	0.011	6.858

Chapter 6

Conclusions

T2K is a long baseline neutrino oscillation experiment located at the J-PARC facility in eastern Japan. The baseline length is 295 km, with both ND280 and Super-Kamiokande located 2.5° off-axis - producing a narrow energy neutrino beam peaked at ∼0.6 GeV.

The analysis presented in this thesis is an NC1π^0 cross-section measurement in which the neutrino interacts in the FGD fiducial volumes (FVs) of ND280 and the resulting π^0 decay photons convert in the Tracker ECals. This is of importance to T2K as NC1π^0 interactions represent the second largest uncertainty in the ν_e appearance measurement - behind only the intrinsic ν_e component of the beam. The reason for this is that the low energy photons resulting from the π^0 decay can mimic the ν_e appearance signal at Super-Kamiokande if either one of the photons is not reconstructed or the decay produces two collinear photons - both scenarios produce an electron-like Cherenkov ring. Hence this measurement can help to constrain such uncertainties. Furthermore, it is also of use to the wider neutrino community as few NC1π^0 measurements exist in the field. Thus such studies are invaluable in testing and improving cross-section models.

The first step in this analysis was to scrutinise each cut of a previous, first generation, analysis. This was necessary as many improvements have been made to ECal reconstruction since the previous analysis was produced. Each cut was reassessed in order to maximise $\epsilon \times \rho$. Thus some cuts remained, others were altered and some removed completely. Furthermore, new cuts were implemented and were again optimised using the $\epsilon \times \rho$ metric. Once complete the relevant systematics were analysed, this included undertaking new studies, particularly those relating to the ECals. Finally, a Bayesian unfolding procedure was used to determine the absolute and differential cross-sections with respect to the π^0 momentum and angle from the z-axis.

The main result of interest is the NC1π^0 inclusive cross-section which was found to be 2.145 ± 0.109 (stat) ± 0.284 (syst) $\times 10^{-40}$ cm^2/nucleon (when unfolded with respect to momentum) and 2.037 ± 0.110 (stat) ± 0.273 (syst) $\times 10^{-40}$ cm^2/nucleon (when unfolded with respect to cosθ), compared to a predicted value of 3.087 $\times 10^{-40}$ cm^2/nucleon - a standard deviation discrepancy of 3.097 and 3.567 respectively. Comparison to the much higher measurement produced by MiniBooNE [63], 4.76 ± 0.05 (stat) ± 0.76 (syst) $\times 10^{-40}$ cm^2/nucleon, suggests much tension between the two. However, although both experiments have a similar neutrino beam energy, $\mathcal{O}(1\text{ GeV})$, to fully understand the discrepancies between both measurements in the future, it is imperative to analyse the differing kinematics between the two experiments - T2K having a narrow-band and MiniBooNE a wide band beam - alongside understanding the effect of the different targets.

Interestingly, for each of the topologies, the measured cross-section showed a deficit compared to that predicted by NEUT. The largest deficit occurred in the exclusive BrlBrl topology, when unfolded with respect to $\cos\theta_{\pi^0}$, in which a cross-section of 0.271 $\times 10^{-40}$ cm^2/Nucleon was expected, whilst a value of 0.165 ± 0.012 (stat) ± 0.022 (syst) $\times 10^{-40}$ cm^2/Nucleon was measured - a 4.230σ discrepancy. Conversely, the smallest discrepancy occurred in the exclusive DsDs topology, when unfolded with respect to $\cos\theta_{\pi^0}$, in which a cross-section of 0.101 $\times 10^{-40}$ cm^2/Nucleon was expected, whilst a value of 0.070 ± 0.019 (stat) ± 0.011 (syst) was measured - a 1.412σ discrepancy. Despite the fact that the significance between the predicted and measured results varies between the topologies, the data/MC ratio for all topologies are consistent with one another - ranging from 0.609 ± 0.092 to 0.762 ± 0.161 - suggesting a normalisation shift. Furthermore, when analysing the $d\sigma/dp_{\pi^0}$ and $d\sigma/dcos\theta_{\pi^0}$ distributions, the deficit again appears to be due to a normalisation shift and is also not a function of momentum or angle.

The reasons for the discrepancies between these results and the NEUT predictions possibly occur due to a wide range of factors. However, an area which could account for a large component of the difference is the FSI modelling; particularly the modelling of pion charge exchange and absorption, alongside the use of 'formation zones'.[1] In the case of charge exchange, an inflated NEUT cross-section prediction would occur in the scenario in which, upon interaction with a nucleus, $\pi^0 \rightarrow \pi^\pm$ charge exchange is under-predicted (similarly, if $\pi^\pm \rightarrow \pi^0$ is over-predicted, the NEUT prediction will also be increased). Furthermore, in the scenario in which π^0 absorption is modelled to occur too rarely, more neutral pions will make it to the final state - again inflating the NEUT cross-section. Finally, the use of formation zones could align the cross-section predictions with the data. However, there are competing effects when varying the size of the zone. If, for example, a charged pion is produced at the interaction vertex, by increasing the size of the formation zone, there is a decreased probability of producing a π^0 in the final state (owing to a reduced chance of re-interaction), hence the NEUT cross-section would decrease. Conversely, if a π^0 is produced at the interaction vertex, a reduction in the size of the formation zone would be required to reduce the NEUT cross-section (as this increases the chance that the π^0 will not make it out of the nucleus).

A further factor that could help account for the difference between the data and the NEUT prediction can be attributed to the cross-section modelling (for a more comprehensive discussion, see Section 1.8). For example, the latest version of NEUT now considers Berger-Sehgal coherent production [75] (whilst previous versions used the Rein-Sehgal model [141]). This is expected to reduce the overall coherent component of the NC1π^0

---

[1] A formation zone is an empirically motivated parameter used by interaction generators. It dictates a region around the interaction vertex in which newly produced particles do not interact.

cross-section prediction [142]. Unfortunately, however, as the coherent component of the NC1π^0 cross-section is expected to be on the order of ∼1%, this alone can not explain the discrepancies with the NEUT prediction.

Furthermore, this thesis presented the scintillator based Trip-t timing calibration procedure. This is the first stage in the timing calibration for the ECals, P0D and SMRD - searching and applying the relevant offsets for the 10 ns timeslips induced during resynchronisation of the clock modules. This procedure is essential for accurate detector reconstruction and effective inter-detector timing, which makes it imperative for the NC1π^0 analysis presented due to the extensive use of the ECals.

Bibliography

[1] W. Pauli, Pauli letter collection: letter to Hans Geiger *et al*, Typed copy, December 1930.

[2] R. Reines and C.L. Cowan, Jr, Phys. Rev. **92**, 830 (1953).

[3] R. Reines and C.L. Cowan, Jr, Nature, **178**, 446 (1956).

[4] G. Danby, J-M. Gillard, K. Goulianos, L.M. Lederman, N. Mistry, M. Schwartz and J. Steinberger, Phys. Rev. Lett. **9**, 36 (1962).

[5] M. L. Perl *et al*, Phys. Rev. Lett. **35**, 1489 (1975).

[6] G.J. Feldman *et al*, Phys. Rev. Lett. **38**, 576 (1977).

[7] The ALEPH Collaboration, The DELPHI Collaboration, The L3 Collaboration, the OPAL Collaboration, The SLD Collaboration, Phys. Rept. **427**, 257 (2006).

[8] The DONUT Collaboration, K. Kodama *et al*, Phys. Lett. B. **503**, 3, (2001).

[9] The Super-Kamiokande Collaboration, Y. Fukuda *et al*, Phys. Rev. Lett. **81**, 1562 (1998).

[10] B. Pontecorvo, Sov. Phys. JETP **6**, 429 (1957).

[11] R. Davis, D.S. Harmer and K.C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).

[12] J.N. Bahcall, M.H. Pinsonneault and G.J. Wasserburg, Rev. Mod. Phys. **67**, 781 (1995).

[13] The SAGE Collaboration, J.N. Abdurashitov *et al*, Phys. Rev. C. **60**, 055801 (1999).

[14] The GALLEX Collaboration, W. Hampel *et al*, Phys. Lett. B. **447**, 127 (1999).

[15] The Kamiokande Collaboration, Y. Fukuda *et al*, Phys. Rev. Lett. **77**, 1683 (1996).

[16] The Super-Kamiokande Collaboration, Y. Fukuda *et al*, Phys, Rev. Lett. **81**, 1158 (1998).

[17] The SNO Collaboration, Q.R. Ahmad *et al*, Phys. Rev. Lett. **87**, 071301 (2001).

[18] The SNO Collaboration, Q.R. Ahmad *et al*, Phys. Rev. Lett. **89**, 011301 (2002).

[19] D. Caspar *et al*, Phys. Rev. Lett. **66**, 2561 (1991).

[20] A. Surdo *et al*, Nucl. Phys. B. **110**, 342 (2002).

[21] W.W.M. Allison *et al*, Phys. Lett. B. **391**, 491 (1997).

[22] The Kamiokande Collaboration, K.S. Hirata *et al*, Phys. Lett. B. **280**, 146 (1992).

[23] The Super-Kamiokande Collaboration, Y. Fukuda *et al*, Phys. Lett. B. **433**, 9 (1998).

[24] T. Kajitia, Nucl. Phys. B. Proceeding Supplements. **77**, 123 (1999).

[25] The Super-Kamiokande Collaboration, Y. Fukuda *et al*, Phys. Rev. Lett. **81**, 1562 (1998).

[26] Z. Maki, M. Nakagawa and S. Sakata, Progress of Theoretical Physics **28**, 870 (1962).

[27] K.A. Olive *et al*, Chin. Phys. C. **38**, 090001 (2014).

[28] The KamLAND Collaboration, A. Gando *et al*, Phys. Rev. D. **88**, 033001 (2013).

[29] The Super-Kamiokande Collaboration, A. Himmel *et al*, AIP Conf. Proc. **1604**, 345 (2014).

[30] The MINOS Collaboration, P. Adamson *et al*, Phys. Rev. Lett. **107**, 181802 (2011).

[31] The T2K Collaboration, K. Abe *et al*, Phys. Rev. D. **91**, 072010 (2015).

[32] The Daya Bay Collaboration, F.P. An *et al*, Phys. Rev. Lett. **112**, 061801 (2014).

[33] The RENO Collaboration, S.B. Kim *et al*, Phys.Rev. Lett. **108**, 191802 (2012).

[34] The NOvA Collaboration, P. Adamson *et al*, Arxiv Physics e-prints **1601.05022**, (2016).

[35] The T2K Collaboration, K. Abe *et al*, Phys. Rev. Lett. **116**, 181801 (2016).

[36] The T2K Collaboration, K. Abe *et al*, Phys. Rev. Lett. **112**, 061802 (2014).

[37] The T2K Collaboration, K. Abe *et al*, Phys. Rev. Lett. **112**, 181801 (2014).

[38] S.F. King and C. Luhn, Rept. Prog. Phys. **76**, 056201 (2013).

[39] T. Nakaya and R.K. Plunked, Arxiv Physics e-prints **1507.08134**, (2015).

[40] The Hyper-Kamiokande Working Group, K. Abe *et al*, Arxiv Physics e-prints **1412.4673**, (2015).

[41] The DUNE Collaboration, R. Acciarri *et al*, Arxiv Physics e-prints **1512.06148**, (2016).

[42] The PINGU Collaboration, S.F. Ge *et al*, JHEP **1409**, 024 (2014).

[43] A.D. Sakharov, JETP. Lett. **5**, 24 (1967).

[44] Y. Takeuchi*et al*, JCAP **1401**, 046 (2014).

[45] S. Bilenky, Introduction to the Physics of Massive and Mixed Neutrinos **817**, 159 (2010).

[46] The EXO Collaboration, M. Auger *et al*, Phys. Rev. Lett. **109**, 032505 (2012).

[47] The GERDA Collaboration, M. Agostini *et al*, Nuclear and Particle Physics Proceedings. **273**, 1876 (2016).

[48] The KamLAND-Zen Collaboration, A. Gando *et al*, Phys. Rev. Lett. **117**, 082503 (2016).

[49] The SNO+ Collaboration, A. Maio *et al*, J. Phys. Conf. Ser. **587**, 012030 (2015).

[50] The MAJORANA Collaboration, M. Abgrall *et al*, Advances in High Energy Physics. **2014**, 365432 (2014).

[51] The SuperNEMO Collaboration, M. Bongrand *et al*, Physics Procedia **61**, 211 (2015).

[52] H.V. Klapdor-Kleingrothaus *et al*, Mod. Phys. Lett. A. **16**, 2409 (2001).

[53] J.M. Conrad *et al*, Annu. Rev. Nucl. Part. Sci. **63**, 45 (2013).

[54] S. Gninenko, Phys. Rev. D. **83**, 015015 (2011).

[55] D. Casper, Nucl. Phys. Proc. Suppl. **112**, 161 (2002).

[56] J.A. Formaggio and G.P. Zeller, Rev. Mod. Phys. **84**, 1307 (2012).

[57] W. Krenz *et al*, Nucl. Phys. B. **135**, 45 (1978).

[58] H. Faissner *et al*, Phys. Lett. B. **125**, 230 (1983).

[59] M. Derrick *et al*, Phys. Rev. D. **23**, 569 (1981).

[60] W. Lee *et al*, Phys. Rev. Lett. **38**, 202 (1977).

[61] P.J. Nienaber, PhD Thesis, University of Illinois at Urbana-Champaign, (1988).

[62] S. Nakayama *et al*, Phys. Lett. B. **619**, 255 (2005).

[63] A.A. Aguilar-Arevalo *et al*, Phys. Rev. D. **81**, 013005 (2010).

[64] Y. Kurimoto *et al*, Phys. Rev. D. **81**, 111102 (2010).

[65] R. Acciarri *et al*, Arxiv Physics e-prints **1511.00941**, (2015).

[66] C.T. Kullenberg *et al*, Phys. Lett. B. **682**, 177 (2009).

[67] J.L. Raaf, PhD Thesis, Fermilab (2005).

[68] E. Isiksal *et al*, Phys. Rev. Lett. **52**, 1096 (1984).

[69] CHARM Collaboration, F. Bergsma *et al*, Phys. Lett. B. **157**, 469 (1985).

[70] SKAT Collaboration, H.J. Grabosch *et al*, Z. Phys. C. **31**, 203 (1986).

[71] C. Baltay *et al*, Phys. Rev. Lett. **57**, 2629 (1986).

[72] MINOS Collaboration, P. Adamson *et al*, Phys. Rev. D. **94**, 072006 (2016).

[73] Y. Hayato, Acta. Phys. Polon. B. **40**, 2477 (2009).

[74] D. Rein and L. M. Sehgal, Ann. Phys., **133**, 79 (1981).

[75] C. Berger and L. M. Sehgal, Phys. Rev. D., **79**, 053003, (2009).

[76] S. L. Adler, Phys. Rev., **135**, B 963, (1964).

[77] T2K Collaboration, K. Abe *et al.*, Phys. Rev. D. **91**, 112010 (2015).

[78] T2K Collaboration, K. Abe *et al.*, Phys. Rev. D. **91**, 112002 (2015).

[79] T2K Collaboration, K. Abe *et al.*, Phys. Rev. D. **92**, 112003 (2015).

[80] T2K Collaboration, K. Abe *et al.*, Phys. Rev. Lett. **113**, 241803 (2014).

[81] T2K Collaboration, K. Abe *et al.*, Phys. Rev. D. **90**, 072012 (2014).

[82] T2K Collaboration, K. Abe *et al.*, Nuclear Instruments and Methods in Physics Research A, **659** 106 (2011).

[83] Official plots produced by the T2K collaboration.

[84] T2K Collaboration, K. Abe *et al*, Phys. Rev. D. **88**, 032002 (2013).

[85] D. Allan *et al*, Journal of Instrumentation, **8**, P10019 (2013).

[86] T2K Collaboration, K. Abe *et al*, Nuclear Instruments and Methods in Physics Research A, **694** 211 (2012).

[87] T2K Collaboration, K. Abe *et al*, Phys. Rev. D. **93**, 072002 (2016).

[88] T2K Collaboration, K. Abe *et al.*, Phys. Rev. D. **89**, 092003 (2014).

[89] S. Assylbekov *et al.*, Nuclear Instruments and Methods in Physics Research A, **686** 48 (2012).

[90] N. Abgrall *et al.*, Nuclear Instruments and Methods in Physics Research A, **637** 25 (2011).

[91] P.-A. Amaudruz *et al.*, Nuclear Instruments and Methods in Physics Research A, **696** 1 (2012).

[92] D. Allan *et al.*, Journal of Instrumentation, **8** P10019 (2013).

[93] S. Aoki *et al.*, Nuclear Instruments and Methods in Physics Research A, **689** 135 (2013).

[94] The T2K Collaboration, K. Abe *et al*, Nuclear Instruments and Methods in Physics Research Section A, **659**, 106 (2011).

[95] S. Ritt, O. Amaudruz and K. Olchanski, http://midas.psi.ch.

[96] S. Agostinelli *et al*, Nuclear Instruments and Methods in Physics Research Section A, **506**, 250 (2003).

[97] R. Brun and F. Rademakers, Nuclear Instruments and Methods in Physics Research Section A, **389**, 81 (1997).

[98] Phys. Rev. D, **87**, 012001 (2013).

[99] R. Brun, F. Carminati and S. Giani, CERN, **W5013**, (1994).

[100] A. Fasso, A. Ferrari, J. Ranft, P.R. Sala, Conf. Proc. **C9309194**, 493 (1993).

[101] C. Andreopoulos *et al*, Nuclear Instruments and Methods in Physics Research Section A, **614**, 87 (2010).

[102] S. Assylbekov *et al*, Nuclear Instruments and Methods in Physics Research Section A, **686**, 48 (2012).

[103] S. Aoki *et al*, Nuclear Instruments and Methods in Physics Research Section A, **698**, 135 (2013).

[104] P.A. Amaudruz *et al*, Nuclear Instruments and Methods in Physics Research Section A, **696**, 1 (2012).

[105] N. Abgrall *et al*, Nuclear Instruments and Methods in Physics Research Section A, **637**, 25 (2011).

[106] The ND280 Calibration Group, T2K Internal Technote (T2K-TN-37), (2012).

[107] P.T. Le, PhD Thesis, Stony Brook University (2009).

[108] A. Cerveira-Villanueva, J.J. Gomez-Cardenas and J.A. Hernando, Nuclear Instruments and Methods in Physics Research Section A, **534**, 180 (2004).

[109] A. Hillairet *et al.*, T2K internal technote (T2K-TN-072), (2011).

[110] G. Barker *et al.*, T2K internal technote (T2K-TN-111), (2012).

[111] Official Plots Provided by the T2K Collaboration.

[112] I. Jolliffe, Principal Component Analysis. Wiley StatsRef: Statistics Reference Online (2014).

[113] M. Lawe, PhD Thesis, The University of Sheffield (2015).

[114] S. Short, PhD Thesis, Imperial College London (2013).

[115] A. Hoecker *et al*, Proceedings, 11th International Workshop on Advanced Computing and Analysis Techniques in Physics Research, (2007).

[116] M. Scott, PhD Thesis, Imperial College London (2013).

[117] P. Martins, T2K internal technote (T2K-TN-270), 2016.

[118] A. Fiorentini, T2K internal technote (T2K-TN-234), 2015.

[119] P. Bartet *et al*, T2K internal technote (T2K-TN-212), 2015.

[120] S. Bordoni *et al*, T2K internal technote (T2K-TN-221), 2015.

[121] F. Sanchez *et al*, T2K internal technote (T2K-TN-229), 2016.

[122] C. Andreooulos *et al*, T2K internal technote (T2K-TN-007), 2012.

[123] A. Bercellie *et al*, T2K internal technote (T2K-TN-193), 2015.

[124] C. H. Llewellyn Smith, Phys. Rep. C., **3**, 261 (1972).

[125] R. A. Smith and E. J. Moniz, Nucl. Phys. B., **43**, 605 (1972).

[126] R. A. Smith and E. J. Moniz, Nucl. Phys. B., **101**, 547 (1975).

[127] D. Rein, Z. Phys. C., **35**, 43 (1987).

[128] E. J. Moniz *et al*, Phys. Rev. Lett. 26., **8**, 445 (1971).

[129] A. M. Ankowski and J. T. Sobczyk, Phys. Rev. C., **77**, 044311 (2008).

[130] A. A. Aguilar-Arevalo *et al*, Phys. Rev. D., **81**, 092005 (2010).

[131] V. Lyubushkin *et al*, Eur. Phys. J., **C63**, 355 (2009).

[132] P. Adamson *et al*, Phys. Rev. D., **81**, 072002 (2010).

[133] P. de Perio *et al.*, T2K internal technote (T2K-TN-033), (2012).

[134] P. de Perio *et al.*, T2K internal technote (T2K-TN-032), (2011).

[135] P. de Perio *et al.*, T2K internal technote (T2K-TN-108), (2012).

[136] N. Abgrall *et al.*, Phys. Rev. C., **84**, 034604, (2011).

[137] N. Abgrall *et al.*, Phys. Rev. C., **85**, 035210, (2012).

[138] T. Eichten *et al.*, Nucl. Phys. B., **44**, 333, (1972).

[139] J. V. Allaby *et al.*, Technical Report 70-12, CERN, (1970).

[140] G. D'Agostini, Nuclear Instruments and Methods in Physics Research A, **362**, 487, (1995).

[141] D. Rein and L. M. Sehgal, Nucl. Phys. B., **223**, 29, (1983).

[142] P. Stowell and C. Wret, "Private communication." (2016).

Appendix A

BDT supplementary information

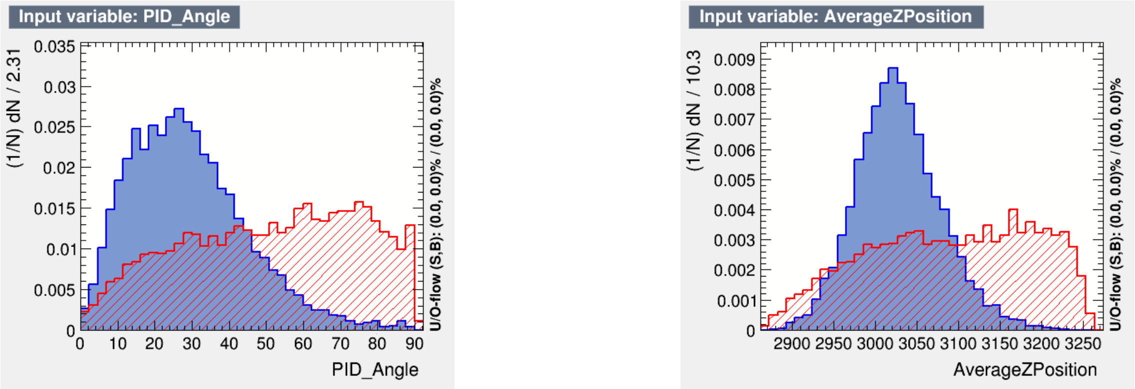

(a) PID Angle (1).

(b) Average Z Position (2).

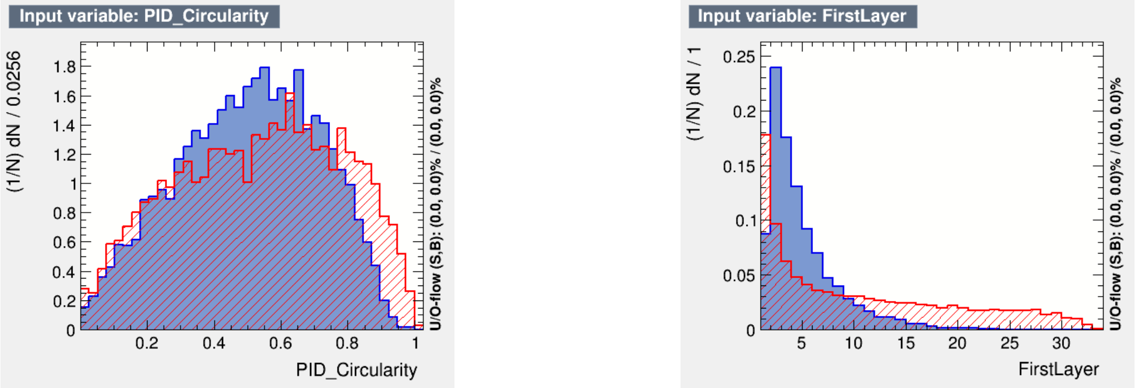

(c) PID Circularity (3).

(d) First Layer (4).

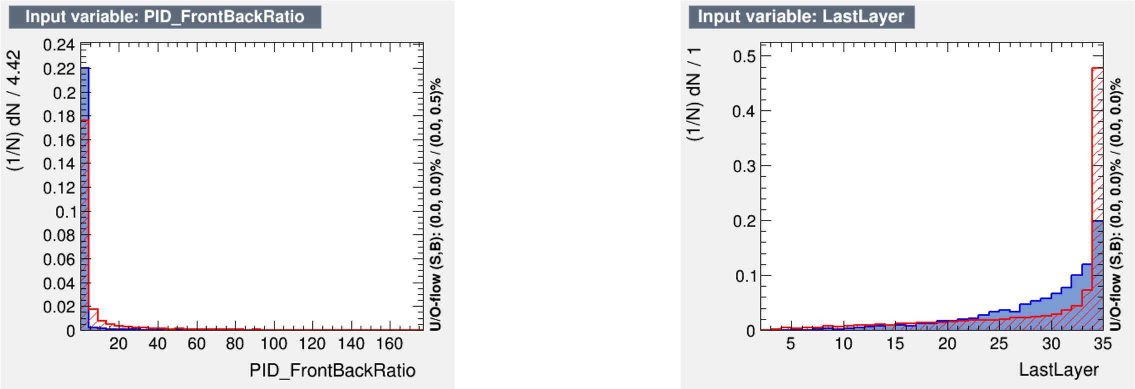

(e) PID Front Back Ratio (5).

(f) Last Layer (6).

Figure A1: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the DsDs topology.

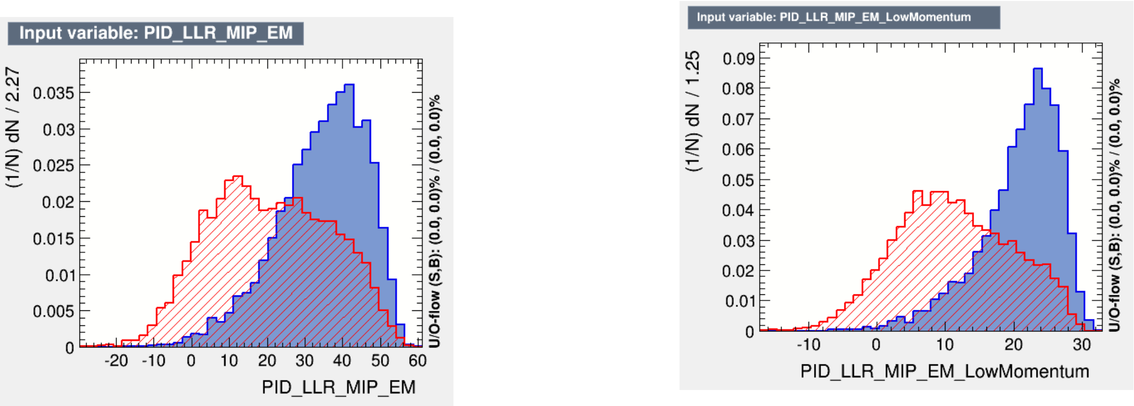

(g) PID LLR MIP EM (7).

(h) PID LLR MIP EM Low Momentum (8).

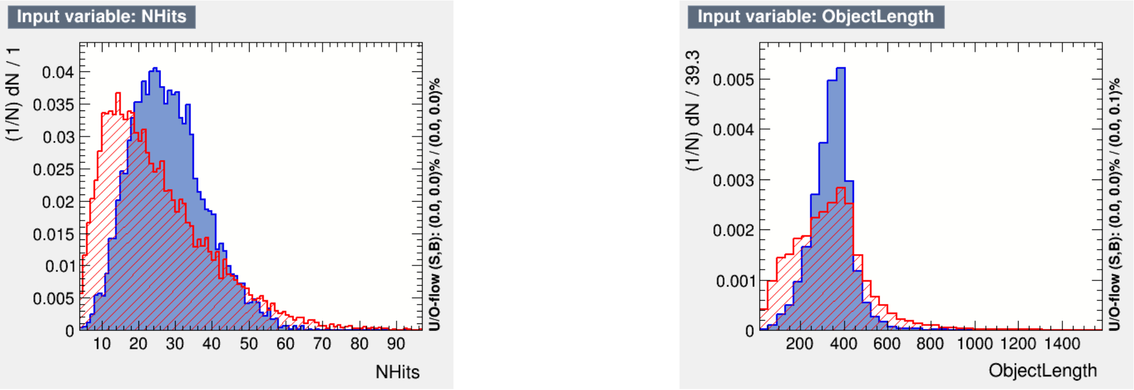

(i) NHits (9).

(j) Object Length (10).

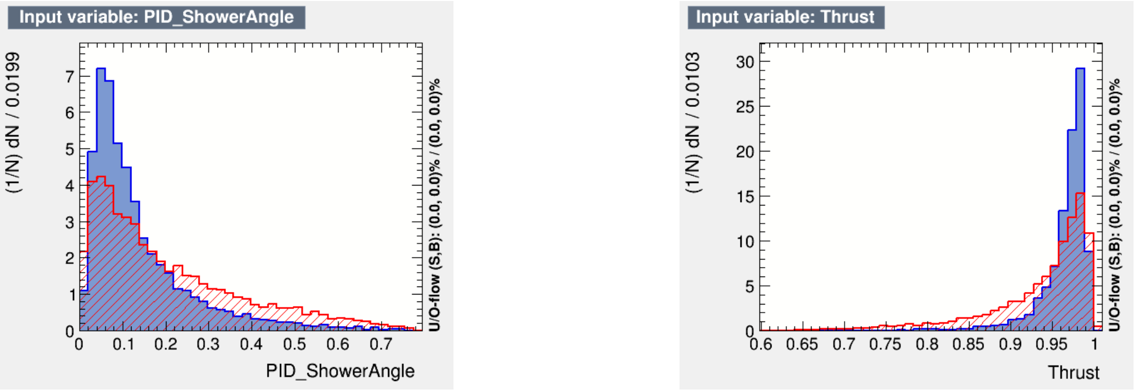

(k) PID Shower Angle (11).

(l) Thrust (12).

Figure A1: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the DsDs topology.

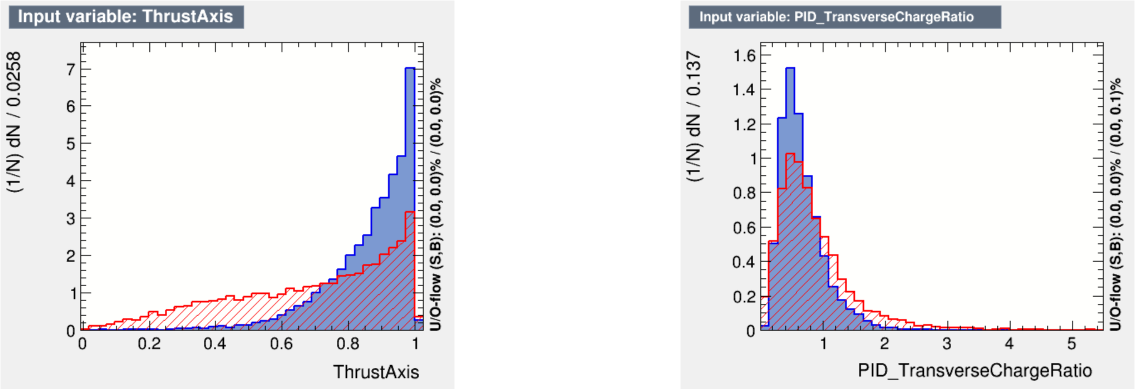

(m) Thrust Axis (13).

(n) PID Transverse Charge Ratio (14).

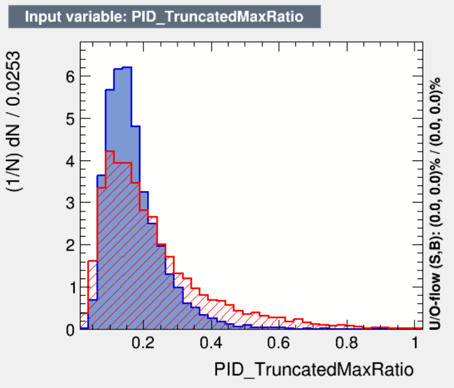

(o) PID Truncated Max Ratio (15).

Figure A1: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the DsDs topology.

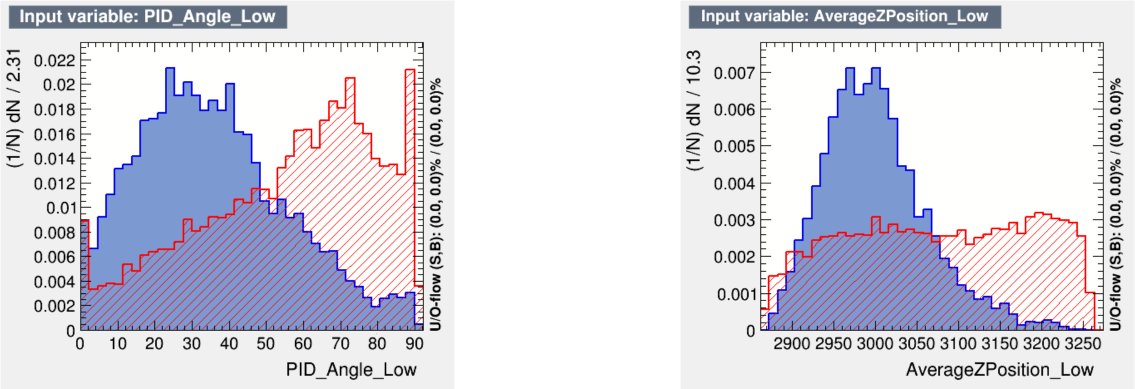

(a) PID Angle Low (16).

(b) Average Z Position Low (17).

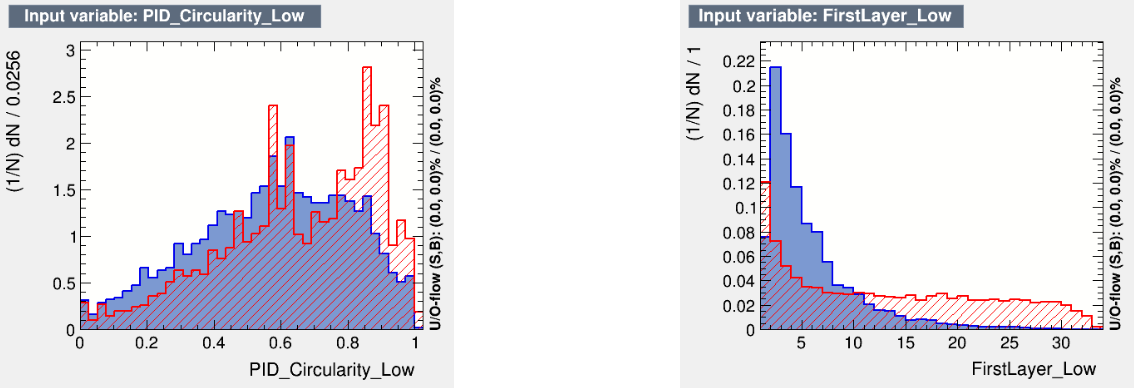

(c) PID Circularity Low (18).

(d) First Layer Low (19).

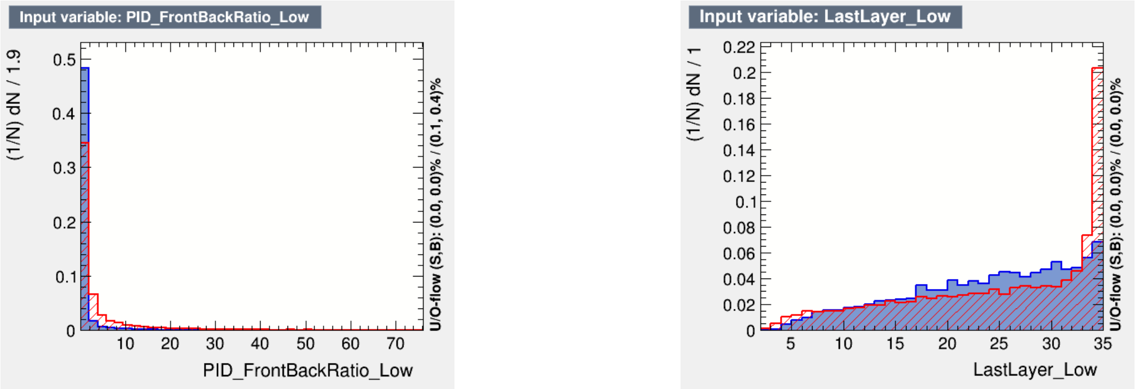

(e) PID Front Back Ratio Low (20).

(f) Last Layer Low (21).

Figure A2: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the DsDs topology.

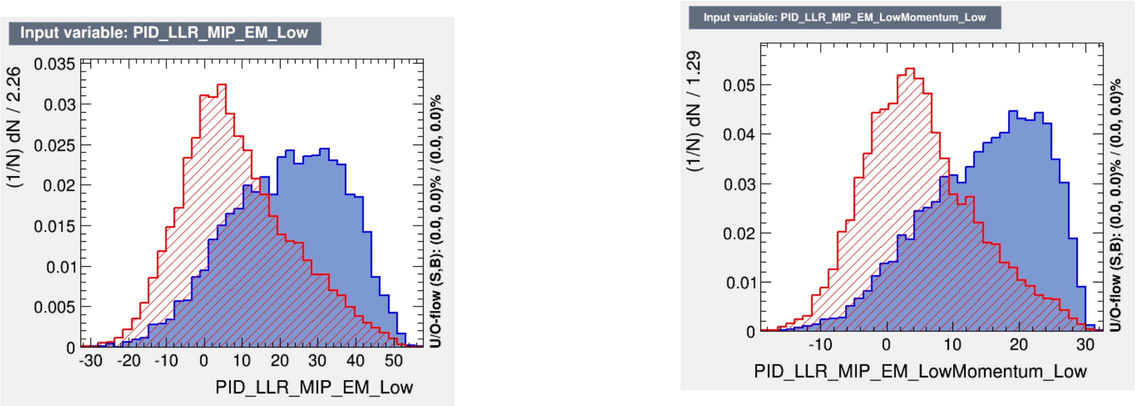

(g) PID LLR MIP EM Low (22).

(h) PID LLR MIP EM Low Momentum Low (23).

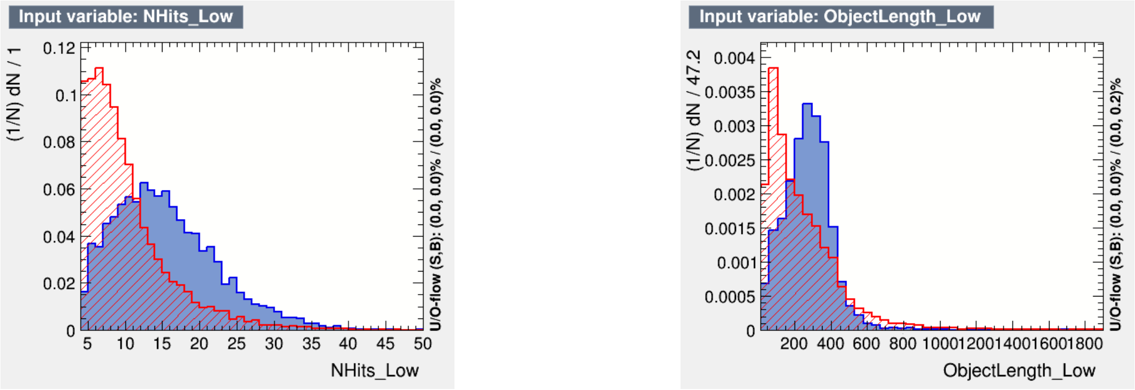

(i) NHits Low (24).

(j) Object Length Low (25).

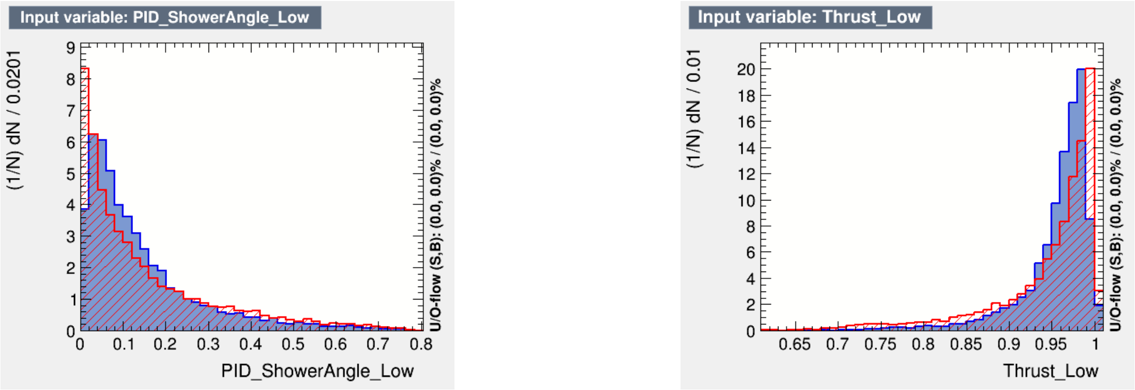

(k) PID Shower Angle Low (26).

(l) Thrust Low (27).

Figure A2: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the DsDs topology.

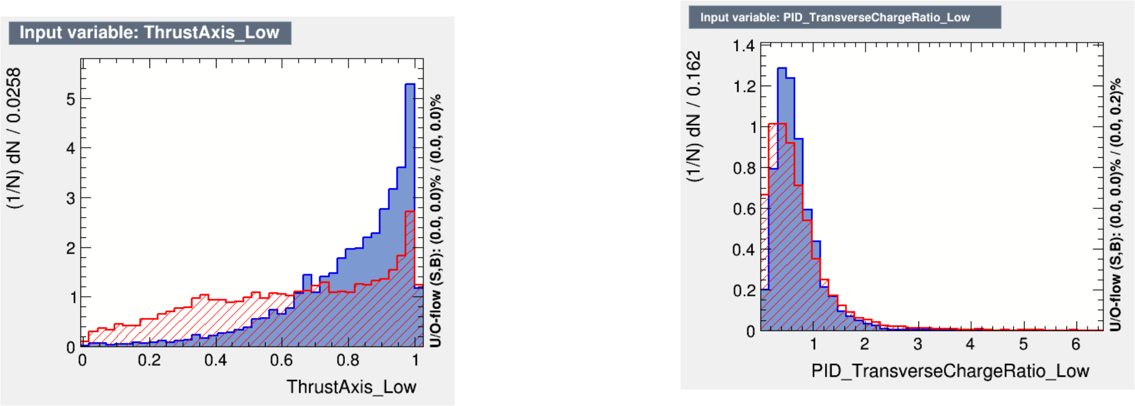

(m) Thrust Axis Low (28).

(n) PID Transverse Charge Ratio Low (29).

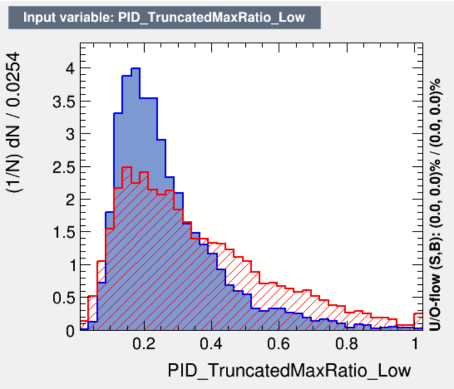

(o) PID Truncated Max Ratio Low (30).

Figure A2: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the DsDs topology.

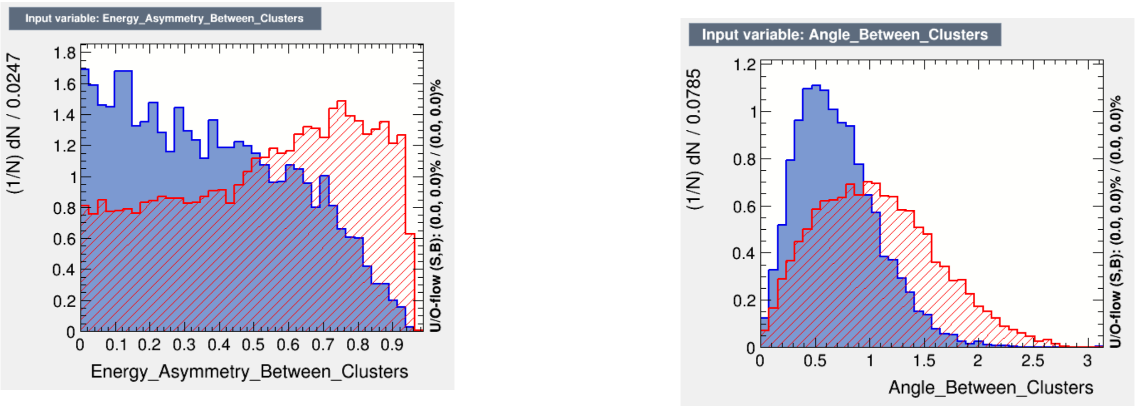

(a) Energy Asymmetry Between Clusters (31).

(b) Angle Between Clusters (32).

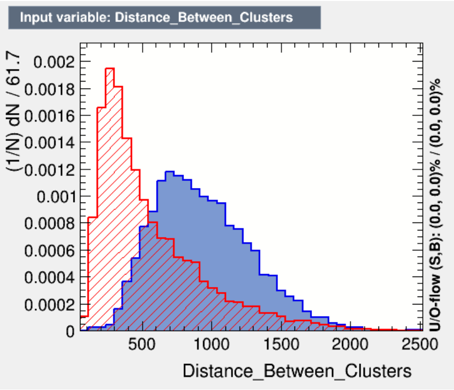

(c) Distance Between Clusters (33).

Figure A3: The signal (blue) and background (red) combined variable distributions of both ECal clusters used to train the BDT for the DsDs topology.

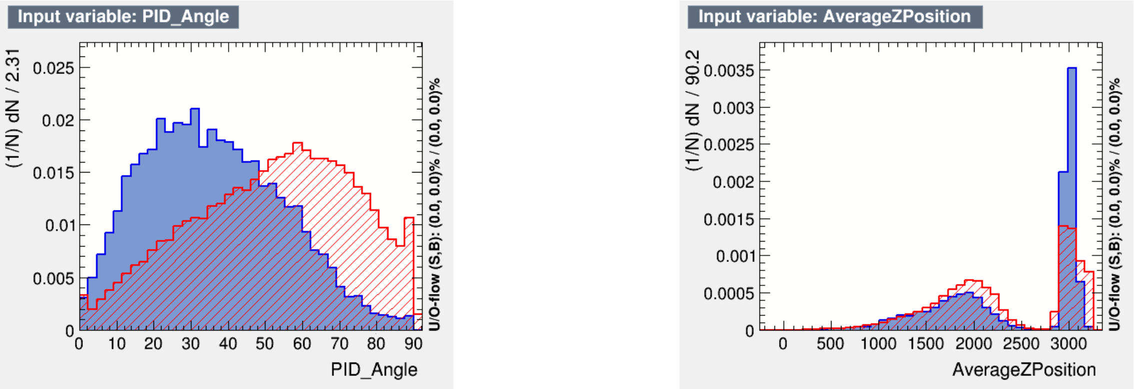

(a) PID Angle (1).

(b) Average Z Position (2).

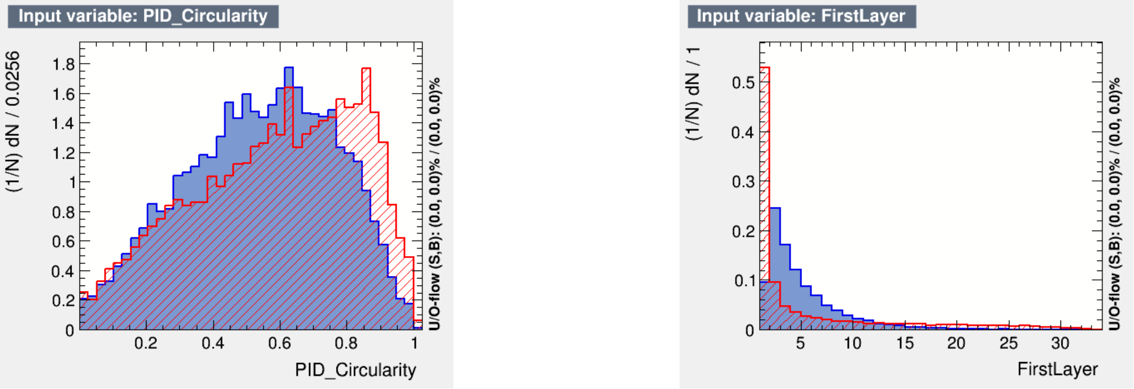

(c) PID Circularity (3).

(d) First Layer (4).

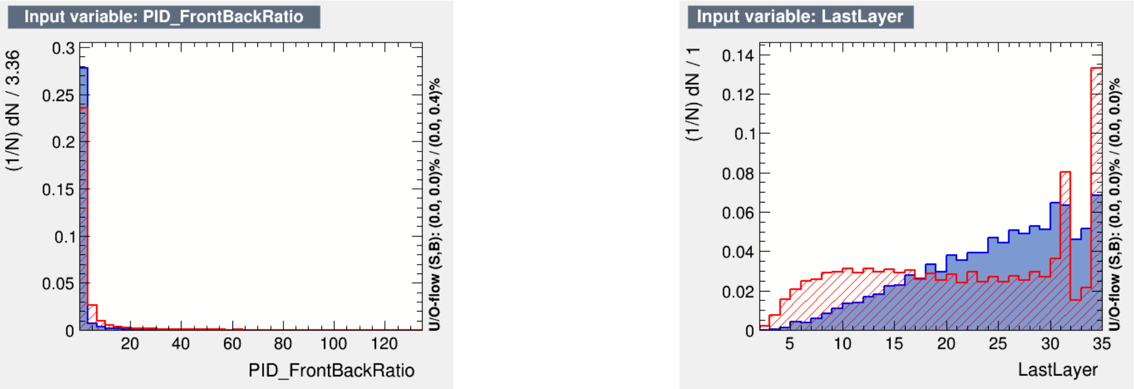

(e) PID Front Back Ratio (5).

(f) Last Layer (6).

Figure A4: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the BrlDs topology.

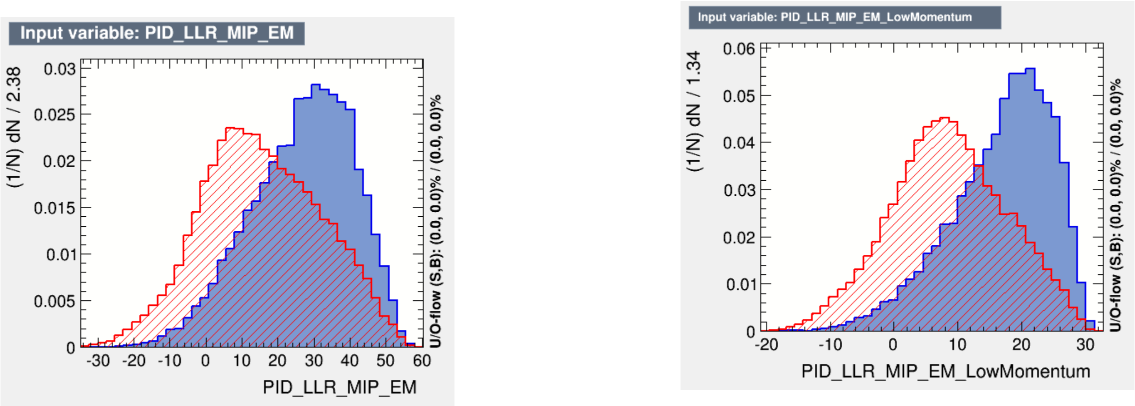

(g) PID LLR MIP EM (7).

(h) PID LLR MIP EM Low Momentum (8).

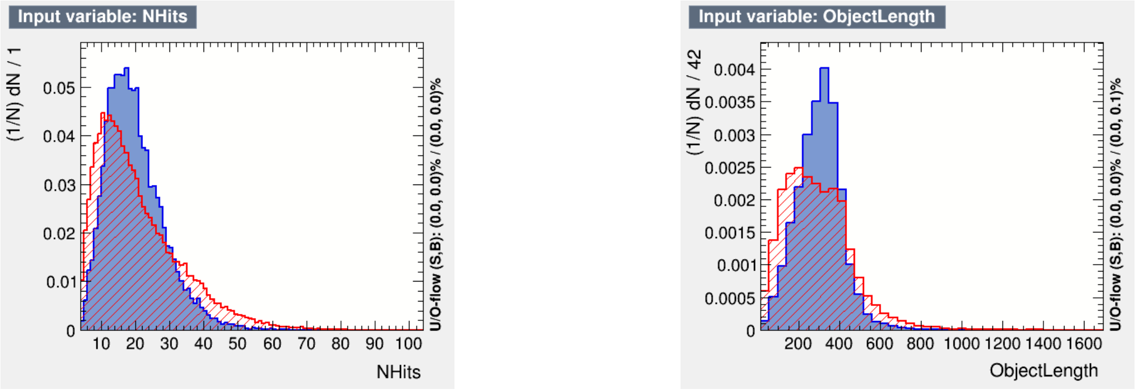

(i) NHits (9).

(j) Object Length (10).

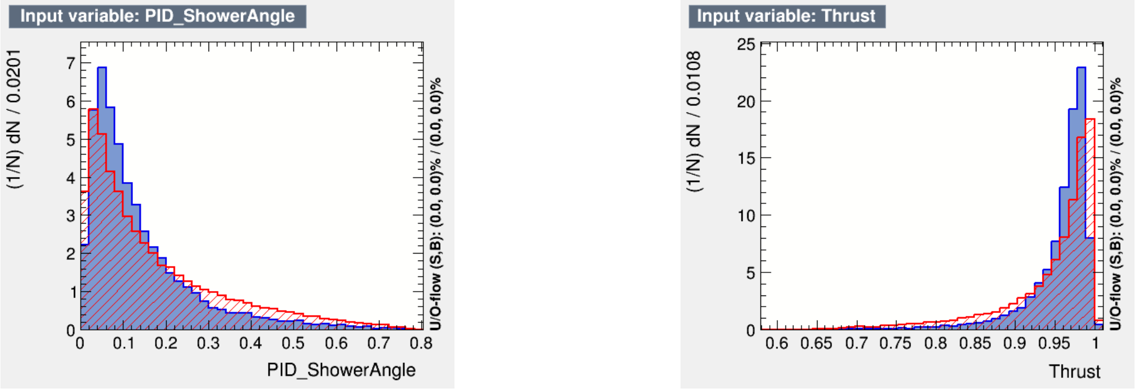

(k) PID Shower Angle (11).

(l) Thrust (12).

Figure A4: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the BrlDs topology.

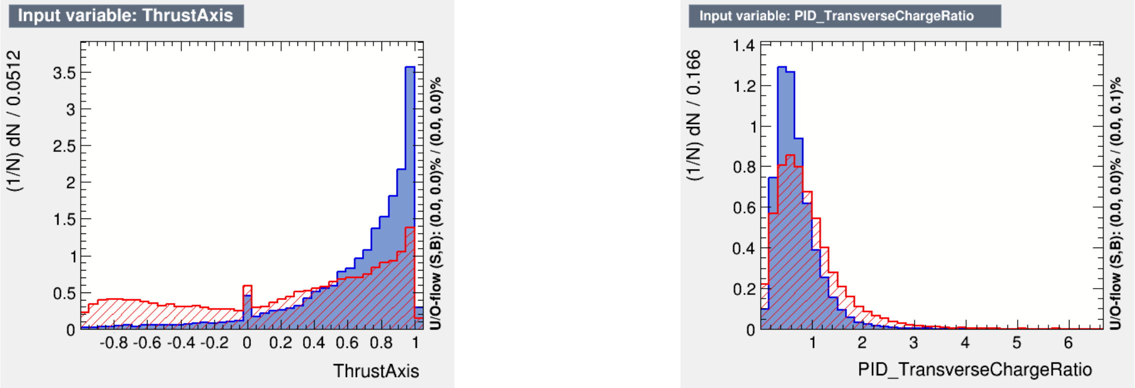

(m) Thrust Axis (13).

(n) PID Transverse Charge Ratio (14).

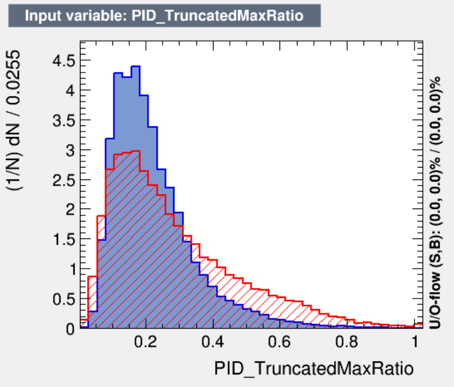

(o) PID Truncated Max Ratio (15).

Figure A4: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the BrlDs topology.

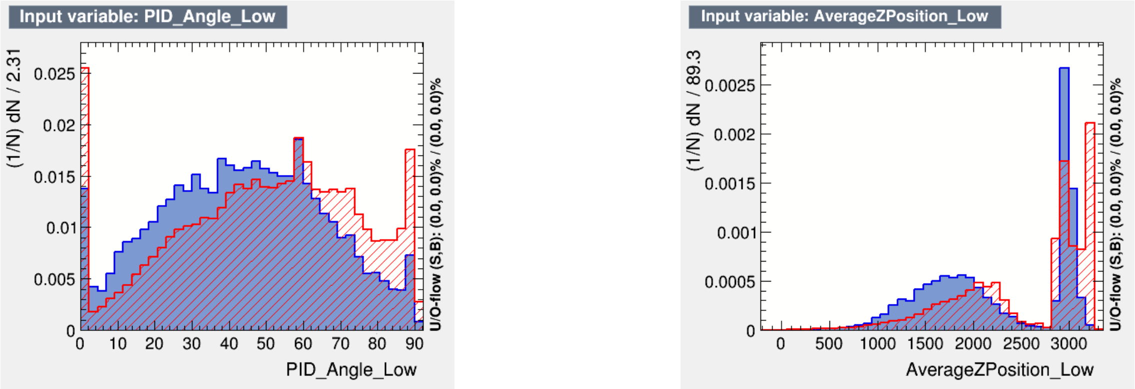

(a) PID Angle Low (16). (b) Average Z Position Low (17).

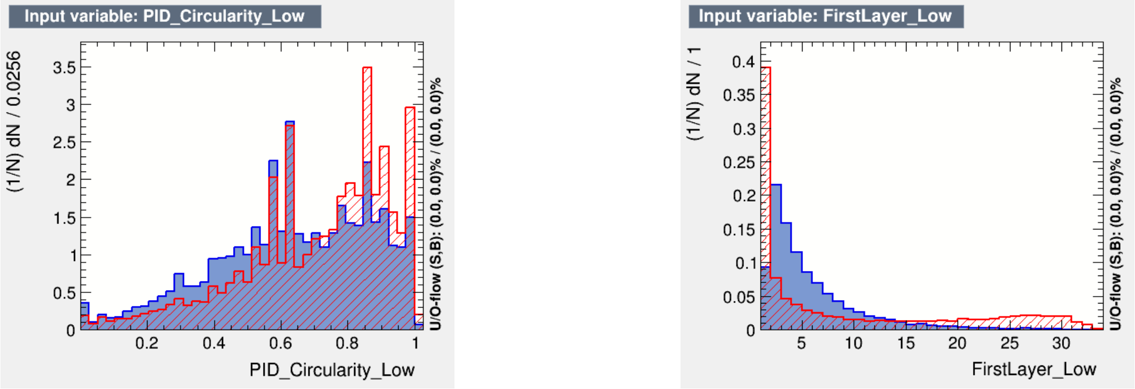

(c) PID Circularity Low (18). (d) First Layer Low (19).

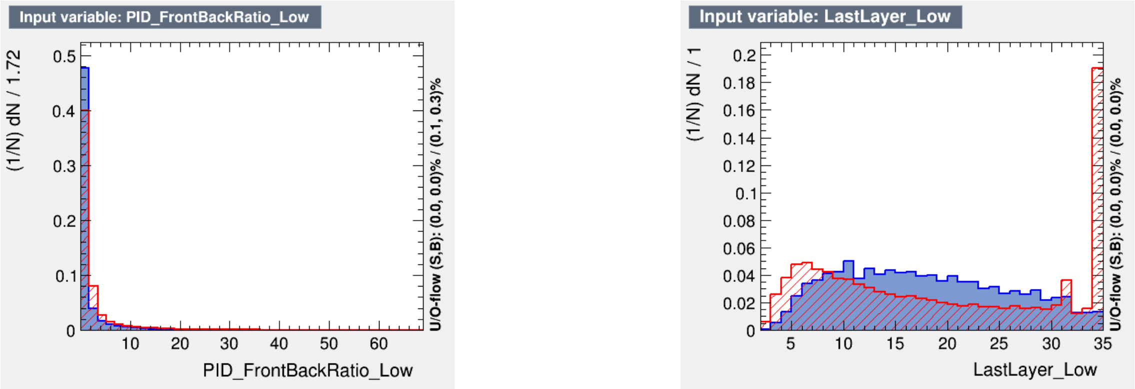

(e) PID Front Back Ratio Low (20). (f) Last Layer Low (21).

Figure A5: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the BrlDs topology.

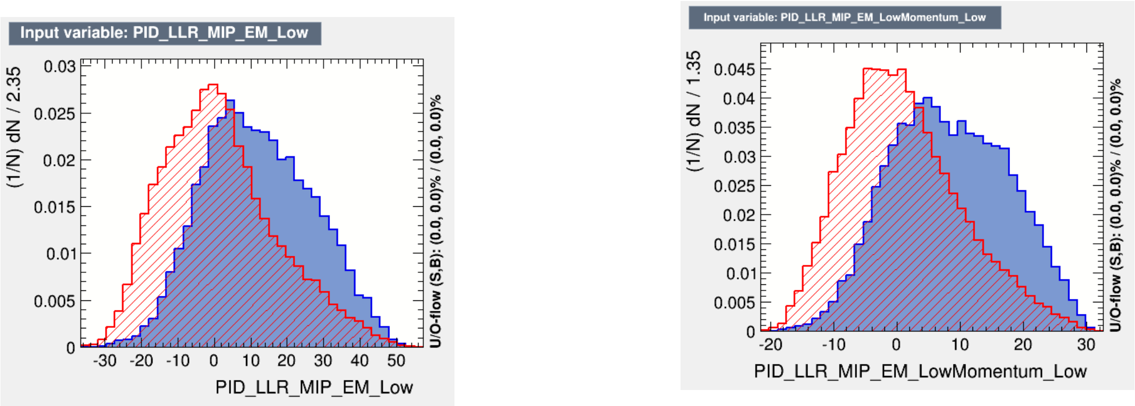

(g) PID LLR MIP EM Low (22).

(h) PID LLR MIP EM Low Momentum Low (23).

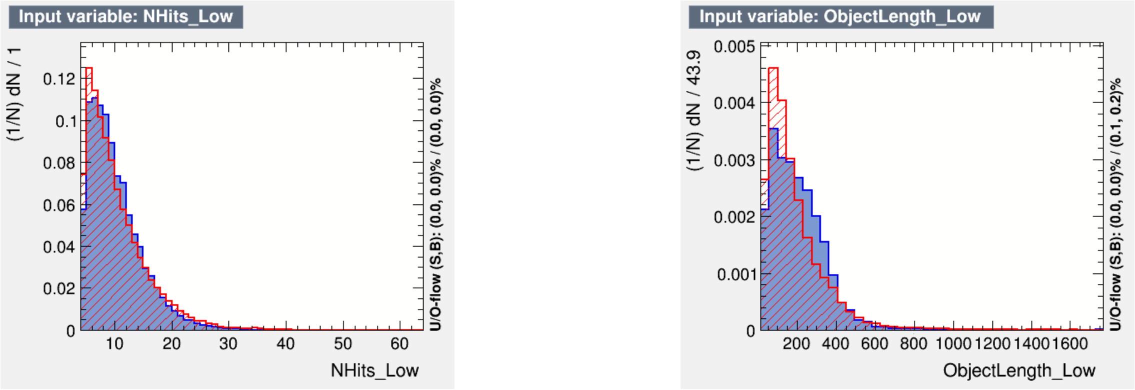

(i) NHits Low (24).

(j) Object Length Low (25).

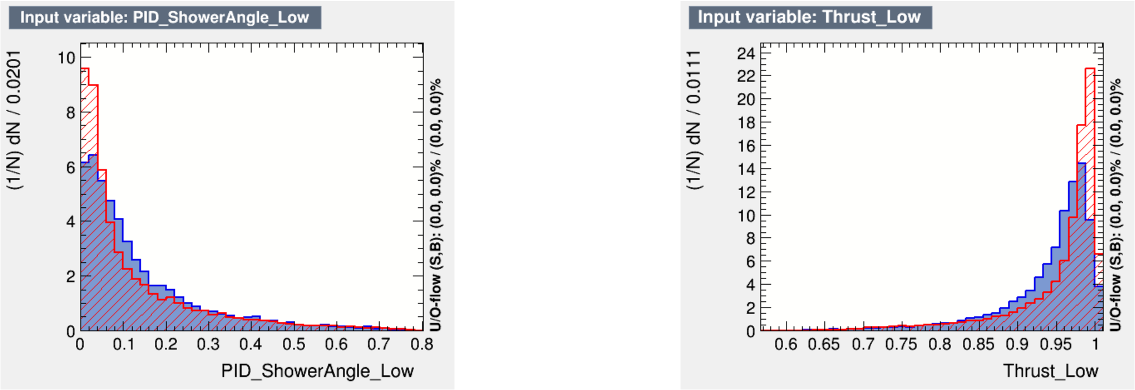

(k) PID Shower Angle Low (26).

(l) Thrust Low (27).

Figure A5: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the BrlDs topology.

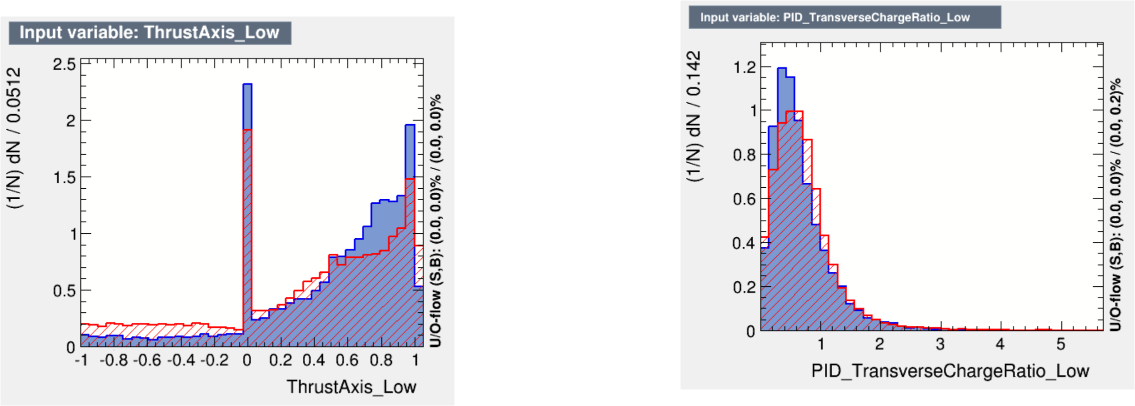

(m) Thrust Axis Low (28).

(n) PID Transverse Charge Ratio Low (29).

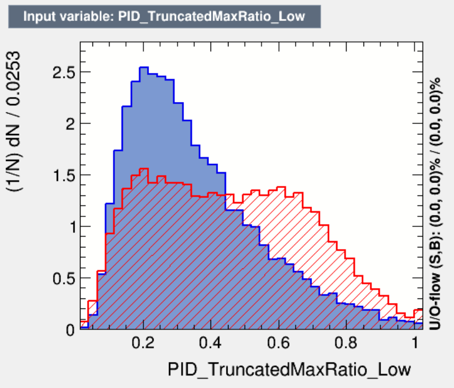

(o) PID Truncated Max Ratio Low (30).

Figure A5: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the BrlDs topology.

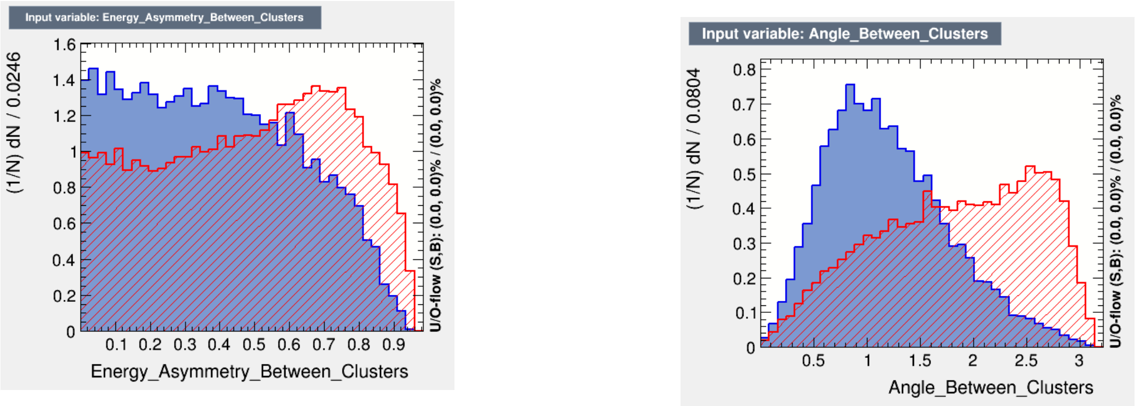

(a) Energy Asymmetry Between Clusters (31).

(b) Angle Between Clusters (32).

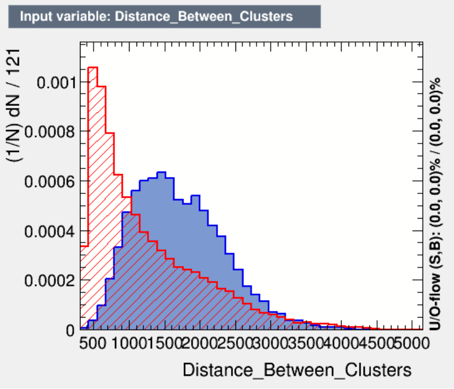

(c) Distance Between Clusters (33).

Figure A6: The signal (blue) and background (red) combined variable distributions of both ECal clusters used to train the BDT for the BrlDs topology.

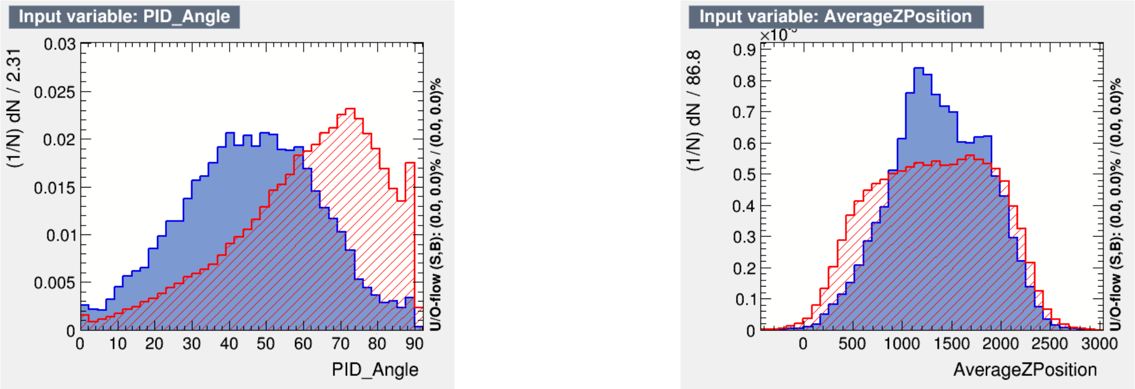

(a) PID Angle (1).

(b) Average Z Position (2).

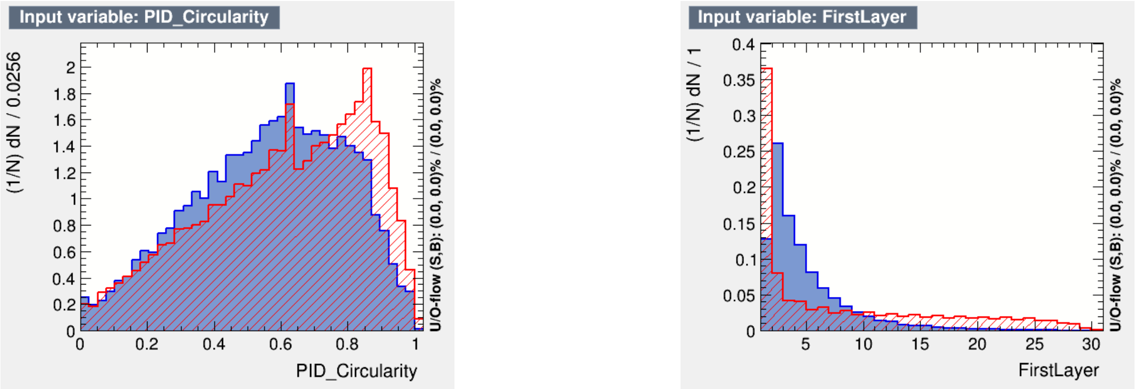

(c) PID Circularity (3).

(d) First Layer (4).

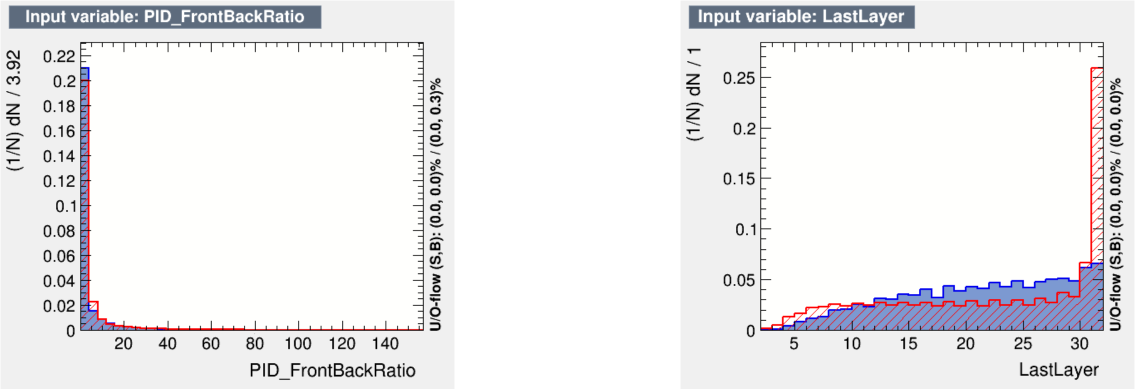

(e) PID Front Back Ratio (5).

(f) Last Layer (6).

Figure A7: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the BrlBrl topology.

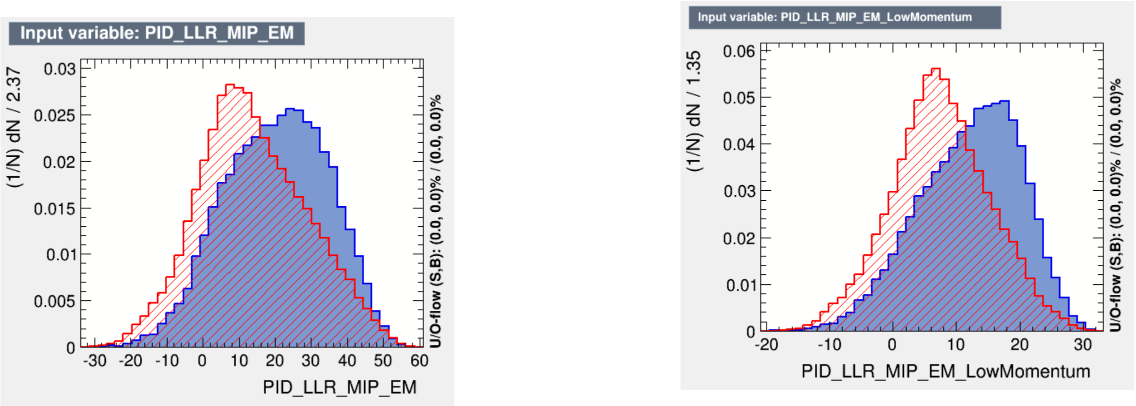

(g) PID LLR MIP EM (7).

(h) PID LLR MIP EM Low Momentum (8).

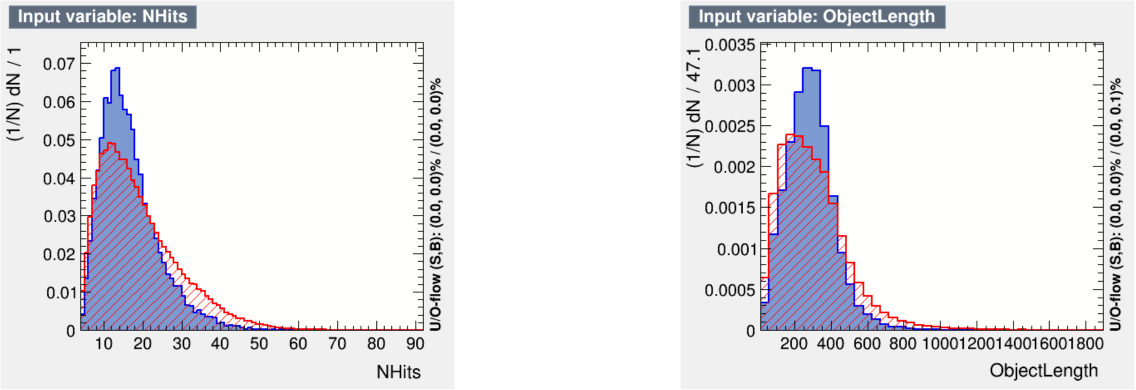

(i) NHits (9).

(j) Object Length (10).

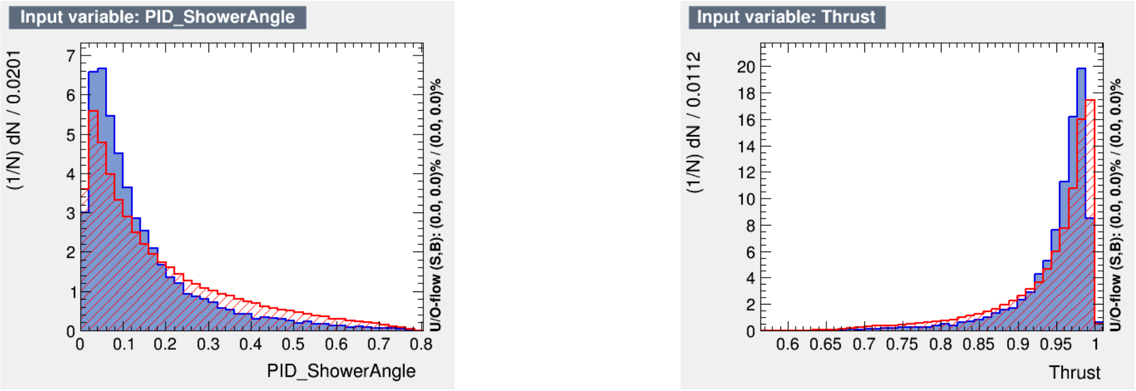

(k) PID Shower Angle (11).

(l) Thrust (12).

Figure A7: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the BrlBrl topology.

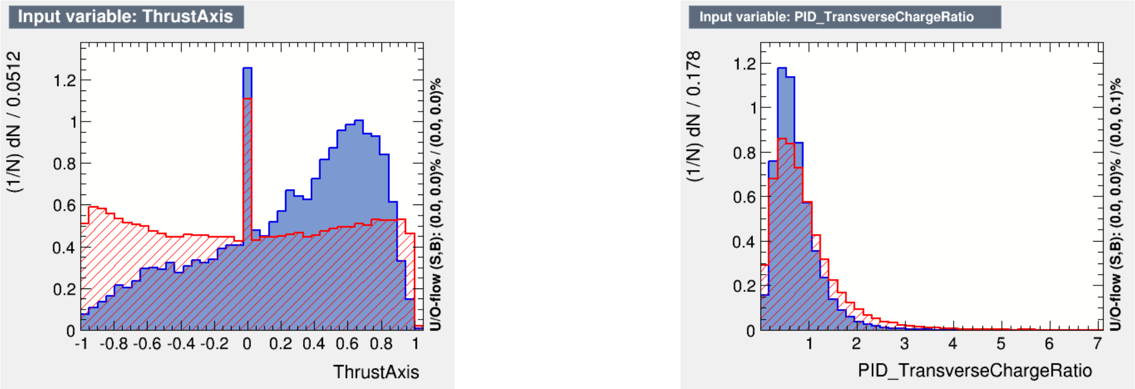

(m) Thrust Axis (13).

(n) PID Transverse Charge Ratio (14).

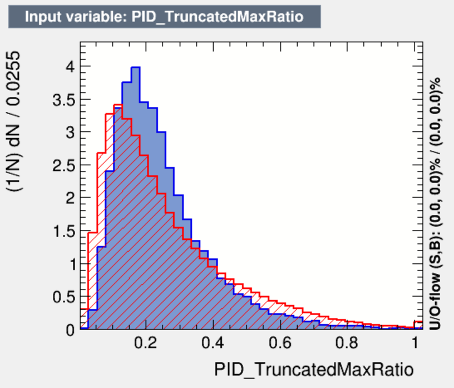

(o) PID Truncated Max Ratio (15).

Figure A7: The signal (blue) and background (red) variable distributions of the high energy ECal cluster used to train the BDT for the BrlBrl topology.

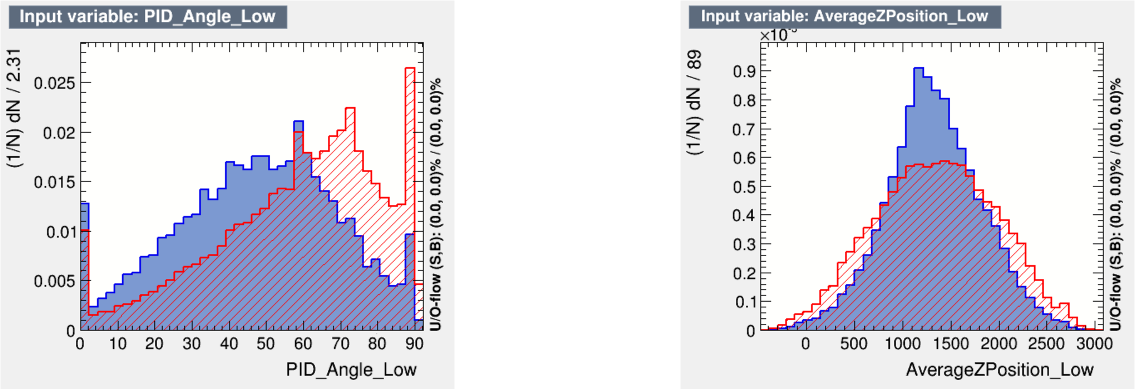

(a) PID Angle Low (16).

(b) Average Z Position Low (17).

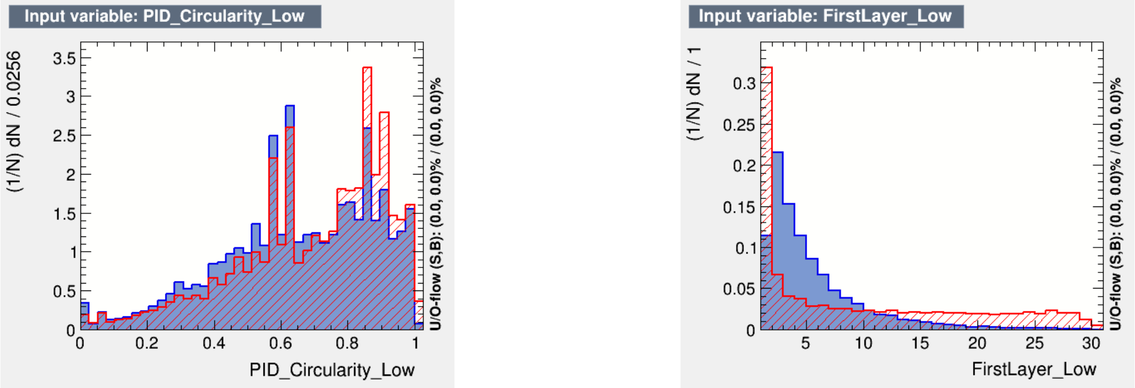

(c) PID Circularity Low (18).

(d) First Layer Low (19).

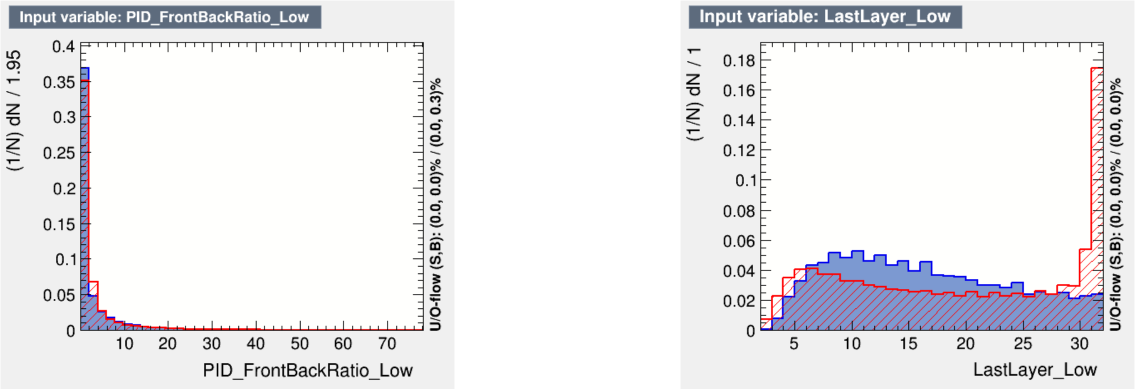

(e) PID Front Back Ratio Low (20).

(f) Last Layer Low (21).

Figure A8: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the BrlBrl topology.

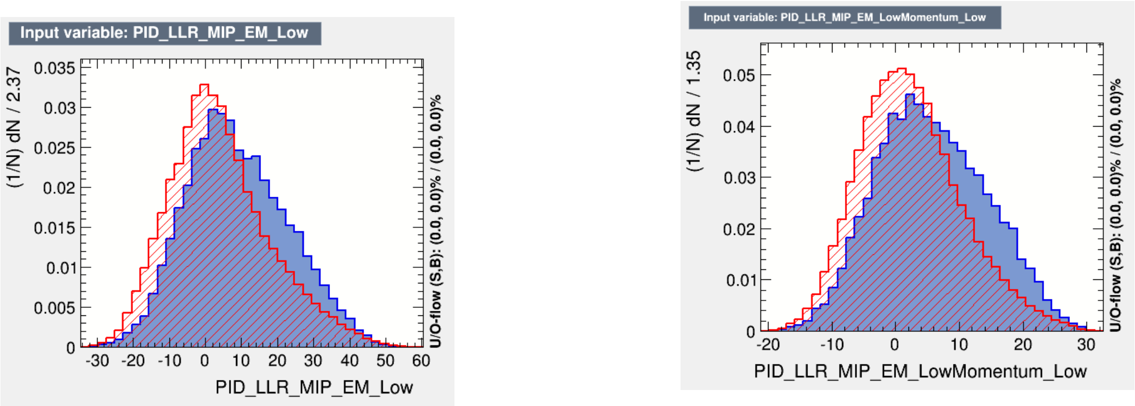

(g) PID LLR MIP EM Low (22).

(h) PID LLR MIP EM Low Momentum Low (23).

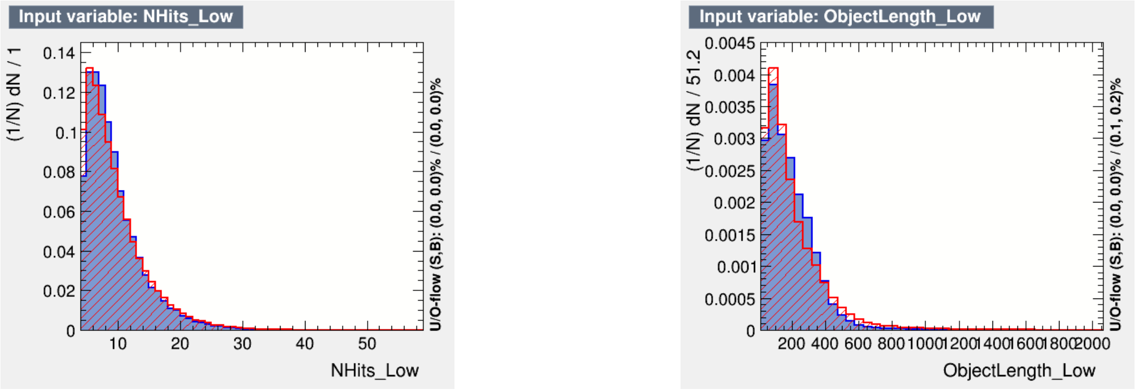

(i) NHits Low (24).

(j) Object Length Low (25).

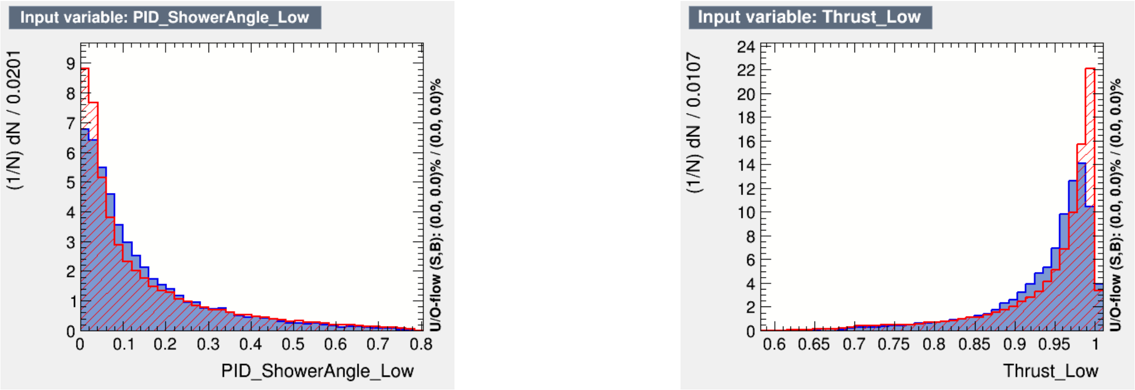

(k) PID Shower Angle Low (26).

(l) Thrust Low (27).

Figure A8: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the BrlBrl topology.

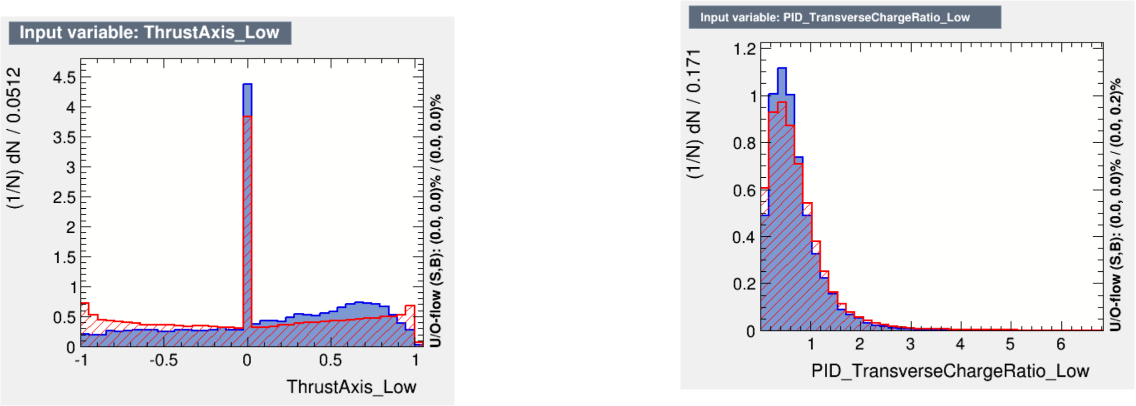

(m) Thrust Axis Low (28).

(n) PID Transverse Charge Ratio Low (29).

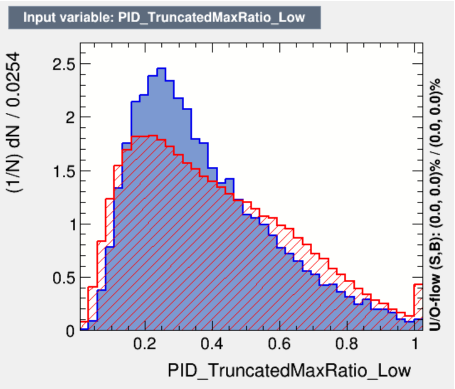

(o) PID Truncated Max Ratio Low (30).

Figure A8: The signal (blue) and background (red) variable distributions of the low energy ECal cluster used to train the BDT for the BrlBrl topology.

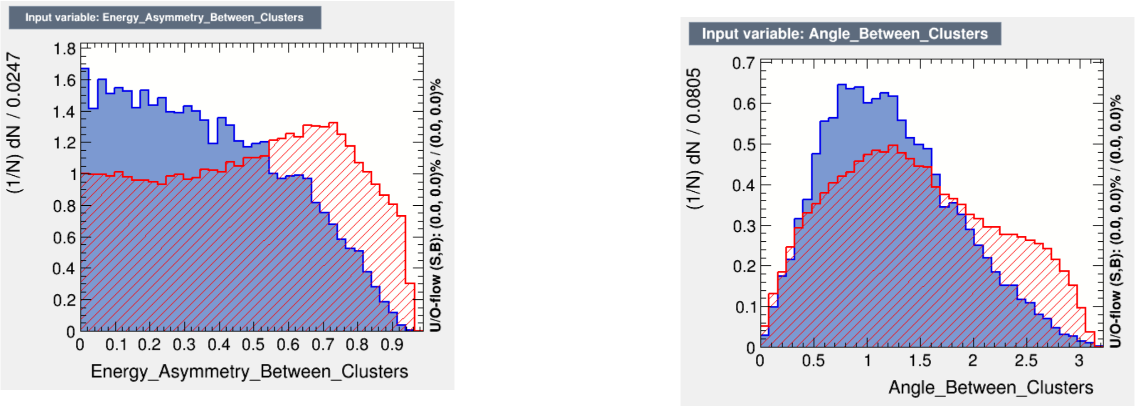

(a) Energy Asymmetry Between Clusters (31).

(b) Angle Between Clusters (32).

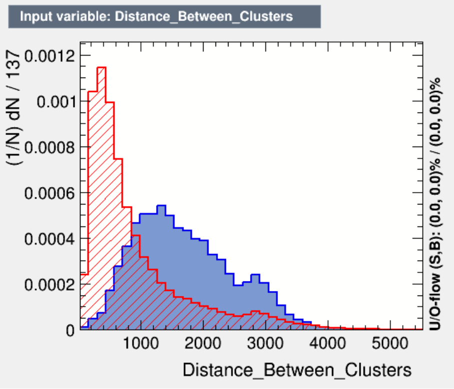

(c) Distance Between Clusters (33).

Figure A9: The signal (blue) and background (red) combined variable distributions of both ECal clusters used to train the BDT for the BrlBrl topology.

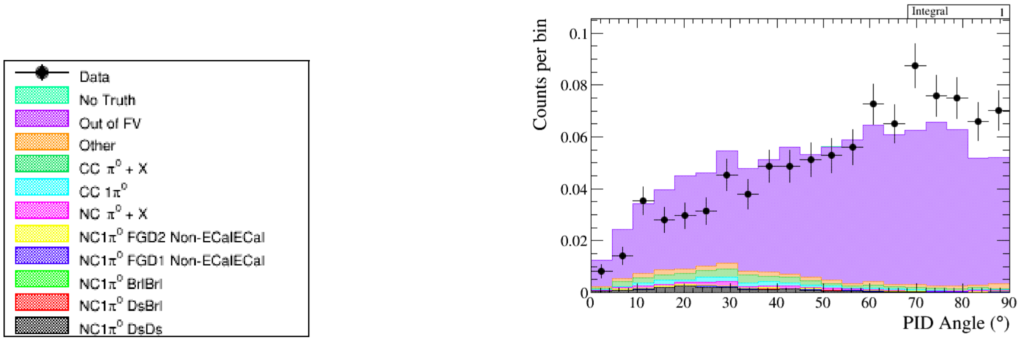

(a) PID Angle (1).

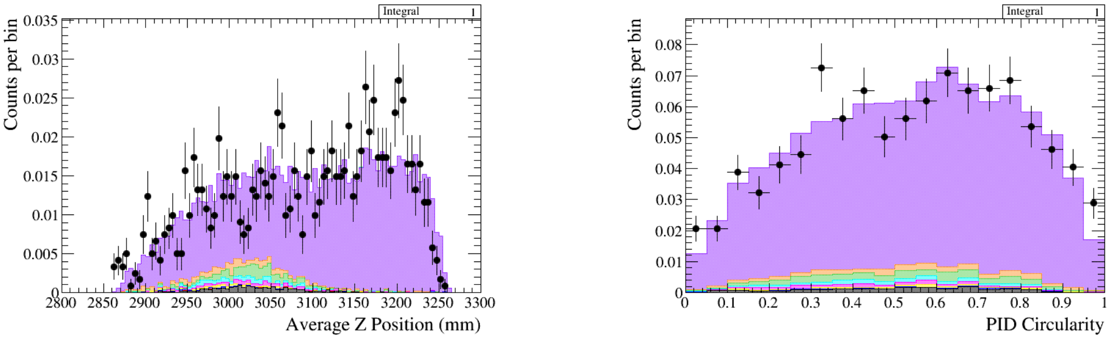

(b) Average Z Position (2).

(c) PID Circularity (3).

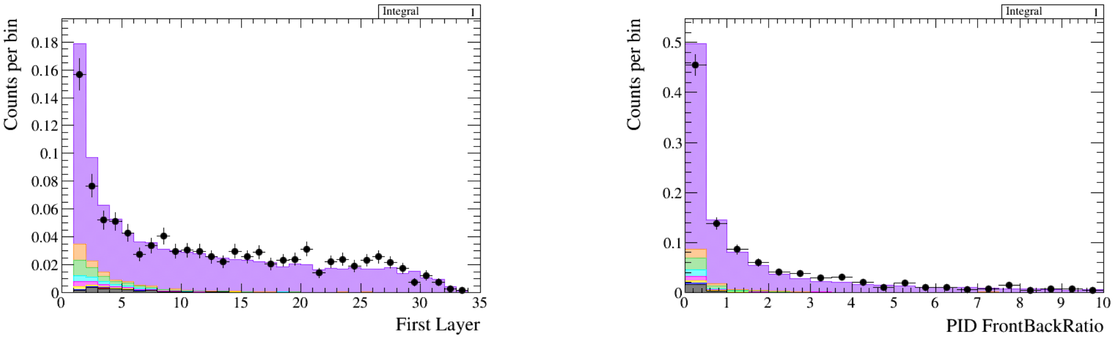

(d) First Layer (4).

(e) PID Front Back Ratio (5).

Figure A10: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the DsDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

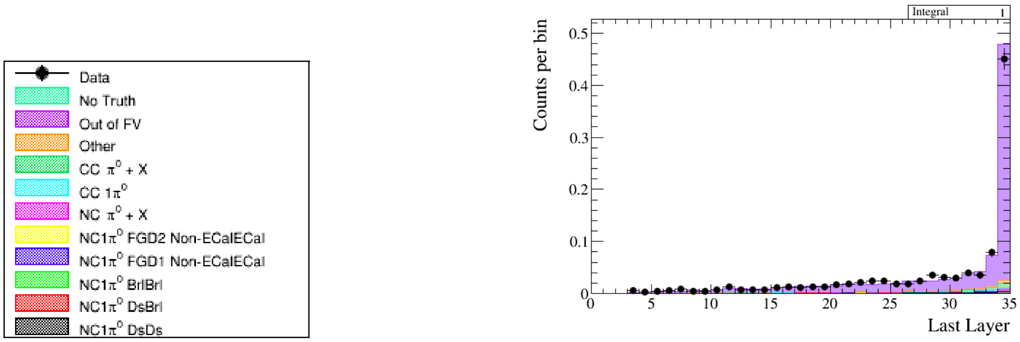

(f) Last Layer (6).

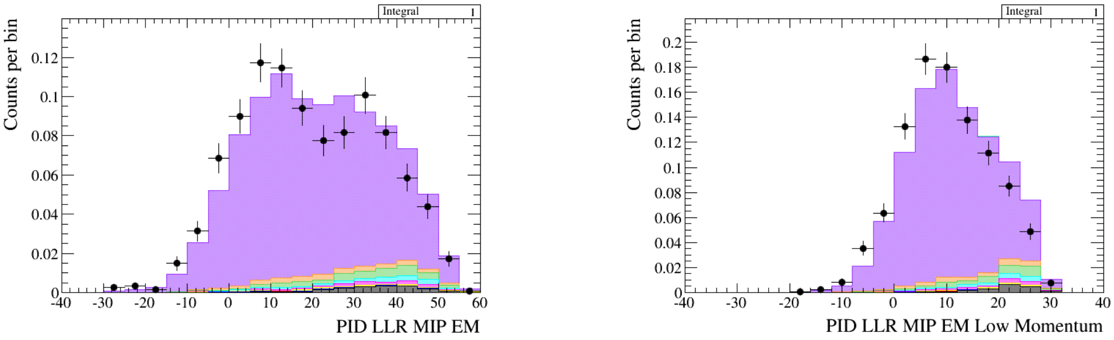

(g) PID LLR MIP EM (7). (h) PID LLR MIP EM Low Momentum (8).

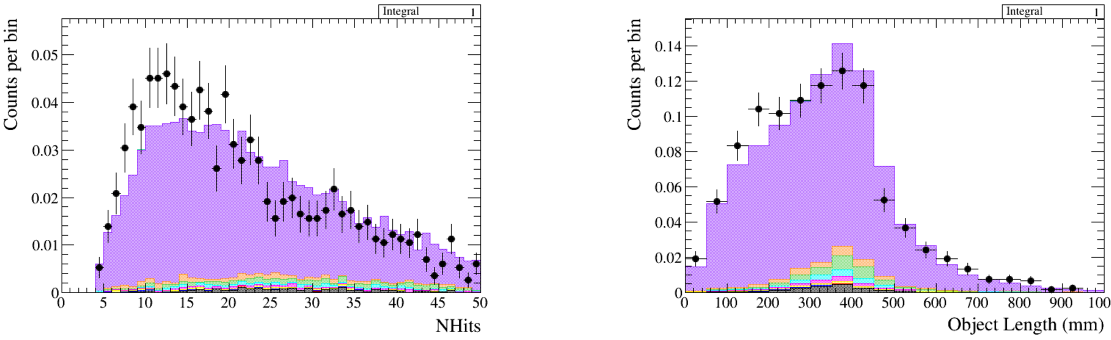

(i) NHits (9). (j) Object Length (10).

Figure A10: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the DsDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

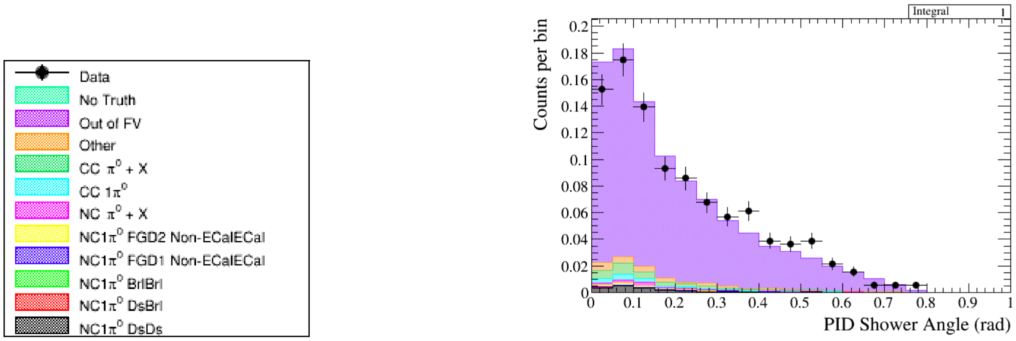

(k) PID Shower Angle (11).

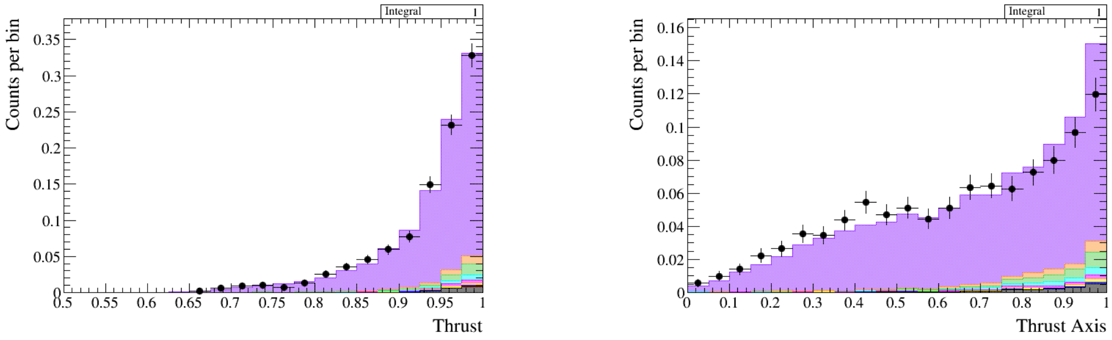

(l) Thrust (12).

(m) Thrust Axis (13).

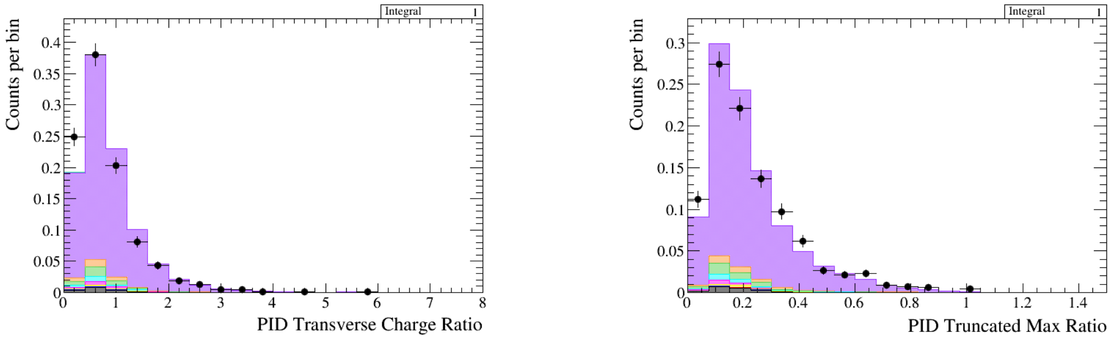

(n) PID Transverse Charge Ratio (14).

(o) PID Truncated Max Ratio (15).

Figure A10: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the DsDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

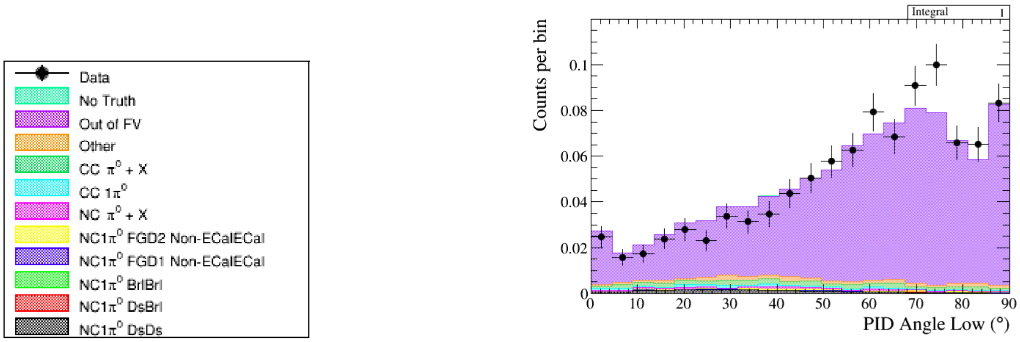

(a) PID Angle Low (16).

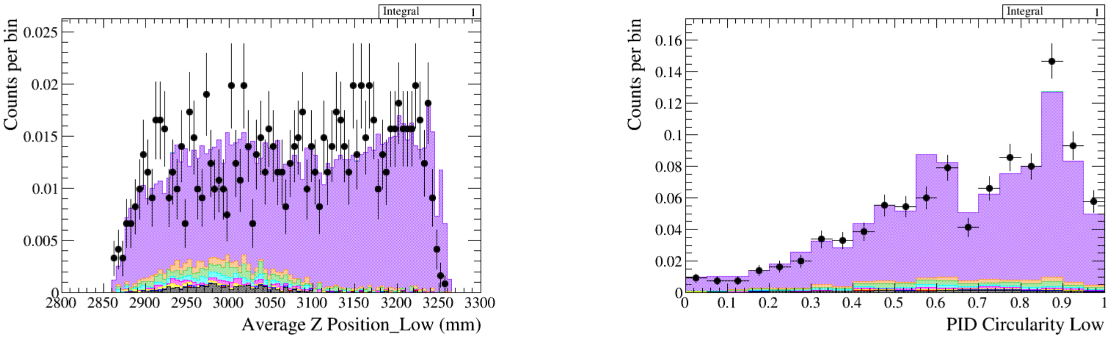

(b) Average Z Position Low (17).

(c) PID Circularity Low (18).

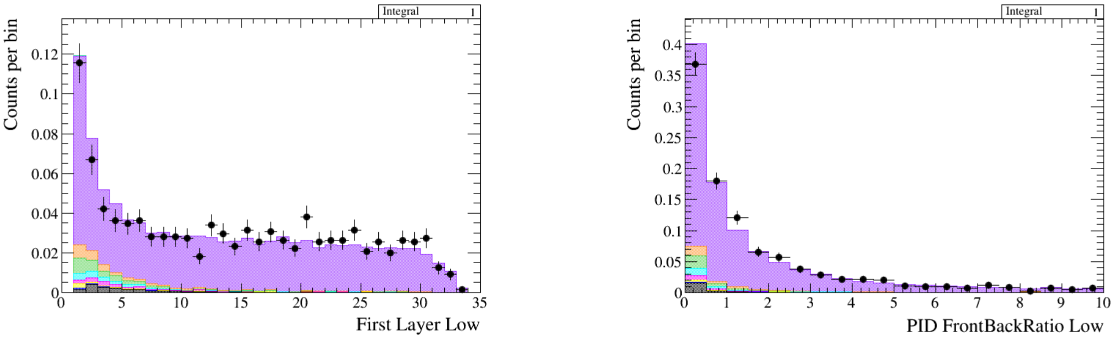

(d) First Layer Low (19).

(e) PID Front Back Ratio Low (20).

Figure A11: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the DsDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

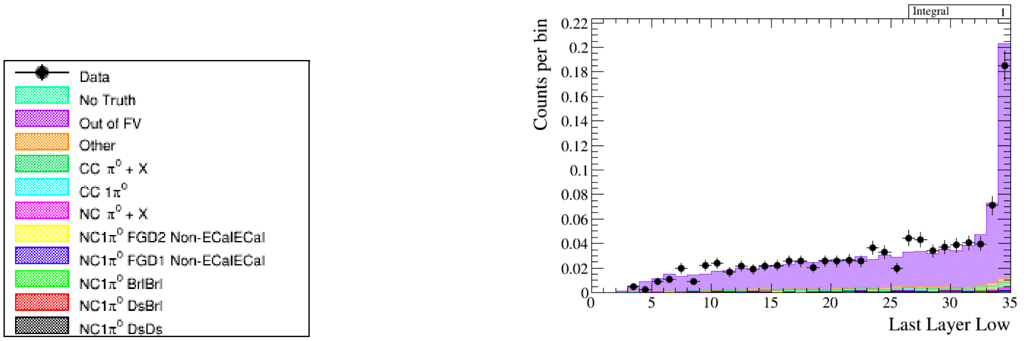

(f) Last Layer Low (21).

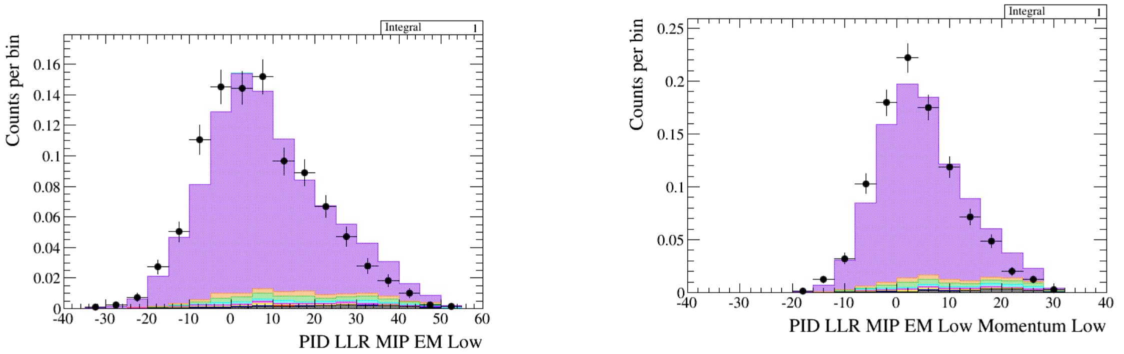

(g) PID LLR MIP EM Low (22).

(h) PID LLR MIP EM Low Momentum Low (23).

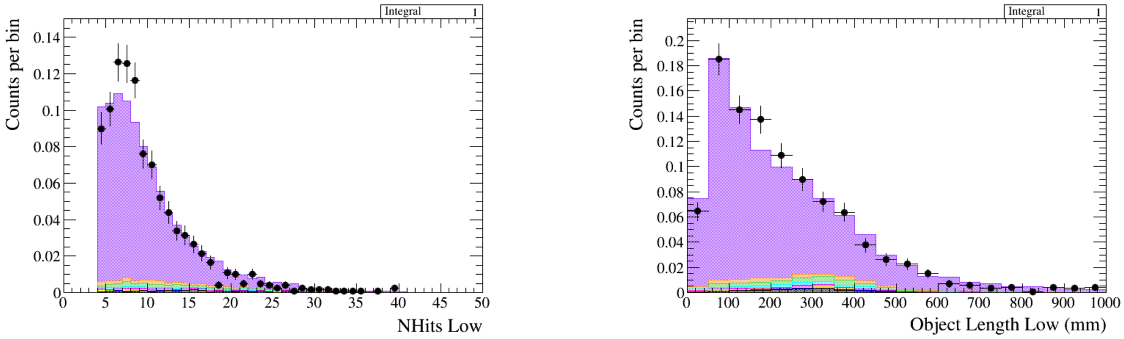

(i) NHits Low (24).

(j) Object Length Low (25).

Figure A11: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the DsDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

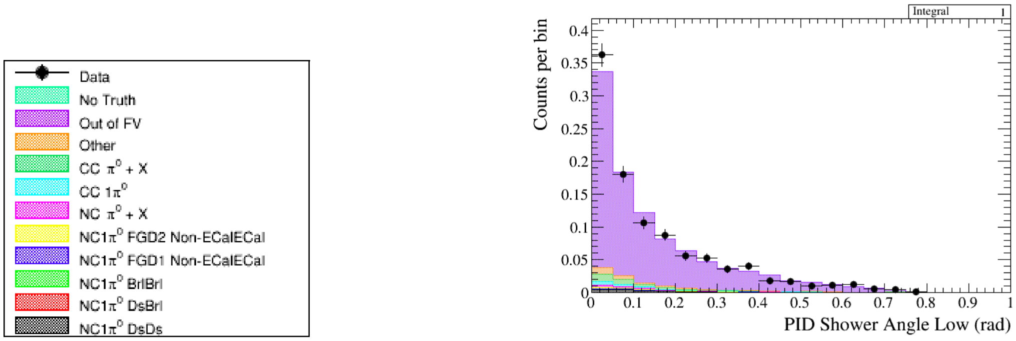

(k) PID Shower Angle Low (26).

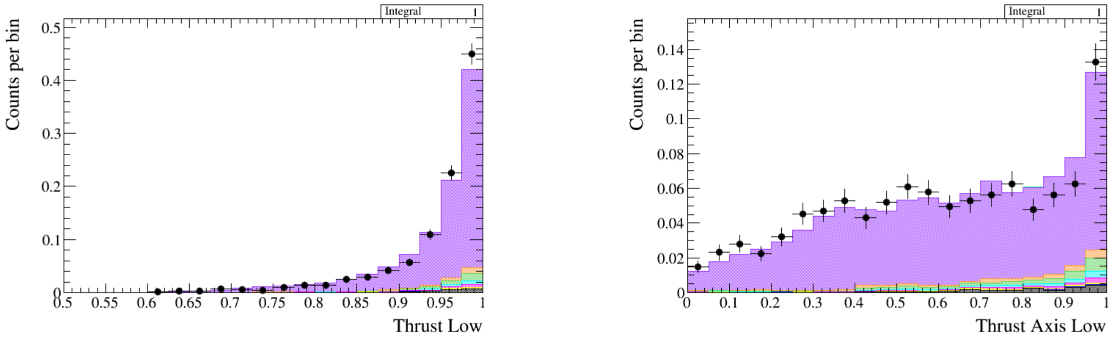

(l) Thrust Low (27).

(m) Thrust Axis Low (28).

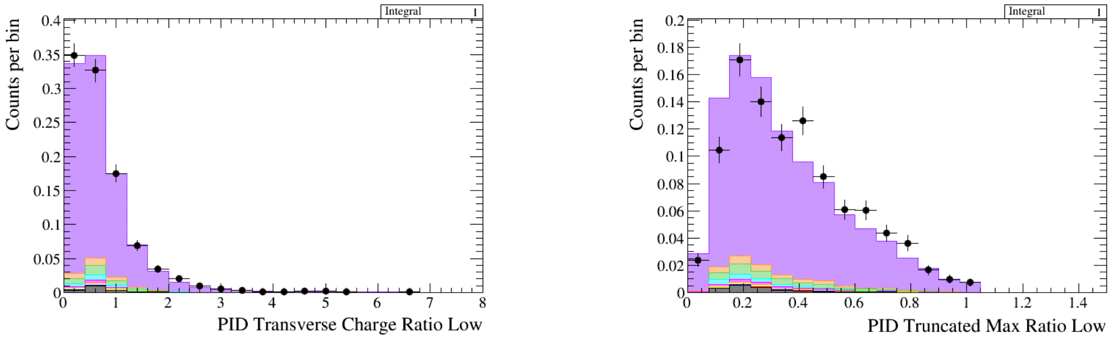

(n) PID Transverse Charge Ratio Low (29).

(o) PID Truncated Max Ratio Low (30).

Figure A11: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the DsDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

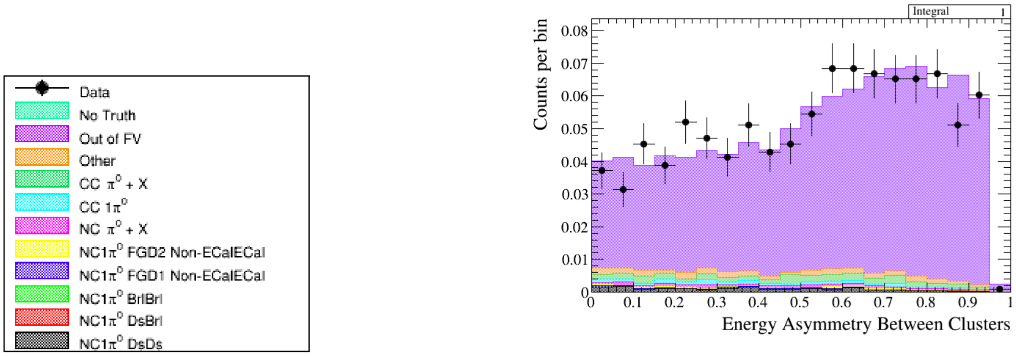

(a) Energy Asymmetry Between Clusters (31).

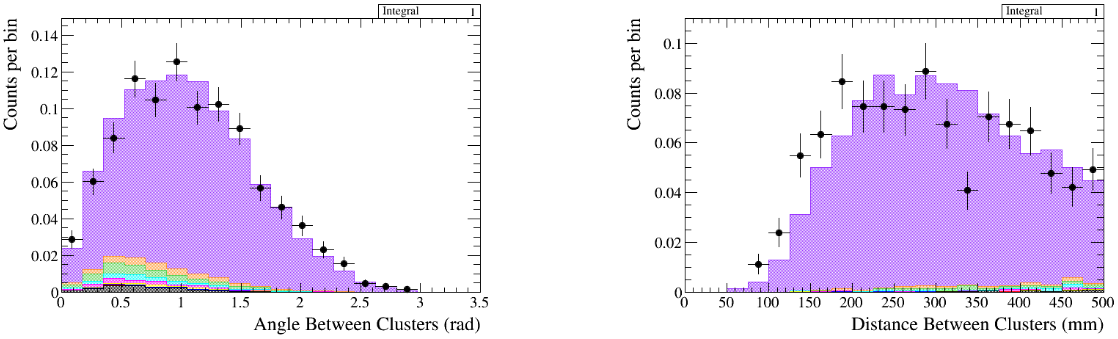

(b) Angle Between Clusters (32).

(c) Distance Between Clusters (33).

Figure A12: Data/MC comparisons of the combined ECal cluster variables used to train the BDT for the DsDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

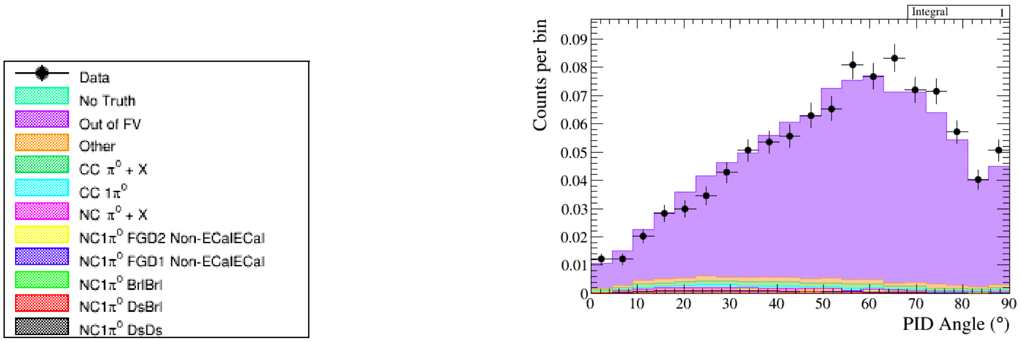

(a) PID Angle (1).

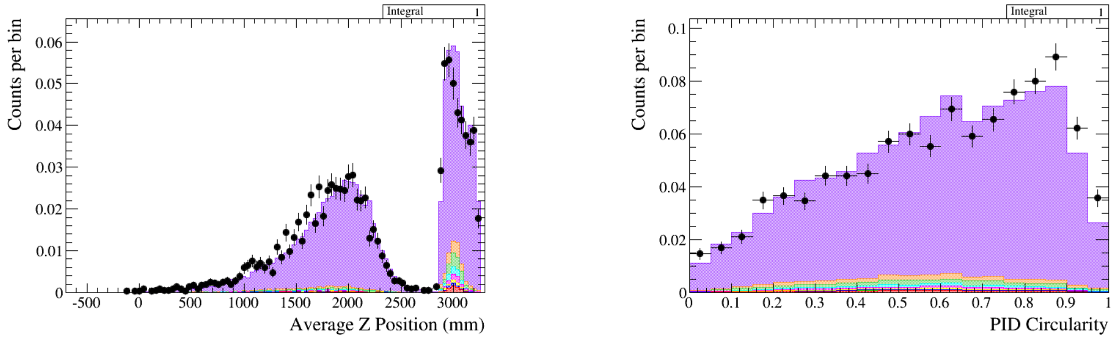

(b) Average Z Position (2).

(c) PIC Circularity (3).

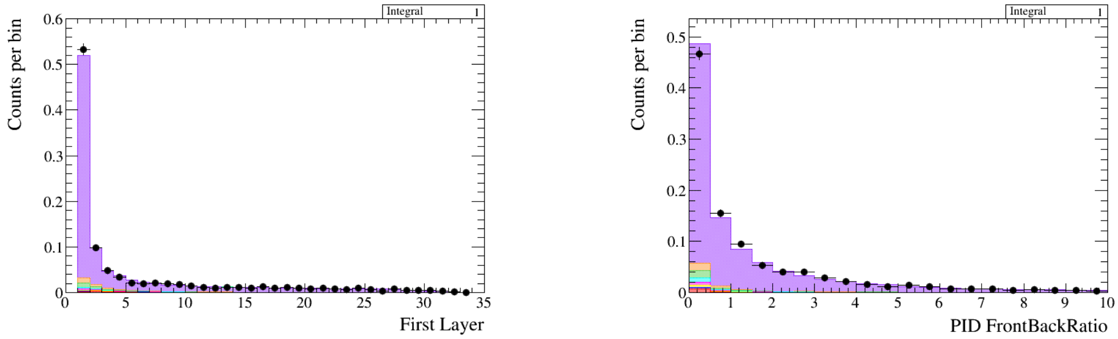

(d) First Layer (4).

(e) PID Front Back Ratio (5).

Figure A13: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the BrlDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

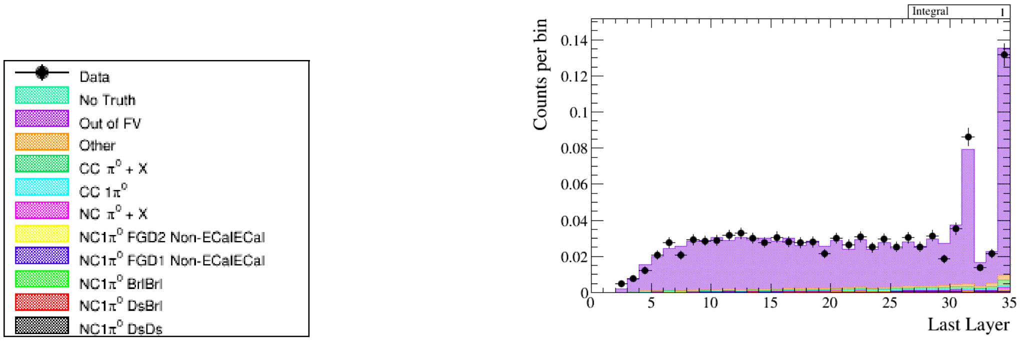

(f) Last Layer (6).

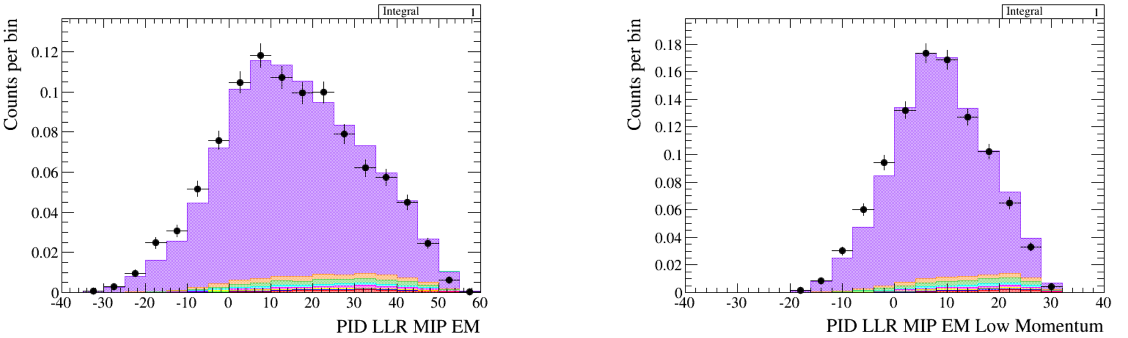

(g) PID LLR MIP EM (7). (h) PID LLR MIP EM Low Momentum (8).

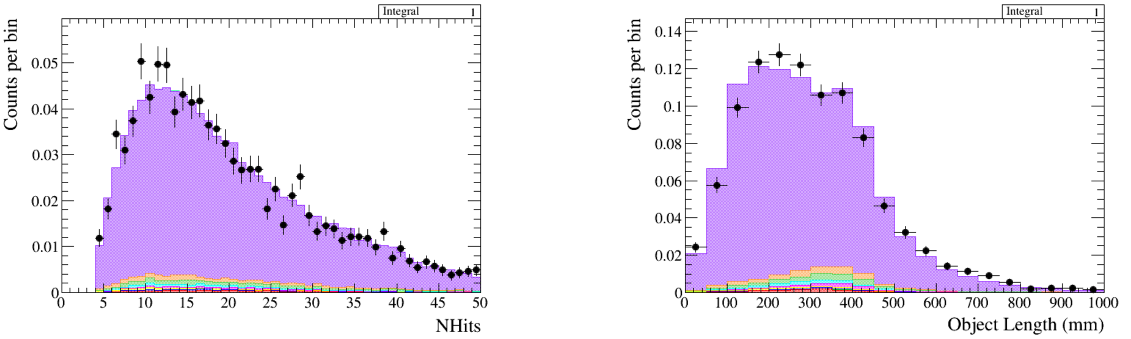

(i) NHits (9). (j) Object Length (10).

Figure A13: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the BrlDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

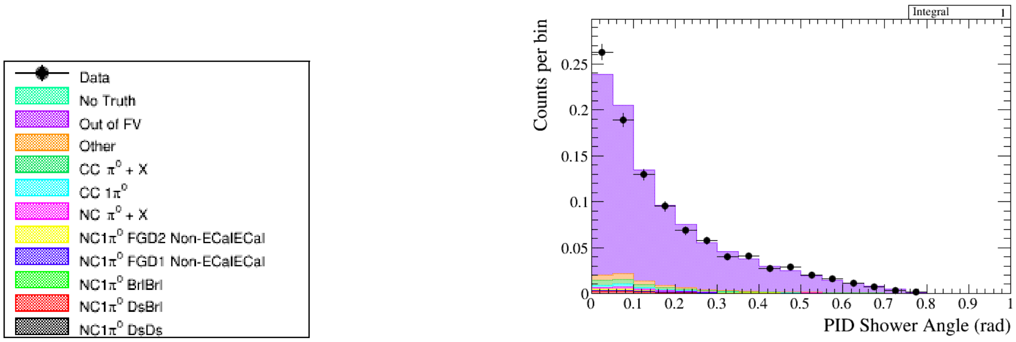

(k) PID Shower Angle (11).

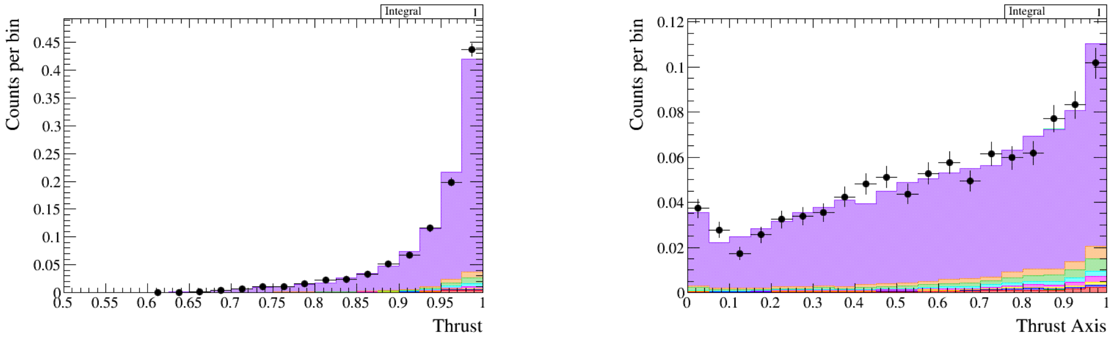

(l) Thrust (12).

(m) Thrust Axis (13).

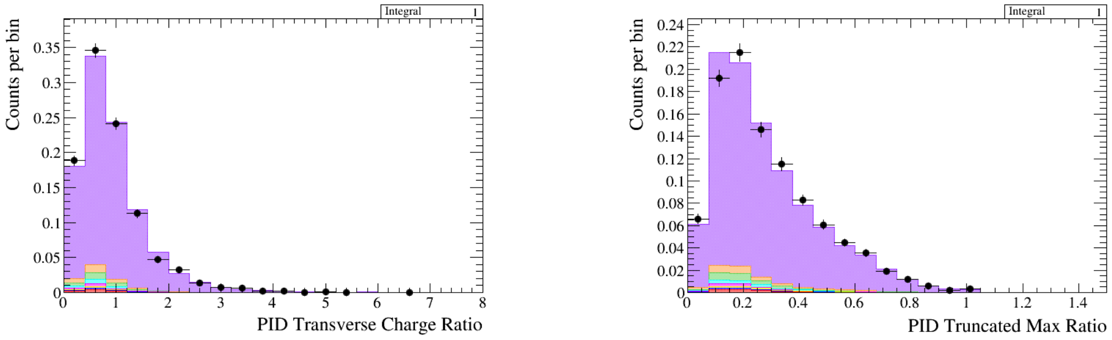

(n) PID Transverse Charge Ratio (14).

(o) PID Truncated Max Ratio (15).

Figure A13: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the BrlDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

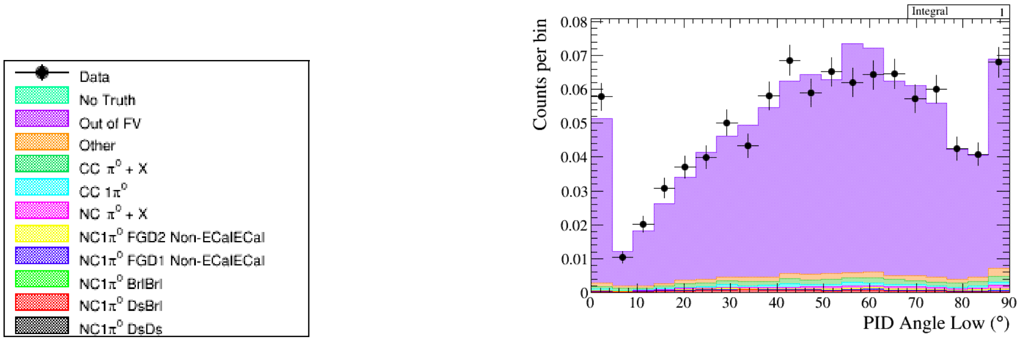

(a) PID Angle Low (16).

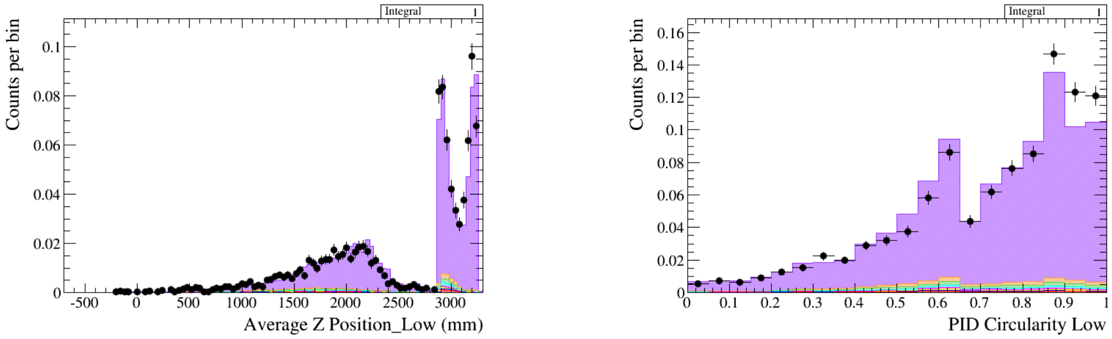

(b) Average Z Position Low (17).

(c) PID Circularity Low (18).

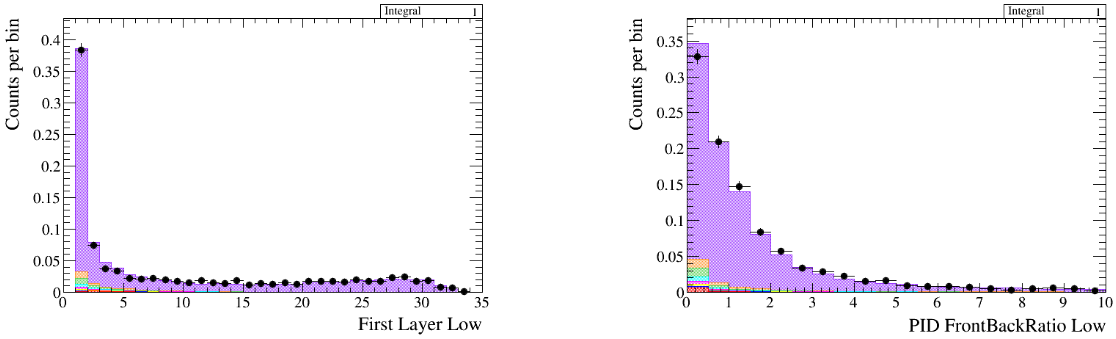

(d) First Layer Low (19).

(e) PID Front Back Ratio Low (20).

Figure A14: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the BrlDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

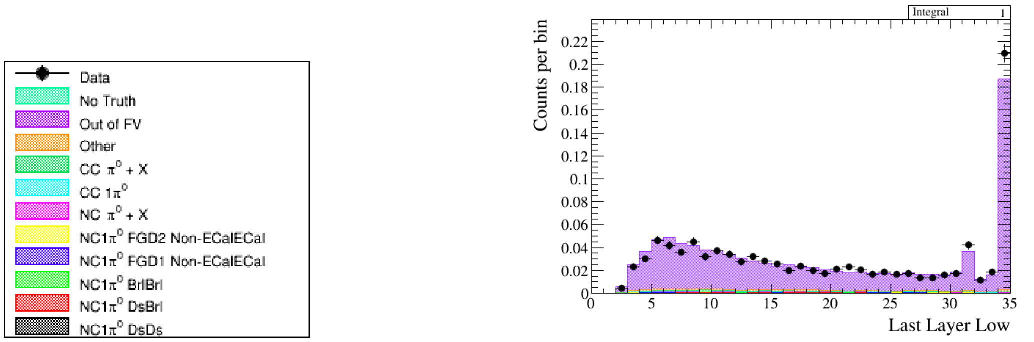

(f) Last Layer Low (21).

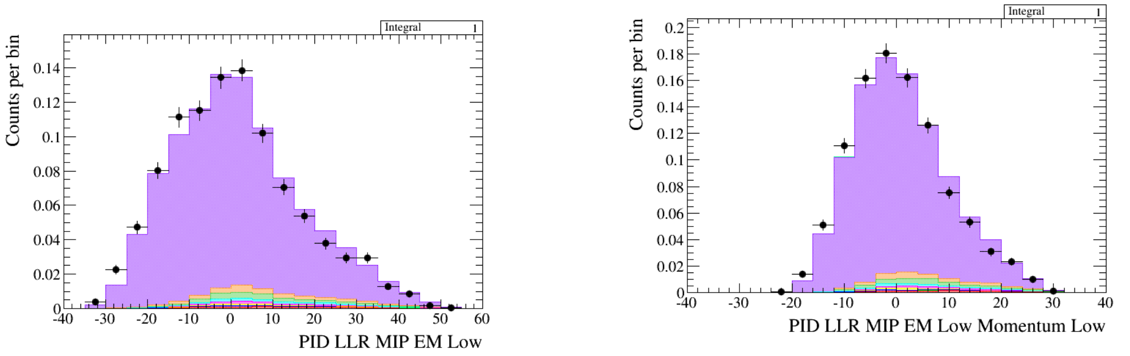

(g) PID LLR MIP EM Low (22).

(h) PID LLR MIP EM Low Momentum Low (23).

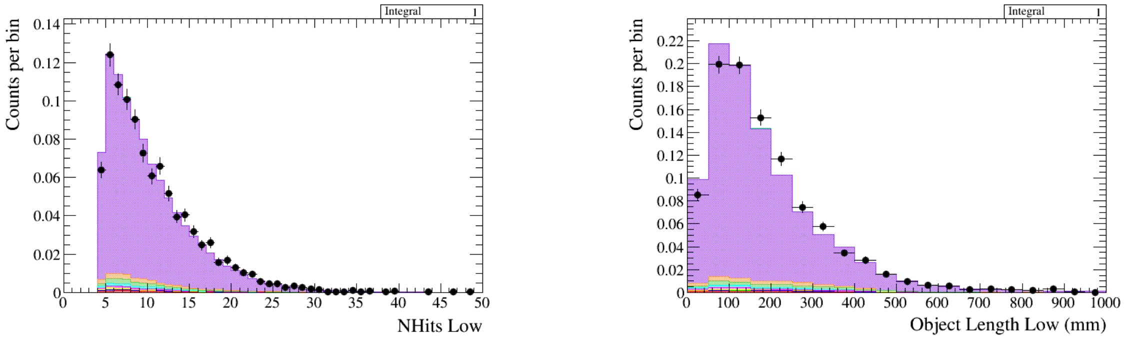

(i) NHits Low (24).

(j) Object Length Low (25).

Figure A14: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the BrlDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

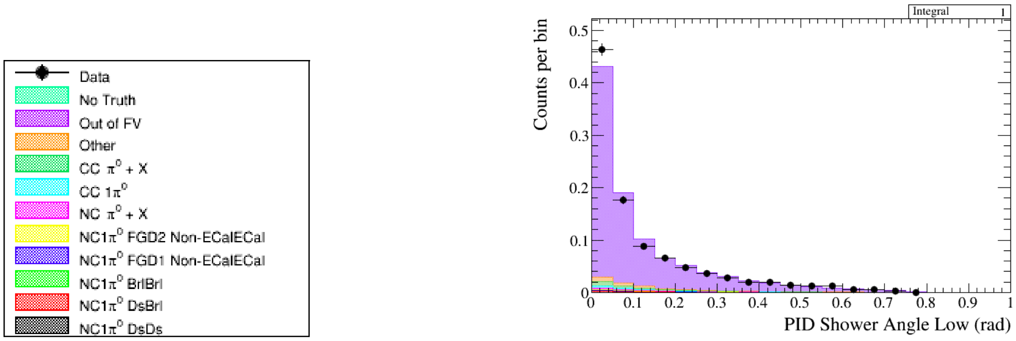

(k) PID Shower Angle Low (26).

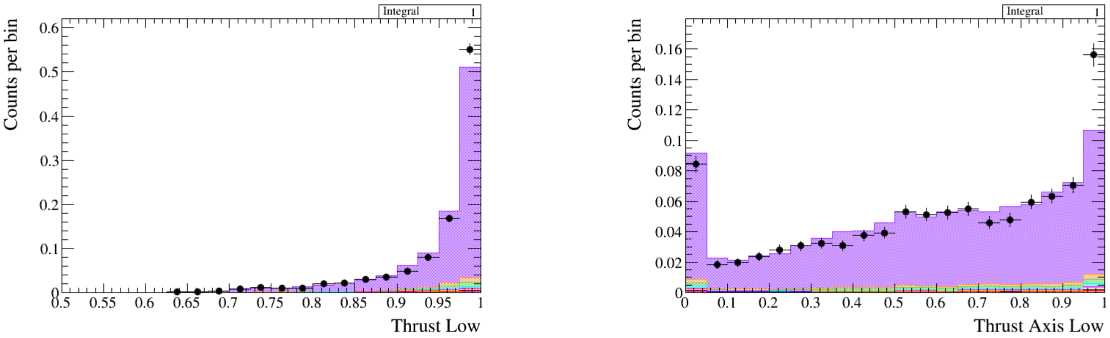

(l) Thrust Low (27).

(m) Thrust Axis Low (28).

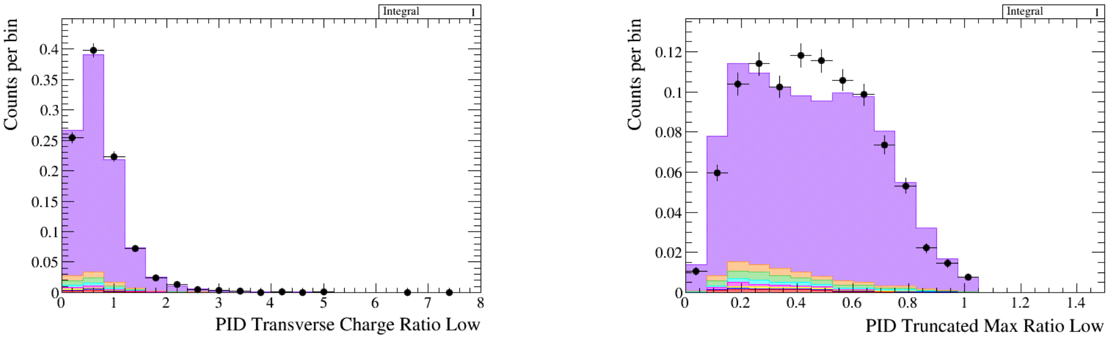

(n) PID Transverse Charge Ratio Low (29).

(o) PID Truncated Max Ratio Low (30).

Figure A14: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the BrlDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

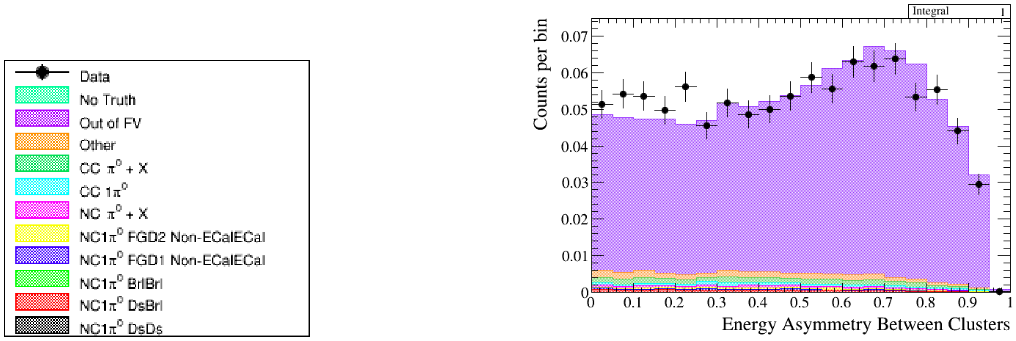

(a) Angle Between Clusters (31).

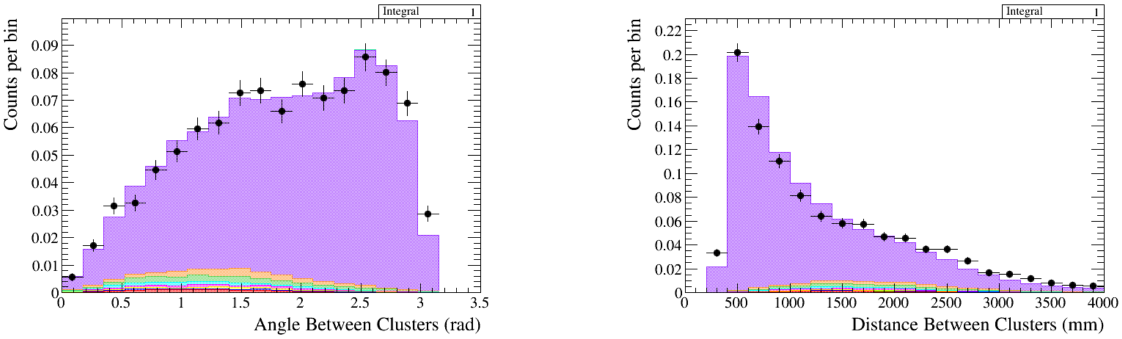

(b) Angle Between Clusters (32).

(c) Distance Between Clusters (33).

Figure A15: Data/MC comparisons of the combined ECal cluster variables used to train the BDT for the BrlDs topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

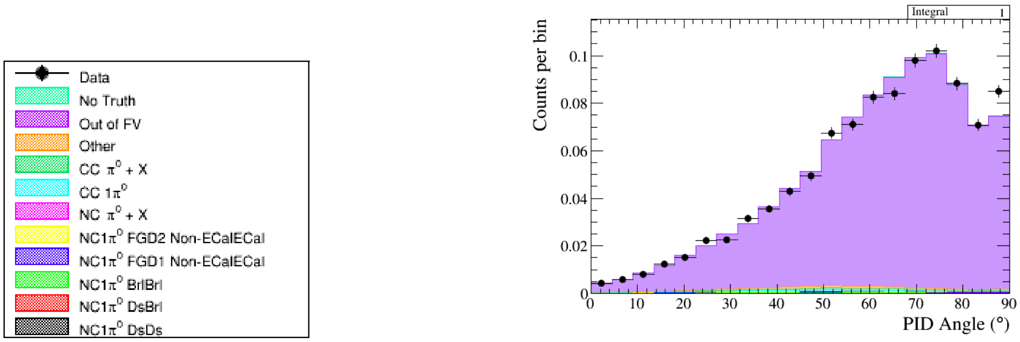

(a) PID Angle (1).

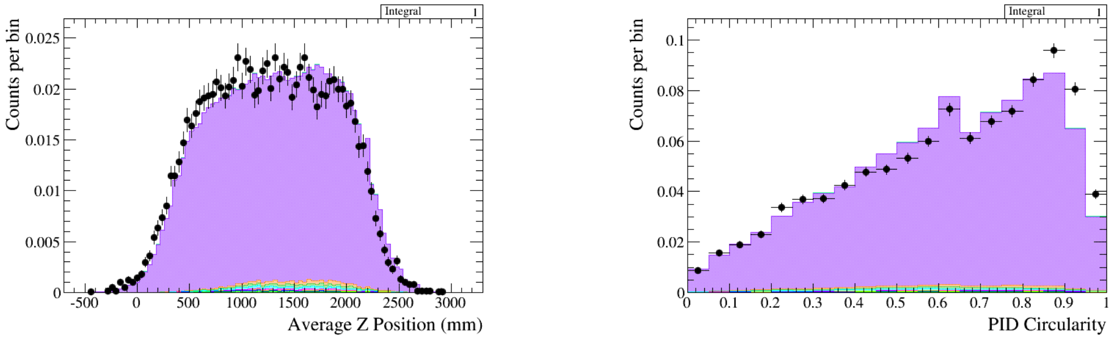

(b) Average Z Position (2).

(c) PID Circularity (3).

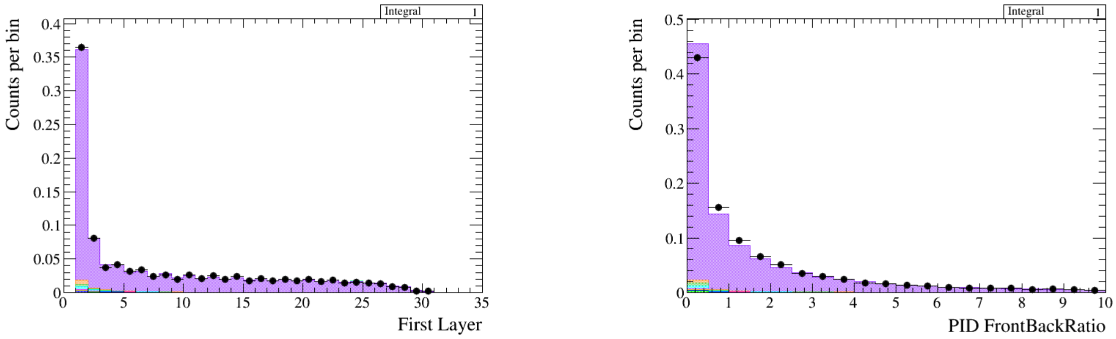

(d) First Layer (4).

(e) PID Front Back Ratio (5).

Figure A16: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the BrlBrl topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

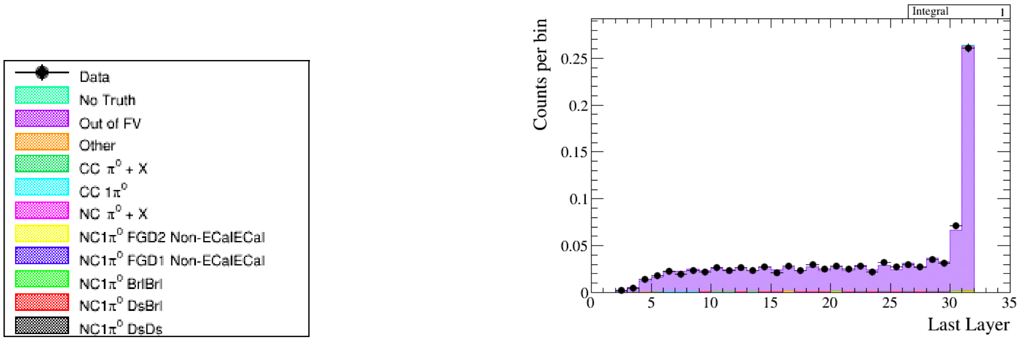

(f) Last Layer (6).

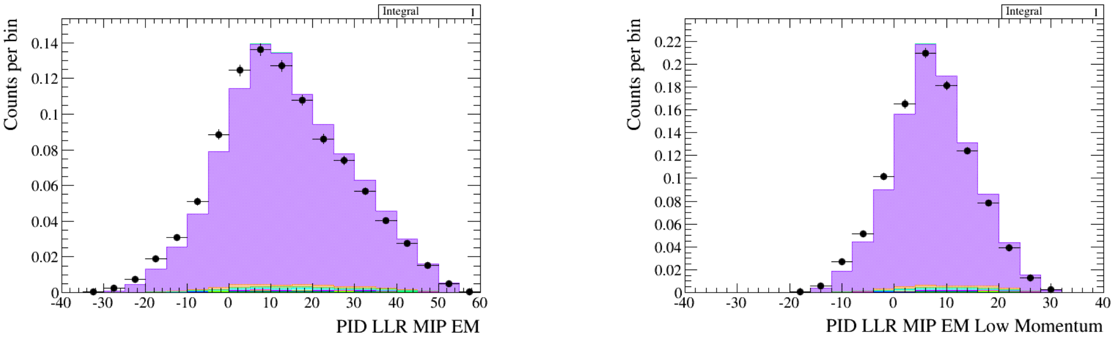

(g) PID LLR MIP EM (7). (h) PID LLR MIP EM Low Momentum (8).

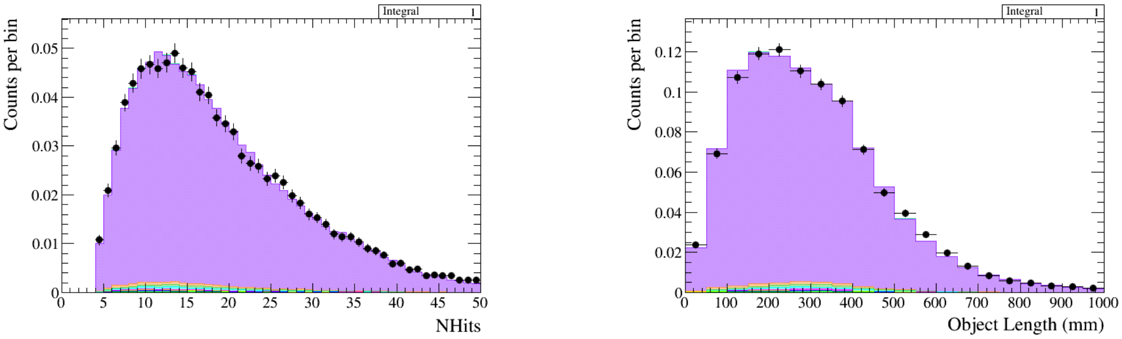

(i) NHits (9). (j) Object Length (10).

Figure A16: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the BrlBrl topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

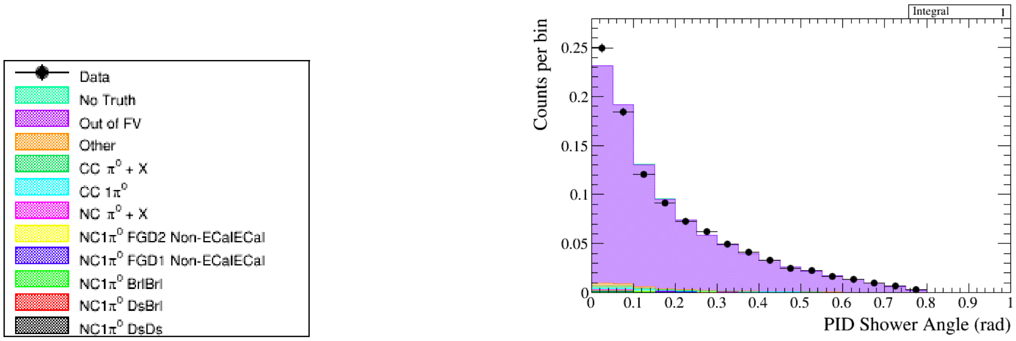

(k) PID Shower Angle (11).

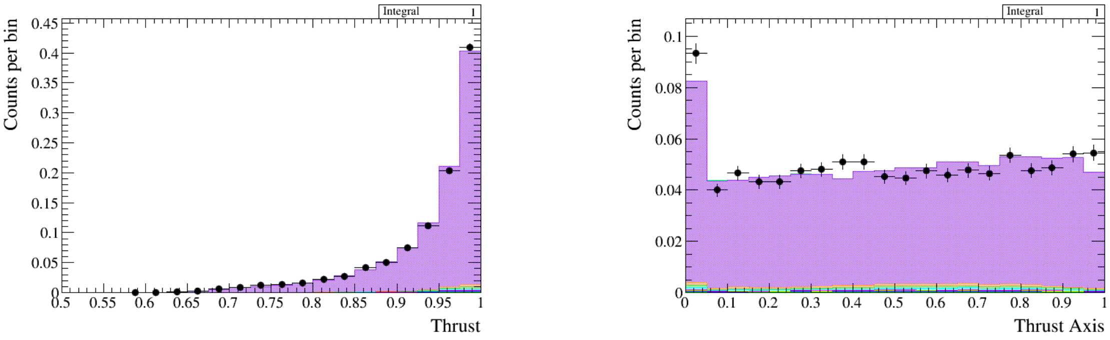

(l) Thrust (12).

(m) Thrust Axis (13).

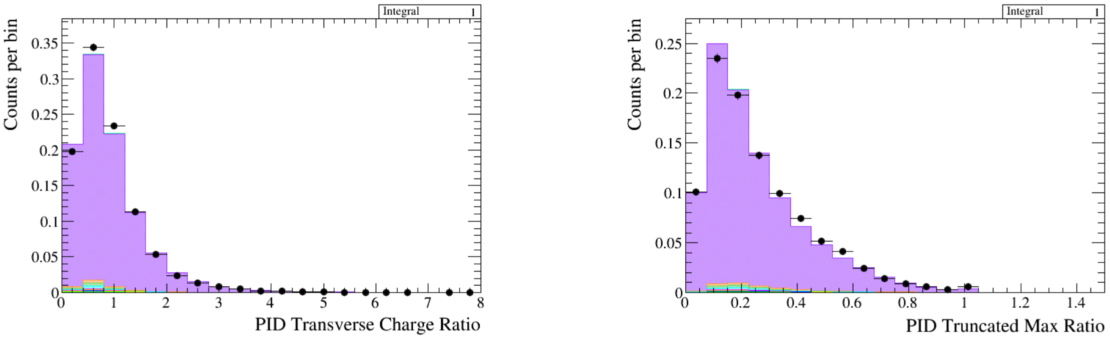

(n) PID Transverse Charge Ratio (14).

(o) PID Truncated Max Ratio (15).

Figure A16: Data/MC comparisons of the high energy ECal cluster variables used to train the BDT for the BrlBrl topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

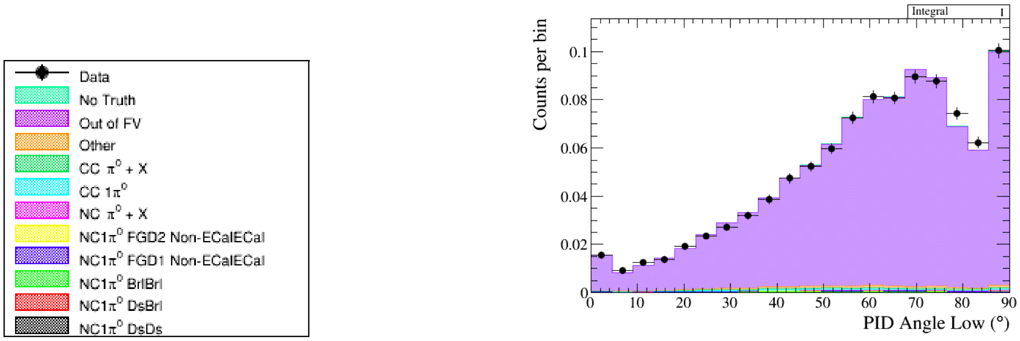

(a) PID Angle Low (16).

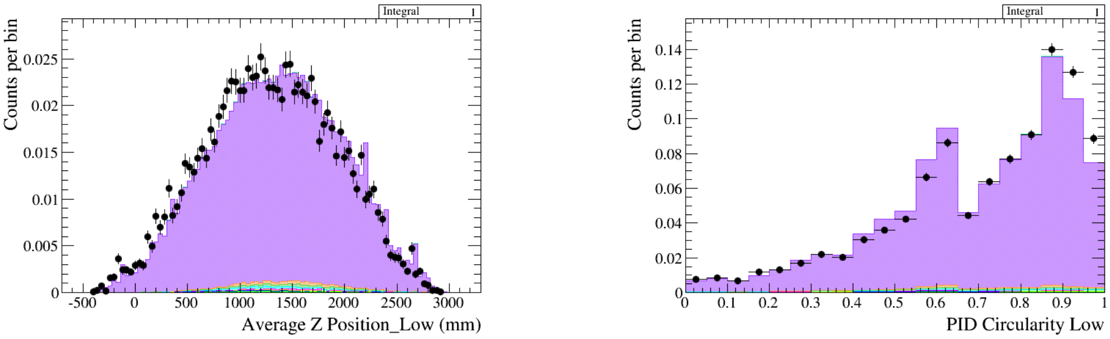

(b) Average Z Position Low (17).

(c) PID Circularity Low (18).

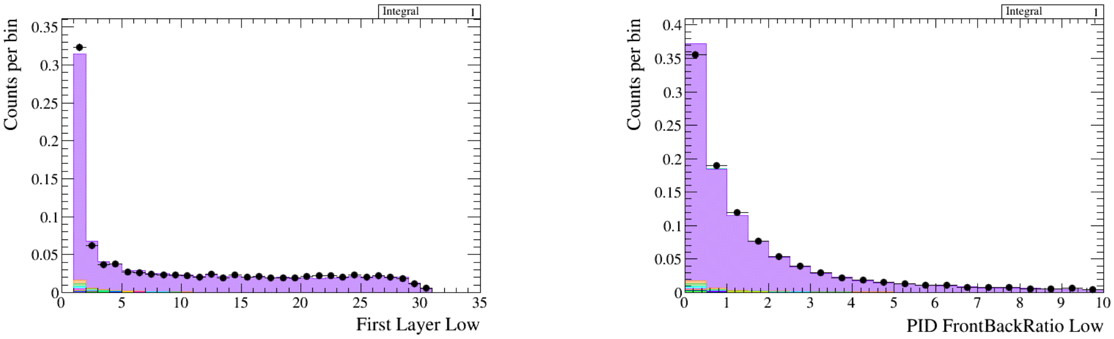

(d) First Layer Low (19).

(e) PID Front Back Ratio Low (20).

Figure A17: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the BrlBrl topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

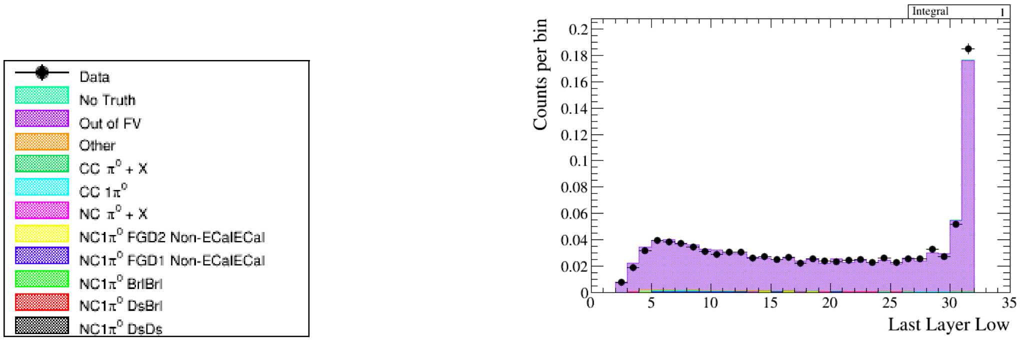

(f) Last Layer Low (21).

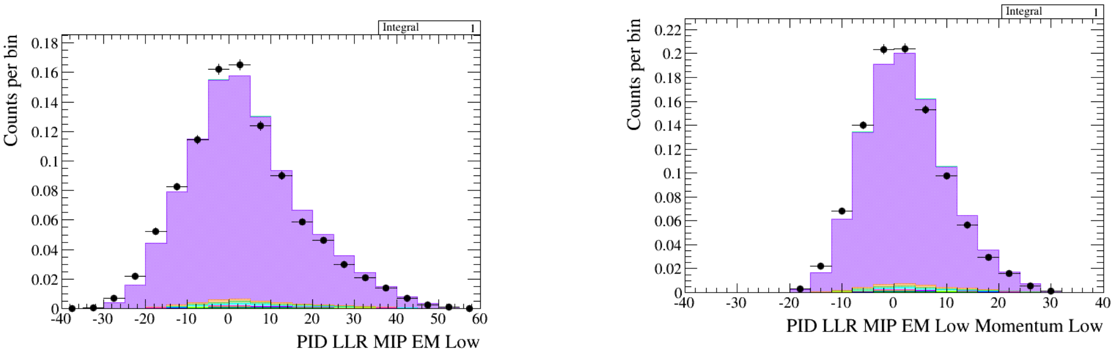

(g) PID LLR MIP EM Low (22).

(h) PID LLR MIP EM Low Momentum Low (23).

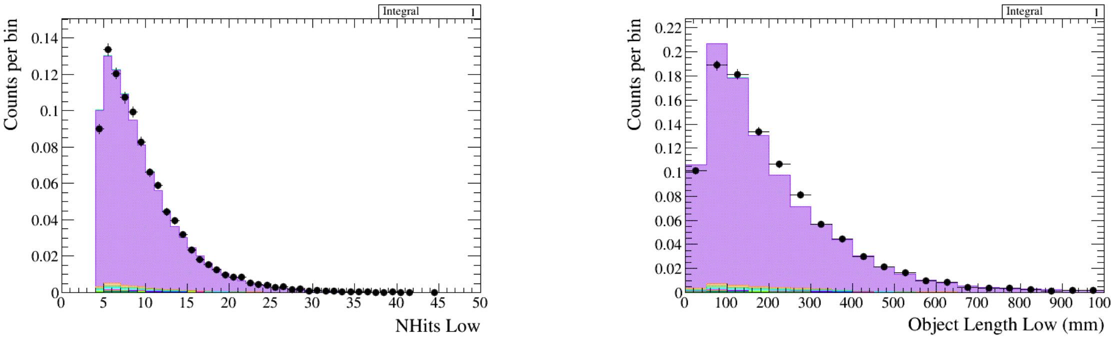

(i) NHits Low (24).

(j) Object Length Low (25).

Figure A17: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the BrlBrl topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

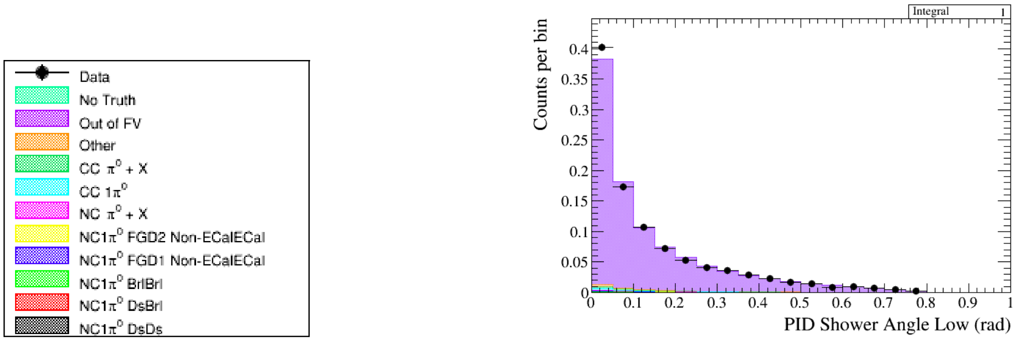

(k) PID Shower Angle Low (26).

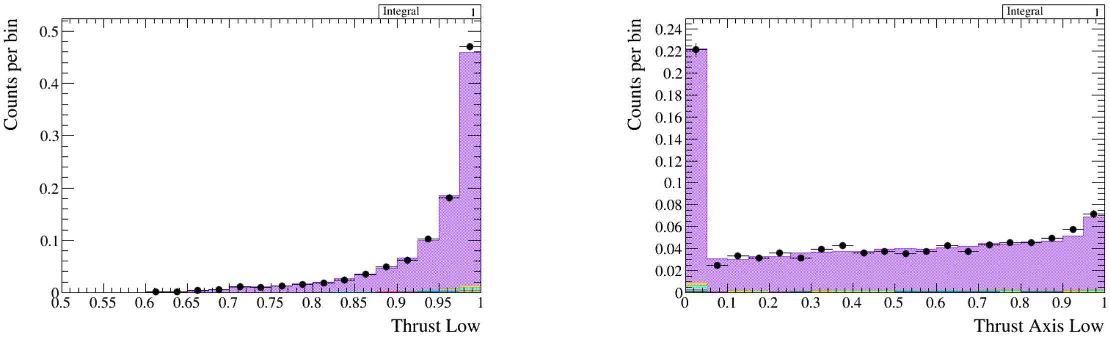

(l) Thrust Low (27).

(m) Thrust Axis Low (28).

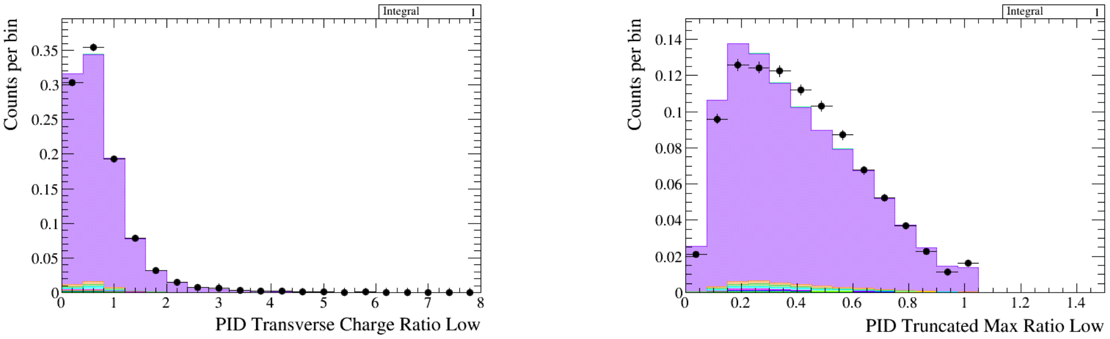

(n) PID Transverse Charge Ratio Low (29).

(o) PID Truncated Max Ratio Low (30).

Figure A17: Data/MC comparisons of the low energy ECal cluster variables used to train the BDT for the BrlBrl topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

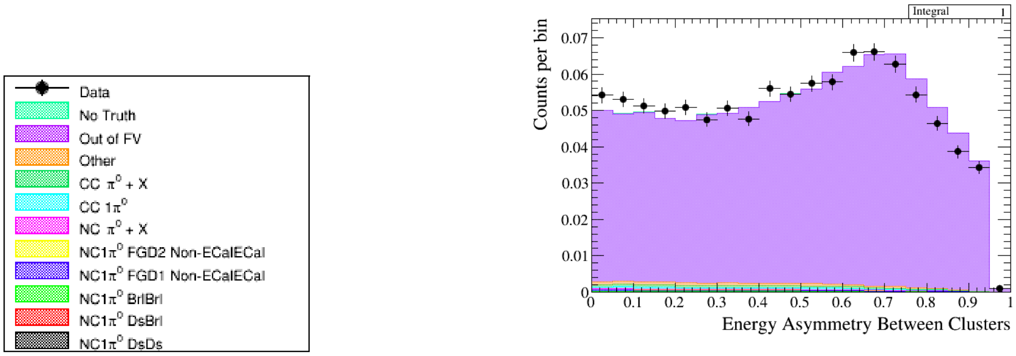

(a) Energy Asymmetry Between Clusters (31).

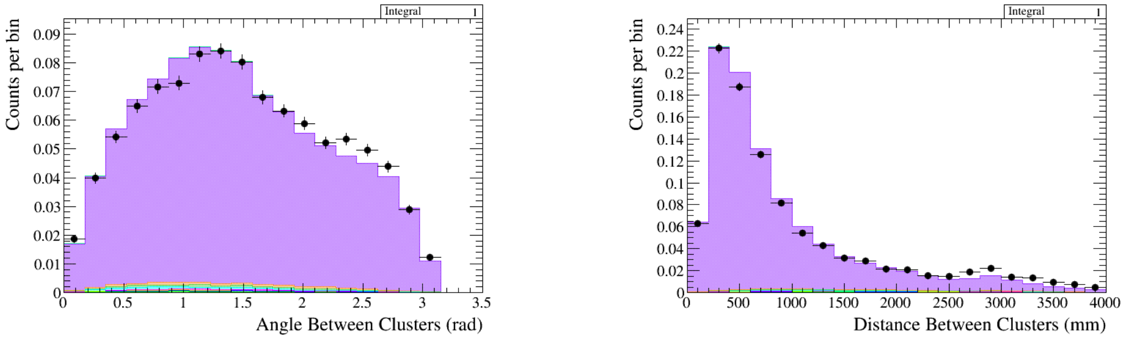

(b) Angle Between Clusters (32).

(c) Distance Between Clusters (33).

Figure A18: Data/MC comparisons of the combined ECal cluster variables used to train the BDT for the BrlBrl topology (Note, Run4 Air and Run4 Water Magnet samples are compared to the Run 4 Data).

Appendix B

Singular value decomposition

In order to use SVD instead of Cholesky decomposition, it is imperative that both methods give the same results. To prove this is the case, consider a matrix which can be factorised into its two Cholesky decomposed triangular matrices:

$$M = CC^{\dagger}. \tag{B.1}$$

This can also be factorised into its SVD counterpart:

$$M = USV^{\dagger}. \tag{B.2}$$

where U and V are unitary matrices and S is diagonal. Now, as M is real, the following relations are true:

$$C^{\dagger} = C^{T}, \tag{B.3}$$

$$M = CC^{T}. \tag{B.4}$$

Furthermore, as M is also diagonalisable, the following relations are also true:

$$V^{\dagger} = V^{T} = V^{-1} \tag{B.5}$$

$$U = V. \tag{B.6}$$

Thus:

$$M = USU^{\dagger} = USU^{T}. \tag{B.7}$$

As S is diagonal, matrix P can be defined as:

$$P = P^{T} = S^{\frac{1}{2}}. \tag{B.8}$$

Hence, the following equality is found:

$$M = USU^{T} = (UP)(P^{T}U^{T}) = CC^{T}, \tag{B.9}$$

provided

$$UP = C. \tag{B.10}$$

Appendix C

Cross-section supplementary information

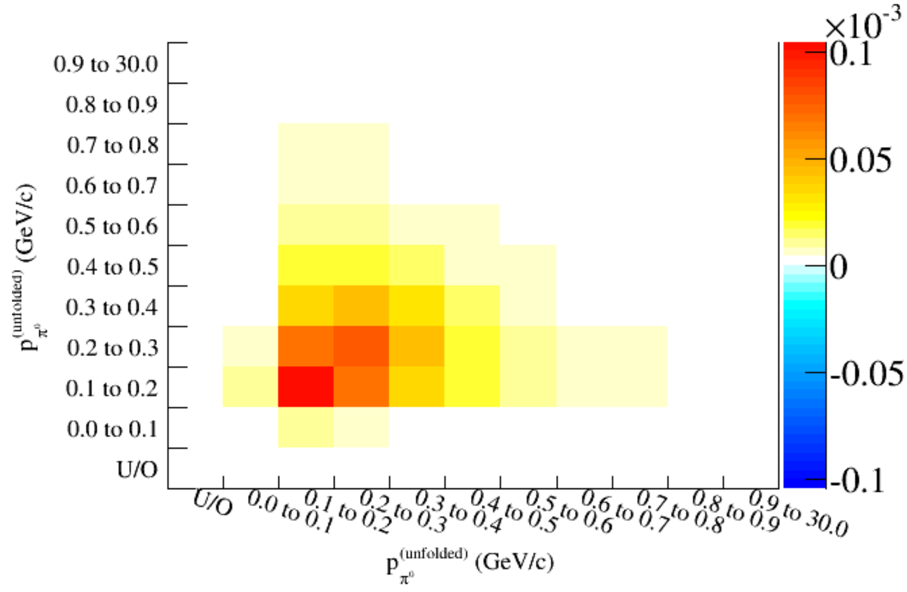

Figure C1: The momentum covariance matrix for all errors combined for the NC1π^0 inclusive measurement.

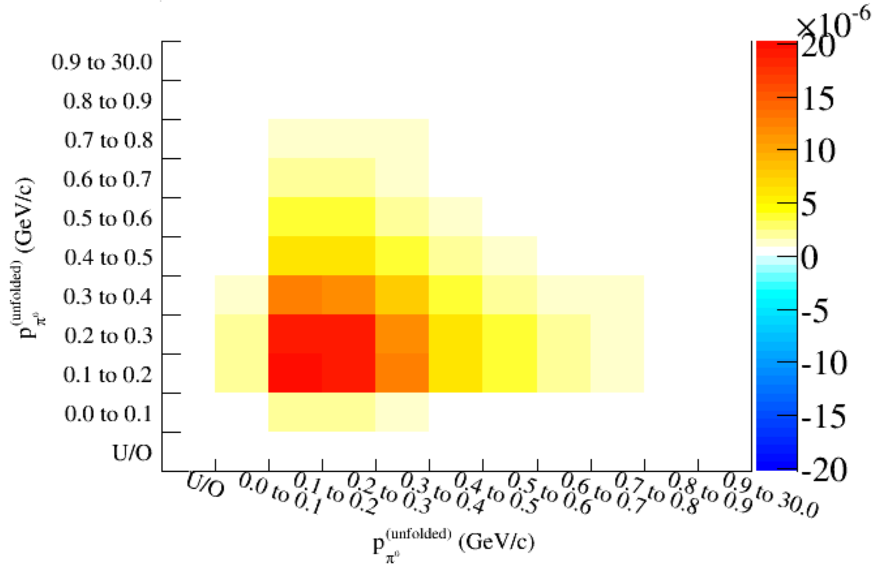

Figure C2: The momentum covariance matrix for all Highland errors for the NC1π^0 inclusive topology.

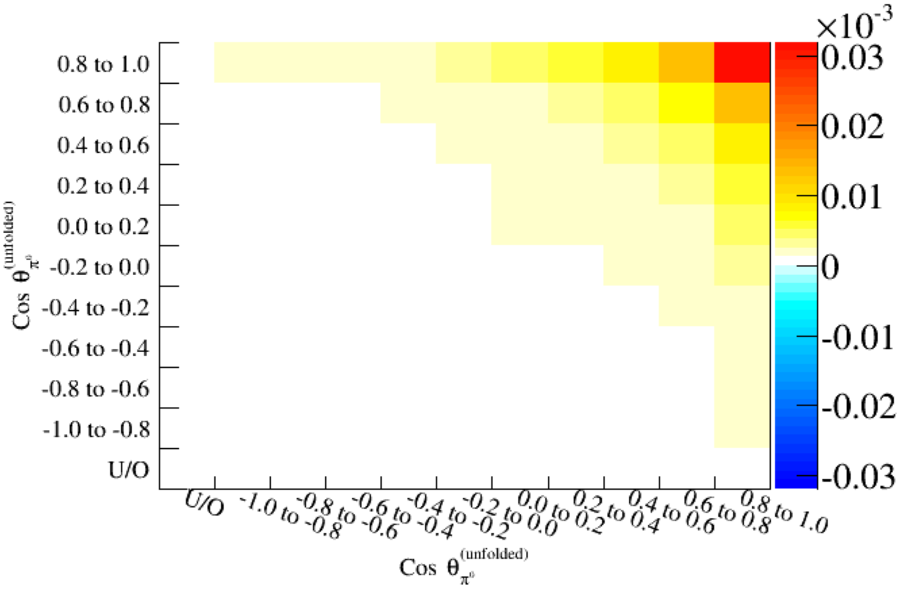

Figure C3: The cosθ covariance matrix for all errors combined for the NC1π^0 inclusive measurement.

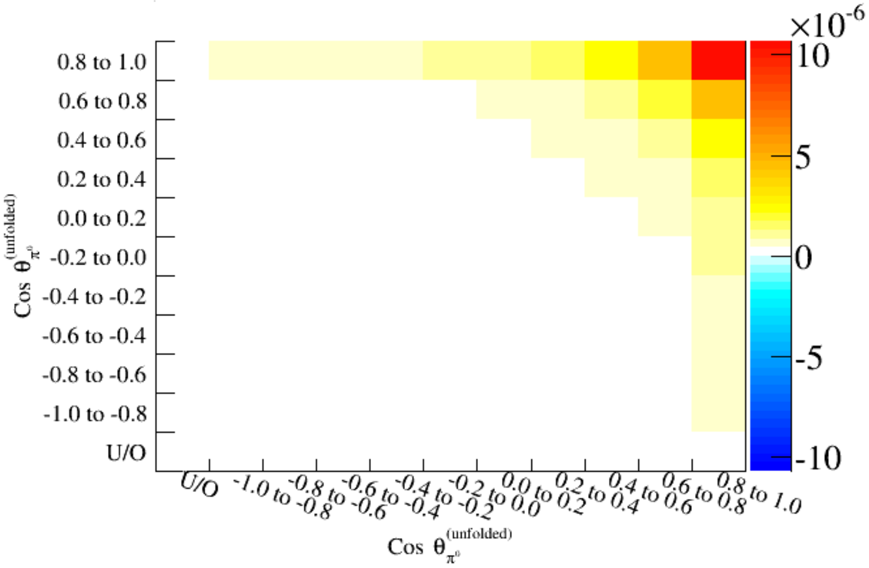

Figure C4: The cosθ covariance matrix for all Highland errors for the NC1π^0 inclusive topology.

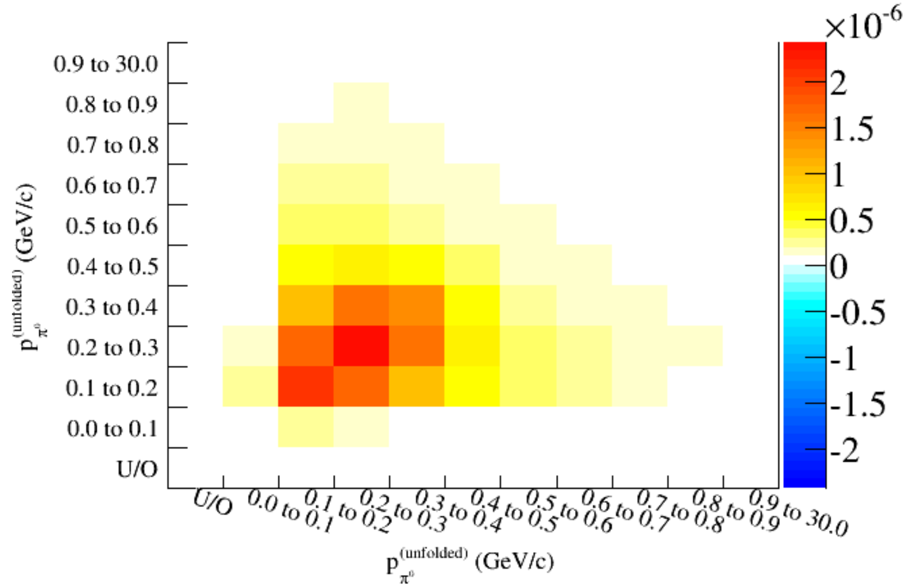

Figure C5: The momentum covariance matrix for all errors combined for the NC1π^0 ECalECal measurement.

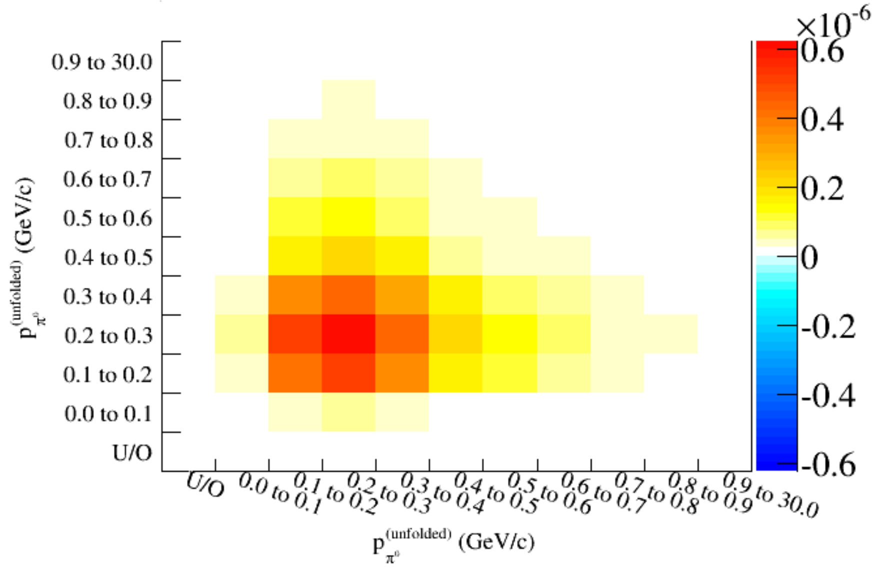

Figure C6: The momentum covariance matrix for all Highland errors for the NC1π^0 ECalECal topology.

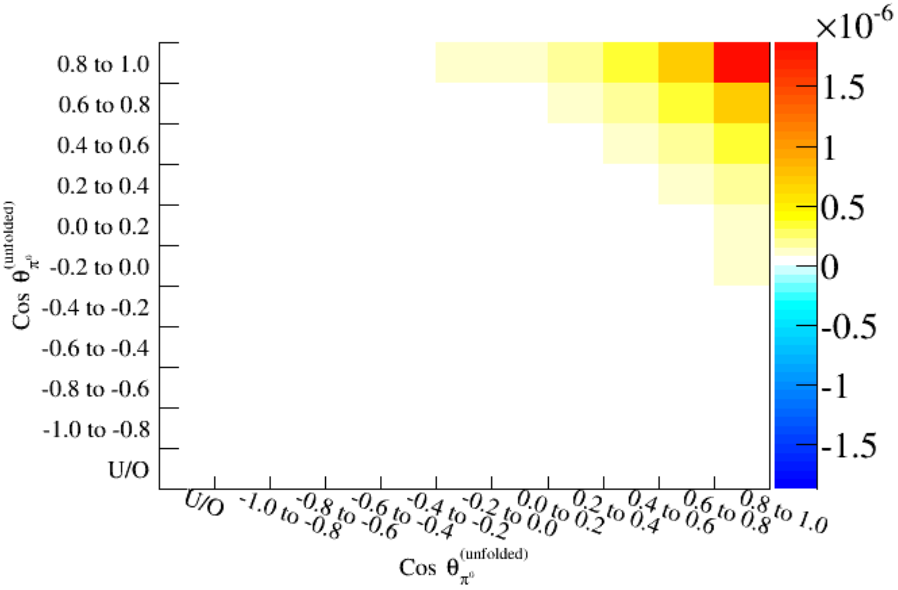

Figure C7: The cosθ covariance matrix for all errors combined for the NC1π^0 ECalECal measurement.

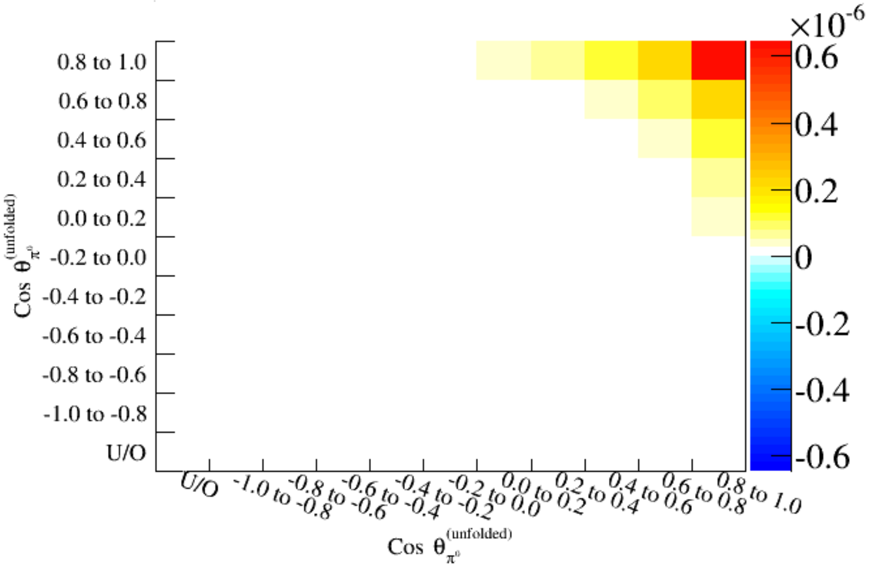

Figure C8: The cosθ covariance matrix for all Highland errors for the NC1π^0 ECalECal topology.

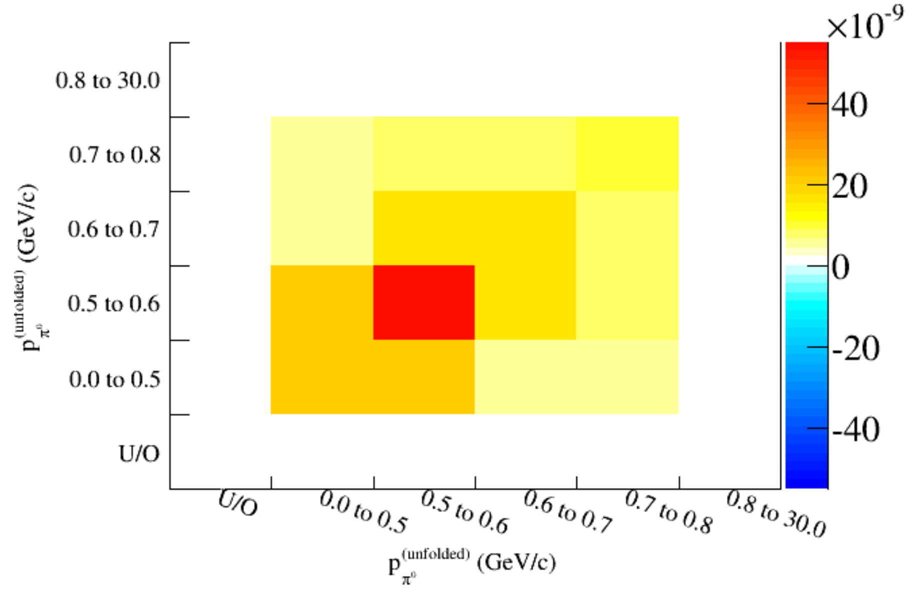

Figure C9: The momentum covariance matrix for all errors combined for the NC1π^0 DsDs measurement.

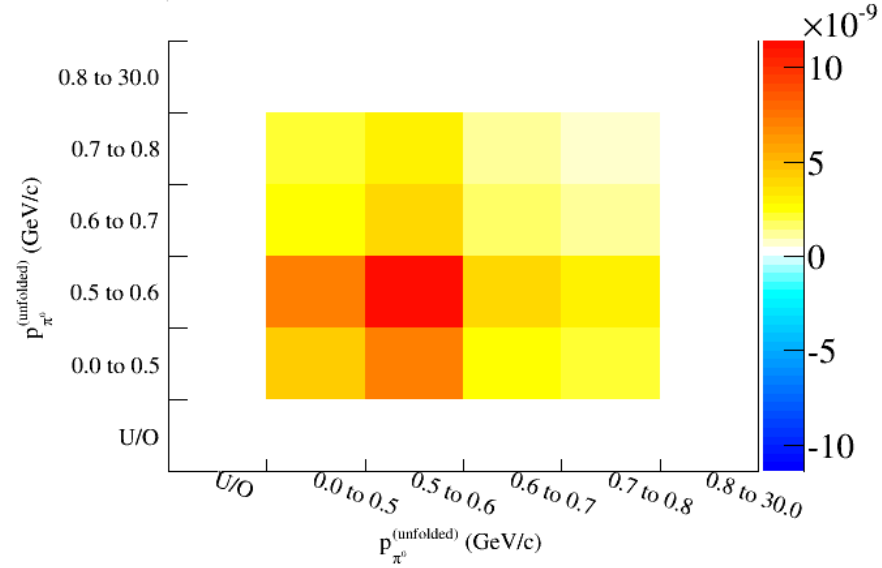

Figure C10: The momentum covariance matrix for all Highland errors for the NC1π^0 DsDs topology.

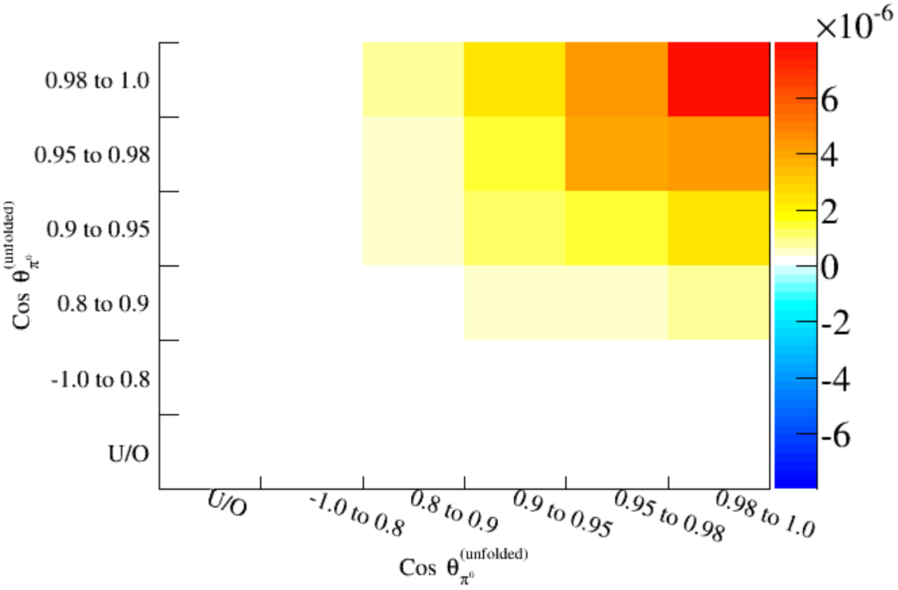

Figure C11: The cosθ covariance matrix for all errors combined for the NC1π^0 ECalECal measurement.

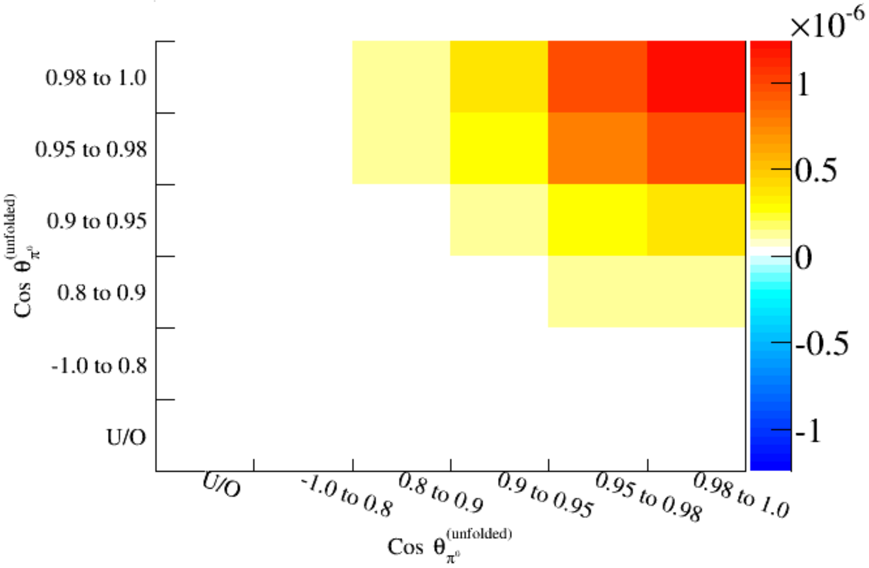

Figure C12: The cosθ covariance matrix for all Highland errors for the NC1π^0 DsDs topology.

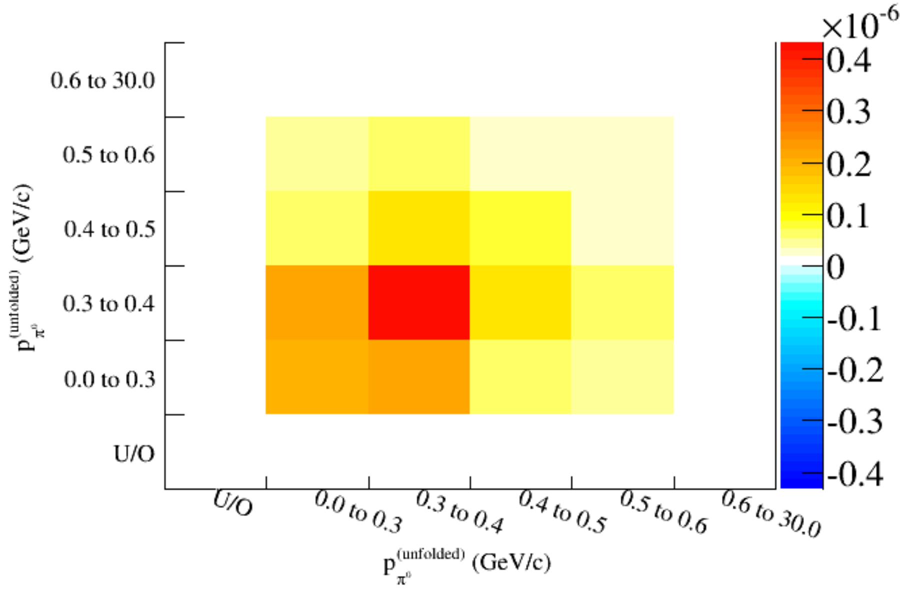

Figure C13: The momentum covariance matrix for all errors combined for the NC1π^0 DsBrl measurement.

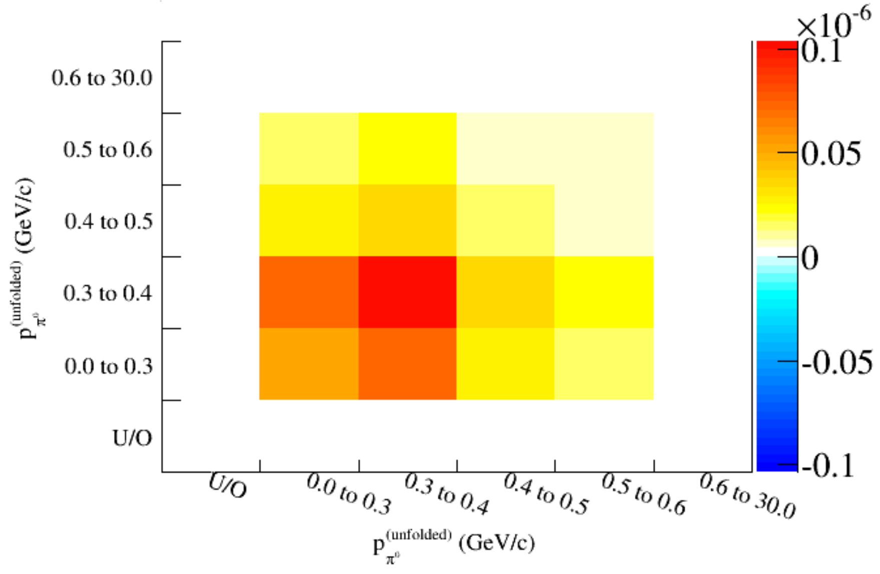

Figure C14: The momentum covariance matrix for all Highland errors for the NC1π^0 DsBrl topology.

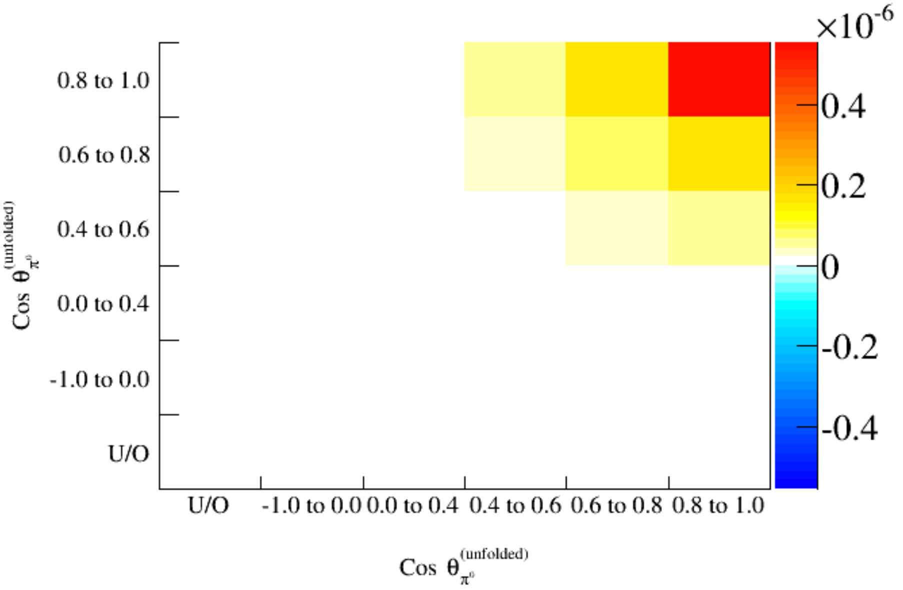

Figure C15: The cosθ covariance matrix for all errors combined for the NC1π^0 DsBrl measurement.

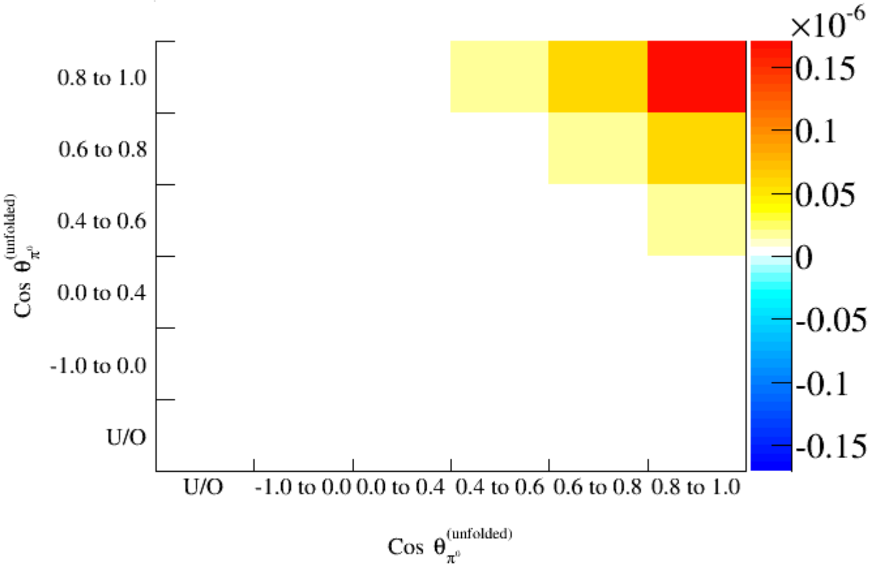

Figure C16: The cosθ covariance matrix for all Highland errors for the NC1π^0 DsBrl topology.

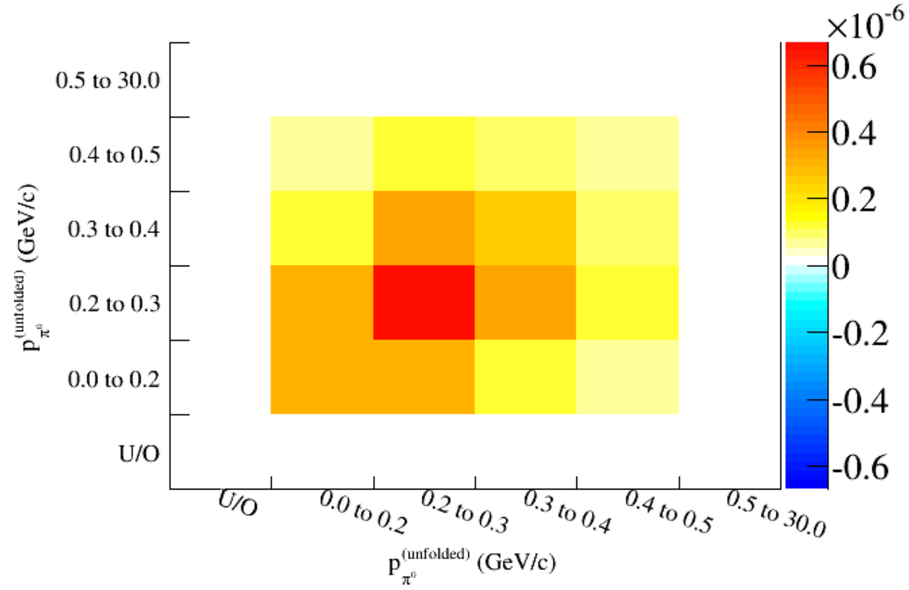

Figure C17: The momentum covariance matrix for all errors combined for the NC1π^0 BrlBrl measurement.

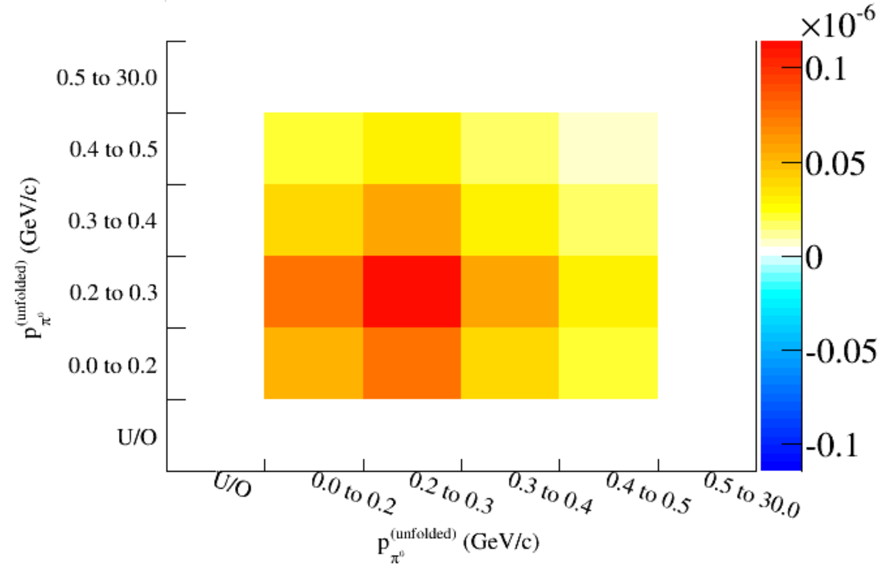

Figure C18: The momentum covariance matrix for all Highland errors for the NC1π^0 BrlBrl topology.

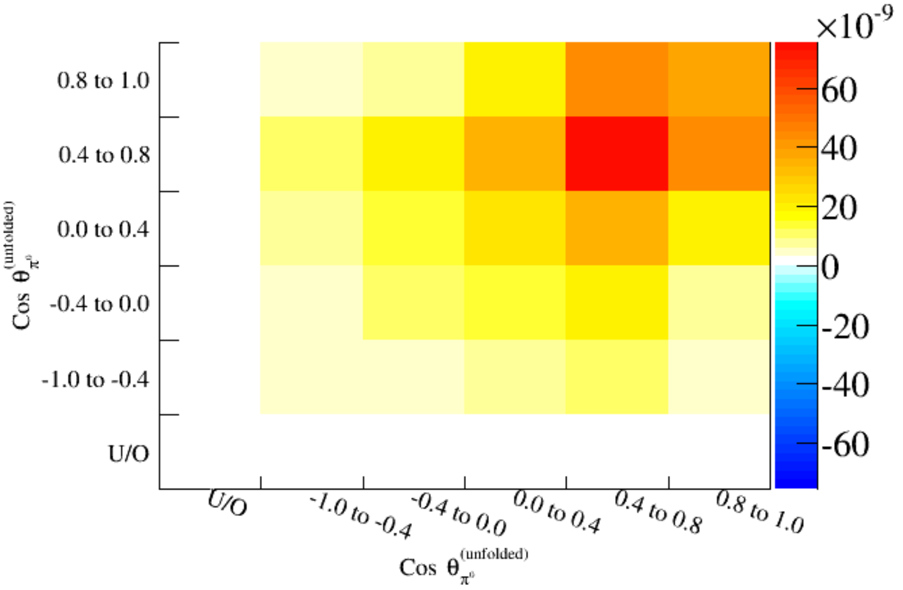

Figure C19: The cosθ covariance matrix for all errors combined for the NC1π^0 BrlBrl measurement.

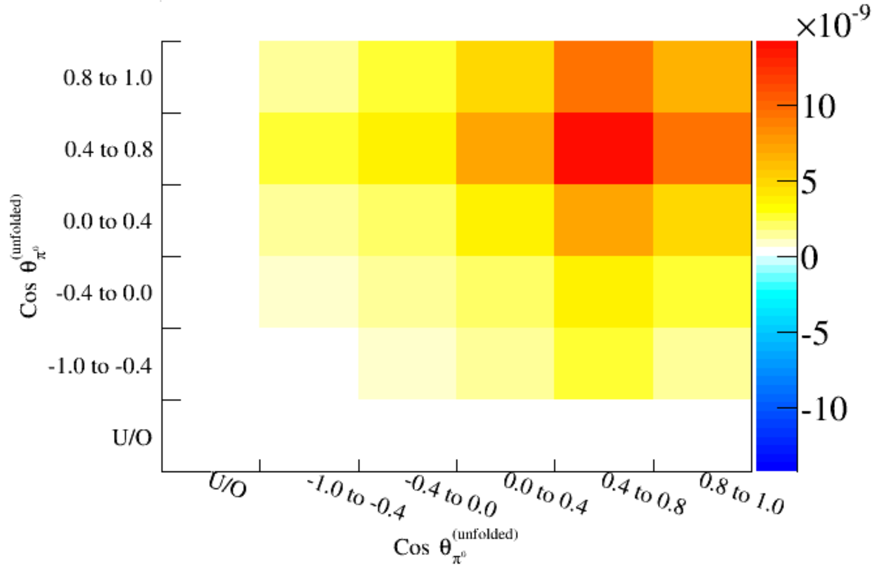

Figure C20: The cosθ covariance matrix for all Highland errors for the NC1π^0 BrlBrl topology.