Measurement of Neutral-Current π^0 production for ν_μ interactions in ND280

Zachary J. Williamson

St. Cross College, Oxford

Thesis submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2014

Measurement of Neutral-Current π^0 production for ν_μ interactions in ND280

Zachary J. Williamson
St. Cross College, Oxford

Thesis submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2014

Abstract

This thesis presents a study of neutral-current π^0 production from ν_μ interactions. The study uses data taken at the Tokai-to-Kamiokande off-axis experiment's near detector (ND280). Monte Carlo studies were used to develop selection criteria to reconstruct event signatures produced from such interactions. This physics analysis focuses on the tracker region of the ND280 subdetector. Other π^0 analyses using different detector regions are either underway or completed. The physics analysis, after being developed from Monte Carlo studies, was then run on real data.

This analysis goes partway towards measuring the neutral-current π^0-producing neutrino interaction cross-section, by measuring the rate of such interactions in the T2K experiment's near detector, ND280. The observed results were:

$$N_{\text{Data}}^{NC\pi^0} = 294.4 \pm 32.2_{\text{stat}} \pm 173.3_{\text{syst}} \quad (1)$$

This rate measurement, when made as a ratio of the predicted Monte Carlo rate, yields a ratio of:

$$\frac{R_{\text{Data}}}{R_{\text{MC}}} = 1.20 \pm 0.13_{\text{stat}} \pm 0.72_{\text{syst}} \quad (2)$$

Declaration

I declare that this thesis and the work presented in it are my own and were produced by me as the result of my own original research. The work was done while a candidate for a degree at the University of Oxford, and has not been submitted for any other qualification. Results and figures from published works have been clearly attributed.

I developed the analysis presented in chapter 4 and was not based on prior work, however the underlying ND280 reconstruction software was developed by others. When calculating cross-section systematic errors in chapter 8, a software package called T2KReWeight was used, which was developed by other collaboration members.

Acknowledgements

I would like to thank my father, Peter, and my brother, Ben, for being so supportive throughout the duration of my thesis.

I would like to thank my partner, Suzie, for always being there for me regardless of the quality of my companionship. I would also like to thank my good friend, Bernard; even in my darkest moments you never once questioned my sanity.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis documents research based at the T2K long-baseline neutrino oscillation experiment. The Tokai-to-Kamioka (T2K) experiment produces a beam of muon neutrinos, which is aimed towards two detectors: a near detector located 280 metres from the beam source and a far detector at 295km from the beam source (the Super-Kamiokande water-Cherenkov detector).

The research goal of this thesis was to use the near detector complex to measure the cross-section for a subset of neutral-current neutrino interactions. Particularly, neutral-current muon neutrino interactions which produce at least one π^0 particle. The two decay photons from the π^0 particle must convert in the tracker region of the near detector, which is described in chapter 3.

When measuring electron-neutrino interactions in the T2K's far detector, neutral-current π^0-producing interactions from muon neutrino interactions represent a potentially large background. The π^0 particle decays into two photons which then emit electrons through Cherenkov radiation. The key event signature of an electron-neutrino electromagnetic scattering interaction is the emission of an electron and the Cherenkov radiation from photons produced by a π^0 decay

can mimic this signal. If the π^0 decays in such a way that one photon has such low momentum that its interaction products emit little Cherenkov radiation, then the interaction is indistinguishable from the event signature of a charged-current quasi-elastic scattering interaction produced from an electron neutrino.

Measuring a cross-section for this event type allows the T2K experiment to more accurately estimate how many of these neutral-current background events one should expect, which will lead to more accurate electron neutrino appearance measurements.

1.1 Overview of Thesis

This section documents the structure of this thesis. A description of neutrino physics and the interaction types neutrinos undergo are detailed in chapter 2. Section 2 also details a brief history of experimental neutrino physics, as well as the key experiments which have led to advancements in the knowledge of neutrino interactions and neutrino oscillations.

A description of the T2K experiment is described in chapter 3. This chapter details the muon-neutrino beam and how muon neutrinos are generated. The chapter also gives an overview of the far detector and the near detector, including the two near detectors ND280 and INGRID.

Chapters 4, 5 and 6 describe in depth the physics analysis, developed by the author, used to investigate neutral-current π^0-producing neutrino interactions in the ND280 near detector. The chapter documents the construction of the analysis, as well as testing and validating methods used to reduce backgrounds whilst maintaining high statistics. This chapter deals exclusively with Monte-

Carlo generated information designed to simulate the performance of the ND280 detector and output of the T2K muon-neutrino beam.

Chapter 7 compares the selected events from the analysis for both data and Monte-Carlo samples to investigate any differences between the data and Monte-Carlo.

Chapter 8 discusses systematic errors associated with the analysis. There are three main classes of systematic errors to be considered: Errors associated with uncertainties in the muon-neutrino beam flux, errors associated with cross-section uncertainties in the Monte-Carlo and errors associated with uncertainties in the properties of the ND280 detector.

Chapter 9 documents the results of the analysis. Finally, chapter 10 discusses the value of the results in chapter 9, and further work that can be done to improve the analysis.

1.1.1 Summary of experimental NCπ^0 results

Experimental knowledge about the neutral-current π^0-producing neutrino interaction cross-section is limited, due to the difficulty in measuring this interaction's event signature. Currently, the best results are given by the MiniBooNE experiment, which measured ν_μ and $\bar{\nu}_\mu$ induced neutral current single π^0 production on mineral oil [16]. Their results give a total cross section of $(4.76 \pm 0.05_{\text{stat}} \pm 0.76_{\text{syst}}) \times 10^{-40}$ cm^2/nucleon at a mean neutrino energy of 808MeV. A measurement was also made at a mean neutrino energy of 664MeV, which gave a total cross section of $(1.48 \pm 0.05_{\text{stat}} \pm 0.23_{\text{syst}}) \times 10^{-40}$ cm^2/nucleon [16].

Source	Target	NC/CC Ratio	Value	Ref
ANL	H_2	$\sigma(\nu_\mu p \to \nu_\mu p\pi^0)/\sigma(\nu_\mu p \to \mu^- p\pi^+)$	0.51±0.25	[11]
ANL	H_2	$\sigma(\nu_\mu p \to \nu_\mu p\pi^0)/\sigma(\nu_\mu p \to \mu^- p\pi^+)$	0.09±0.05	[12]

Table 1.1: Measurement of neutral-current π^0 production as a ratio of charged-current π^+ production [11][12].

Measurements of the neutral-current π^0-production cross section have also been made as a ratio of charged-current π^0-production and charged-current π^+ production, by the Argonne National Laboratory in 1981 [12], the Brookhaven National Laboratory [17] [18] and the Gargamelle experiment at CERN [19]. However, these measurements vary between each other by up to a factor of 3 [20]. The two Argonne National Laboratory measurements are shown in table 1.1.

1.1.2 Summary of Experimental Measurements of θ_{13}

The primary physics goal of the T2K experiment is to accurately measure the neutrino oscillation parameter θ_{13}. Neutrino oscillations are described in section 2.2.1. This measurement requires measuring the appearence of electrons from ν_e scattering interactions in the Super-Kamiokande water Cherenkov detector 3.2. π^0 production from neutral current muon neutrino interactions is a key source of background and understanding the cross section of this interaction would improve the accuracy of future measurements of θ_{13}.

In 2011, T2K published its first measurement results, which suggested θ_{13} was not zero, to an accuracy of 2.5σ [21]. In 2013 this result was updated to give a confidence of 7.3σ and measured θ_{13} to be $\sin^2 2\theta_{13} = 0.140^{+0.003}_{-0.032}$ [22].

The first non-zero measurement of θ_{13} whose accuracy exceeded 3σ was made by the Daya Bay collaboration in 2012 [23]. Their current most accurate mea-

surement of θ_{13}, made in 2014, is $\sin^2(2\theta_{13}) = 0.083 \pm 0.018$ [24]. Further discussion of key neutrino experimental results is in section 2.3.

Chapter 2

Neutrino Oscillation Theory

This chapter provides an overview of the theory of neutrino oscillations and milestone experimental measurements of neutrino properties, as well as the latest measurements of key parameters such as the neutrino oscillation parameter θ_{13} 2.3.3. In order to further the development of neutrino physics it is necessary to have a detailed understanding of potential sources of background in order to reduce the uncertainty of experimental measurements. Specifically in the context of this thesis, neutral-current π^0-producing muon neutrino interactions can mimic charged-current electron neutrino interactions in water-Cherenkov detectors such as kamiokande 2.3.1 and super-kamiokande 2.3.1. The effect of π^0 production in water-Cherenkov detectors was further discussed in chapter 1.

2.1 Neutrino Theory: The Origin of Neutrino Physics

The neutrino was first postulated in 1936, by Pauli, to explain beta decay. Beta decay, as it was understood at the time, appeared to violate conservation of energy because the neutrino was not observable. Pauli hypothesised the existence of the neutrino to explain this paradox in a letter to his colleagues.

The existence of this hypothetical particle was resolved in 1953, when Fred Reines, Clyde Cowan, H. Kruse and A. McGuire detected positrons produced by inverse beta decay: $\bar{\nu}_e + p \rightarrow n + e^+$. Their source of neutrinos was a nuclear reactor at the Savannah River Plant in South Carolina [25].

This chapter discusses the theory behind neutrino interactions and the experiments whose measurements advanced our knowledge of neutrino physics.

2.1.1 Electroweak Interactions

What follows is a description of the theory of electroweak interactions. Electroweak interactions were first hypothesised in the 1960s by S. Glashow [26], S Weinberg [27] and B Salam [28], although each developed their own version of the theory independently. They were jointly honoured with the Nobel Prize in Physics for their work on electroweak interactions in 1979[29].

The observed results of μ decay and π^- decay led to the first theories of the weak interaction by Enrico Fermi in 1933 [30]. Neither of these decays can be explained through strong or electromagnetic interactions because of their unusual longevity (the μ has a half-life of 2.6×10^{-8} seconds, the π^- 2.2×10^{-6} seconds).

The weak interaction can be described as an analogue to the electromagnetic interaction. The electromagnetic scattering amplitude is a combination of the electromagnetic currents of the two interacting particles and a propagator term to represent the exchanged photon:

$$M = (\overline{u}_p \gamma^\mu u_p) \frac{-e^2}{q^2} (-\overline{u}_e \gamma_\mu u_e) \tag{2.1}$$

Fermi hypothesised that the amplitude for β decay, $n \rightarrow pe^-\overline{\nu}_e$, was given

by:

$$M = G_F(\overline{u}_n\gamma^\mu u_p)(-\overline{u}_{\nu_e}\gamma_\mu u_e) \tag{2.2}$$

G_F is the weak coupling constant. The formalism in equation 2.2 lacks a propagator term as seen in equation 2.1, and therefore does not really provide any physical insight into what mediates the weak interaction. It also implies that the weak interaction is inherently charge-raising or charge-lowering (neutral current weak interactions were not discovered until 1973) [31].

Further insight into the weak interaction can be acquired by examining the weak coupling constant G and whether it is truly a constant. By comparison to the electromagnetic interaction, G replaces the propagator $\frac{e^2}{q^2}$. The electromagnetic propagator is dimensionless, so G has units of GeV^{-2}. In order for equation 2.2 to represent a probability amplitude, it must overall be dimensionless. This suggests that G contains within it a propagator term where the associated vector boson is massive (the full weak interaction propagator is given in equation 2.6).

Symmetry arguments can be used to deduce the properties of the vector bosons associated with the weak interactions. Specifically, one can construct fermion doublets which are invariant under the weak interaction:

$$\begin{pmatrix} u \\ d \end{pmatrix} \begin{pmatrix} e \\ \nu_e \end{pmatrix} \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix} \tag{2.3}$$

The particle names in the doublets represent Dirac spinors. Group theory can be used to relate symmetries to conserved quantities, by linking these conserved quantities to experimental observables the underlying theory can be val-

idated or rejected.

Because these doublets are invariant under the weak interaction one can propose that the fermion in each doublet possesses a quantity that is conserved under the weak interaction. This symmetry is similar to the conservation of spin-angular momentum and a SU(2) symmetry can be applied in this context.

Under a 2x2 unitary rotation matrix the fermion doublets are invariant, which agrees with experimental observations. A U(2) formalism, however, treats the phase of the Dirac spinors as an extra degree of freedom. In this context a choice of phase is irrelevant. Removing this degree of freedom leaves us with the special unitary group SU(2).

There are three generators associated with the SU(2) group; the three Dirac matrices.

$$\tau_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tau_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \tau_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{2.4}$$

The generating set of a group is a subset, from which any member of the group can be reconstructed by using a finite combination of generating group members and their inverses. This leads one to postulate that for each generating group there exists a physical gauge boson which mediates the weak interaction.

The probability amplitudes of the weak interaction, acting upon the fermion doublets, is then given by:

$$(\overline{\nu}_\mu, \overline{\mu})(g\overrightarrow{W} \cdot \frac{\overrightarrow{\tau}}{2}) \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix} = (\overline{\nu}_\mu, \overline{\mu})\frac{g}{2} \begin{pmatrix} W_3 & W_1 - iW_2 \\ W_1 + iW_2 & -W_3 \end{pmatrix} \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix} \tag{2.5}$$

Using the relations $W^+ = \frac{(W_1 - iW_2)}{\sqrt{2}}$ and $W^- = \frac{W_1 + iW_2}{\sqrt{2}}$, we can see that the charge raising and charge lowering properties of the weak interaction are

mediated by two charged W bosons. The propagator in the weak interaction is contained within the Fermi Constant, and is given by equation 2.6. Due to the q^2 term this is only constant at low energies.

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8(m_W^2 - q^2)} \tag{2.6}$$

Each weak interaction requires a virtual W-boson, and the subsequent probability amplitude is suppressed by the W boson mass, which explains its relative weakness compared to the electromagnetic and strong interactions.

A problem with the formalism as it is currently presented in this chapter is that we know from experiment that the charged-current weak interaction only couples with left-handed fermions.

Parity Violation in the Weak Interaction

In 1956, T.D. Lee and C.N. Yang [32] examined experimental data on K-meson decays to examine parity conservation in the weak interaction. K^+ mesons have two decay channels of interest, $K^+ \rightarrow 2\pi$ and $K^+ \rightarrow 3\pi$. The two-pion and three-pion final states have different parities and so previously it was postulated that the K^+ meson was in fact two different particles, as otherwise the K^+ meson decay channels violate parity conservation. Lee and Yang were the first to propose that the two pion and three pion final states came from the same particle, and experiments began in earnest to verify this.

Later in 1956, Chien-Shiung Wu discovered parity violation in the weak interaction by examining Cobalt-60 decays [33]. Cobalt-60 undergoes beta decay to produce:

$$^{60}_{27}Co \rightarrow ^{60}_{28}Ni + e^- + \overline{\nu}_e + 2\gamma \tag{2.7}$$

If parity is conserved in the weak interaction, then the emitted electron

should not have a directional preference along the Cobalt's axis of spin. By using a magnetic field to align the spin of the Cobalt-60 atoms, Wu measured the directionality of the emitted electrons. She then inverted the magnetic field to change the Cobalt's spin-alignment, and repeated the experiment. The two emitted gamma–rays acted as a control in the experiment. As the electromagnetic interaction conserves parity, Wu could compare the directionality of the photons with the directionality of the electrons, and a significant discrepancy would prove parity violation.

She found that the electrons *did* have a directional preference and therefore parity is not conserved in the weak interaction.

The physical interpretation of this discovery is that the weak interaction only interacts with fermions that posses a left-handed chirality (combined spin and parity).

One can propose that the symmetry group laid out earlier in equation 2.3 only acts upon left-handed fermion doublets. Returning to the Fermi's weak probability amplitude, the vector component is replaced by a vector-axial-vector component.

However this formalism is still not yet a valid representation of reality. The third proposed boson, W_3, does not posses the same properties as any experimentally known gauge boson. W_3 does not have charge-raising or charge-lowering properties and so one can propose that it mediates a neutral-current weak interaction. However W_3 only couples to left-handed fermions, and we know from experiment that the neutral-current weak interactions couple to both left-handed and right-handed fermions.

SU(2) symmetry arguments alone cannot completely explain the weak interaction because of the non-physical W_3 boson. The solution is to construct a symmetry which is a combination of both weak isospin and another conserved quantity, electric charge. A spin-1 boson, B, is added to equation 2.5. The prob-

ability amplitude then becomes:

$$(\overline{\nu}_\mu, \overline{\mu})(g\overrightarrow{W} \cdot \frac{\overrightarrow{\tau}}{2} + \frac{g'B}{2}) \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix} = (\overline{\nu}_\mu, \overline{\mu}) \frac{g}{2} \begin{pmatrix} W_3 + \frac{Bg'}{g} & W_1 - iW_2 \\ W_1 + iW_2 & -W_3 + \frac{Bg'}{g} \end{pmatrix} \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix} \tag{2.8}$$

The field W_3 does not represent a physical state, but W_3 and the new spin-1 field B can mix to create observable states, represented by the bosons Z and A, shown in equation 2.9, where g and g′ are coupling constants.

$$\begin{aligned} \sqrt{g^2 + g'^2} Z &= gW_3 - g'B \\ \sqrt{g^2 + g'^2} A &= g'W_3 + gB \end{aligned} \tag{2.9}$$

Using the expression in equation 2.10, we can rewrite equation 2.9 as equation 2.11:

$$\theta_W = sin^{-1} \frac{g'}{\sqrt{g'^2 + g^2}} \tag{2.10}$$

$$\begin{aligned} Z &= W_3 \cos\theta_W - B \sin\theta_W \\ A &= W_3 \sin\theta_W + B \cos\theta_W \end{aligned} \tag{2.11}$$

The Z-boson mediates the neutral current weak interaction and A represents the electromagnetic photon. The overall coupling strength of the electroweak interaction must be equal to e, therefore g and g′ are related by the relation $g \sin(\theta_W) = g' \cos(\theta_W) = e$. θ_W is the weak mixing angle.

This weak hypercharge symmetry unifies both the weak and electromagnetic interaction into the electroweak interaction and accurately explains experimental observations [34]. We can use this formalism to construct a picture

of neutrino interactions.

2.1.2 Neutrino Interactions

This subsection details the interactions that neutrinos can undergo in T2K (approximate energy 600 MeV). Neutrinos interact only with the weak force and gravity (which is ignored due to its relative strength). Neutrino interactions are broadly classed as charged-current (mediated via a W boson) or neutral-current (mediated via a Z boson). Only neutrino interactions with up and down quarks in protons and neutrons are considered, because the overwhelming quantity of interacting matter in neutrino experiments are composed from these two quark flavours. Neutrino-electron scattering is not considered due to the interaction rate being several orders of magnitude lower than neutrino-quark interactions.

Charged Current

The most common neutrino interaction type is charged-current quasi-elastic scattering. Figure 2.1 shows a Feynman diagram of the interaction; a muon neutrino scatters off of a neutron, generating a proton and a muon. The original neutron which the neutrino scatters off of is converted to a proton. There is also enough energy in the centre-of-mass frame to produce resonances, such as a Δ^{++}. Figure 2.2 shows a Feynman diagram of such a resonance. These resonances decay predominantly into protons/neutrons and pions.

Charged-current neutrino interactions can also directly produce pions, however, the branching ratio of such interactions is reduced at the low energies of the T2K neutrino beam, as it requires quark-antiquark pair production. Charged-current muon neutrino interactions can only produce neutral and positively-charged pions due to charge and lepton number conservation. Anti-neutrinos can produce neutral and negatively-charged pions. Figure 2.3 shows a Feyn-

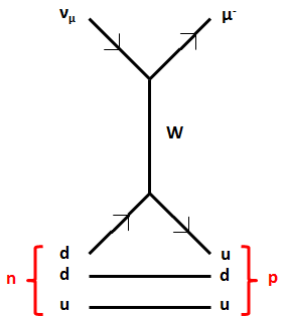

Figure 2.1: Charged-current quasi-elastic neutrino scattering.

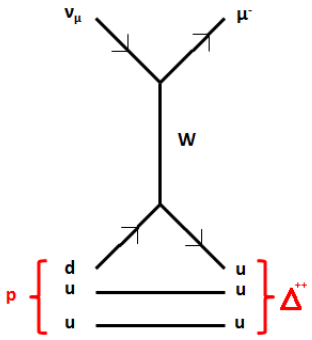

Figure 2.2: Charged-current quasi-elastic neutrino scattering where a resonance is produced.

man diagram of charged-current π^+-producing interaction, and figure 2.4 shows a Feynman diagram of a charged-current π^0-producing interaction.

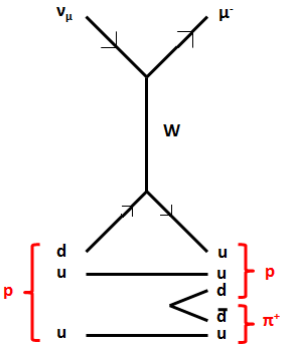

Figure 2.3: Charged-current neutrino interaction where a π^+ is produced.

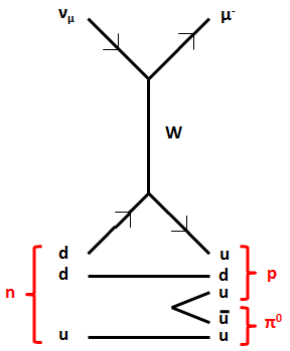

Figure 2.4: Charged-current neutrino interaction where a π^0 is produced.

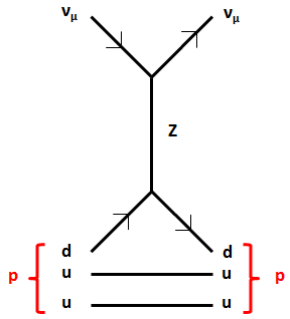

Figure 2.5: Neutral-current elastic neutrino scattering.

Neutral Current

Figure 2.5 shows a neutral-current elastic scattering interaction. The initial and final particle states are identical, making this interaction type all but impossible to detect. However it is possible for such an interaction to produce a resonance, where a proton with J=1/2 is converted to a Δ^+ with J=3/2. The resonance will then decay, in this case to a π^0 and a proton, or a π^+ and a neutron.

These resonant neutral-current interactions contribute to the overall neutral-current π^0-producing cross-section, from a Δ^+ or Δ^0 decay.

Another way pions can be produced in neutral-current interactions is when neutrino scattering is coupled with quark-antiquark pair production. Figures 2.6, 2.7, 2.8 show such interactions for π^+ creation, π^0 creation and π^- creation.

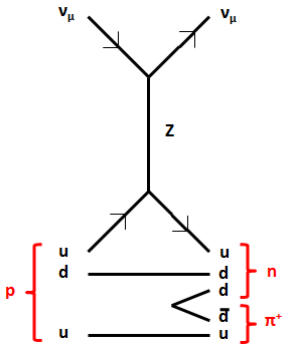

Figure 2.6: Neutral-current neutrino interaction where a π^+ is produced.

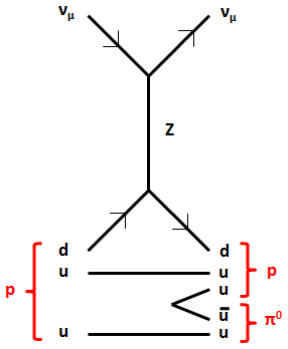

Figure 2.7: Neutral-current neutrino interaction where a π^0 is produced.

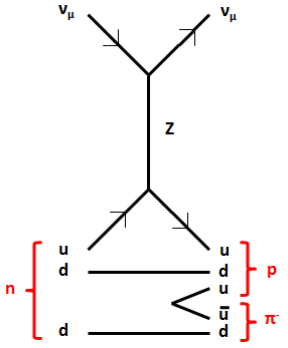

Figure 2.8: Charged-current neutrino interaction where a π^- is produced.

2.1.3 Neutrino mass under the Standard Model

In 1958 M. Goldhaber, L. Grodzins and A. W. Sunyar discovered that neutrinos are exclusively left-handed through experimental observation of Europium-152 decays [35].

A major consequence of left-handed neutrinos is that they *cannot* couple to the Higgs Boson, as the Higgs field can only interact with fermions that posses both left and right-handed helicities [36].

2.2 Beyond the Standard Model: Neutrino Mass and Neutrino Oscillations

The Standard Model of Particle Physics, and in particular the massless neutrino, did not completely agree with experimental observations dating as far back as the 1960s. In 1965, R. Davis constructed the Homestake experiment [37], designed to measure the solar neutrino flux. The experiment ran continuously from 1970 to 1994, although results were first published in 1968 [38].

The main component of the experiment was a tank of Tetrachloroethylene, located 4.2 km underground. Figure 2.9 shows the layout of the experiment. Electron neutrinos could interact with the chlorine to produce argon:

$$\nu_e + {}^{37}Cl \rightarrow {}^{37}Ar + e^- \tag{2.12}$$

The experiment then measured the amount of produced argon and compared it with expected results. The experiment recorded significant deficit of argon atoms, and Davis concluded that he was detecting only a third of the expected solar neutrino flux. Today, we understand that the reason behind this is that the majority of the solar neutrinos oscillate into muon neutrinos and tau neutrinos, both of which his experiment could not detect.

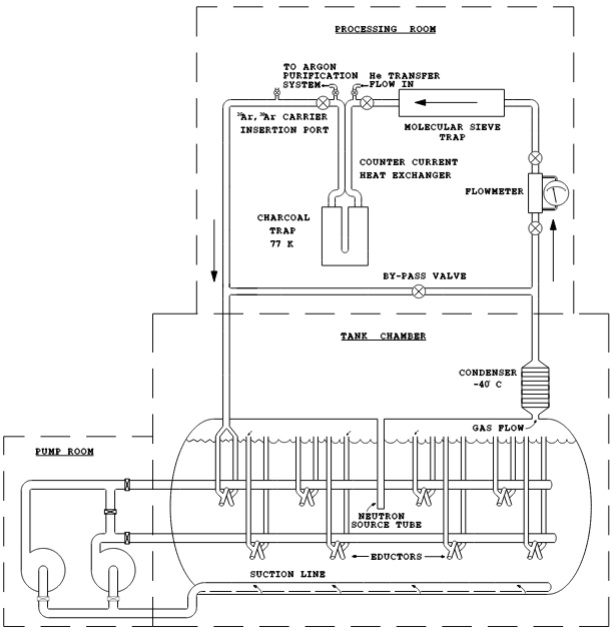

Figure 2.9: Overview of the Homestake experiment

The detector required the ability to detect a very small number of expected argon atoms; daily measurements averaged at 0.478 ± 0.030 (stat) ± 0.029 (syst) argon atoms per day [38]. For Homestake's discovery to be confirmed, the solar neutrino flux would need to be measured by another detector utilising a different detector technology, to rule out systematic effects. This would not happen until 2001, when the Sudbury Neutrino Observatory conclusively discovered a deficit in the solar neutrino flux [39].

2.2.1 Neutrino Mass and Neutrino Oscillations

The consequence of having massive neutrinos, namely neutrino oscillations, was first hypothesised by Pontecorvo in 1957 [40]. There are two main theories behind massive neutrinos. One hypothesises that neutrinos and anti-neutrinos are their own antiparticle. This theory was first advocated by M. Majorana in 1937 [41]. The second theory invokes the concept of the right-handed neutrino. Such a neutrino is termed a Dirac neutrino; if left and right-handed neutrinos

exist in nature, then neutrinos can couple with the Higgs field, generating a neutrino mass term.

Under the Standard Model, the Lagrangian that governs Higgs-lepton coupling is given by:

$$\mathcal{L}_{H,L} = -(\frac{\nu + H}{\sqrt{2}})[\overline{\boldsymbol{l}'_L} Y'^{l} \mathbf{l}'_R] + H.c. \tag{2.13}$$

Where $\boldsymbol{l}'_L = \begin{pmatrix} l'_{eL} \\ l'_{\mu L} \\ l'_{\tau L} \end{pmatrix}$, where each term is the left-handed Dirac spinor for each lepton flavour. $\boldsymbol{l}'_R$ and $\boldsymbol{\nu}'_L$ serve identical functions for their lepton counterparts. Y'^{l} is the matrix of Yukuwa couplings, which couples fermion fields to the scalar Higgs field.

Adding the concept of right-handed neutrinos adds a second term to represent left and right-handed neutrinos interacting with a new set of neutrino Yukawa couplings:

$$\mathcal{L}_{H,L} = -(\frac{\nu + H}{\sqrt{2}})[\overline{\boldsymbol{l}'_L} Y'^{l} \mathbf{l}'_R + \overline{\boldsymbol{\nu}'_L} Y'^{\nu} \boldsymbol{\nu}'_R] + H.c. \tag{2.14}$$

Y'^{ν} is not a diagonal matrix, however it can be diagonalised by multiplying it by unitary matrices. The resulting eigenvalues in the diagonal matrix are the physical mass terms of particles which couple to the Higgs field.

By defining unitary matrices $V_L^{\nu\dagger}$ and V_R^{ν}, we can diagonalise Y'^{ν}: $V_L^{\nu\dagger} Y'^{\nu} V_R^{\nu} = Y^{\nu}$, where $Y^{\nu}_{kj}(k, j = 1, 2, 3)$.

Using these unitary matrices, $V_L^{\nu\dagger}$ and V_R^{ν}, it is now possible to define massive neutrino arrays for left and right-handed neutrinos:

$$\boldsymbol{n}_L = V_L^{\nu\dagger}\boldsymbol{\nu}_L^{'} = \begin{pmatrix} \nu_{1L} \\ \nu_{2L} \\ \nu_{3L} \end{pmatrix} \tag{2.15}$$

The resulting neutrino mass of the Dirac neutrino is given by equation 2.16:

$$m_{1,2,3} = \frac{y_{1,2,3}^{\nu}\nu}{\sqrt{2}} \tag{2.16}$$

What is of considerable consequence is that, similarly to quarks, the mass eigenstates do not map onto the flavour eigenstates. To see the relationship between the two, the weak charged current can be described by:

$$j_{W,L}^{p} = 2\overline{\boldsymbol{\nu}_L^{'}}\gamma^{\rho}\boldsymbol{l}_L^{'} = 2\overline{\boldsymbol{n}_L}V_{L\dagger}^{\nu}V_L^{l}\gamma^{\rho}\boldsymbol{l}_L \tag{2.17}$$

The current mixes mass eigenstates $\overline{\boldsymbol{n}_L}$ and flavour eigenstates $\boldsymbol{l}_L^{'}$. The current is dependent on the unitary matrix:

$$U = V_L^{l\dagger}V_L^{\nu} \tag{2.18}$$

This matrix is similar to the quark Cabibbo-Kobayashi-Masakawa matrix [42][43].

2.2.2 Neutrino Oscillation Formulae

From section 2.2.1, neutrino flavour states can be described by a unitary oscillation matrix acting on a vector of neutrino mass states:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.19}$$

Or more concisely,

$$|\nu_\alpha >= \sum_k U_{\alpha,k}|\nu_k > \tag{2.20}$$

The neutrino mass states propagation through space is described by the Shrödinger equation:

$$i\frac{\partial}{\partial t}|\nu_k(x,t) >= -\frac{1}{2m_k}\frac{\partial^2}{\partial x^2}|\nu_k(x,t) > \tag{2.21}$$

The solution implies that the neutrino mass states are described by plane waves:

$$|\nu_k(t) >= e^{-iE_k t}|\nu_k > \tag{2.22}$$

Equation 2.22 can be used to describe the time-evolution of flavour states:

$$|\nu_\alpha(t) >= \sum_k U_{\alpha,k}e^{-iE_k t}|\nu_k > \tag{2.23}$$

Since mass states are not directly observable, it is desirable to describe equation 2.23 in terms of flavour states only. Because $U_{\alpha k}$ is a unitary matrix by definition, it can be inverted to describe the evolution of mass states in terms of flavour states:

$$|\nu_k >= \sum_\alpha U^*_{\alpha k}|\nu_\alpha > \tag{2.24}$$

Merging equation 2.24 and equation 2.23 yields the time evolution of neutrino flavour states without any neutrino mass state terms:

$$|\nu_\alpha(t) >= \sum_{\beta=e,\mu,\tau}(\sum_k U_{\alpha k}e^{-iE_k t}U_{\beta k})|\nu_\beta > \tag{2.25}$$

The transition probability between flavour states α and β is given by the

square of $< \nu_\beta | \nu_\alpha >$:

$$P_{\nu_\alpha \to \nu_\beta} = | < \nu_\beta | \nu_\alpha > |^2 = (\sum_k U_{\alpha k} e^{-iE_k t} U_{\beta k})^2 = \sum_{k,j} U_{\alpha k} U^*_{\beta k} U^*_{\alpha j} U_{\beta j} e^{-i(E_k - E_j)t} \tag{2.26}$$

The energy eigenvalues, E_k and E_j, are given by $E_k = \sqrt{\overrightarrow{\boldsymbol{p}}^2 + m_k^2}$. It is safe to assume that, for neutrinos, $v \approx c$ and therefore the energy values can be described by the following relation:

$$E_k = \sqrt{\overrightarrow{\boldsymbol{p}}^2 + m_k^2} = \overrightarrow{\boldsymbol{p}} \sqrt{1 + \frac{m_k^2}{\overrightarrow{\boldsymbol{p}}^2}} = \overrightarrow{\boldsymbol{p}} (1 + \frac{m_k^2}{2\overrightarrow{p}^2}) = E_k + \frac{m_k^2}{2E_k} \tag{2.27}$$

Since the difference between the energy of mass states is described entirely by their mass, which is small in comparison to their momentum, equation 2.27 can be further simplified to:

$$E_k = E + \frac{m_k^2}{2E} \tag{2.28}$$

The $E_k - E_j$ term in equation 2.26 can now be written as:

$$E_k - E_j = \frac{m_k^2}{2E} - \frac{m_j^2}{2E} = \frac{\Delta m_{kj}^2}{2E} \tag{2.29}$$

Equation 2.29 is not, however, described entirely by terms which are observable, namely t. What is known, however, is the distance between the neutrino source and the detector, L. Since neutrinos propagate at a speed which is almost equivalent to c, t is proportional to L. The final oscillation probability can be written as:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{13} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{-1\alpha_1} & 0 & 0 \\ 0 & e^{-i\alpha_2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.31}$$

Figure 2.10: PMNS Mixing Matrix. s and c refer to sine and cosine.

$$P_{\nu_\alpha \to \nu_\beta}(L,E) = |<\nu_\beta|\nu_\alpha>|^2 = (\sum_k U_{\alpha k} e^{-E_k t} U_{\beta k})^2 = \sum_{k,j} U_{\alpha k} U^*_{\beta k} U^*_{\alpha j} U_{\beta j} exp(-i\frac{\Delta m^2_{kj} L}{2E}) \tag{2.30}$$

The oscillation probability is now described entirely by the unitary mixing matrix U, and a phase factor $\phi_{kj} = \frac{\Delta m^2_{kj} L}{2E}$. The elements of U, as well as the mass-squared difference between flavour states, are fixed quantities. The values these quantities take can be determined by measuring flavour oscillation probabilities with respect to a known E and L.

2.2.3 PMNS Mixing Matrix

In 1962 a parametrisation of the unitary oscillation matrix from equation 2.30 was developed by Ziro Maki, Masami Nakagawa and Shoichi Sakata [44], which describes three-flavour neutrino oscillation. This chapter adopts the notation used by Giunti and Kim [45]. The parametrisation of this 3x3 mixing matrix is shown in figure 2.2.3.

This matrix is similar to the CKM matrix [43] in that the parametrisation can be described in terms of mixing angles, in this case θ_{12}, θ_{13} and θ_{23}. There is also an additional phase factor, δ in the parametrisation. The PMNS matrix is especially convenient because it splits the overall unitary matrix into three rotational matrices, which roughly correspond to different neutrino sources that are experimentally observable. Long-baseline beam neutrinos are used to measure θ_{13}, alongside reactor-based experiments. Beam neutrinos, as well as at-

mospheric neutrino measurements, are also used to measure θ_{23}. The mixing angle θ_{12} is measured through solar neutrino observations.

δ and CP violation

The phase factor term, δ, in the above parametrisation leads to the possibility of CP violation in the neutrino sector. A non-zero δ would imply that the mixing matrix for neutrinos has different values compared with the anti-neutrino mixing matrix. This would result in CP violation. As of writing, there is no compelling evidence to either support or reject a non-zero δ hypothesis.

Sterile Neutrinos

From looking at the cross-section of Z-boson decays at LEP, it is possible to determine the number of neutrino flavours that participate in the weak interaction.

The total decay width of the Z-boson is:

$$\Gamma_Z = \sum_{l=e,\mu,\tau} \Gamma_{Z\to l\bar{l}} + \sum_{q\neq t} \Gamma_{Z\to q\bar{q}} + \Gamma_{inv} \tag{2.32}$$

Γ_{inv} represents the decay width from neutrino production, which cannot be directly measured, but can be inferred by measuring the Z-boson total width Γ_z from the cross section as a function of beam energy measured from its visible constituents. Γ_{inv} can be written as $N_\nu \Gamma_{Z\to\nu\bar{\nu}}$, where N_ν is the number of neutrino flavours.

Data from the LEP experiment was collated and fitted against the number of neutrino flavours as a free parameter, with the experimental evidence overwhelmingly supporting a 3-flavour hypothesis, shown in Fig 2.11 [1].

The way more neutrinos flavours can be introduced, given this experimental evidence, is if new neutrinos which don't participate in the weak interaction

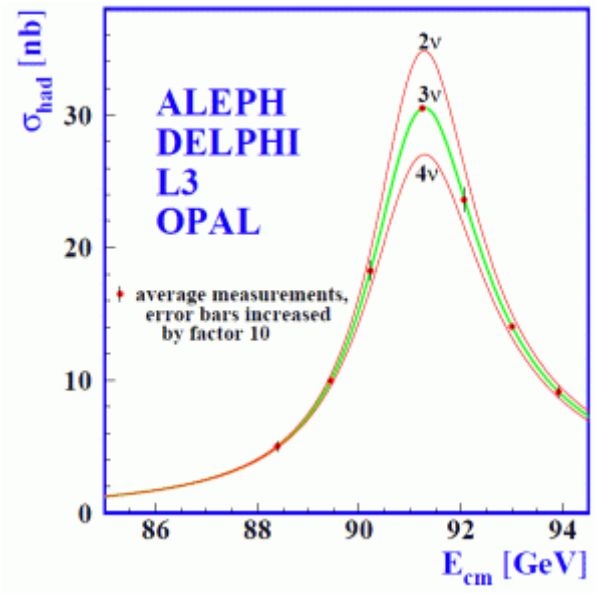

Figure 2.11: Hadron production cross-section at the Z resonance [1].

exist. These are referred to as sterile neutrinos.

2.3 Measuring neutrino properties

This section documents landmark experiments which measured important properties of neutrinos, including the parameters of the mixing matrix. The first experiment which discovered the solar neutrino problem, Homestake, has been discussed in section 2.2.

2.3.1 Solar Neutrino Experiments

Kamiokande

The Kamiokande detector was a water-Cherenkov detector initially built to investigate proton decay. It was located at the Kamioka mine in Japan, and was situated 1000 m underground to shield the detector from cosmic rays [46]. Aside from proton decay, one of Kamiokande's primary physics goals was to

investigate the solar neutrino problem by detecting solar neutrinos produced through the decay of boron in the sun through the fusion of beryllium and a proton (shown in equation 2.33). ^8B neutrinos were preferred due to their higher energies.

$$^8B = {}^8Be + e^+ + \nu_e \tag{2.33}$$

The Kamiokande experiment measured the solar neutrino flux and measured roughly half the amount of electron neutrinos predicted by solar neutrino models [46]. This confirmed the discoveries of the Homestake experiment and added further legitimacy to the solar neutrino problem.

In 1987 the Kamiokande experiment measured eleven neutrinos produced by Supernova 1987A, the first time astrophysical neutrinos had been definitively measured [47].

Super-Kamiokande

The Super-Kamiokande experiment is the successor to the Kamiokande experiment, and consists of a 50 kiloton water-Cherenkov detector located in the Kamioka mine. This Super-Kamiokande detector is used in the Tokai-to-Kamiokande experiment, and a more in-depth description of the detector is given in section 3.2.

The primary physics goal of the Super-Kamiokande experiment, in the context of neutrino physics, was to measure solar and atmospheric neutrinos and compare the measured fluxes with prior predictions, testing the theory of neutrino oscillations.

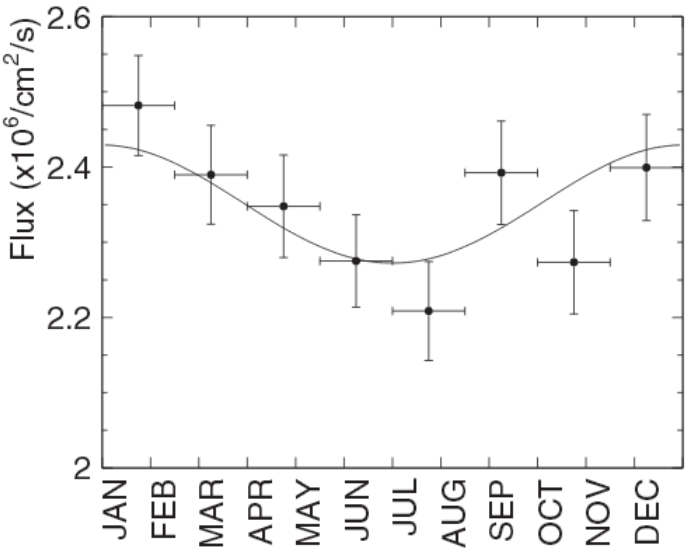

Figure 2.12: Seasonal variation in neutrino flux. Solid line is predictions based on the Earth's eccentricity, data points are from Super-Kamiokande data [2].

For solar neutrinos, the Super-Kamiokande experiment measured the solar neutrino flux and observed how the flux changed as a function of time. A seasonal change in neutrino flux was predicted due to the differing distances between the Super-Kamiokande detector and the Sun. An additional oscillation factor due to vacuum oscillations was hypothesised, but not detected [2]. The seasonal flux changes are shown in figure 2.12.

SNO

The Sudbury Neutrino Observatory is an underground heavy-water Cherenkov detector in Ontario, Canada (Fig 2.13) [3]. Its main component is a 12 m acrylic sphere which contains one kiloton of D_2O. Leptons produced by neutrino interactions on the deuterium can produce rings of Cherenkov radiation. The Cherenkov radiation is then recorded by Photo-multiplier Tubes which line the walls of the detector.

Three major neutrino interaction types can occur within SNO, charged-current, neutral-current and electron-scattering interactions.

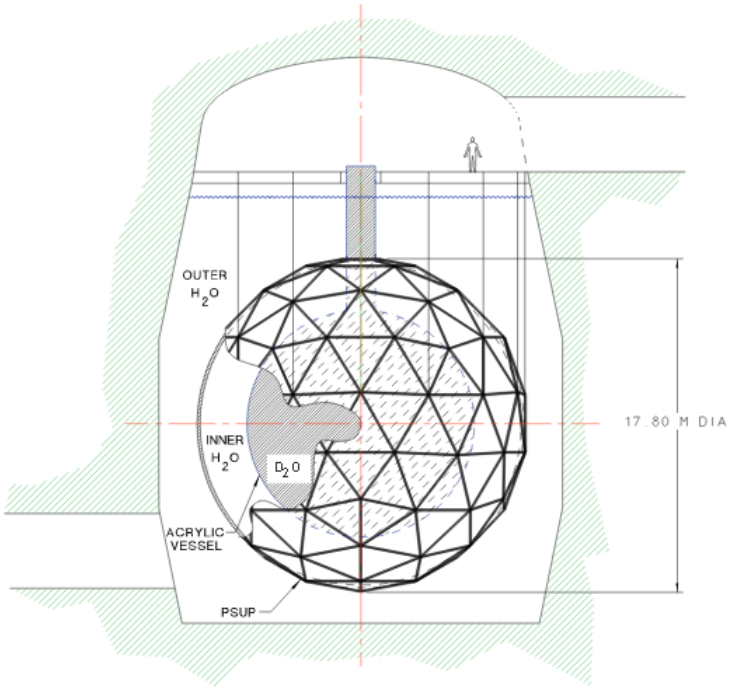

Figure 2.13: Overview of the SNO experiment. Figure taken from [3]

$$\text{Charged Current: } \nu_e + d \rightarrow 2p + e^- \tag{2.34}$$

$$\text{Neutral Current: } \nu_x + d \rightarrow p + n + \nu_x \tag{2.35}$$

$$\text{Electron Scattering: } \nu_x + e^- \rightarrow e^- + \nu_x \tag{2.36}$$

Neutral Current and Electron Scattering interactions are sensitive to all three neutrino flavour types, whilst Charged Current interactions can only be produced through electron neutrinos due to the limited centre-of-mass energy available. This enabled the experiment to discriminate between electron neutrinos and other types.

All neutrinos produced by the Sun are initially electron-flavoured. SNO not only had the ability to measure a deficit of electron neutrinos, but could also compare the electron-neutrino flux with an all-flavour neutrino flux.

The SNO experiment discovered, in 2001 [39], a 17 standard-deviation dis-

crepancy between the measured charged-current neutrino flux and both the neutral-current and electron-scattering neutrino flux. This proved beyond all reasonable doubt that neutrinos oscillate between flavour states and therefore that neutrinos have mass, and confirmed the previous evidence of neutrino oscillations put forward by the Super-Kamiokande experiment [4]. Super-Kamiokande's atmospheric neutrino oscillation result is discussed in section 2.3.2.

2.3.2 Atmospheric Neutrino Experiments

Kamiokande

The Kamiokande experiment has already been mentioned in section 2.3.1 in the context of solar neutrino observations. However, Kamiokande also performed pioneering work in the atmospheric neutrino sector.

Kamiokande was the first detector to uncover evidence behind the atmospheric neutrino anomaly. In 1988, the experiment found only 59% ($\pm$7%) of the ν_μ flux predicted by Monte Carlo simulations, whilst the ν_e flux matched the Monte Carlo predictions, leading to a net deficit of neutrinos [48].

This ν_μ disappearance was due to $\nu_\mu \rightarrow \nu_\tau$ oscillations, however the Kamiokande experiment could not distinguish between $\nu_\mu \rightarrow \nu_e$ and $\nu_\mu \rightarrow \nu_\tau$ oscillations. The CHOOZ experiment, mentioned in section 2.3.3 ruled out $\nu_\mu \rightarrow \nu_e$ oscillations [49], leading to the conclusion that the $\nu_\mu \rightarrow \nu_\tau$ mixing angle is non-zero. Figure 2.14 shows the allowed regions of mass/mixing-angle combinations using Kamiokande data and Super-Kamiokande data. The best fit puts $\sin^2 2\theta$ at 0.95.

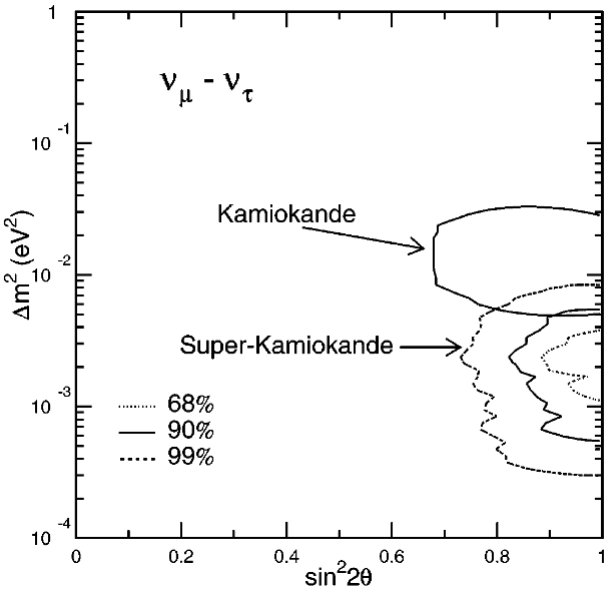

Figure 2.14: 68%, 90% and 99% confidence intervals for $\sin^2 2\theta$ and Δm^2 for $\nu_\mu \rightarrow \nu_\tau$ oscillations based on data from Super-Kamiokande. The 90% confidence interval for Kamiokande data is also shown. Figure taken from [4]

Super-Kamiokande

The Super-Kamiokande experiment has been described, in the context of solar neutrinos, in section 2.3.1.

In the context of atmospheric neutrinos, Super-Kamiokande's major contribution to neutrino physics was measuring the up-down ν_μ asymmetry. The up-down asymmetry refers to the ratio of neutrino flux for neutrinos traveling from the top of the Super-Kamiokande detector, compared with the flux for neutrinos traveling through the Earth and entering the bottom of the detector. Without neutrino oscillations, one would expect no up-down asymmetry. However, when considering neutrino oscillations, the baseline for up-neutrinos and down-neutrinos is different, leading to different oscillation rates and therefore a flux asymmetry. In 1998, Super-Kamiokande published the first definitive evidence pointing towards neutrino oscillations when it measured an up-down asymmetry of $A_\mu^{up-down} = -0.296 \pm 0.048 \pm 0.01$ [4].

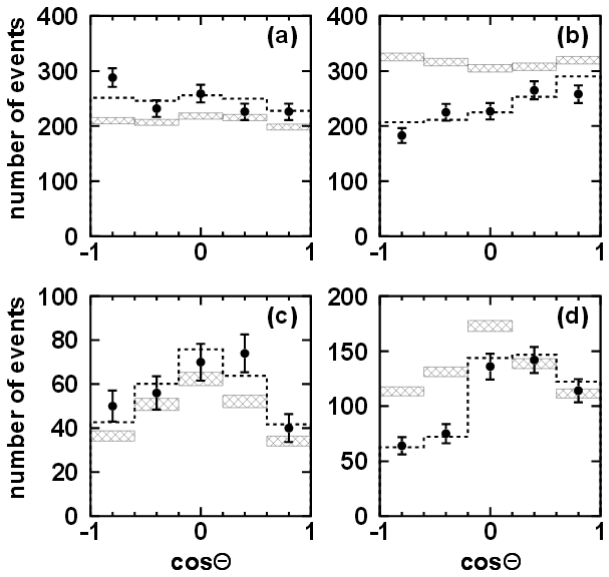

Figure 2.15: Zenith angle distributions of Super-Kamiokande events. cosθ = 1 means downwards-going particles. Dotted line represents Monte Carlo prediction if neutrino mixing is maximal. Shaded bar represents Monte Carlo prediction for the no-oscillations hypothesis [4].

Super-Kamiokande also measured the number of neutrino events relative to their zenith angle, shown in figure 2.15. A cosθ of 1 means a neutrino is entering from the top of the detector, a cosθ of -1 means a neutrino is entering from the bottom of the detector. The dotted line represents the Monte Carlo prediction for a maximal mixing angle, the shaded bar represents the Monte Carlo prediction for a no-oscillation hypothesis. The data clearly fits the oscillation hypothesis and not the no-oscillation hypothesis.

Soudan 2

The Soudan 2 experiment was located in the Soudan mine in Northern Minnesota, and was based around a 960 tonne iron tracking calorimeter [50]. Its primary physics goal was measuring proton decay, however it was also capable of measuring neutrino interactions. The tracking calorimeter could measure

the directionality of neutrino interactions, and therefore the Soudan 2 experiment was capable of measuring the up-down ν_μ asymmetry. Soudan 2 therefore could confirm Super-Kamiokande's result with a completely different type of detector. This type of confirmation is important because Soudan 2 had a different set of systematic errors compared to water-Chereknov detectors.

In 2005 the Soudan 2 detector published results showing that the muon-neutrino flux coming up through the Earth was lower than atmospheric models, whilst the downward flux matched predictions [50]. This result confirmed Super-Kamiokande's discovery of neutrino oscillations [4].

2.3.3 Reactor Neutrino Experiments

Reactor experiments provide valuable insight into neutrino physics as they are a very numerous source of anti-electron neutrinos, produced via β decays in nuclear material. The energy of these neutrinos originates from the binding energy of the fissile material in the reactor and is of the order of 1 MeV. This enables relatively short-baseline neutrino oscillation maxima to be studied. The drawback to reactor experiments is that the neutrinos are produced isotropically and therefore the neutrino flux drops off rapidly ($\frac{1}{r^2}$) with distance.

CHOOZ and Double-CHOOZ

The CHOOZ reactor experiment (pronounced as show) is a reactor experiment with a baseline of roughly 1 km. The atmospheric neutrino anomaly, a lower than expected muon-neutrino flux from atmospheric muons, had been previously identified by the Kamiokande experiment [51], discussed in section 2.3.2, and one hypothesis was that the reduced muon neutrino flux was due to $\nu_\mu \rightarrow \nu_e$ oscillations. The primary research goal of CHOOZ was to investigate this hypothesis, by determining the oscillation parameter θ_{13} by measuring

anti-electron-neutrino disappearance. [49].

The CHOOZ experiment was based around two reactors in the Ardennes region of France. CHOOZ's measurement of the ratio between the expected and observed anti-electron-neutrino flux was:

$$R = 1.01 \pm 0.023 \pm 0.027 \tag{2.37}$$

This result [49] agreed with the hypothesis that there were no oscillations at this distance scale, rejecting the $\nu_\mu \rightarrow \nu_e$ explanation for the atmospheric neutrino anomaly.

This result, combined with the result of Kamiokande, enabled finer limits to be placed on both θ_{23} and the mass hierarchy.

Double Chooz was the successor to CHOOZ and is located in the same laboratory space[52]. Its primary physics goal was to measure the oscillation parameter θ_{13} through anti-electron-neutrino disappearance, with a greater sensitivity than CHOOZ. In 2011, Double Chooz published its disappearance measurement [53]. It hinted towards a non-zero value of θ_{13}, although with only a statistical sensitivity of 1-sigma further measurements were required to make a definitive judgment. Their most recent measurement of θ_{13} is $\sin^2 2\theta_{13} = 0.090^{+0.032}_{-0.029}$ [54].

KamLAND

The Kamioka Liquid scintillator AntiNeutrino Detector, KamLAND for short, is a very-long baseline reactor neutrino experiment. The average baseline is 180km, however the neutrino sources come from 53 nuclear reactors in Japan with distances ranging from 80 km to 800 km [55]. Its primary physics goal was

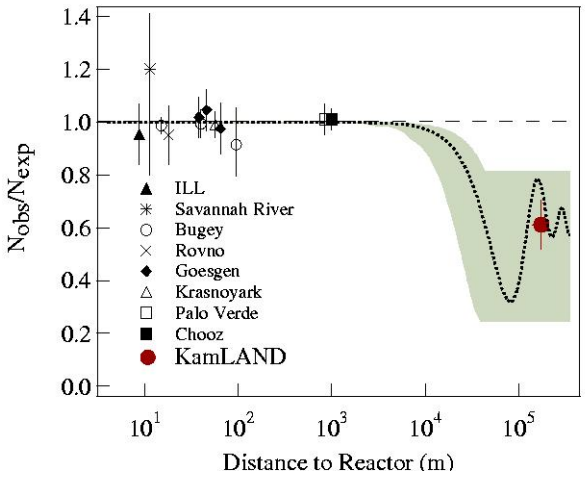

Figure 2.16: Ratio of measured to predicted ν_e flux as a function of source-detector distance. The dotted line represents the expected ratio if θ_{12} is large. Figure taken from the KamLAND collaboration. [5]

to check the large-mixing-angle hypothesis (LMA) to solving the solar neutrino problem. The LMA theory hypothesises that the mixing angle θ_{12} is large, causing the Sun's electron neutrinos to oscillate into muon neutrinos and tau neutrinos. Figure 2.16 shows the ratio of measured to expected $\bar{\nu_e}$ flux for reactor neutrino experiments as a function of their source-detector distance. Only experiments with a long baseline are sensitive to changes predicted by the LMA hypothesis.

KamLAND published, in 2004 [5], its measurement of the ratio between observed and expected $\bar{\nu_e}$ fluxes. This value deviated from unity by 5σ. These deviations also matched up with predictions made by the LMA model. This result effectively confirmed the large-mixing-angle theory as a valid solution to the solar neutrino problem.

Daya Bay

The Daya Bay neutrino experiment is located at the Daya Bay nuclear reactor facility in China. Its key research goal is to accurately measure the neutrino mixing angle θ_{13} by measuring the anti-electron neutrino flux output of six nuclear reactors, using eight separate neutrino detectors. Each detector consists of 20 tons of liquid scintillator and the $\bar{\nu}_e$ flux is measured through inverse-beta decay. The neutrino detectors are situated by choosing an L/E value where the value of θ_{13} is the greatest. In addition, systematic errors associated with the detectors are reduced by having identical near detectors to more accurately measure the neutrino flux near the reactors.

In 2012, the Daya Bay experiment published the first measurement of θ_{13} whose accuracy exceeded 3σ [23]. They measured θ_{13} to be:

$$\sin^2(2\theta_{13}) = 0.092 \pm 0.016\,(\text{stat}) \pm 0.005\,(\text{syst}) \tag{2.38}$$

Their latest published result is the most accurate measurement of θ_{13} to data [24]:

$$\sin^2(2\theta_{13}) = 0.083 \pm 0.018 \tag{2.39}$$

2.3.4 Long-Baseline Accelerator Neutrino Oscillation Experiments

Accelerator-based experiments have the advantage of being able to produce a relatively focused beam of neutrinos, enabling large neutrino fluxes at large distances. Accelerator experiments also have the ability to control the energy-range of the neutrinos produced. Together this enables the construction of experiments which have less constraints over which $\frac{L}{E}$ values are chosen.

Figure 2.17: Overview of the Daya Bay reactor experiment, showing three of eight detectors and the nuclear reactors. Image taken from Daya Bay's public website http://dayabay.ihep.ac.cn/docs/experiment.html

The disadvantages of accelerator neutrino experiments include the fact that short-baseline reactor experiments have a considerably larger flux than accelerator-based experiments. Accelerator-based experiments are also more expensive than reactor, solar or atmospheric neutrino experiments, as a neutrino beam must be built.

K2K

The KEK to Kamioka experiment (K2K) was located in Japan. It has a baseline of 250 km, and uses the Super-Kamiokande facility as its detector. Its primary physics goal was to measure atmospheric neutrino oscillations [56] via ν_μ disappearance measurements. The optimal baseline for such an experiment, using KEK's neutrino beam energy of 1.3 GeV, would have been 650 km, however the KEK facility and Super-Kamiokande detector were both pre-existing locations,

Figure 2.18: MINOS far detector. Picture taken from Fermilab: http://www-numi.fnal.gov .

so the 250 km distance was a necessary compromise.

The final results of the K2K experiment measured a net disappearance of muon neutrinos with an accuracy of 4.3 σ, with an associated mass squared difference between muon and tau neutrinos of $\Delta m^2 = 2.8 \times 10^{-3}\text{eV}^2$ [51].

MINOS

The Main Injector Neutrino Oscillation Search (MINOS) is a long baseline neutrino experiment whose detector is based in Northern Minnesota. The beam originates at Fermilab, and the baseline is 735km [57]. There is also a near detector based at Fermilab. The primary physics goals of MINOS was to measure θ_{23}, as well as the relative mass difference between the 2nd and 3rd neutrino mass states: Δm^2_{23}. Figure 2.18 shows a the MINOS far detector.

In 2011 MINOS published updated results, which found that $\theta_{23} > 0.90$ with 90% confidence, and that $\Delta m^2_{23} = 2.43^{+0.12}_{-0.18}$ [58].

In 2013, the MINOS experiment started running (as MINOS+) using the higher energy beam for the NOVA experiment. The MINOS far-detector can be used to measure neutrinos in the 4-10GeV energy range and investigate sterile neutrinos and non-standard neutrino interactions [59].

T2K

For a detailed description of the T2K experiment, see section 3.1. The T2K experiment is located in Japan, and uses a muon neutrino beam with a baseline of 295 km [7]. It has two detectors; a near detector 280 m from the beam source and a far detector at 295 km. T2K's primary physics goals are measuring the mixing angles θ_{13} and θ_{23}. In 2011, T2K published its first measurement of electron neutrino appearance from its muon neutrino beam, which suggested θ_{13} was not zero, to an accuracy of 2.5σ [21]. In 2013 this result was updated to give a confidence of 7.3σ, and a measurement of θ_{13} of $\sin^2 2\theta_{13} = 0.140^{+0.003}_{-0.032}$ [22].

NOvA

The NuMi Off-Axis ν_e Appearance experiment is a neutrino experiment based at Fermilab, and was designed as the successor to MINOS. The far detector is also in Northern Minnesota, this time with a baseline of 810 km [60].

NOvA's two primary physics goals are to put further constraints on the mixing angle θ_{13}, and to investigate the mass hierarchy of neutrinos. The mass ordering of neutrinos can be investigated by comparing the oscillation probabilities of neutrinos in vacuum, with the oscillation probabilities through matter. However, a very long baseline is required to get an observable difference between the two scenarios. With an 810 km baseline and high luminosity neutrino beam, NOvA is the only current-generation experiment that has the sensitivity to measure the mass hierarchy [60].

As of writing, the NOvA experiment is still under construction, with first measurements expected to be taken in 2015.

OPERA

The Oscillation Project with Emulsion-tRacking Apparatus is a neutrino experiment designed to detect tau neutrinos produced from a muon neutrino beam. The beam is situated at CERN, with the detector in Gran Sasso in Italy [61].

Its primary physics goal was to observe tau neutrino oscillations from a neutrino beam. It is one of few neutrino experiments sensitive to tau neutrinos; taus produced in charged-current tau-neutrino interactions can be identified using OPERA's emulsion-based detector. In 2010, OPERA reported its first discovery of a tau neutrino candidate from the muon neutrino beam [62]. Its latest result, in 2013, documents three tau neutrino candidates [63].

The OPERA experiment's most talked-about publication was made in September 2011. The experiment published a paper documenting a neutrino travel-time from the beam at CERN to Gran-Sasso which was inconsistent with the speed of light [64]. In 2012, the OPERA experiment published a paper explaining that the 2011 result was likely due to instrumental effects, and that its updated time-of-flight measurements were consistent with the speed of light [65].

2.3.5 Sterile Neutrino Studies

LSND Experiment

The Liquid Scintillator Neutrino Experiment was build in 1987 to investigate neutrino oscillations. The experiment was based around a liquid scintillation counter at the Los Alamos National Laboratory, measuring neutrinos produced from an anti-muon neutrino beam [66]. The primary physics goal was to investigate neutrino oscillations by measuring anti-electron neutrinos produced from the anti-muon neutrino beam.

Results from LSND, published in 2001, found an excess of anti-electron neutrinos which rejected the no-oscillation hypothesis with a confidence level of 5σ [67]. However, the measured oscillation probability conflicted with solar neutrino oscillation results published by SNO [39] and atmospheric neutrino oscillation results published by Super-Kamiokande [4]. The three results cannot be explained by a three-flavour theory of neutrino oscillations. This suggests a fourth neutrino flavour, a sterile neutrino, could exist, which would explain the results. However, errors across the three detectors are too large for the sterile-neutrino hypothesis to be confirmed by this data.

MiniBooNE

The Mini Booster Neutrino Experiment (MiniBooNE) is an experiment based at Fermilab. It uses a muon-neutrino beam to investigate neutrino oscillations, by measuring the appearance of electron neutrinos. One of MiniBooNE's primary physics goals was to investigate the results produced by the LSND experiment [68].

In 2006 MiniBooNE published results which refuted a basic 2-flavour explanation of the LSND anomaly, however more complex models could not be ruled out [68].

2.3.6 Future Neutrino Experiments

Hyper Kamiokande

Hyper Kamiokande is a proposed successor to the T2K experiment. It would use the pre-existing neutrino beam at Tokai as its source, however a new far detector would be built at the Kamiokande site. This detector would be a water-Cherenkov detector consisting of 900,000 tons of water. This greatly increased

detector volume would allow for much more precise measurements of ν_μ and ν_e appearance. The primary physics goals of such a facility would be to measure the value of the CP-violating phase factor δ, and solve the mass hierarchy problem.

LBNE

The Long-Baseline-Neutrino-Experiment is a proposed experiment based at the Homestake Mine, using a new muon neutrino beam at Fermilab. The baseline will be 800 km. LBNE is proposed to pioneer a new type of neutrino detector; a 10 kt liquid argon scintillator detector. Gaseous argon detectors are already in use, and currently form a key component of the T2K experiment's near detector [7]. Charged particles ionise the argon and the liberated electrons can be detected via charged-sensitive planes, allowing for precise 3D tracking. The advantage of a liquid argon detector over a gaseous one is that the detector mass in considerable, increasing the number of neutrino interactions. LBNE's physics goals would include the measurement of δ and resolving the mass hierarchy problem. [69].

Chapter 3

The T2K Experiment

The analysis laid out in chapters 4, 5 and 6 are based around neutrino interactions in the near detector complex (ND280 for short) of the Tokai to Kamiokande (T2K for short) neutrino experiment. In order to give the analysis context, this chapter provides a description of the T2K experiment and its major components. Section 3.1 provides an overview of the experiment. Section 3.1.2 describes the off-axis technique used by this experiment. Sections 3.2, 3.2.1 and 3.2.2 describe the three detectors used in T2K: the Super-Kamiokande far detectors, and the INGRID and ND280 near detectors. Technical information about the experiment was largely taken from T2K's initial publication in Nuclear Instruments and Methods [7].

3.1 Overview of the T2K Experiment

The Tokai to Kamiokande Long Baseline Neutrino Oscillation Experiment is located in Japan and has three major components. Figure 3.1 shows an overview of the experiment. A muon neutrino beam produced by a 30 GeV proton synchrotron at the JPARC physics research facility, a near detector complex 280 metres from the beam source, and the far detector Super-Kamiokande: A 50,000

tonne water Cherenkov neutrino detector [4].

The T2K experiment's major physics goal is to accurately measure the neutrino oscillation parameters θ_{13} and θ_{23} by measuring ν_e appearance and ν_μ disappearance probabilities at Super-Kamiokande.

The Super-Kamiokande detector is displaced 2.5° away from the neutrino beam axis. This is done to displace the peak neutrino flux to an energy that maximises the θ_{13} oscillation probability. An off-axis detector also reduces the high energy tail of the beam, which can reduce backgrounds. A description of the off-axis technique is given in section 3.1.2

The near detector complex consists of two detectors: INGRID and ND280. ND280 is also located off-axis, and was constructed to measure neutrino interaction cross section parameters. Some of these, in particular the neutral-current single π^0 production interaction cross-section, form significant backgrounds at Super-Kamiokande, and reducing the cross-section uncertainties complements the major physics goals of T2K.

3.1.1 T2K Beam

The T2K neutrino beam is designed to produce muon neutrinos at an average energy of 600 MeV. The beam is created at the JPARC high energy physics complex in Tokai, Japan (figure 3.1) [7]. The neutrino beam flux prediction at T2K's near detector is shown in figure 3.2.

A beam of 30 GeV protons is created by a main accelerating ring (which is fed by a small synchrotron, which in turn is fed by a linear accelerator). This proton beam then enters the neutrino beam-line. The beam-line is made up of six sections, described below and shown in figure 3.3.

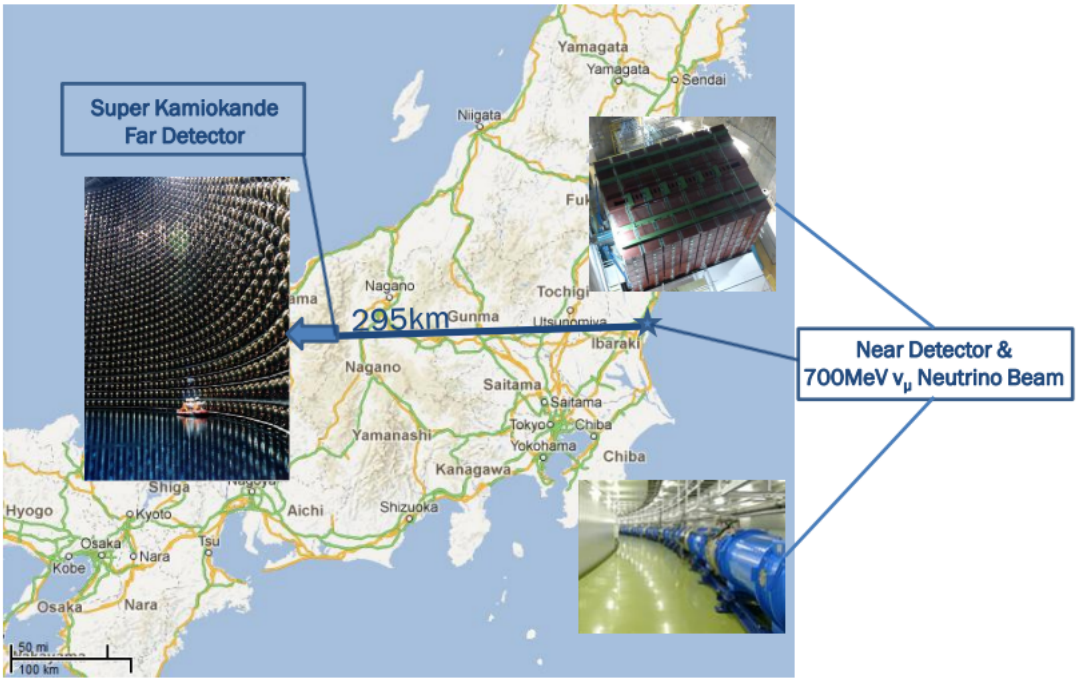

Figure 3.1: Overview of the T2K experiment

The first section of the beam-line, the preparation section, is 54 m long and this is where the beam is focused via eleven normal conducting magnets, so that it can be fed into the next part of the beam-line, the arc section.

This second section is 104 m long and bends the proton beam towards the direction of Kamioka.

The third section, the final focusing section, focuses the proton beam to a point, so that it can accurately hit the beam target.

The proton beam is fired into the target station, where interactions on carbon produce pions, amongst other interaction products.

The pions then enter the decay volume, a 96 m long steel tunnel, and decay

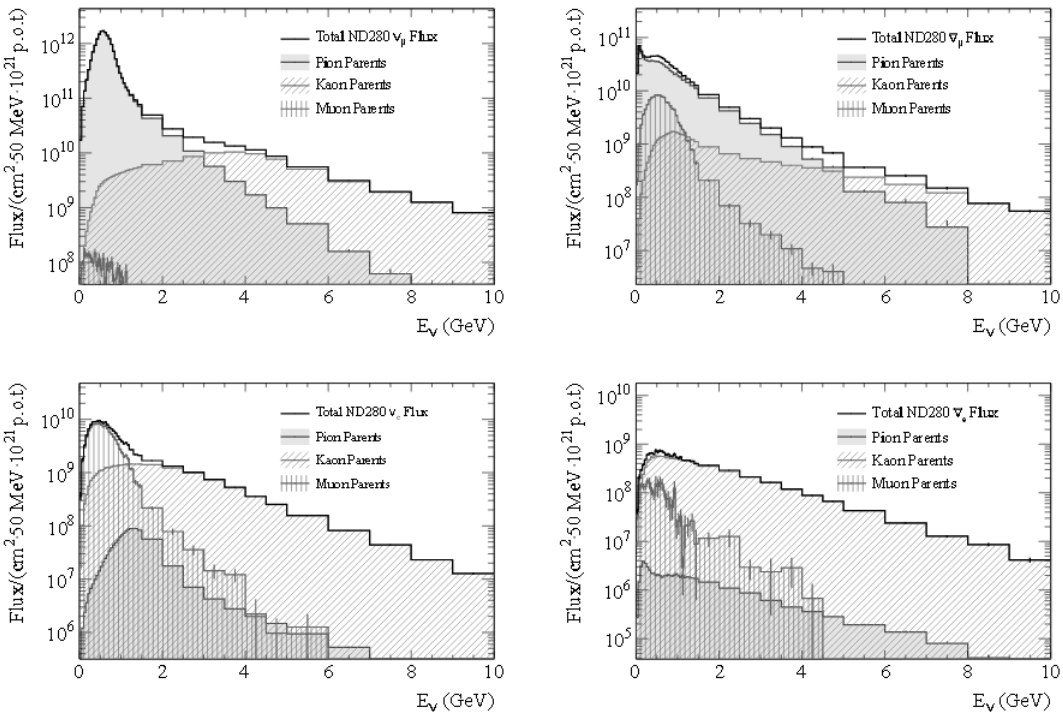

Figure 3.2: The flux prediction at the ND280 near detector, broken down by neutrino parent particle type. Going clockwise starting at the top-left, the neutrino types are: ν_μ, $\bar{\nu}_\mu$, ν_e, $\bar{\nu}_e$. Errors shown are Monte Carlo statistical errors. Image taken from T2Ks published result on the neutrino flux prediction [6].

primarily into muons and muon neutrinos. There are three sets of horn-shaped magnets in the decay volume which focus the pions so that neutrinos produced from pion decays are forward-going.

At the end of the decay volume is the beam dump is composed of 75 tonnes of graphite, and intercepts any hadrons, as well as muons below 5 GeV. After the beam dump, what is left of the initial proton beam is mostly neutrinos. Any escaping muons are monitored by an electromagnetic calorimeter placed after the beam dump. The neutrino beam consists approximately of 94% muon neutrinos, and 6% electron neutrinos. Approximately 1% of all neutrinos produced are anti-neutrinos [7].

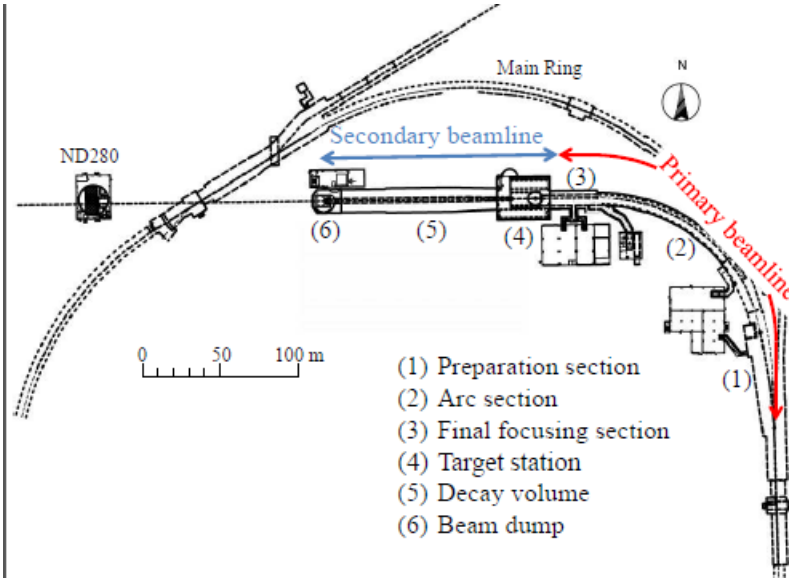

Figure 3.3: The JPARC beam. Image taken from the T2K NIM paper [7].

3.1.2 The Off-Axis Technique

The idea of placing a neutrino detector off-axis from the beam source was put forward by Beavis *et al.* in 1995 [8]. Using an off-axis neutrino detector to narrow the energy spectrum of neutrinos entering a detector. Consider the neutrino-producing π^+ decay: $\pi^+ \rightarrow \nu_\mu \mu^+$. It is a two-body decay and we can choose the direction of travel of the π^+ to be along the z-axis, and the decay to occur in the xz plane. The 4-momenta of the ν can be written as:

$$p_\nu = E^*_\nu \begin{pmatrix} 1 \\ \sin\theta^* \\ 0 \\ \cos\theta^* \end{pmatrix} = \begin{pmatrix} E^*_\nu \\ p^*_{\nu_x} \\ p^*_{\nu_y} \\ p^*_{\nu_z} \end{pmatrix} \tag{3.1}$$

where θ^* represents the opening angle between the neutrino and the muon in the centre of mass frame. Applying a Lorentz boost along the z-axis yields the neutrino's 4-vector using centre of mass quantities:

$$p_\nu = E_\nu \begin{pmatrix} \gamma_\pi(1 + \beta_\pi \cos\theta) \\ \sin\theta \\ 0 \\ \gamma_\pi(\beta_\pi + \cos\theta) \end{pmatrix} \tag{3.2}$$

Taking $\frac{p^*_{\nu_x}}{p^*_{\nu_z}}$ yields

$$\tan\theta = \frac{E_\nu \sin\theta}{\gamma_\pi E_\nu(\beta_\pi + \cos\theta)} \tag{3.3}$$

Since the pion is produced from a 30 GeV accelerator it is safe to assume that $\beta_\pi \approx 1$. Using $E^*_\nu = \gamma_\pi E_\nu(\beta_\pi + \cos\theta)$, and rearranging to yield the neutrino energy in the lab frame gives:

$$E^*_\nu = \frac{E_\nu \sin\theta}{\tan\theta} \tag{3.4}$$

The energy of the neutrino in the centre of mass frame is given by the two body decay kinematics. From conservation of momentum and energy, the 4-momentum of the muon can be written as:

$$p_\mu = p_\pi - p_\nu \tag{3.5}$$

Taking the square and expanding yields:

$$m^2_\mu = m^2_\pi - 2\boldsymbol{p}_\pi \cdot \boldsymbol{p}_\nu \tag{3.6}$$

In the centre-of-mass frame, the pion has no kinetic energy and so equation 3.6 can be written as:

$$\begin{aligned} m_\mu^2 &= m_\pi^2 - 2m_\pi E_\nu^* \\ E_\nu^* &= \frac{m_\pi^2 - m_\mu^2}{2m_\mu} \end{aligned} \tag{3.7}$$

The neutrino energy in the centre-of-mass frame is therefore a constant. Considering $\sin\theta$ can never exceed one, this shows that the neutrino energy in the laboratory frame is inversely proportional to $\tan\theta$, where θ is the opening angle of the decay in the laboratory frame.

For a detector that is directly in line with the neutrino beam, the detector is only sensitive to decays where $\theta \approx 0$. At $\theta = 0$, $\tan\theta = 0$ and no bounds are set on the neutrino energy.

However, for an off-axis detector, the detector is only sensitive to neutrinos where the opening angle is non-zero, therefore equation 3.4 limits the maximum neutrino energy that can be observed.

This is useful in neutrino physics, because of the difficulty of creating a beam with a narrow range of neutrino energies. The neutrinos produced from the JPARC accelerator complex cover a broad energy range: by putting the detector off-axis neutrinos with undesirable energies can be suppressed. This is shown in figure 3.4, which shows the distribution between neutrino energies and parent pion energies, for a given θ.

By using the off-axis technique, it is possible to 'tune in' to the neutrino energy spectrum, which provides an oscillation maximum at T2K, as shown in figure 3.5

Both Super-Kamiokande and ND280 are located 2.5° away from the beam axis to maximise the oscillation probability.

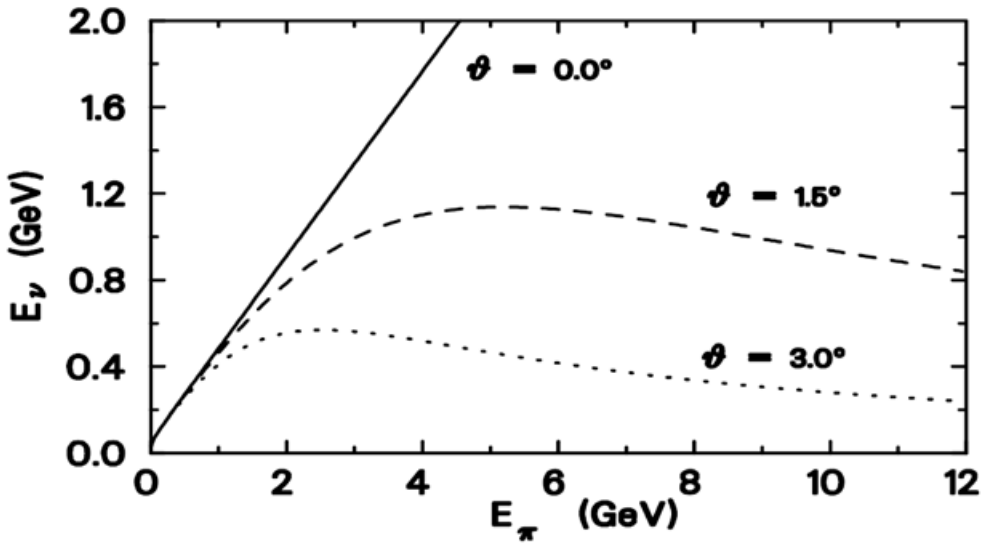

Figure 3.4: How opening angle affects neutrino energy distributions [8].

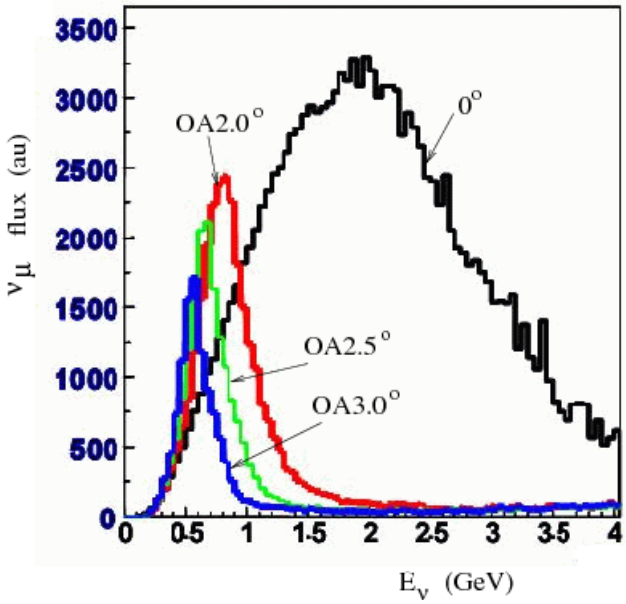

Figure 3.5: How varying opening angle affects neutrino flux at T2K. The y-axis shows predicted number of ν_μ interactions [9].

3.2 The Super-Kamiokande Detector

The Super-Kamiokande Far Detector, shown in figure 3.6 is a water Cherenkov detector, which has been in operation since 1996. Its main components are two concentric chambers which contain 50 kilotons of water. The walls of the chamber are lined with Photo-Multiplying-Tubes, which measure cones of light emitted from Cherenkov showers when a neutrino interaction occurs. The detector is located 1km inside Mount Kamioka. Its depth serves to limit cosmic backgrounds.

Super-Kamiokande's detector volume is split into an inner and outer cylinder. The former has a 36.2m height and 33.8m diameter and is lined by 11,129 inward-facing Photo-Multiplier tubes (PMTs).

The outer detector contains 1885 outward-facing PMTs along its inner wall. The boundary of the two chambers is delineated by a stainless steel scaffold covered in plastic sheets, which keep the two chambers optically isolated. In order to compensate for the relatively small number of PMTs lining the outer detector, the walls facing the outer chamber are covered in reflective Tyvek(®).

The outer detector's main purpose is to act as a veto for cosmic rays entering the inner chamber.

Both ν_e and ν_μ interactions can be detected by Super-Kamiokande. Charged-Current Quasi Elastic (CCQE) neutrino interactions are of specific interest due to their large branching ratio and relative ease of reconstruction.

Muons produced by a CCQE ν_μ interaction undergo relatively little scattering compared to their lighter electron-neutrino counterparts. This distinction, combined with the likelihood of electrons to shower, causes CCQE ν_μ interactions to produce much more well defined Cherenkov light rings than CCQE ν_e interactions. This difference is how Super-Kamiokande distinguishes electron and muon neutrinos.

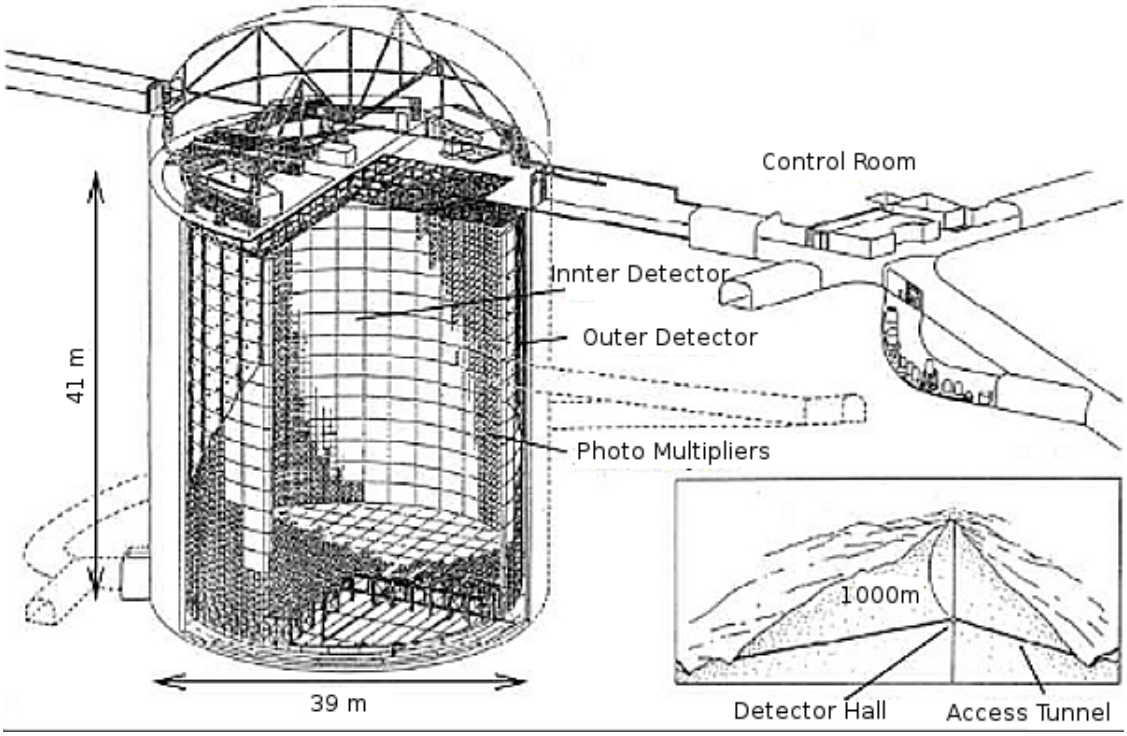

Figure 3.6: The Super-Kamiokande water cherenkov detector [7].

A complication arises when other interaction types are considered. Specifically, π^0s produced by neutral-current interactions (figure 3.8) of both neutrino types can mimic CCQE ν_e interactions in Super-Kamiokande (figure 3.7). Photon and electron Cherenkov cones are visually similar. Due to the low angular separation between the two photons produced in an energetic π^0 decay, the photon pair can imitate a single photon, which in turn mimics an electron. On top of this, one of the two photons can have a boosted laboratory momentum, which is low enough for any electrons produced to produce a limited amount of Cherenkov radiation. Understanding the neutral current π^0 production neutrino interaction cross-section will serve to constrain backgrounds introduced by this interaction type.

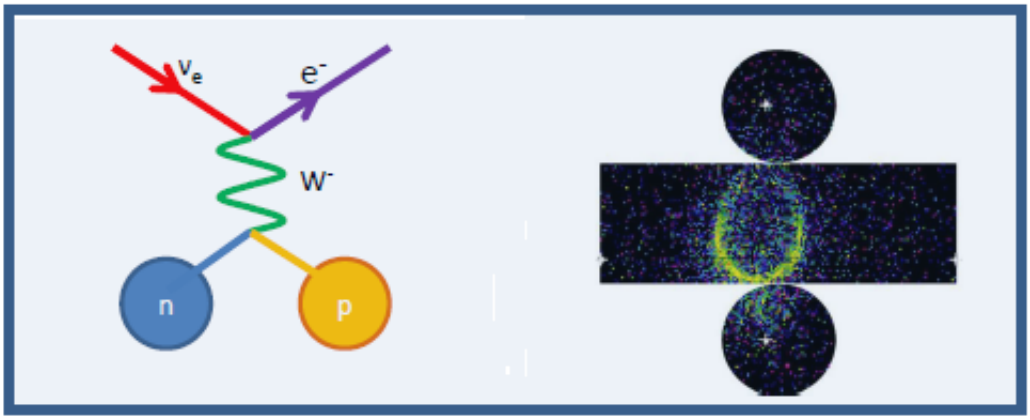

Figure 3.7: A CCQE ν_e interaction and the image produced by such an interaction in Super-Kamiokande.

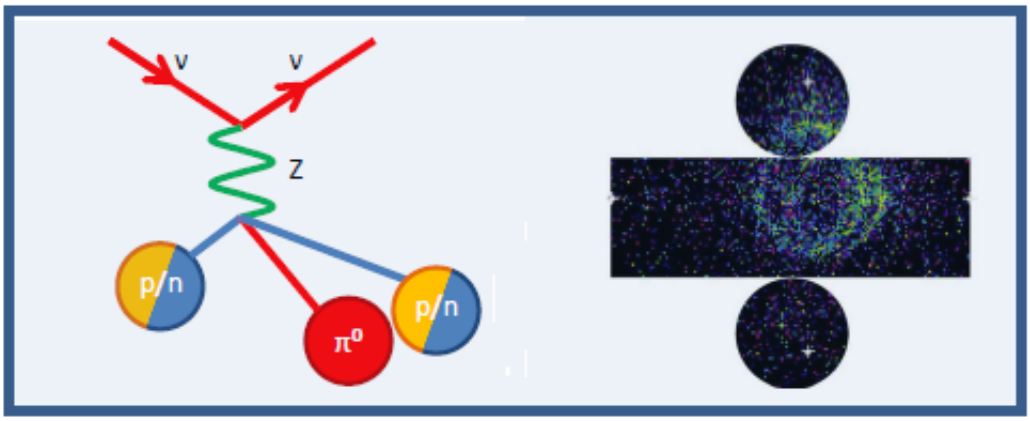

Figure 3.8: A NCπ^0 interaction and the image produced by such an interaction in Super-Kamiokande.

3.2.1 The INGRID Detector

The Interactive Neutrino GRID (INGRID) detector is centered on the neutrino beam axis and is localed 280 m away from the neutrino beam. Its primary function is to accurately determine the neutrino beam's direction and intensity. As the ND280 detector relies on the off-axis technique [7], understanding the directionality of the beam is crucial to understanding the beam's energy profile at ND280. INGRID can measure the beam center down to a precision of 0.4 mrad at the near detector.

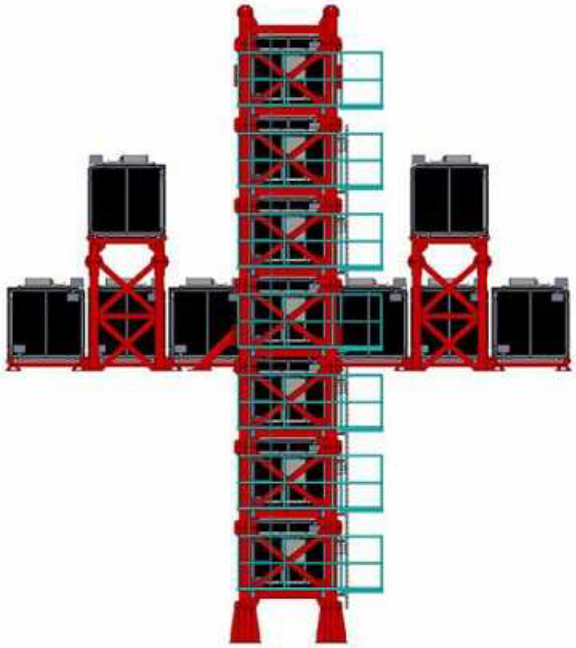

Figure 3.9: Frontal cross-section of the INGRID Detector [7].

INGRID is constructed from 14 modules and forms a cross shape, seen in figure 3.9, as well as 2 extra modules located off-axis. The off-axis modules measure the axial symmetry of the beam. The cross-section of the detector is 10 m x 10 m. The centre of INGRID is aligned to be at the centre of the neutrino beam. Each INGRID module contains 11 polystyrene scintillator planes, with 9 iron plates interleaved between the planes. The scintillator planes are made up of scintillator bars, with dimensions 1.0 cm x 5.0 cm x 120.3 cm.

3.2.2 The ND280 Detector

The ND280 detector is a composite neutrino detector located 280 m away from T2K's neutrino beam [7]. It consists of 8 subdetectors, each with its own strengths. As a whole ND280 can accurately reconstruct the majority of neutrino interaction types. ND280 was designed to perform cross section measurements on all common neutrino interaction types, as accurate information on many is cur-

rently lacking.

ND280 is designed to accurately measure the ν_e and ν_μ components of the beam. As ν_e appearance and ν_μ disappearance measurements at Super-Kamiokande form the major physics goals of the T2K Experiment, understanding the beam profile is critical.

Given the scarcity of statistics in the ν_e appearance measurement (T2K's first published appearance result measured six events [21]), it is also crucial to understand ν_μ interactions that create backgrounds in the ν_e search in order to reduce systematic errors these backgrounds introduce. The background is dominated by neutral current single π^0 production.

Figure 3.10 depicts an overview of ND280. Figures 3.11, 3.12 and 3.13 show the layout of the ND280 detector. Figure 3.11 shows the cross-section of ND280 as viewed from the side. Figure 3.12 shows ND280 from the front, and figure 3.13 shows ND280 from the back. The front of the detector is defined to be where the neutrino beam enters ND280, the back is defined to be where the neutrino beam exits ND280. The subdetectors which make up ND280 consist of a π^0 detector (P0D), two Fine Grain Detectors (FGDs), three Time Projection Chambers (TPCs), an Electromagnetic Calorimeter (ECal) which surrounds the inner detectors, and a Side Range Muon Detector (SMRD). The SMRD is embedded within ND280's focusing magnet, which was acquired from the UA1 experiment.

3.2.3 UA1 Magnet

The UA1 magnet started life as the bending magnet for the UA1 experiment at CERN, and later at NOMAD [70]. It produces a dipole magnetic field of 0.2 T and is designed to bend charged particles as they travel through ND280, allowing the various subdetectors to accurately measure particle momentum

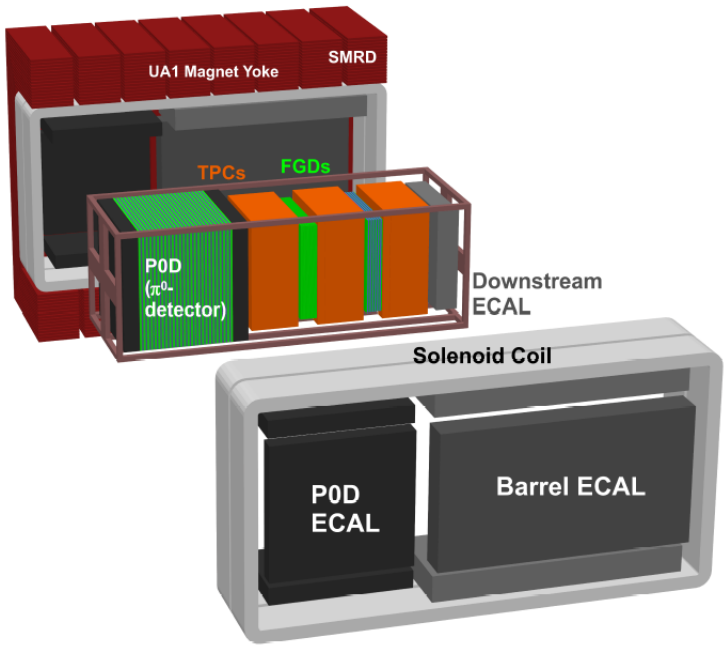

Figure 3.10: Exploded image of ND280. The middle subdetector components are surrounded by the P0D ECal, barrel ECal and solenoid coil. These in turn surround the inner subdetector components in the centre of the image. Neutrino beam enters from the left [7]

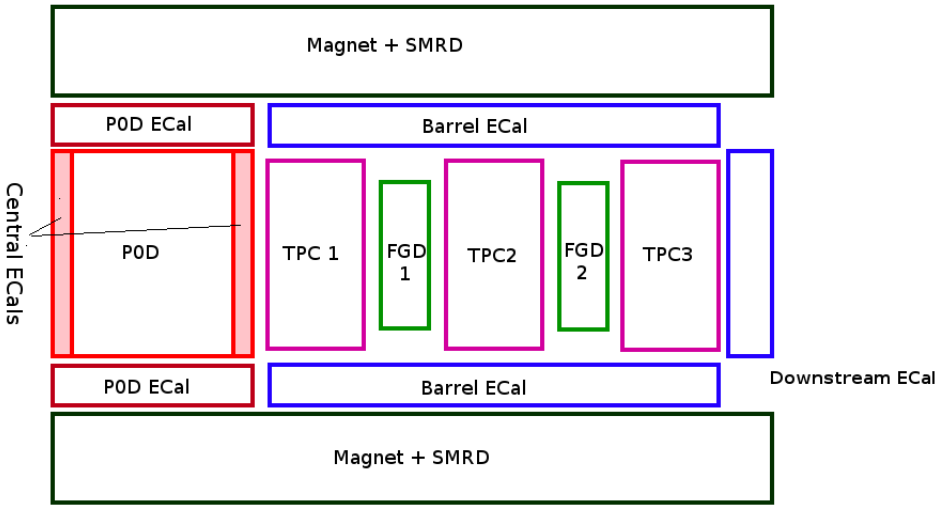

Figure 3.11: Cross-section of the side of ND280. The neutrino beam enters the detector on the left and exists on the right.

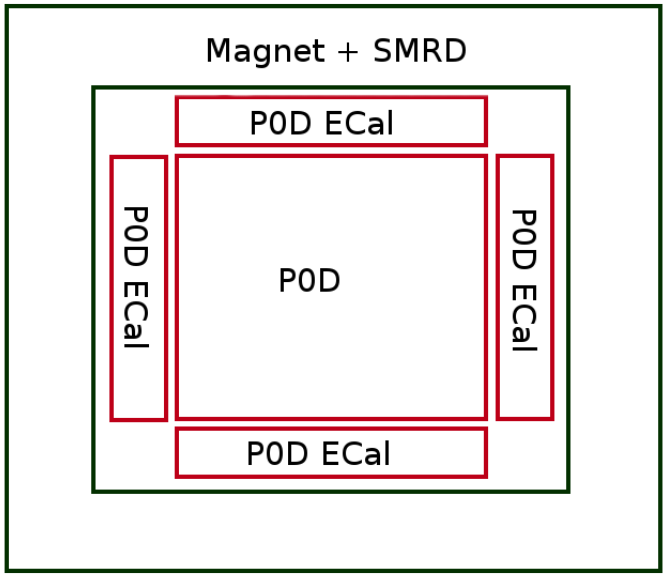

Figure 3.12: Cross-section of the front of ND280, where the neutrino beam enters the detector.

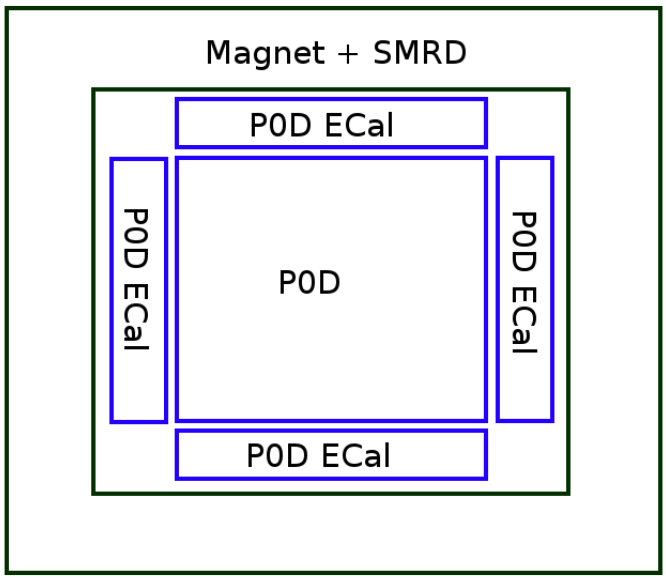

Figure 3.13: Cross-section of the rear of ND280, where the neutrino beam exits the detector.

and determine the sign of charged particles.

The inner volume of the magnet is 7.0 m x 3.5 m x 3.6 m and is composed of aluminum coils. It weights 18 tonnes and is water-cooled.

3.2.4 π^0 Detector

The π^0 detector (P0D) is most upstream subdetector of ND280 and is primarily designed to detect π^0s. It consists of planes of triangular-shaped polystyrene scintillator bars, interleaved with bags of water and sheets of lead and brass.

The detector can be configured to run with either the water bags full or empty. This allows for cross-section measurements to be made in 'water in' and 'water out' modes, allowing neutrino interaction cross-sections on water to be measured. This is one of the main goals of the π^0 detector, as π^0 backgrounds in Super-Kamiokande are considerable. The mass of the detector with water is 16.1 tonnes. Without water the detector mass is 13.3 tonnes. The mass of water inside the P0D is measured to be 1902 ± 26 kg

The scintillator bars were designed to fulfill the twin roles of accurately tracking charged particles, and measuring the total energy deposition of electromagnetic showers. The cross-section of a triangular scintillator bar is 33 mm x 17 mm. A 1.5 mm hole passes through the centre of each bar to accommodate a wavelength shifting fiber. The fibers then feed into MPPCs.

Figure 3.14 shows the construction of the P0D, and the layering of brass, lead, water and scintillator. The two ends of the P0D are capped by layers of scintillator and lead, without any brass or water. These two end-regions form two electromagnetic calorimeters (P0D Ecals).

The P0D ECals serve the twin purpose of containing electromagnetic showers and enabling analysts to discriminate between neutrino interactions inside the water region of the P0D, and particles entering the P0D from external interac-

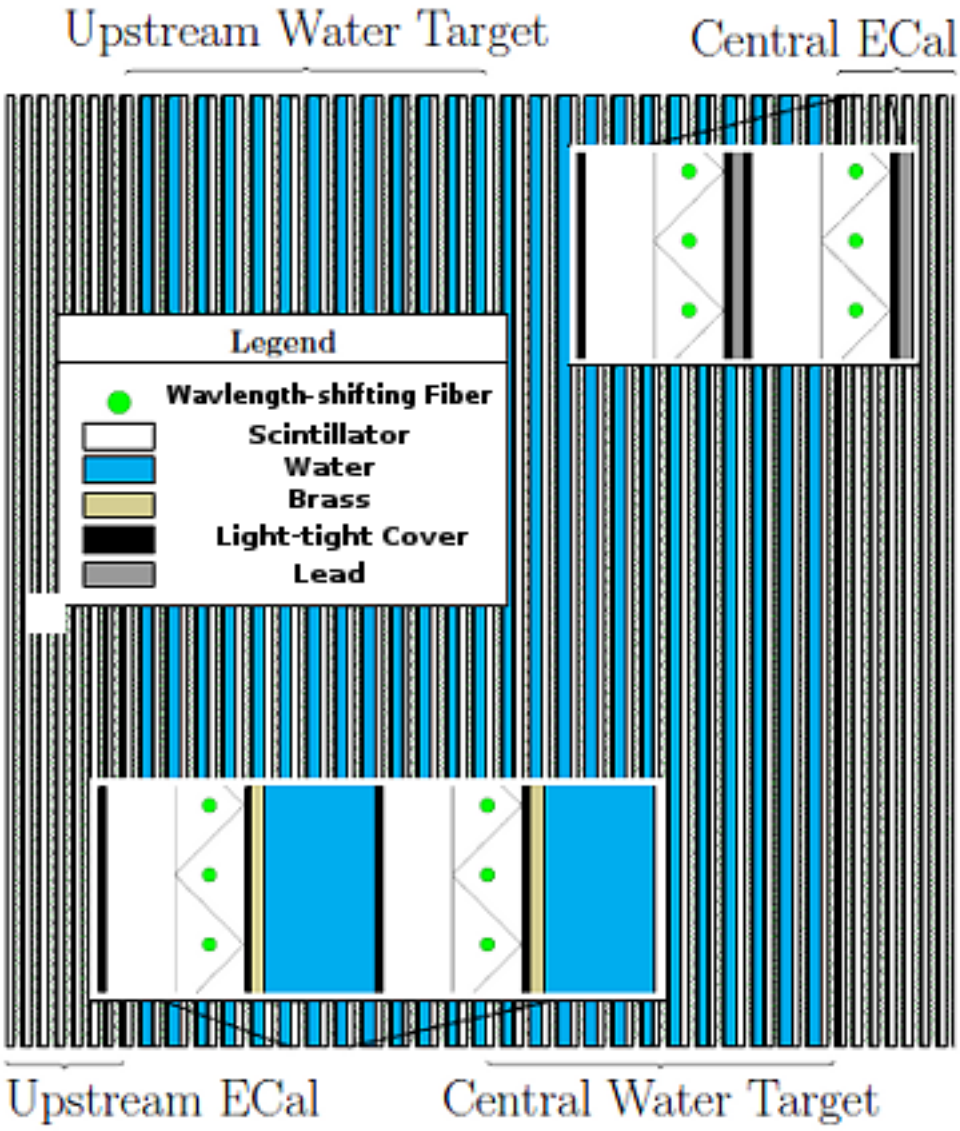

Figure 3.14: Side view of the distribution of brass, lead, scintillator and water layers in the P0D. Neutrino beam enters from the left [7].

tions.

3.2.5 Time Projection Chambers

There are three time projection chambers (referred to as TPCs) in ND280. Each TPC has dimensions of 2.3 m by 2.4 m by 1.0 m [7] The TPCs fulfill the role of having very high tracking efficiency and accuracy with low detector mass. The

TPCs are also used to accurately measure the rate of energy loss of a particle, a property which is crucial for particle identification.

The TPCs main two components are a chamber of argon-based gas and a uniform electric field applied across the TPC. The walls of the argon container are made of copper, formed into a copper strip pattern. This forms the anode of the electric field; the cathode is created from a central panel in the TPC (figure 3.15). When a charged particle travels through a TPC, the argon gas is ionised and the ionised electrons drift under the influence of the electric field into readout planes.

The readout planes consist of 12 micromegas detectors covering an area of $9m^2$, which amplify and record the ionization electrons. The resolution of the detectors is 7.0mm x 9.8mm (vertical/horizontal) [7]. These detectors measure the position of the particle in the axis parallel to the detector. The arrival time of the electrons is also logged, allowing accurate positioning of the electrons in the axis perpendicular to the micromegas detector and subsequently provide a complete 3D reconstruction of the path of the original charged particle.

An outer container filled with carbon dioxide gas is used to insulate the argon gas. There are six front-end boards of electronics, used to sample and digitise the analog signals coming from the readout pads. A final electronics board is used to collate the TPC data and send it to the data acquisition module (DAQ).

The TPCs are uniquely suited to measuring the rate of energy loss of particles because of its very high granularity when measuring particle positions. As a charged particle travels through the TPC, it will lose energy through ionising the argon gas. This energy loss leads to a corresponding loss in momentum, which can be measured via a change in the particle's curvature. This energy-loss rate is predicted by the Bethe-Bloch formula [71]. The energy-loss rate, $\frac{dE}{dx}$, is different depending on the particle mass, and this information can be used to infer what kind of particle is traveling through the TPC.

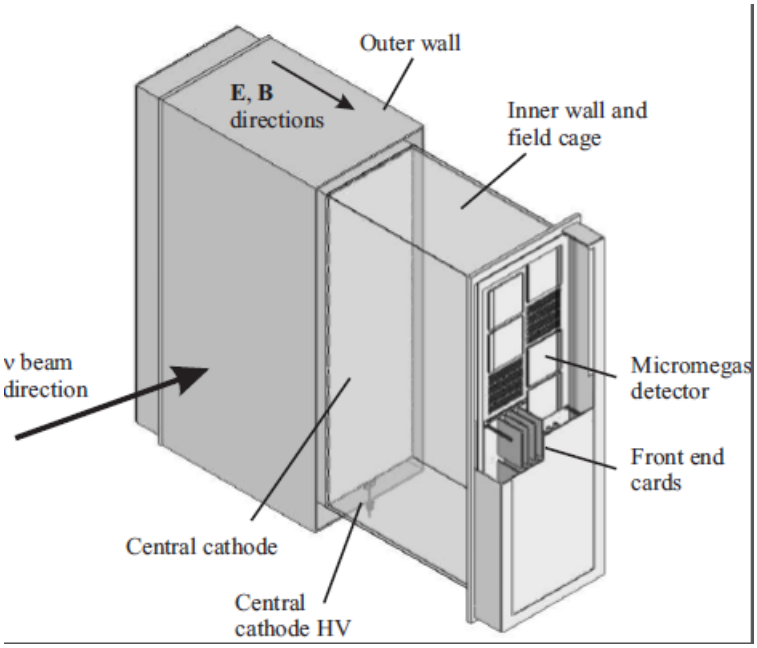

Figure 3.15: Simplified cutaway of a time projection chamber [7].

The measured energy loss rate for a given particle is compared against the 'ideal' energy-loss rate for various particle candidates: electrons, protons, muons, pions and kaons. The difference between the expected and measured value for each case is referred to as a 'pull'. For example, the electron pull of a particle is the difference between the particle's energy loss rate compared with an ideal electron. A pull value of 0 denotes perfect agreement. Positive pulls imply the energy loss is greater than expected, negative pulls imply the energy loss is less than expected. An in-depth formalism of the TPC pulls is described in an internal T2K technical note [72].

The pull, δ_E, for a given particle hypothesis i, is defined as:

$$\delta_E(i) = \frac{\bar{C}_T - C_E(i)}{\sigma_o(i)} \tag{3.8}$$

$\bar{C}_T$ is the measured energy loss, and $C_E(i)$ is the expected energy loss for the

particle hypothesis. $\sigma_o(i)$ represents a width which is a function of the variance of C_T, σ_T and the uncertainty of the momentum measurement of the particle σ_p. $\sigma_o(i)$ is given as:

$$\sigma_o(i) = \sqrt{\sigma_T(i)^2 + \left(\frac{dC_E}{dp}\sigma_p\right)^2} \tag{3.9}$$

3.2.6 Fine Grain Detectors

The two Fine Grain Detectors (referred to as FGDs, shown in figure 3.16) in ND280 provide detector mass for neutrino interactions, and complement the adjacent TPCs. Each FGD is constructed from planes of extruded polystyrene scintillator bars with dimensions 9.61 mm x 9.61 mm x 1864.3 mm. Each bar is oriented perpendicular to either the x or y direction, with an MPPC (multi-pixel-photon-counter) at one end. The overall volume of each FGD is 2.3 m x 2.4 m x 0.365 m, with 1.1 tonnes of target mass. The tracking resolution of each FGD is given by the dimensions of each scintillator bar, allowing for sub-1 cm resolution .

The second, downstream FGD differs from the first in that it has water containers embedded inside the detector. There are six 2.5 cm thick water modules sandwiched between seven XY modules of scintillator bars. This contrasts wtih fifteen XY modules of scintillator bars in the first FGD. Each water module consists of a thin-walled hollowed polycarbonate casing. The water modules are kept at a pressure of slightly less than 1 atmosphere: any leaks will suck air into the module instead of expelling water. By comparing interaction rates between FGD1 and FGD2, interaction cross-sections on water can be calculated. This was one of the motivations for performing the NCπ^0 analysis in this thesis, however significantly more statistics are needed before such a comparison is viable.

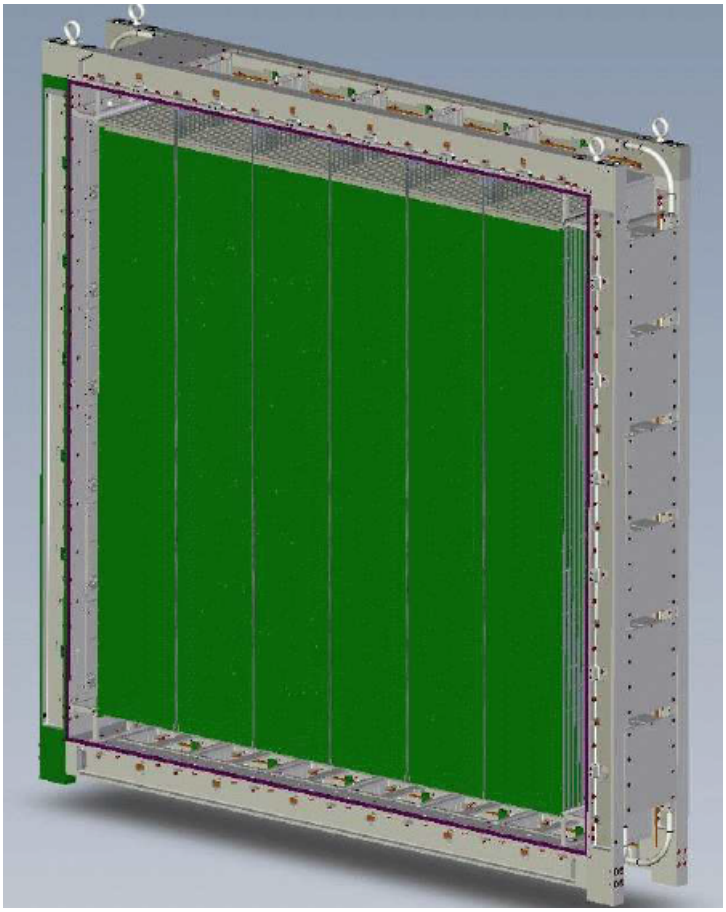

Figure 3.16: Image of the FGD with the front cover removed. Image taken from [7]

The entire FGD is enclosed within an aluminium dark-box, preventing external light sources from interfering with the scintillator bars and their attached MPPCs.

3.2.7 Electromagnetic Calorimeter

The electromagnetic calorimeter (referred to as the ECal, shown in figure 3.17) is designed to provide significant detector mass for neutrino interactions, and it surrounds all of the inner subdetectors (P0D, FGDs, TPCs). It consists of doped polystyrene scintillator bars with cross-section 4.0 cm x 1.0 cm, and varying length which are arranged in planes. Sandwiched between the scintillator

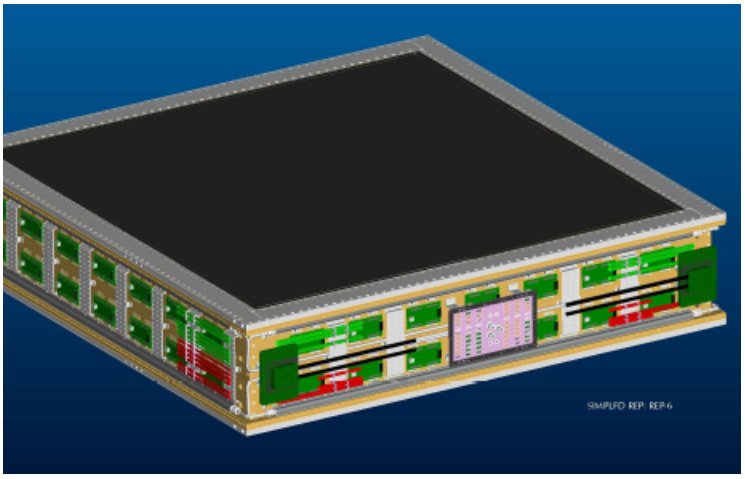

Figure 3.17: An ECal module. The readout electronics are mounted on aluminium plates at the edge of the module. The top of the module is a carbon fiber sandwich plate [7].

planes are sheets of lead.

Each ECal scintillator bar has a 2.0 mm x 3.0 mm hole in the centre, running across its length. A 1 mm fiber [7] runs through the hole and an MPPC is connected to one or both ends.

The ECals considerable detector mass complements the inner components, specifically the TPCs and FGDs which are considerably lighter. The ECal is also well suited to measuring the total energy deposition of electromagnetic showers, something neither the TPCs or FGDs are well-suited to do.

The ECal is split into 13 subsections: an ECal at the downstream exit of the detector (referred to as the DsECal). Six ECal modules surrounding the inner detectors on the four sides parallel to the beam's z-axis. Six ECAL modules similarly surround the P0D detector. This is illustrated in figure 3.17.

3.2.8 Side Range Muon Detectors

The Side Range Muon Detector (referred to as the SMRD, components of which are shown in figure 3.18) is primarily used as a trigger for cosmic ray events. It

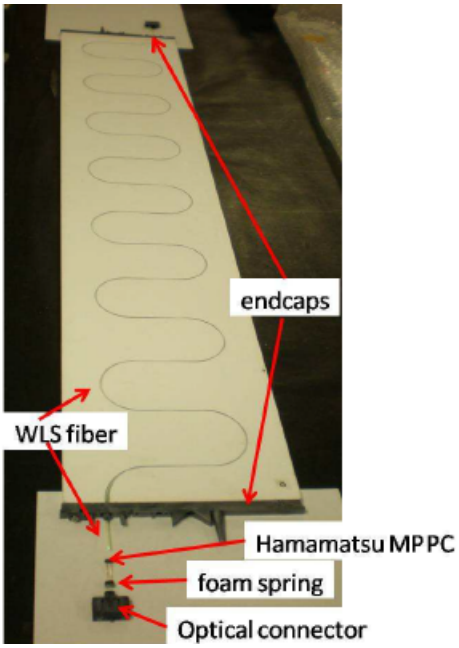

Figure 3.18: SMRD scintillator components. Image taken from [7]

is built into the UA1 magnet that surrounds ND280 and consists of 440 scintillator modules. The scintillator modules are placed within the 1.7 cm air gaps that exist between the steel plates of UA1's return yolk. There are 15 such air gaps. Due to the UA1's considerable mass and orientation, the SMRD can detect charged particles which are entering the detector but not in the direction of the neutrino beam.

3.2.9 Detector Electronics

For all the detectors which use MPPCs (P0D, ECal, SMRD, INGRID), the output signals are fed to custom front-end boards (called TFBs) which digitise the analog signals and measures the signal pulse height and time. The FGD and TPC subdetectors use front-end boards based on different electronics. The output for these subdetectors is digitised into a waveform, similar to an oscilloscope output. This waveform is then processed offline, with pulses identified and the pulse height and pulse times determined.

The output of the TFBs is then fed into a Readout Merger Module (RMM), which logs received data and saves a record when activated by a trigger.
There are three different types of trigger that can activate the RMM. The first is a beam trigger, sent by the neutrino beam whenever a beam spill is created. The second is a cosmic ray trigger, which can be sent by either the SMRD, ECal and FGDs when a cosmic ray is detected entering the top of ND280. The final trigger type is the calibration trigger, used when calibrating individual subdetectors. An example is when the signal baseline of a subdetector needs to be measured.

Chapter 4

Analysis Techniques and Signatures

4.1 NCπ^0 Interactions in ND280

This analysis specifically searches for neutral-current neutrino interactions within ND280, which produce at least one π^0 in the final interaction state. Neutral current π^0 interactions can produce an event signature which mimics electron-neutrino scattering in T2K's water Cherenkov far detector, Super-Kamiokande. A diagram of this interaction is shown in figure 4.1. Due to the fact that the ND280 detector contains a large number of subsystems, it was decided to split up the neutral-current π^0 analysis into several analyses which focused on specific sub-detector combinations, to better optimise overall π^0 reconstruction quality. These analyses are as follows:

1. NC interactions within the P0D and P0D ECal. This was carried out by Dr Karen Giilje, the results of which are documented in her PhD thesis, which is not yet published.

2. NC interactions occurring within the tracker region, where the decay photons convert in the barrel and downstream ECals. This was performed by Dr Sam Short [73], following for Dr Pawel Guzowski [74].

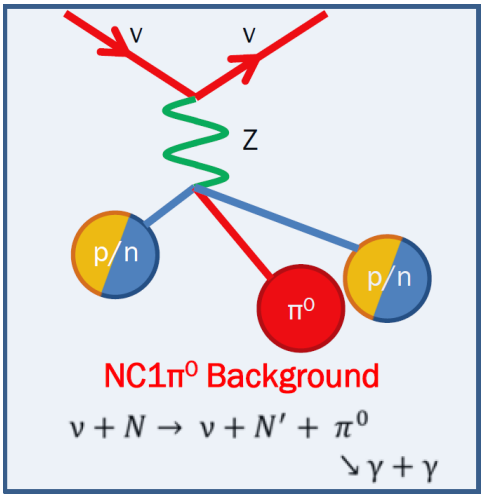

Figure 4.1: A neutral-current π^0-producing neutrino interaction

3. NC interactions occurring within the tracker region, where one decay photon converts in the tracker region and the second decay photon converts in the barrel or downstream ECals. This is being performed by Dr Helen O'Keeffe and is ongoing.

4. NC interactions occurring within the tracker region, where both decay photons convert in the tracker. This analysis is the subject of this document.

Decay photons refer to the two photons produced when a π^0 decays.

4.1.1 Definitions of Signal and Background

When searching for a set of specific event profiles, it is necessary to split all the events in ND280 into two categories: signal and background. The signal consists of neutrino interactions this analysis is trying to reconstruct and the background consists of particle interactions in ND280 that this analysis has no interest in and should therefore be rejected. In the context of this analysis, the signal and background are defined as follows:

- **The Signal consists of neutral-current neutrino interactions which produce at least 1 π^0 meson. This interaction must occur within the FGD1 or FGD2 subdetector volumes. The two decay photons produced from the π^0 decay must convert within the tracker volume of ND280.**
- **The Background consists of all interactions, excluding the Signal.**

The definition of signal is consistent with the fourth list item in subsection 4.1. The analysis focuses on decay photons converting in the tracker region, because the criteria to reconstruct decay photons which convert in the ECal is very different to tracker-based conversions (the ECal-based neutral-current π^0 analyses are being carried out by Sam Short and Helen O'Keeffe, mentioned in subsection 4.1).

The number of all signal events combined with all background events is equal to the number of all interactions that can be seen in the ND280 detector. In an ideal scenario the analysis would reconstruct 100% of the signal and 0% of the background. This is not the case, but these definitions of signal and background can be used to define metrics which measure how effective the analysis is at achieving its goal, namely reconstructing neutral-current π^0-producing neutrino interactions in ND280. This is further discussed in subsection 4.1.11.

4.1.2 Searching for π^0s in the Near Detector

The π^0 decays into two photons with a branching ratio of ($98.78 \pm 0.032\%$). The next most likely decay mode, the Dalitz decay, is into a photon and an electron/positron pair, with a ratio of ($1.198 \pm 0.032\%$), which this analysis is sensitive to. All other decay modes are not included in this analysis as they have small branching ratios and so are statistically irrelevant (table 4.1 shows the decay modes of the π^0 and their branching ratios). The eventual goal of the analysis is to measure the cross section of neutral-current π^0 interactions on water via

π^0 **decay modes**	Fraction (Γ_i/Γ)
2γ	$(98.823 \pm 0.034)\%$
$e^+e^-\gamma$	$(1.174 \pm 0.034)\%$
$e^+e^+e^-e^-$	$(3.34 \pm 0.16) \times 10^{-5}\%$
e^+e^-	$(6.46 \pm 0.33) \times 10^{-8}\%$
γ positronium	$(1.82 \pm 0.29) \times 10^{-9}\%$

Table 4.1: Decay modes of the π^0 and their branching ratios [10].

a rate subtraction between FGD1 and FGD2, therefore only pions produced in the FGD's are targetted for this analysis. This, combined with the scope of the analysis being constrained to the tracker, leads to the following event profile:

1. The interaction vertex occurs in an FGD.
2. The π^0 instantly decays into two photons.
3. The two photons both convert somewhere in the tracker region.

Due to the nature of the neutral-current interaction (figure 4.1), either a proton or a neutron is produced at the interaction vertex. In the latter case, such neutral particles are undetectable in ND280. In either case, the nucleon kinematic information is not used in this analysis. This is because there is very little knowledge about the relative cross-sections between producing a $\pi^0 + n$ and a $\pi^0 + p$ (or between the relative probabilities of the nucleus remaining intact or breaking up) [10]. Incorporating the nucleon into the reconstruction algorithm would introduce a bias and associated significant systematic error, which is difficult to quantify.

The only reconstructable evidence of a neutral-current π^0 interaction is the two e^+/e^- pairs produced by the two converting primary decay photons. For the rest of the reconstruction chapter, electron-positron pairs are also referred to as 'lepton pairs'. The term 'lepton pair' in this context strictly refers to this specific combination of electrons and positrons and does not include muons or taus.

4.1.3 Searching for converting e^+/e^- pairs

Figure 4.3 shows the distribution of momenta for e^+/e^- pairs produced by primary decay photons. What is of interest and concern, is the quantity of low-momenta leptons produced by this interaction. Despite the heavy Lorentz-boost these particles experience, there is no symmetry between the momenta of the two leptons in any produced pair. The two leptons have equal and opposite momenta in their centre of mass frame, however their momenta components relative to the z-axis of the detector are not correlated. Therefore, a Lorentz boost along the z-axis creates the unequal momenta distribution seen in fig 4.3.

Leptons with a momenta of less than 50 MeV/c^2 cannot be reliably reconstructed by the TPC [75]. This leads to a significant majority of the primary photon conversions producing only one reconstructable lepton. For the remainder of the text these types of leptons are called isolated leptons. Due to the limited quantity of statistics available for this type of interaction[1], it was necessary to develop an analysis that attempted to reconstruct the isolated lepton scenarios, as well as situations where both photon interactions produced reconstructable lepton pairs. This naturally lead to the idea of developing a topological analysis, with the following topology definitions:

1. Twin pair topology: Both converting primary photons produce reconstructable lepton pairs.

2. Pair and isolated electron topology: One photon produces a lepton pair, the other an isolated electron.

3. Pair and isolated positron topology: One photon produces a lepton pair, the other an isolated positron.

[1]Roughly 1% of neutrino interactions in ND280 are neutral-current π^0 interactions and only 15% of these are contained within the tracker region. These numbers are the result of initial investigations of the Monte Carlo sets 4.1.6 by the author.

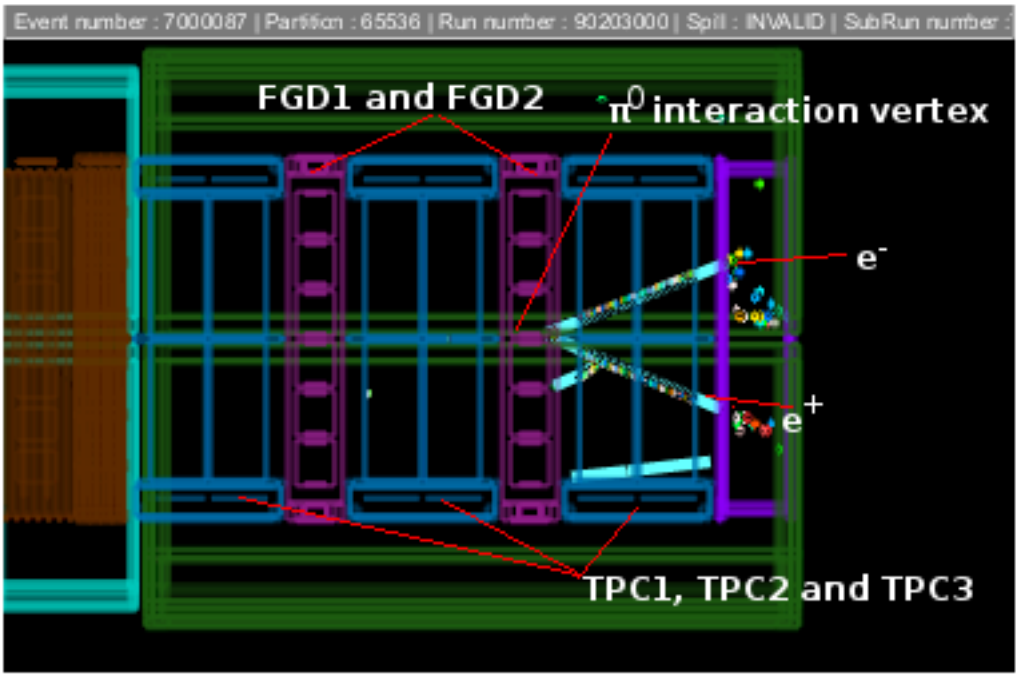

Figure 4.2: A reconstructed event from the Monte Carlo simulation, showing an isolated electron and isolated positron topology

4. Twin isolated electron topology: Both photons only produce reconstructable electrons.

5. Twin isolated positron topology: Both photons only produce reconstructable positrons.

6. Isolated electron and isolated positron topology: One photon produces an isolated electron, the other an isolated positron. An ND280 event display of such an interaction is shown in figure 4.2.

Each of these topology types require different reconstruction criteria to maximise the reconstructed signal and reduce the background. The actual kinematic variables used to reconstruct each topology are kept the same among the 6 topologies, to simplify systematic error evaluation.

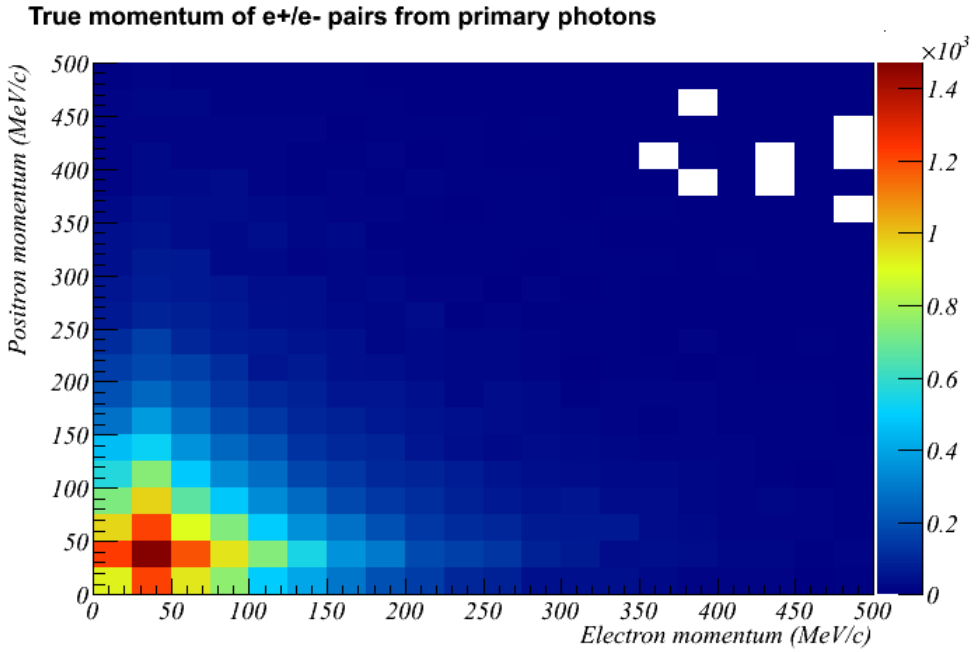

Figure 4.3: Distribution of momenta for e^+/e^- pairs from converting primary decay photons in neutrino interactions. All plots in this chapter are derived from Monte Carlo information.

4.1.4 Overview of the ND280 Reconstruction Software

For a given particle traveling through ND280, the reconstruction software attempts to amalgamate reconstructed objects from each subdetector into a single object, called a Global Particle Identity (referred to as PID). Each PID contains global kinematic information as well as information specific to each subdetector.

This analysis relies heavily on the Time Projection Chambers to provide accurate momentum reconstruction: the Fine Grain Detectors are limited by their 1 cm tracking resolution combined with their small size. However, the TPCs cannot reconstruct the momenta of an object without timing information from the FGD. As a result, each lepton must have both an FGD and a TPC component in order to be reconstructed. Scenarios where the primary photon converts in

the TPC, not the FGD, will not be reconstructed. However because the detector mass of the TPC is low this does not greatly affect the available statistics.

The latter case is unavoidable because, at this stage in time, the reconstruction software is unable to calculate the momenta of objects isolated within an FGD.

4.1.5 Truth Studies and Monte Carlo generation

The core of the analysis begins at the bunch level. Data from T2K's neutrino beam is recorded in event spills. Each spill consists of eight bursts of neutrinos produced by the beam from one filling of the accelerator, with all eight bursts being emitted within five micro-seconds [7]. Beam spills are referred to as 'events', with each burst referred to as a 'bunch'. Each bunch can contain multiple neutrino interaction vertices. The Monte Carlo keeps a record of every particle that travels through the detector and groups the particles according to which bunch they belong to.

The Monte Carlo information contains two sets of information: truth information and reconstructed information. The truth information contains accurate descriptions of the particles produced by the Monte Carlo from information provided by the NEUT event generator, described in section 4.1.6. The reconstructed information contains descriptions of particles produced by the Monte Carlo after they have been run through ND280's reconstruction software. i.e. the truth information is an accurate record of what happens, whilst the reconstructed information is a simulation of what ND280 will report from real data interactions.

The truth information is used to conduct a truth study on every neutrino interaction, in order to class the interaction as 'signal' or 'background'. Definitions

of signal and background are given in subsection 4.1.1.
It is important to classify the entire bunch as either signal or background, because when analysing data the number of neutrino interaction vertices cannot be accurately known. Any bunch with signal interactions is classed as signal, to do otherwise would cause the real NC1π^0 cross-section to diverge from the measured NC1π^0 cross-section.

4.1.6 Data and Monte Carlo Samples

This section of the analysis was performed on a set of Monte Carlo events generated to mimic neutrino interactions in ND280. The Monte Carlo information was generated by the NEUT event generator. The NEUT event generator simulates all possible neutrino-on-nuclei interactions in the kinematic region covered by T2Ks neutrino beam (600MeV peak [7]). A list of the physics interactions simulated by NEUT is shown below in table 4.2.
Two Monte Carlo sets were used to reflect the two different beam profiles used during data taking. These are labeled Beam B samples and Beam C samples. The results of the analysis for the Beam B and Beam C samples were scaled to their respective protons-on-target. Protons-on-target refers to the number of protons fired into the T2K neutrino beam's target (the T2K neutrino beam is further described in section 3.1.1).

The finished analysis was run on data measured in ND280 from real neutrino interactions. The data is split up into different run periods denoted as Run 1, Run 2, Run 3 and Run 4. Run 1 data is not used in this analysis because the barrel ECal modules were not installed for this run and the amount of protons

Interaction Name	Interaction Formula
Charged-current quasi-elastic scattering	$\nu_l n \rightarrow l^- p$
Charged-current pi+ production (from proton)	$\nu_l p \rightarrow l^- p \pi^+$
Charged-current pi+ production (from neutron)	$\nu_l n \rightarrow l^- n \pi^+$
Charged-current π^0 production	$\nu_l n \rightarrow l^- p \pi^0$
Charged-current coherent pion production	$\nu_l A \rightarrow l^- A \pi^+$
Charged-current resonant gamma production	$\nu_l n \rightarrow l^- p \gamma$
Charged-current multiple pion production	$\nu_l N \rightarrow l^- N' n \pi$
Charged-current resonant eta production	$\nu_l n \rightarrow l^- p \eta$
Charged-current resonant kaon production	$\nu_l n \rightarrow l^- \Delta K^+$
Charged-current deep inelastic scattering	$\nu_l N \rightarrow l^- N'$mesons
Neutral-current π^0 production (from neutron)	$\nu_l n \rightarrow \nu_l n \pi^0$
Neutral-current π^0 production (from proton)	$\nu_l p \rightarrow \nu_l p \pi^0$
Neutral-current coherent π^0 production	$\nu_l A \rightarrow \nu_l A \pi^0$
Neutral-current π^- production	$\nu_l n \rightarrow \nu_l p \pi^-$
Neutral-current π^+ production	$\nu_l p \rightarrow \nu_l n \pi^+$
Neutral-current γ production (from neutron)	$\nu_l n \rightarrow \nu_l n \gamma$
Neutral-current γ production (from proton)	$\nu_l p \rightarrow \nu_l p \gamma$
Neutral-current multiple pion production	$\nu_l N \rightarrow \nu_l N n \pi$
Neutral-current eta production (from neutron)	$\nu_l n \rightarrow \nu_l n \eta$
Neutral-current eta production (from proton)	$\nu_l p \rightarrow \nu_l p \eta$
Neutral-current K^0 production (from neutron)	$\nu_l p \rightarrow \nu_l \Lambda K^0$
Neutral-current K^+ production	$\nu_l p \rightarrow \nu_l \Lambda K^+$
Neutral-current deep inelastic scattering	$\nu_l N \rightarrow \nu_l N'$mesons
Neutral-current elastic scattering (from neutron)	$\nu_l n \rightarrow \nu_l n$
Neutral-current elastic scattering (from proton)	$\nu_l p \rightarrow \nu_l p$

Table 4.2: List of physics processes simulated by the NEUT Monte Carlo event simulator. Sources taken from [13]

on target is very small compared to the other runs.

The protons-on-target used in the Monte Carlo samples as follows:

1. Beam B Monte Carlo: 9.465×10^{20} pot
2. Beam C Monte Carlo: 3.009×10^{21} pot

The protons-on-target used when generating data is as follows:

1. Run 2 = 0.78377×10^{20} pot
2. Run 3 = 1.34821×10^{20} pot
3. Run 4 = 1.20428×10^{20} pot

The Monte Carlo sets came in two forms: basket and magnet. Magnet Monte Carlo simulates the entire ND280 detector, whilst basket Monte Carlo excludes the magnet. The magnet, due to its mass, is a major source of neutrino interactions and serves to add significant backgrounds, whilst the basket Monte Carlo is cleaner. The magnet Monte Carlo was used for this analysis, as it is a more authentic representation of neutrino interactions within ND280. Monte Carlo describing sand muons was also used to determine the predicted amount of background. Neutrinos from the beam can interact with the sand and rock surrounding the ND280 detector pit, producing muons. These muons can then enter ND280 and mimic neutrino interactions. Monte Carlo describing this process is referred to as the sand muon Monte Carlo.

All of the available magnet Monte Carlo was used. However, due to the limited statistics of this analysis the statistical errors from the Monte Carlo samples are noticeable. The errors shown on the efficiency and purity histograms in chapter 5 and 6 are the statistical error on the Monte Carlo.

Definition of FGD volumes

The signal definition requires a neutrino interaction to occur within an FGD. What 'within an FGD' actually constitutes needs to be rigidly defined. This analysis uses two FGD volume delimiters: the fiducial volume and the extended volume.

The x and y fiducial volume coordinates for FGD1 and FGD2 are -874.51 $<$ x $<$ 874.51, -819.51 $<$ y $<$ 929.51. The z-coordinates for FGD1 are 136.88 $<$ z $<$ 446.96. For FGD2 the z-coordinates are 1474.05 $<$ z $<$ 1806.96. [76].

The extended volume x and y coordinates are -874.51 $<$ x $<$ 874.51, -867.56 $<$ y

< 977.56. For FGD1, the z-coordinates are $116.045 < z < 620.00$ and for FGD2 they are $1474.05 < z < 2000.0$

The difference between fiducial and extended volume arises because particles need to travel through at least 4 FGD planes before their directionality can be accurately defined. This is because each FGD plane is oriented in either the XZ or YZ axis. A given track has to pass through both an XZ and a YZ plane for its 3D position to be calculated. For a given track, two 3D coordinates are required to discern the track's direction, which requires 4 planes in total. The extended volume is the maximal volume of every FGD scintillator bar. The fiducial volume is a subset of the detector volume which excludes the exterior scintillator bars, to ensure that any particle originating in the fiducial volume can be accurately reconstructed. The extended volume also includes some detector dead-space between subdetector modules and extends slightly into the FGD's corresponding downstream TPC to better enable particle identification.

The signal definition requires the neutrino interaction to occur within the FGD extended volume. To do otherwise would cause a subset of the background to be neutral-current π^0 producing interactions occurring in the exterior FGD planes. This background would have an identical kinematic profile to the signal and would corrupt the definition of 'signal', because the background would contain the exact processes that this analysis is trying to measure.

4.1.7 Lepton Cataloging

For a given bunch, a list of every electron-like candidate is created, to be later classified into pairs or isolated leptons. A basic set of selection cuts is applied at this stage:

A position cut: The candidate must have a starting position in either an FGD or TPC2 or TPC3. Hypothetically, any π^0 decay photons could propagate into a TPC before converting into an e^+/e^- pair. Events which propagate into the ECals are beyond the scope of this analysis, they are described in other analyses listed in section 4.1.

A direction cut: The candidate must have a forward-going direction. Signal events which produce leptons which propagate upstream are beyond the scope of this analysis due to the limited ability to accurately reconstruct backwards-going tracks.

A charge cut: The candidate must have a charge of +1 to be tagged as a positron and -1 to be tagged as an electron.

This list is then used to construct pairs and isolated leptons. This measure saves on computer processing time, as the pair and isolated reconstruction algorithms need to process a greatly reduced list of particles.

4.1.8 Kinematic Variables

The following kinematic variables are used in this analysis:

1. The starting position of tracks beginning in an FGD
2. The entry coordinates when a track enters the TPC from an FGD
3. The momenta of a track as it enters a TPC
4. The charge of FGD and TPC tracks. The charge is determined by whether the track curvature is positive or negative

5. The 'pull' values of tracks as they travel through the TPC. The pull values are proportional to dE/dx and provides a crude form of particle identification. Pulls are described below in section 4.1.9 and 3.2.5

6. The fiducial volume of the FGD.

7. The number of objects in other detectors for any given event. This is used as a veto.

4.1.9 Electron Pull

The electron pull is a quantity derived from the rate of energy loss of a particle as it travels through the tracker region of ND280, as described in section 3.2.5.

The reconstruction software measures the rate of energy loss for a given track and compares it to the energy loss of an idea electron. An electron pull of zero means the track in question matches the ideal scenario perfectly. Larger pull values correspond to a worse alignment. There are similarly muon pulls, charged pion pulls and kaon pulls.

Figure 4.4 shows the measured rate of energy loss through the PEP4/TPC-9 detector, an early time-projection-chamber [10], for electrons, muons and pions. The electron pull, defined in section 3.2.5, is a measure of how far the measured $\frac{dE}{dx}$ of the particle is from the expected $\frac{dE}{dx}$ for an ideal electron at the measured momentum of the track. This analysis focuses on tracks with momenta between roughly 50 MeV and 400 MeV. At the lower and upper end of this momentum region, the energy loss of electrons is distinct compared with pions and muons, enabling the electron pull to be a potentially powerful discriminator between electrons and muons and pions.

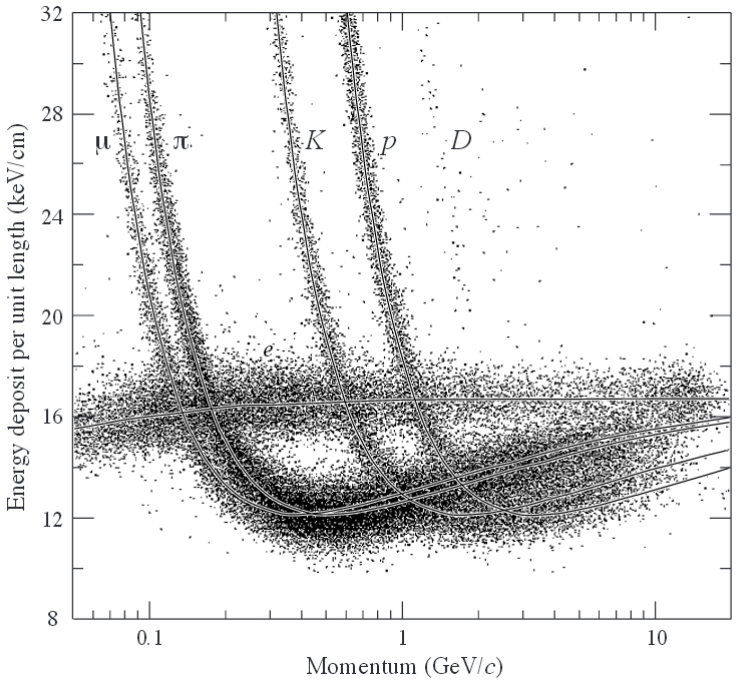

Figure 4.4: A histogram the measured energy loss for electrons, muons and pions traveling through the PEP4/9-TPC detector [10]

4.1.10 First and Second Generation Cuts

The analysis was performed in two stages. Pairs, isolated electrons and isolated positrons were reconstructed in the first stage and selection cuts specific to these three topologies were applied. These cuts are referred to as first generation cuts.

Next, π^0 candidates were constructed from the six different topology types: twin pair, twin isolated electron, twin isolated positron, pair and isolated electron and pair and isolated positron. Selection cuts specific to these six topologies were then applied. These cuts are referred to as second generation cuts.

When beginning this analysis, the first set of selection cuts were loosely defined by looking extensively at ND280 event displays depicting Monte Carlo information. The signal π^0 events were defined as being relatively clean events, with

low track multiplicities. Aside from tracks produced through π^0 decay photons, there was little activity in the detector. Tracks were very forward-going.

In contrast the background frequently had high track multiplicities, lacking the well defined electron/positron pairs exhibited in the signal. Muons were very prevalent in the background.

4.1.11 Evaluating Selection Cuts: Efficiency, Purity and Quality

This analysis reconstructs π^0 candidates by applying selection cuts to reconstructed quantities in order to remove background interactions. The value of cutting on a kinematic variable could be clearly defined and π^0 vertices have a relatively unique event topology, so the analysis lends itself to a cuts-based approach.

It is necessary to quantify the value of each of the selection cuts being considered in the overall analysis. Simply put; is the amount of background events removed from the analysis worth the amount of signal events which are lost? To answer this, the effects of a cut on the signal sample and background sample must be quantified. To this end, the terms **efficiency** and **purity** are defined. The efficiency represents how effective the analysis is at reconstructing signal events, whilst the purity represents how effective the analysis is at rejecting background events whilst retaining signal events. The definitions are as follows:

- **The efficiency is the probability of a signal event being selected by the analysis.**

- **The purity is the probability of a selected event also being a signal event (as opposed to a background event).**

To further clarify what efficiency and purity are, efficiency is represented by $\frac{n}{N}$, where n is the number of reconstructed signal events and N is the total number of signal events in the data set. Purity is represented by $\frac{n}{n+b}$, where b is the number of reconstructed background events. The efficiency and purity parameters are derived and further described, along with their associated errors, in subsections 4.1.12 and 4.1.13 respectively.
A figure of merit is required to evaluate the perceived value of any selection cut, defined as the **quality**. This is a function of the efficiency and purity, designed to minimise statistical and systematic errors. Quality is further described in subsection 4.1.14.

The efficiency, purity and quality histograms are designed to show the effects on the *overall selection* of a given selection cut, including events which cannot be affected by the cut. For example, when investigating the effect of track multiplicity on twin pair topologies, the efficiency, purity and quality values on the histogram represent the values of the entire selection. This includes, for example, pair and isolated electron events which aren't directly affected by a selection cut applied to the twin pair topology.

The design of the efficiency, purity and quality histograms was chosen to assess the effect of event migration. For example, imagine the scenario where an event where both a twin pair event and a pair and isolated electron event are reconstructed. Applying a twin pair tracker multiplicity cut removes the twin pair from the selection, however because the event is also a pair and isolated electron candidate, the overall number of reconstructed π^0s does not change.

A histogram showing the effectiveness of the selection cut should reflect this, which this design does.

4.1.12 Calculating Errors on Selection Cuts: Efficiency

The efficiency is defined as the probability of selecting an event, when the set of events being considered is limited to N, the total number of signal events. The variables and their definitions used in this subsection are defined in table 4.3. Because the efficiency is the probability of a 'success' when considering a population of events, it is a binomially distributed variable. The best estimate of the parameters of the efficiency function can be determined by considering the maximum likelihood of the efficiency function.

The likelihood function is the product of the probability of each event being observed. The true efficiency function is unknown, but by finding the parameters which maximise the likelihood function one can find the best estimate using available data. In this chapter the log-likelihood function is used instead of the likelihood function because it is easier to manipulate. The conclusions drawn, however, are the same [10, pp. 303–304].

The likelihood of the efficiency function is given by equation 4.1.

$$\mathcal{L} = \epsilon^n \left(1 - \epsilon\right)^{N-n} \tag{4.1}$$

In equation 4.2, we maximise the log-likelihood function and determine the efficiency to be defined as $\epsilon = \frac{n}{N}$. This serves as a useful check to justify the validity of the approach taken as this is the common definition of an efficiency.

$$\frac{\partial \ln(\mathcal{L})}{\partial \epsilon} = \frac{\partial}{\partial \epsilon}\left(n \ln \epsilon + (N-n)\ln(1-\epsilon)\right) = 0 \tag{4.2}$$

$$\frac{\partial \ln(\mathcal{L})}{\partial \epsilon} = \left(\frac{n}{\epsilon} + (N-n)\frac{-1}{1-\epsilon}\right) = 0 \tag{4.3}$$

$$n(1-\epsilon) = (N-n)\epsilon \tag{4.4}$$

$$\epsilon = \frac{n}{N} \tag{4.5}$$

Because the log-likelihood efficiency function has an asymptotic chi-squared distribution, the error on the efficiency can be defined as the difference between the true efficiency and the value which is one standard deviation from the true efficiency [10, pp. 303–304]. Consider the ratio between the log-likelihood function for an 'estimated' efficiency with a free-floating parameter ϵ' and the log-likelihood function for the 'true' efficiency which is considered a fixed value. This ratio only has one degree of freedom, ϵ'. Because the ratio is a chi-squared distribution the standard deviation is 2. We therefore want to find the values ϵ' can take for the log-likelihood ratio to equal -2, shown in equation 4.8. The negative sign is due to the fact that we are maximising the log-likelihood, which is the same as minimising the negative log-likelihood because the log-likelihood function is monotonic.

$$-\ln \frac{\mathcal{L}}{\mathcal{L}_{\text{estimated}}} = 2 \tag{4.6}$$

$$-1/2 = n\ln\epsilon' + (N-n)\ln(1-\epsilon') - n\ln\frac{n}{N} - (N-n)\ln\left(\frac{N-n}{n}\right) \tag{4.7}$$

$$n\ln\epsilon' + (N-n)\ln(1-\epsilon') = -\frac{1}{2} + n\ln\frac{n}{N} + (N-n)\ln\left(\frac{N-n}{n}\right) = C \tag{4.8}$$

N and n are, for a given selection cut value, fixed. Therefore the right hand side of equation 4.8 is a constant and, for easier manipulation, represented by the

constant C. This equation can be simplified by taking exponentials, shown in equation 4.11.

$$\ln \epsilon'^n \cdot (1-\epsilon')^{N-n} = C \tag{4.9}$$

$$\epsilon'^n (1-\epsilon')^{N-n} = e^C \tag{4.10}$$

$$\epsilon'^n (1-\epsilon')^{N-n} - e^C = 0 \tag{4.11}$$

In equation 4.11, values for ϵ' which satisfy the equation represent the upper and lower error bounds on the efficiency. There is no analytical solution to find ϵ' and so a numerical method must be implemented. The equation has two roots, representing the lower and upper error bounds. As e^C is always positive, ϵ' cannot take a value either greater than 1 or less than 0, or else the LHS of the equation would be negative and therefore not provide a valid solution to the equation. These limits are what one would intuitively expect for the upper and lower bounds of the efficiency.

To find the roots of equation 4.11, numerical methods are applied. Initially a bisection method is used. As the roots of the equation lie either side of the measured efficiency, $\epsilon = \frac{n}{N}$, a small value is continuously added (or subtracted to find the lower root value) to ϵ. The best value of ϵ' is taken and further refined via the secant method.

Table 4.3 gives an overview of the variables used in this subsection.

4.1.13 Calculating Errors on Selection Cuts: Purity

The derivation to find the uncertainty on the measured purity, p, is very similar to the efficiency. The purity is defined as the probability of measuring a reconstructed signal event, from a combined population of reconstructed signal and

Symbol	Definition
N	total number of signal events
n	number of reconstructed signal events
ϵ	efficiency
$\delta\epsilon$	the difference between the mean and the half-width maxima of the maximum likelihood distribution of the efficiency function. This is considered the uncertainty.
ϵ'	$\epsilon + \delta\epsilon$

Table 4.3: List of variables used to define the uncertainty on the efficiency.

reconstructed background events. This definition provides a measure of how effective a selection cut is at eliminating background events. Ideally, only signal events would be reconstructed by our analysis, giving a purity of 1. In the worst case scenario, only background events are reconstructed and no signal events are reconstructed, giving a purity of 0. This definition of purity provides both reasonable limits and directly reflects the effectiveness of a selection cut at improving the analysis in the context of increasing signal and rejecting background.

The derivation to find the uncertainty is almost identical to the efficiency derivation, the only difference is a different definition of N is used. Previously, N was the population of total signal events. For the purity, N is the combined population of reconstructed signal and reconstructed background events. $N = n + b$. Therefore, the purity is parametrised as $p = \frac{n}{n+b}$. The uncertainty on the purity is found using equation 4.12. p' represents an estimated purity which is exactly one standard deviation from the true purity, p.

$$p'^{n}\left(1-p'\right)^{N-n} - e^{C} = 0 \tag{4.12}$$

Where $C = -\frac{1}{2} + n\ln\frac{n}{N} + (N-n)\ln\left(\frac{N-n}{n}\right)$. Again the roots of the equation for a given selection cut value are found numerically. A rough estimation of the root is found using the bisection method, which is then improved using the secant

method.

4.1.14 Evaluating Selection Cuts: Quality

This cuts-based analysis requires a metric which is capable of evaluating how effective the analysis is at achieving its goal: reconstructing neutral-current π^0-producing neutrino interactions. One desires both a high efficiency and purity, but in practise efficiency is often increased only at the expense of purity and vice-versa. The question that needs to be answered is, how should efficiency and purity be optimised?

This analysis was expected from the start to be very statistics-limited. Initial truth studies indicated that approximately 1% of the Monte-Carlo neutrino interactions were neutral-current π^0-producing interactions. Of these, approximately 15% were produced in the tracker region of ND280. Considering the very limited statistics available, the analysis was initially optimised to minimise statistical errors when measuring the interaction cross-section.

This is addressed in a paper published by the DESY laboratory [77], the derivation which follows is based on this publication. Consider a total number of expected measured events, $< T >$. $< T >=< s > + < b >$, where $< s >$ is the expected number of measured signal events and $< b >$ is the expected number of measured background events. The signal cross-section, σ_s, is expressed in terms of the efficiency in equation 4.13.

$$< s >= \epsilon \cdot \int \mathcal{L}\mathrm{d}t \cdot \sigma_s + < b > \tag{4.13}$$

When T events are observed, the total measured cross-section, σ_m, can now be expressed as:

$$\sigma_m = \frac{T - < b >}{\epsilon \cdot \int \mathcal{L}\mathrm{d}t} \tag{4.14}$$

In this context, $\mathcal{L}$ is the luminosity, *not* the likelihood. The uncertainty on the measured cross section is therefore:

$$\delta\sigma_m = \frac{\delta T}{\epsilon \cdot \int \mathcal{L}\mathrm{d}t} = \frac{\sqrt{T}}{\epsilon \cdot \int \mathcal{L}\mathrm{d}t} \tag{4.15}$$

This uncertainty assumes that the background is well understood and therefore has no associated statistical error.

The purity, defined in subsection 4.1.13, is $p = \frac{<s>}{<s>+<b>}$. The total expected number of events, $< T >$ can now be written as:

$$< T >= \frac{< s >}{p} = \frac{\epsilon \cdot \int \mathcal{L}\mathrm{d}t}{p} \tag{4.16}$$

Therefore the uncertainty on the *expected* measured cross-section is:

$$< \delta\sigma_m >= \frac{\sqrt{\frac{\epsilon \cdot \int \mathcal{L}\mathrm{d}t \cdot \sigma_s}{p}}}{\epsilon \cdot \int \mathcal{L}\mathrm{d}t} = \sqrt{\frac{\sigma_s}{\epsilon \cdot p \cdot \int \mathcal{L}\mathrm{d}t}} \tag{4.17}$$

Therefore the ratio of the measured cross section to the uncertainty is:

$$\frac{< \sigma_m >}{< \delta\sigma_m >} = \sqrt{\frac{\epsilon \cdot p \cdot \int \mathcal{L}\mathrm{d}t}{\sigma_s}} \tag{4.18}$$

From equation 4.18, one can minimise the statistical uncertainty on the cross section by *maximizing* the product of the efficiency and the purity. Therefore the quality, q, was initially defined as $q = \epsilon \cdot p$.

Initially this definition was used to evaluate the selection cuts used in this analysis. Selection cuts that maximised the quality were used in the final analysis.

However, this derivation does not factor in the potential for large systematic errors on the selected background. Initial results from the analysis using this quality derivation returned a low purity, initially 19%. It became apparent that

the systematic errors on the background would be at least 15%, due to uncertainties in the neutrino flux combined with uncertainties on the neutrino interaction cross-sections used in the Monte Carlo samples. Because of this large systematic error, a decision was made to change the quality definition to reflect a desire for a higher purity. Enough progress was made that statistical errors were not going to dominate the analysis, however with a low purity, systematic errors on a large background became a concern.

After qualitative experimenting with different metrics for quality, a new definition of $q = \epsilon \cdot p^2$ was chosen. With this new definition, the overall efficiency of the analysis was only lowered by a few percent, with noticeable improvements in the purity (shown in chapter 9). Any further biasing of purity did not increase the overall purity of the analysis by more than 1%, but at a cost of a significantly lowered efficiency (which dropped significantly when purity was more heavily weighted. e.g. when quality was defined as $q = e \cdot p^3$ the efficiency dropped by a factor of three).

The uncertainties on the quality is not shown on the histograms. When evaluating the effect of a given set of selection cuts using the quality, the effect of the uncertainty on the quality is negligible. This is because what is being considered is the relative difference between two different quality values. Consider a selection cut, e.g. cutting on the number of tracks in the event. Minor alterations in the cut value will either add or remove a small number of events to the overall selection. The extra level of uncertainty added by this very small set of events will be small when considered as a ratio to the overall quality value. For this reason the uncertainties on the quality are not displayed.

4.2 N-1 Histograms

Showing the effect of a selection cut when no other cuts are applied can be misleading, because the effect of the selection cut in question could be already covered by another selection cut used in the analysis. i.e. the selection cuts being applied are not independent of each other.

In order to assess the actual effect a selection cut will have on the final analysis, N-1 histograms are used. These are histograms that show the effect of a given selection cut, when all other selection cuts used in the analysis are applied. The advantage of this approach is that any increase in efficiency, purity or quality shown by the histograms will lead to a direct increase in the corresponding quantity in the final analysis. The disadvantage is that these histograms can be statistics-limited, as large amounts of information are already removed by prior selection cuts.

The histograms detailing selection cuts can show the effects of selection cuts under two scenarios: using the 'base' selection and using the 'final' selection. The 'base' selection shows the effect of a selection cut when no other selection cuts are applied. The 'final' selection shows the effect of a selection cut when all selection cuts are applied. When evaluating the value of a selection cut, the final selection is preferred as it directly measures the effect of the selection cut being investigated and factors in the effects of all of the other selection cuts. However, when this selection is statistics limited, the base selection is used to highlight important trends shown by the final selection. Histograms using the final selection are N-1 histograms.
All of the first generation cuts are evaluated using the final selection. All second generation cuts are evaluated using the final selection, with the exception

of the cuts associated with the pair and isolated electron topology. This topology only contained seven signal events from all Monte Carlo samples. Given the extremely limited number of statistics for this topology, the base selection was used instead.

When showing distributions of signal and background for a given selection cut, the numbers shown on the histogram refer to the number of signal/background candidates reconstructed from the Monte Carlo samples for the specific topology in question. For example, figure 5.1 in section 5.1.1 shows the invariant mass of electron/positron pair candidates.

When showing the effect of a given selection cut, the efficiency, purity and quality of the selection cut are plotted. The efficiency, purity and quality values plotted refer to the overall efficiency and purity of the entire analysis for a given selection cut (e.g. an efficiency of 0.3 implies the efficiency of the analysis with that cut value is 30%). The absolute values of the efficiency and purity are plotted on separate axis. The absolute quality values are not plotted because what matters is not the absolute value of the quality, but rather the selection cut which maximises the quality. The quality is defined as $e \times p^2$ and the absolute value returned by the quality does not convey any new information. Therefore in order to simplify the histograms the absolute values of the quality are not plotted.

All histograms in chapter 5 and 6 display the final selection cut chosen for the specific variable being investigated. This is represented by a dotted vertical line. An arrow is also drawn onto the histogram which reflects the range of values which *pass* the selection cut and are not rejected from the analysis.

4.2.1 Muon Veto

By definition, an NCπ^0 interaction will not produce a muon and so removing events containing muons is a useful way of removing backgrounds. The Charged-Current Quasi Elastic analysis [78] created an optimal set of criteria for reconstructing muons in ND280. This analysis modifies the criteria laid out in their work and uses it as a veto. This veto also removes π^+ and π^- events because their pull signatures are similar to muons at the energy scales of this analysis. This is beneficial because neither of these particles are included in the signal definition. The muon reconstruction is as follows:

1. The muon candidate must have at least 18 hits in a TPC.
2. The muon candidate must have a negative charge.
3. The muon candidate must have a muon pull between -2 and 2.
4. The muon candidate must have an electron pull of less than -2.

4.2.2 Summary

This chapter has described the approach that was taken when reconstructing neutral-current π^0-producing neutrino interaction vertices. Selection cuts are first applied with the goal of reconstructing photon candidates, through first-generation cuts. These photon candidates are combined to form π^0 candidates, with a second generation of selection cuts applied to improve the purity of the analysis. The first generation and second generation selection cuts are described in detail in chapters 5 and 6 respectively.

Chapter 5

Evaluating First Generation Selection Cuts: γ Reconstruction

The first generation of cuts refer to the selection cuts used when reconstructing individual photon candidates, i.e. (lepton pairs, isolated electrons and isolated positrons). These photon candidates are then paired together in chapter 6 to create π^0 candidates. This chapter analyses each selection cut that was applied to reconstruct the photon candidates.

The final set of first generation cuts are described at the end of the chapter in table 5.1 and table 5.2 for pairs and isolated leptons respectively.

The definitions of signal and background used in this chapter are described in section 4.1.1. The definitions of efficiency, purity and quality used in this chapter are described in sections 4.1.12, 4.1.13 and 4.1.14.

All first generation cuts are analysed using the final selection of cut parameters (the 'final' and 'base' selections are described in section 4.2).

5.1 e^+/e^- Pair Reconstruction

Reconstructing lepton pairs is a key step in reconstructing NC π^0 neutrino interactions. Initially, all candidate π^0s in this analysis were constructed using lepton pairs, however this approach yielded unacceptably low statistics. The first attempt of the analysis had a predicted number of reconstructed π^0s using the known number of protons on target (referred to as pot) of only 2. Nevertheless, pair reconstruction forms a vital part of the overall analysis, and the following sections detail the tuning of selection cuts used to maximise the quality of this reconstruction algorithm.

5.1.1 Invariant Mass Cut

The invariant mass of a potential electron/positron pair has the potential to be a key discriminator between signal and background. The invariant mass of an electron/positron pair being produced from a photon is at least 1.02 MeV/c^2. Despite being produced from a photon, the invariant mass of the lepton pair is not zero MeV/c^2, because the pair is created from the interaction between a photon and a nucleon. Therefore the invariant mass is equal to $(\mathrm{p}_\gamma + \mathrm{p}_{p/n})^2$, where p_γ and p_n are the 4-momenta of the photon and nucleon respectively. The 4-momenta of the nucleon is not known, however its energy will be minimal compared to the photon.

The invariant mass of the pair is equal to:

$$(p_{e^+} + p_{e^-})^2 = (p_\gamma + p_{p/n})^2 \tag{5.1}$$

The LHS of this equation can be expanded to:

$$p_{e^+}^2 + p_{e^-}^2 + 2p_{e^+}p_{e^-} \tag{5.2}$$

$$= 2m_e + 2E_{e^-}E_{e^+} - 2\mathbf{p_{e^+}} \cdot \mathbf{p_{e^-}} \tag{5.3}$$

What is the minimum value this equation can take? The leptons are produced from a photon with a large amount of kinetic energy, and so will be highly boosted. The consequence of this is that the opening angle between the leptons will be very small, and $m_e << E_e$. Taking the opening angle to be 0, the invariant mass is equal to:

$$2m_e + 2E_{e^-}E_{e^+} - 2\sqrt{(E_{e^+}^2 - m_e^2)(E_{e^-}^2 - m_e^2)} \tag{5.4}$$

As $m_e << E_e$, the last two terms in the equation effectively cancel, giving a minimum momentum mass of $2m_e = 1.02$ MeV/c^2. This suggests that cutting on the invariant mass (excluding pair candidates with high invariant mass values) could improve the quality of the analysis.

Figure 5.1 shows the invariant mass for signal and background samples. The difference between the signal and background distributions is somewhat marginal. The effect of cutting on the invariant mass is shown in figure 5.2. There is a slight quality maximum at an invariant mass cut of 110 MeV/c^2, and that was the settled upon cut value (the cut is a maximum cut i.e. pair candidates with an invariant mass of 110 MeV/c^2 or higher are rejected).

An invariant mass cut was thought to be promising at first glance, however on inspection, the gains in purity were lower than expected. Whilst there is still a clear benefit in cutting on the invariant mass, there are similarities between the signal and background distributions. There were only two assumptions made

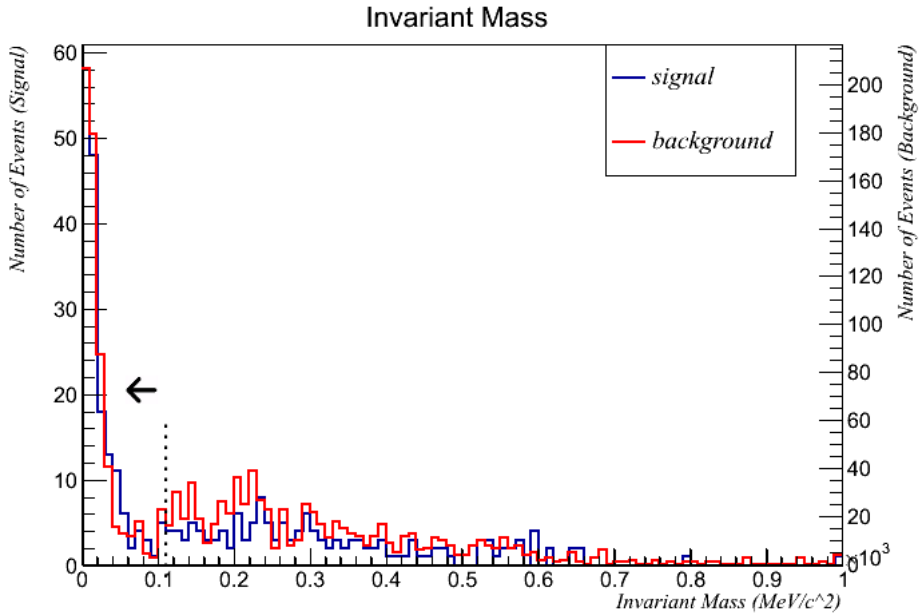

Figure 5.1: Invariant masses of lepton pair candidates. The signal and background event numbers refer to the number of candidates reconstructed in the Monte Carlo samples. The dotted line represents the chosen cut value. The direction of the arrow indicates the range of values which pass the cut.

to reach equation 5.4. That the invariant mass of the pair candidate was constructed from two electrons and/or positrons, and that the leptons are highly boosted. These two assumptions are likely to hold for **any** two selected tracks in the tracker region and so the background distribution will have similar characteristics to the signal distribution. The former assumption will hold because the majority of muons have been removed via the muon veto, and the majority of tracks remaining will be electrons and positrons due to their quantity. Particles are likely to be boosted because the majority of particles in the tracker regions will be one or two interactions away from the initial neutrino interaction vertex, and will carry a sizable fraction of the neutrinos initial kinetic energy.

As a result, the mass cut did not yield expected gains, however a cut to remove pair candidates above 110 MeV/c^2 does provide some benefit.

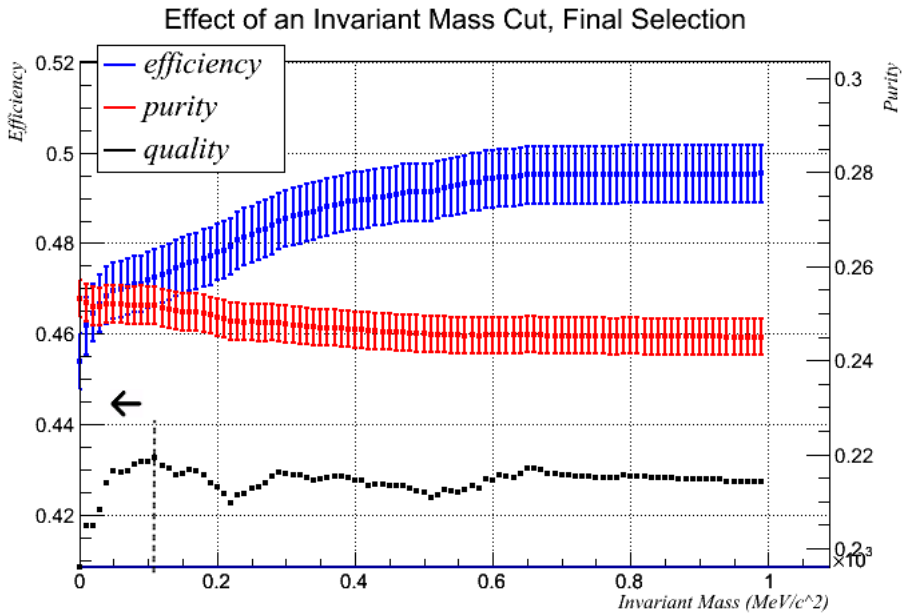

Figure 5.2: Effect of an invariant mass cut on pair reconstruction. The dotted line represents the chosen cut value. The direction of the arrow indicates the range of values which pass the cut.

5.1.2 Electron Pull

The definition of electron pull is described in section 4.1.9. Applying an electron pull cut has the potential to improve quality because it would prevent any non-electron-like particles from being misreconstructed. Two potential cuts can be applied; a maximum electron pull cut and a minimum electron pull cut.

For a given pair candidate, the electron pull of the lepton with the lowest electron pull is plotted in figure 5.3 for signal and background. When applying a minimum electron pull cut, it is the lowest electron pull value of the pair that matters, because if the lowest pull value passes the cut, the higher pull value will also pass by default.

The effect of a minimum electron pull cut is shown in figure 5.4. The effect of a maximum electron pull cut is shown in figure 5.5. For both histograms, the

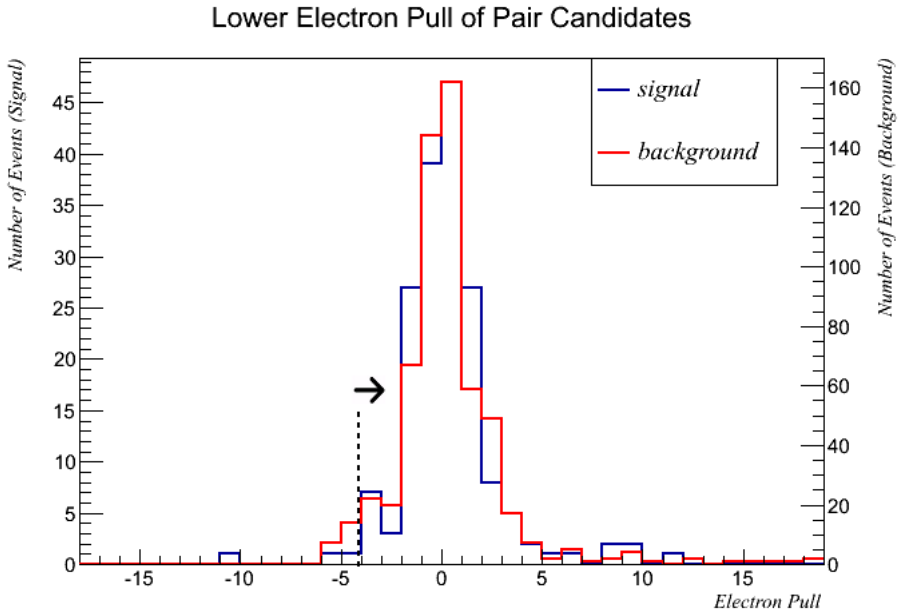

Figure 5.3: Electron pull of the lepton with the highest electron pull of the pair candidate.

quality plateaus somewhat, however there are marginal gains in quality when applying a minimum electron pull cut of -4, and a maximum electron pull cut of +5. The reason why the electron pull cuts only minimally increase quality is likely due to the fact that the background distribution is also made up largely of electrons and positrons. Applying tight cuts on the electron pull value would therefore not discriminate between signal and background.

Although the gains are marginal, the isolated lepton reconstruction discussed later uses a pull cut for more substantial gains. Introducing a pull cut for pair reconstruction therefore does not add any extra systematic errors to this analysis.

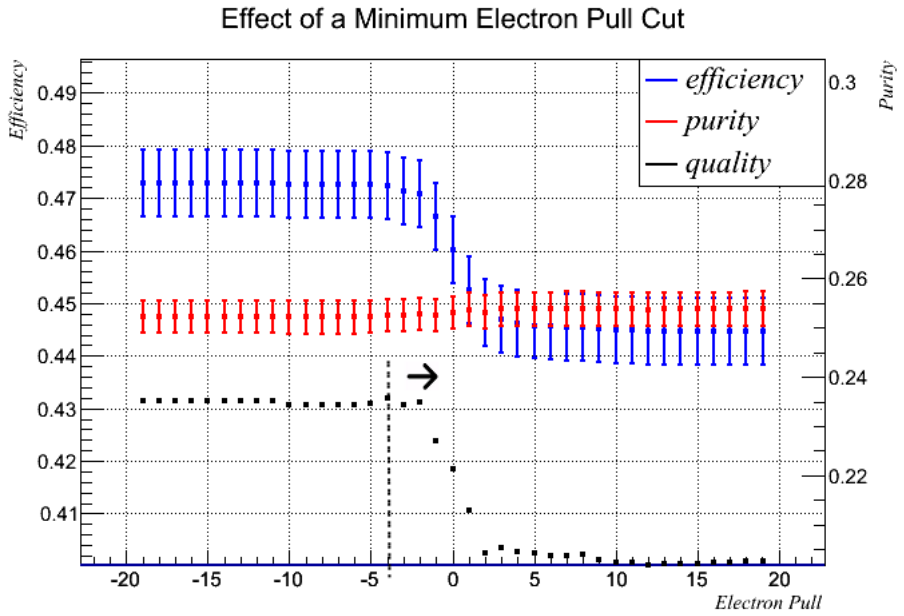

Figure 5.4: Effect of a minimum electron pull cut on pair reconstruction.

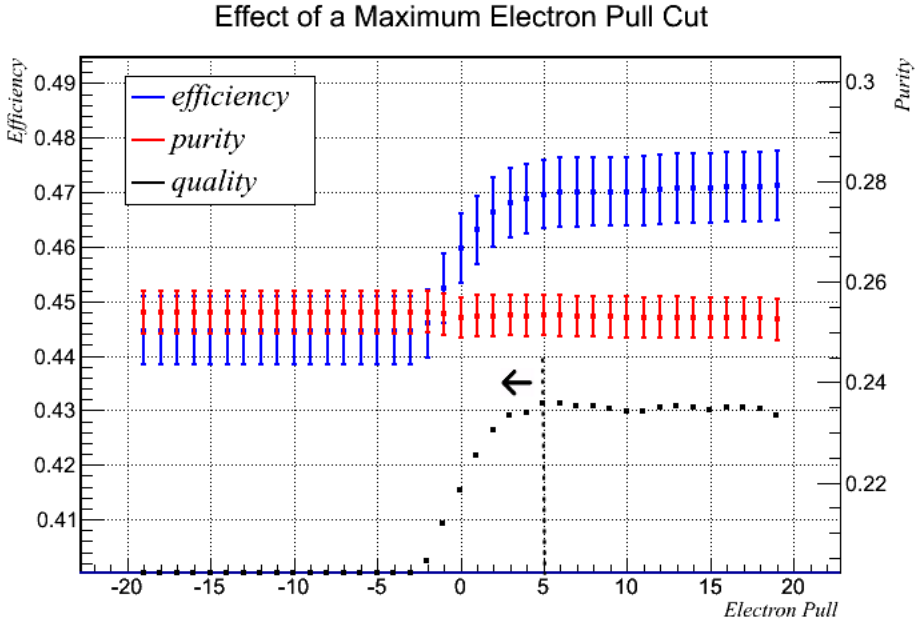

Figure 5.5: Effect of a maximum electron pull cut on pair reconstruction.

5.1.3 Momentum Cut

Applying a momentum cut could yield benefits for pair reconstruction, because the backgrounds are expected to be dominated with low-momentum tracks. The momentum of the lepton with the highest momentum per pair candidate is shown in figure 5.6, with the effects of a cut shown in figure 5.7.

Contrary to this expectation, there is no peak to this distribution. Both signal and background samples produced large numbers of leptons with very low momentum values. With a centre of mass energy of only 141 MeV/c^2, leptons produced from π^0 decays often have a momentum that is at the lower end of what ND280 can reconstruct. A minimum cut of 50 MeV/c was applied because 50 MeV/c is the limit of what the TPCs can reliably reconstruct. This is because at 50 MeV/c, the radius of curvature of the track is the same as the extent of the TPC in the z-direction, and so the tracks bend back on themselves. The TPCs reconstruct a 3-dimensional track by using an electric field to make electrons, ionised by the reconstruc;;.ted track, drift onto charge-sensitive pads. The time delay between successive TPC hits is then used to reconstruct the track. For a particle which bends back on itself, the time delay between successive TPC hits is no longer correlated, and the TPC reconstruction will reconstruct multiple tracks instead of a single track.

The effects of a maximum momentum cut are shown in figure 5.8. There is a slight maximum at around 1000 MeV/c. The selection cut effectively functions as a cut which removes extreme outlier events where the momentum of tracks in the detector is very high and therefore highly unlikely to be produced by a π^0 decay.

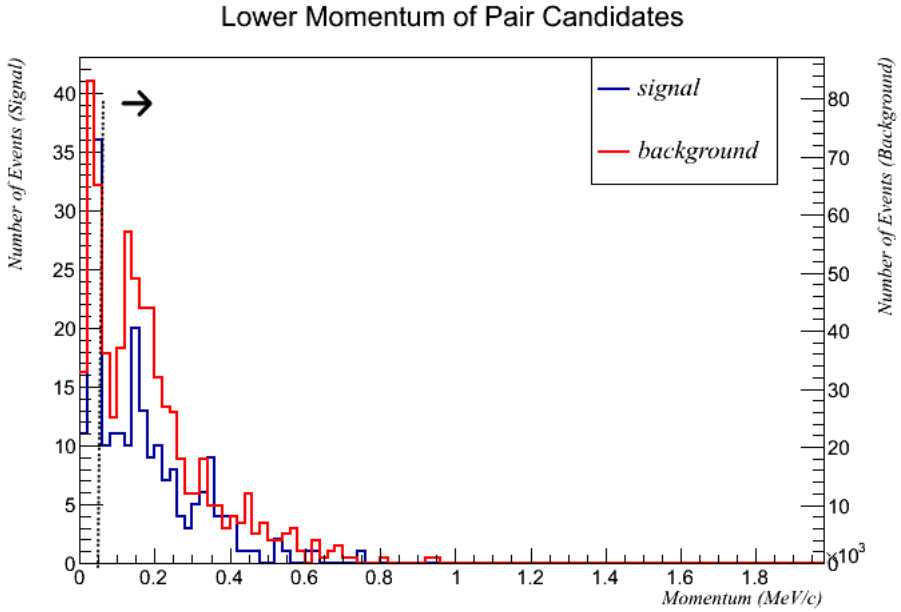

Figure 5.6: Momentum of the lepton with the highest electron pull of the pair candidate.

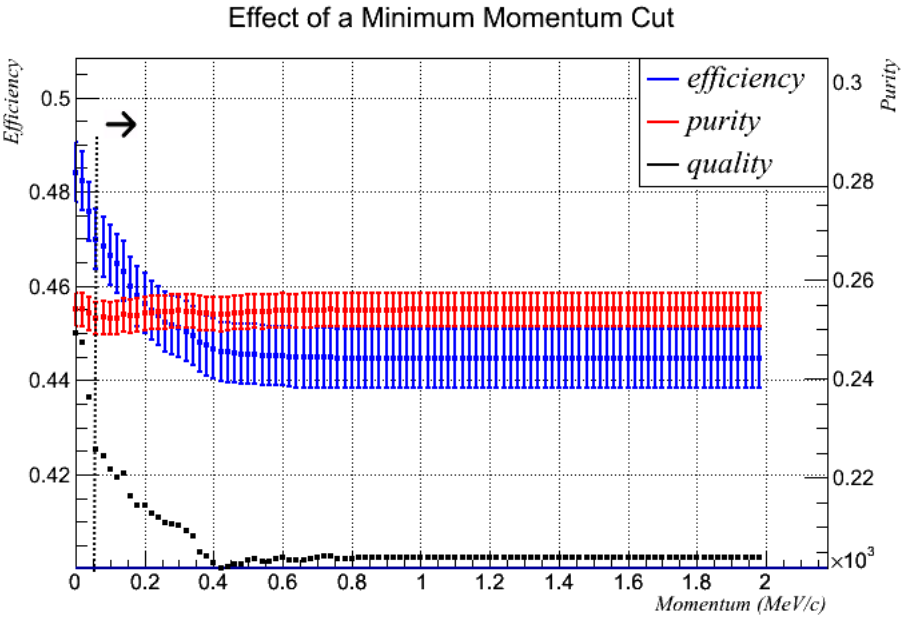

Figure 5.7: Effect of a minimum momentum cut on pair reconstruction.

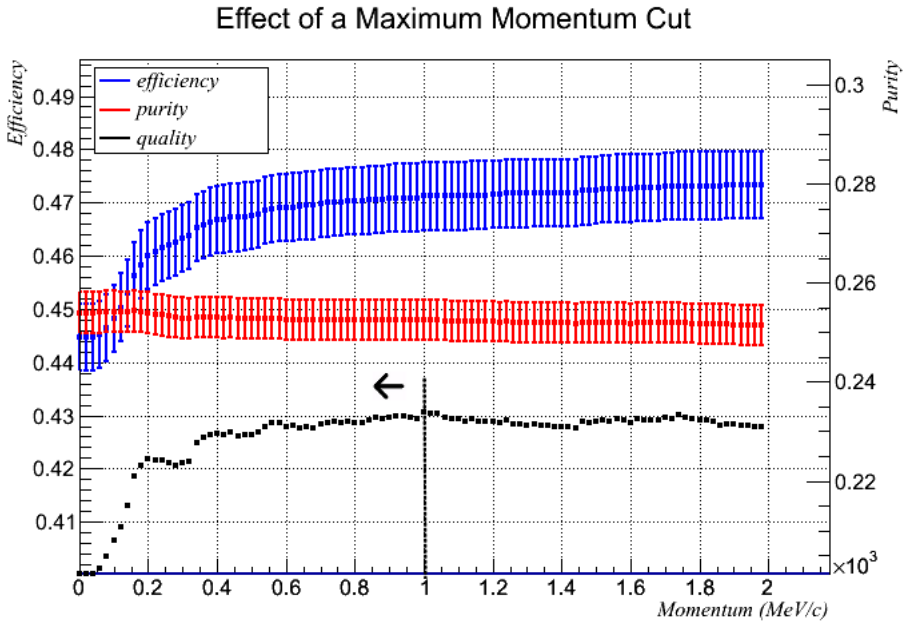

Figure 5.8: Effect of a maximum electron pull cut on pair reconstruction.

5.1.4 Extrapolated Separation Distance Cut

Ideally, any pair produced from a converting photon would have near-identical starting positions. This can cause the reconstruction software to attribute the detector 'hits' to one track only, as the software cannot attribute a detector hits to multiple tracks. To account for this, the starting positions for each lepton are taken, with the most-downstream lepton being projected upstream to the z-position of the most-upstream lepton. This is depicted in figure 5.9. The square root of the sum of the square of the distances between the leptons' X and Y positions are then compared. This value provides a measure of how close the lepton's origin points are to each other. This is plotted for signal and background in Fig 5.10.

This distribution has the majority of its signal and background entries between 0 and 10 mm. The spatial resolution of the FGDs is only 1 cm [7], and so the affect of cutting on distances <1 cm is questionable. Given the reconstruction

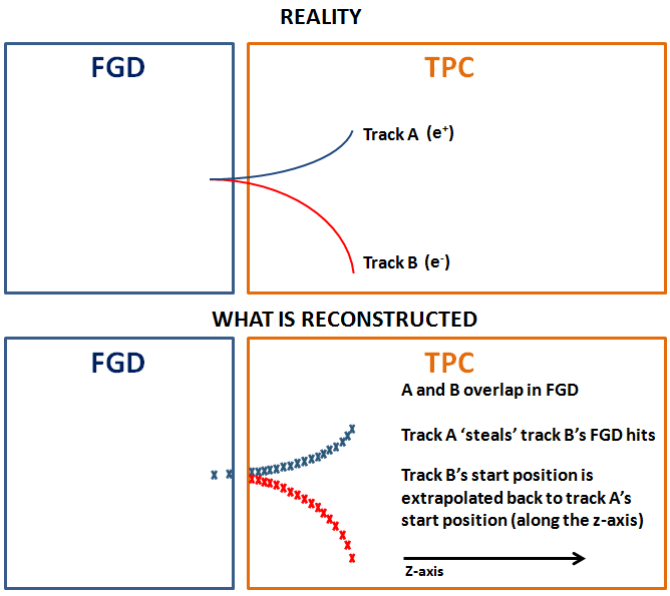

Figure 5.9: Scenario where, for an electron positron pair, one of the tracks is missing its FGD component. After extrapolation, the x-y distance between track A and track B should be small. If track A and track B do not originate from the same point, the x-y distance should be large.

limitations, the signal and background distributions are almost identical, and no cut value was applied.

5.1.5 Separation Distance Cut

As well as an extrapolated distance cut, it is worth considering the actual distance between lepton starting positions, especially considering the limited effect of an extrapolated distance cut.

The distance between lepton starting position is plotted in figure 5.11, with the effect of a cut in figure 5.12.

The results of this histogram were unexpected. Some of the pairs being reconstructed have a distance of 1200 mm between their reconstructed starting positions. This is likely due to the tracks from these pair candidates originating within different FGDs. Such pair candidates are highly unlikely to be genuine

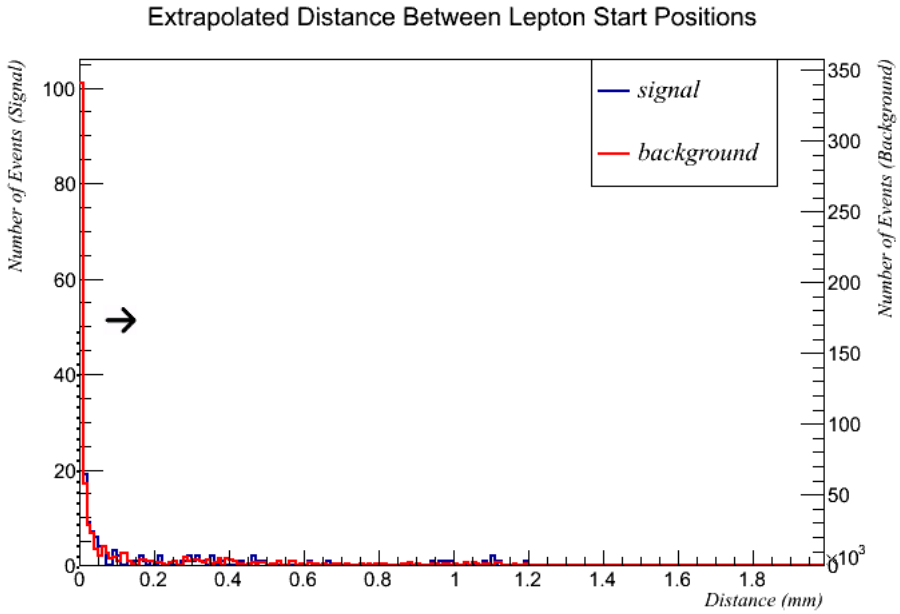

Figure 5.10: Extrapolated distance between lepton start positions.

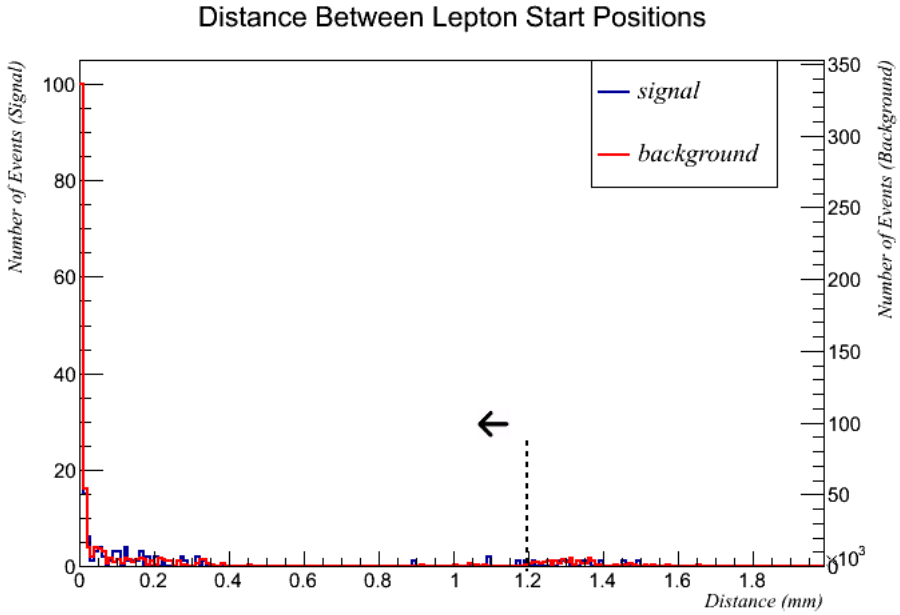

Figure 5.11: Absolute distance between between lepton start positions.

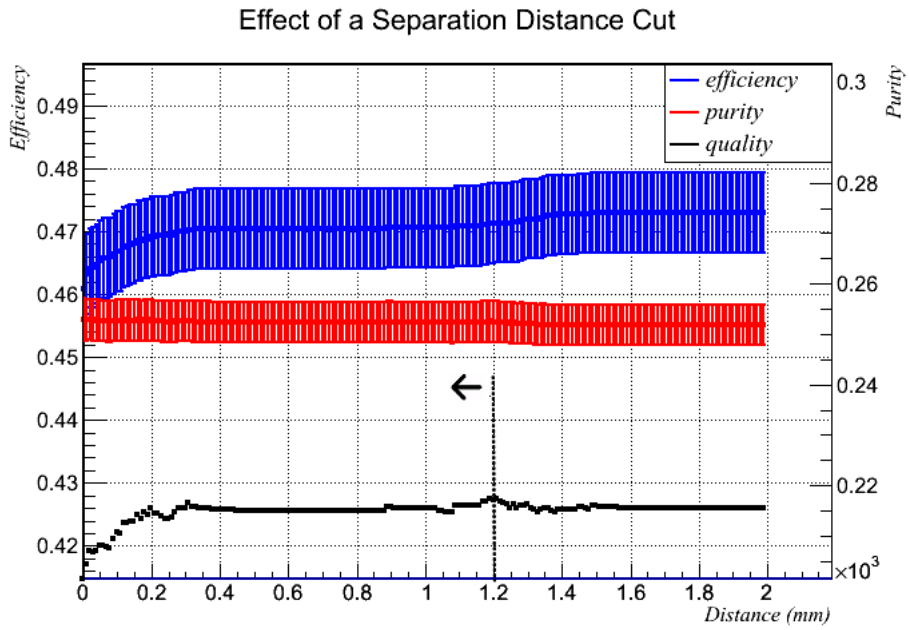

Figure 5.12: Effect of a minimum separation distance cut on pair reconstruction.

electron/positron pairs originating from a single photon. Placing a cut value of 1200mm eliminates these events and raises the quality.

5.2 Isolated Lepton Reconstruction

For isolated leptons (i.e. cases where only one lepton from the pair is detectable), there is a less information available than for the e^+/e^- pair scenario. Particularly an invariant mass selection cut cannot be performed.

Instead, a process of exclusion is applied. For a given charged track, an attempt is made to pair it with every other charged track in the tracker region of the detector. The pairing algorithm used is identical to the pair reconstruction technique, but with different cut parameters to maximise quality. Any leftover, unpaired tracks are considered isolated charged tracks. Tracks are not exclusively considered as part of either a pair or an isolated lepton, but can under certain circumstances be considered candidates for both scenarios. This is ex-

panded upon in the following subsection.

5.2.1 Overlapping candidates

Because the selection cuts for isolated lepton reconstruction and pair reconstruction are different, it is possible for a track to be simultaneously considered a pair candidate and an isolated lepton candidate. Whilst not intuitive, there is no reason not to simultaneously entertain both possibilities: to do otherwise would cause valuable information about the event to be lost. Consider a borderline case where a solidly-reconstructed electron is loosely paired to a low-momentum positron (for example, the invariant mass of the pair could be questionably high). At this stage of the analysis it is not known if either the pair or isolated-lepton scenario can be used to reconstruct a π^0 vertex, so to abandon one scenario in favour of another is premature.

This philosophy is extended to candidate π^0 vertices. It is entirely possible for a single charged track to be included as a candidate for all six topologies. For any given bunch, only one reconstructed π^0 vertex can be returned, however that is the only stage where any discrimination occurs between candidate types.

5.2.2 Electron Pull

Figure 5.13 plots the electron pull of isolated electron candidates when considering a minimum electron pull cut. Figure 5.14 plots the electron pull of isolated electron candidates when considering a maximum electron pull cut. These histograms are different to each other because, like all the plots in this chapter, they are N-1 histograms. N-1 histograms show the distribution of signal and background for a given cut under investigation, after all other plots have been

applied. In the case of figure 5.14, when considering the maximum electron pull the *minimum* electron pull cut has already applied leading to a sharp cutoff at -2, the chosen value for the minimum cut.

In figure 5.13 there is signficantly more background than signal at large electron pull values indicating an electron pull cut can signficantly improve the purity of the analysis with little effects on the efficiency. Figure 5.15 shows the effect of minimum and maximum electron pull cuts on isolated electron candidates. A minimum electron pull of -2 improves the quality of the analysis the most and this was the settled upon cut value. For isolated positrons, the quality also peaked at a pull value of -2 as the distribution of signal and background was highly similar to that of the isolated electrons.

In figure 5.14 the signal and background distributions are very similar and there is much less value in a maximum electron pull cut than the minimum electron pull cut for isolated electron candidates. The effect of a maximum electron pull cut for isolated electron candidates is shown in figure 5.16, where the quality plateaus instead of peaking. An eventual cut value of +2 was settled upon. Whilst the quality differences between having a cut of +2 and having no cut at all are very slim, the purity of the analysis is noticeably greater. Improving the purity of the analysis is of critical importance when considering large beam and cross-section systematic errors.

For isolated positrons, there was a much clearer gain in a maximum electron pull cut. Figure 5.17 shows the electron pull for isolated positron candidates when considering a maximum electron pull cut. For high electron pull values there is significantly more background and signal. This is reflected in figure 5.17, which shows the effect of a maximum electron pull cut for isolated positrons. The quality peaks at a maximum pull value of +1.

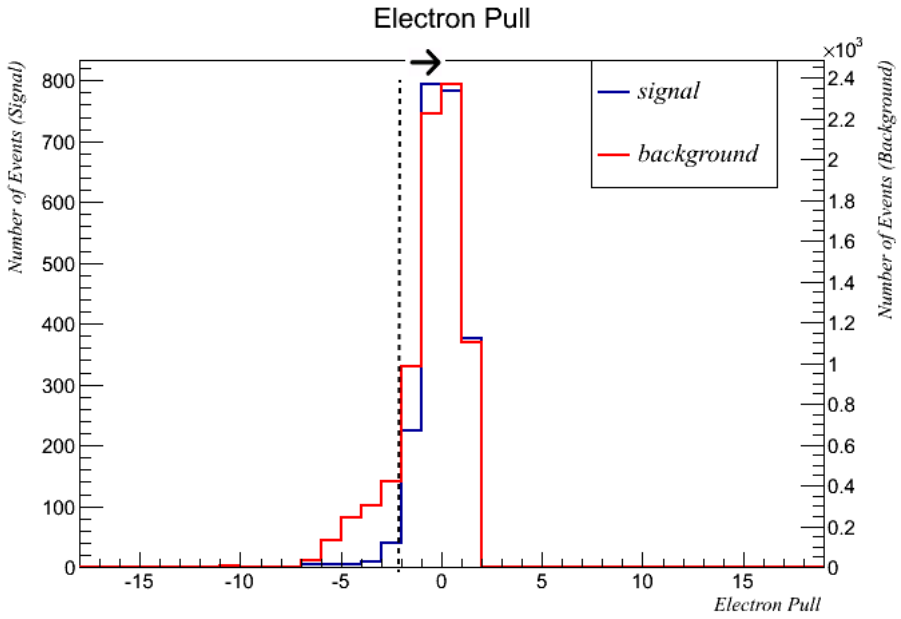

Figure 5.13: Electron pull of isolated electron candidates when considering a minimum electron pull cut.

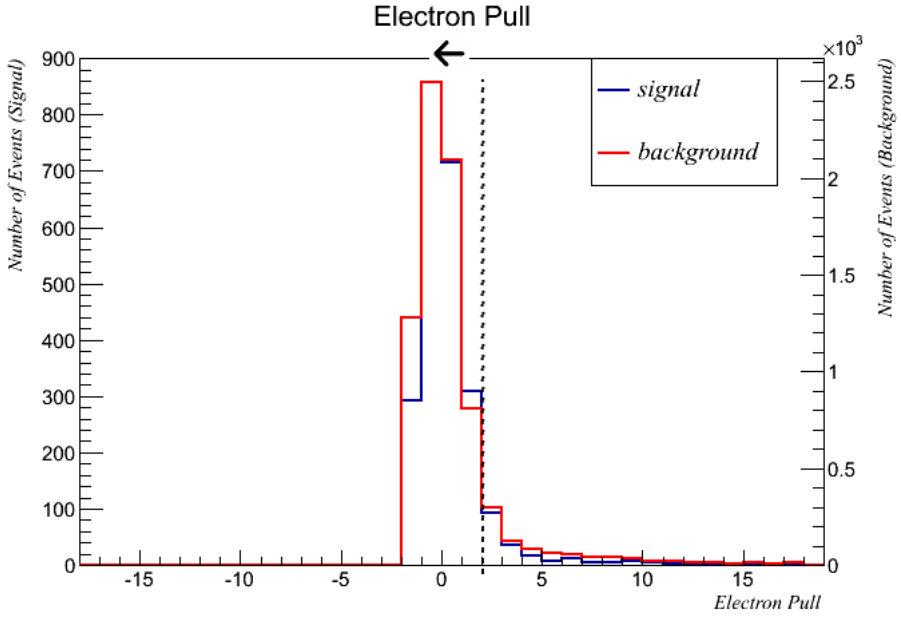

Figure 5.14: Electron pull of isolated electron candidates when considering a maximum electron pull cut.

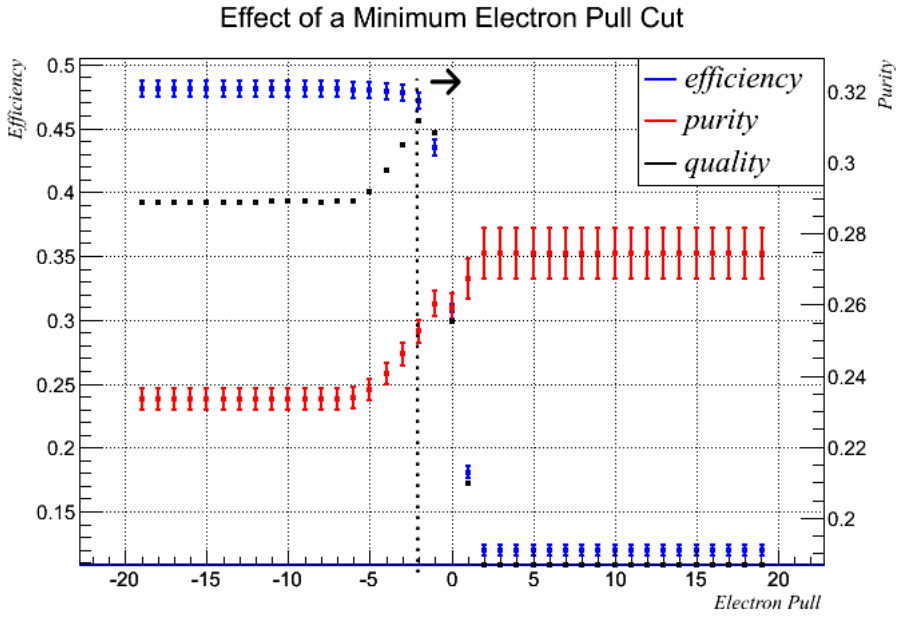

Figure 5.15: Effect of a minimum electron pull on isolated electron candidates.

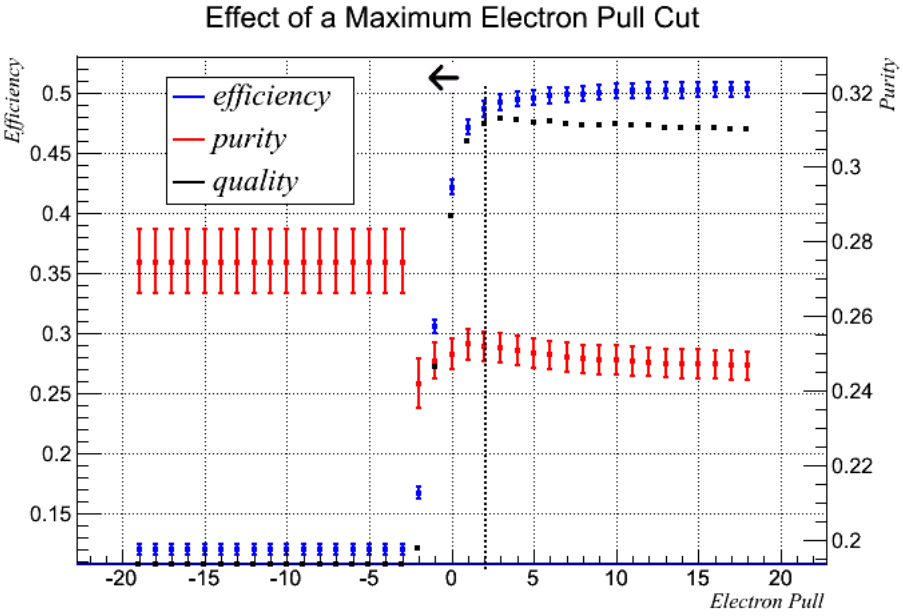

Figure 5.16: Effect of a maximum electron pull on isolated electron candidates.

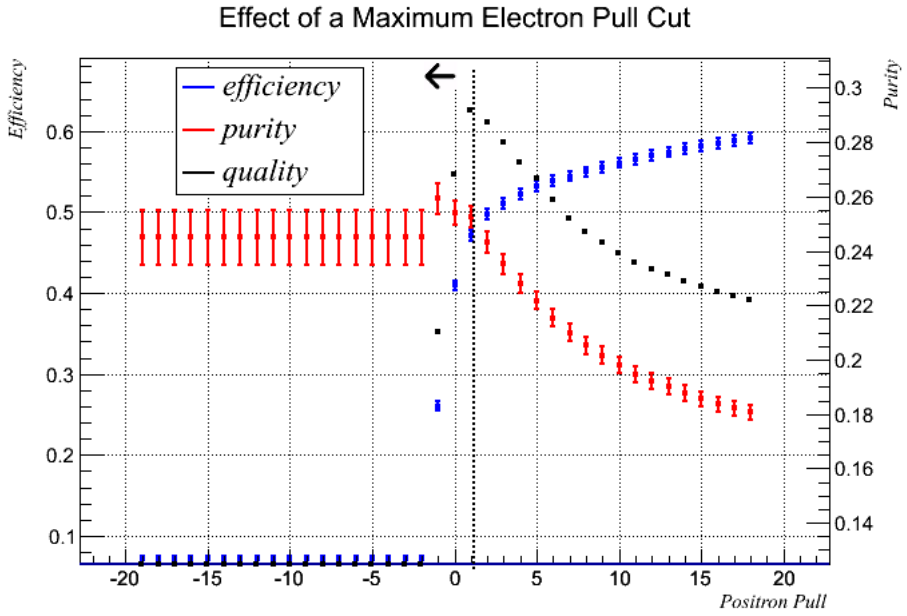

Figure 5.17: Effect of a maximum electron pull on isolated positron candidates.

5.2.3 Momentum Cut

The momentum of isolated positron candidates, when considering a minimum momentum cut is shown in figure 5.18. There is significantly greater amounts of signal when compared to background after 50 MeV/c. This difference between signal and background is not present when looking at the momentum distribution *before* any cuts are applied. This is shown in figure 5.18. This excess of signal is likely due to another selection cut removing background from this momentum region whilst leaving the signal unaffected.

The effect of a minimum momentum cut for isolated positron candidates is shown in figure 5.19. There is no value in applying a minimum momentum cut at any level, however a minimum cut of 50 MeV/c is applied for both isolated electron and isolated positron candidates due to the added systematic errors using very low-momentum tracks would introduce.

The momentum of isolated electron candidates when considering a maximum momentum cut is shown in figure 5.20. The effect of a maximum momentum

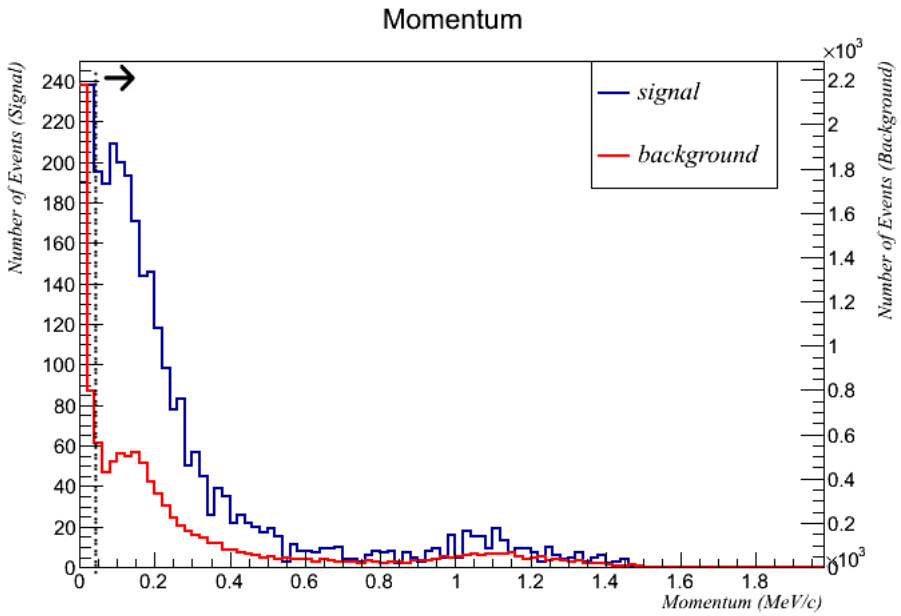

Figure 5.18: Momentum of isolated positron candidates when considering a minimum momentum cut.

cut for isolated electrons is shown in figure 5.21. Similarly to the pair reconstruction, there is value in applying a maximum momentum cut as it removes high-energy outlier events. The quality peaks at 700 MeV/c, which is the cut value used. The isolated positron quality distribution also peaks at 700 MeV/c.

5.2.4 Pair Exclusion Cuts

One of the factors involved in classifying isolated leptons is to ensure that they are, in fact, isolated. To this effect, exclusion selection cuts were investigated. For a given isolated lepton candidate, every other track in the event was considered a potential pair partner. The properties of potential pair partners were taken.

- The distance between origin points of the isolated lepton candidate and its pair partners
- The extrapolated distance between the isolated lepton candidate and its

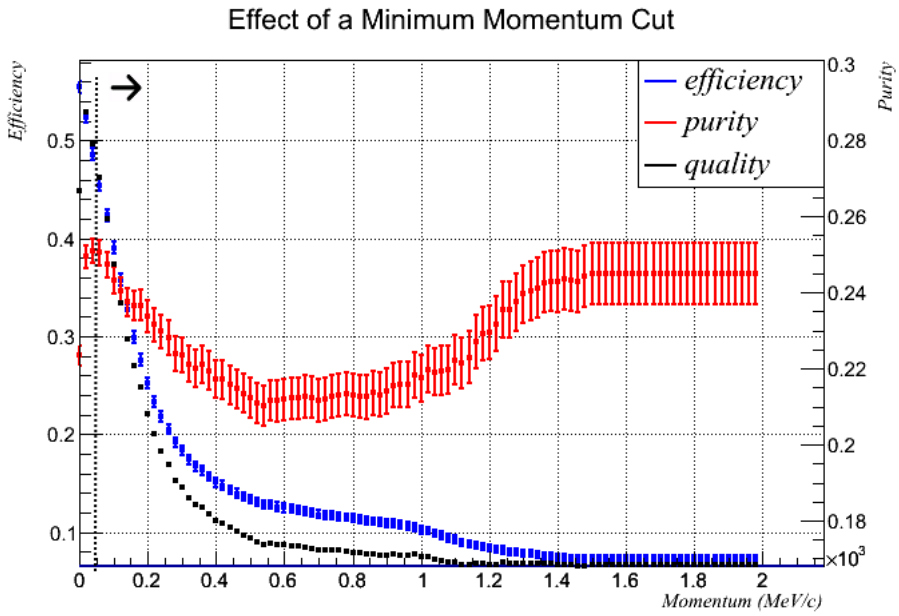

Figure 5.19: Effect of a minimum momentum cut on isolated positron candidates.

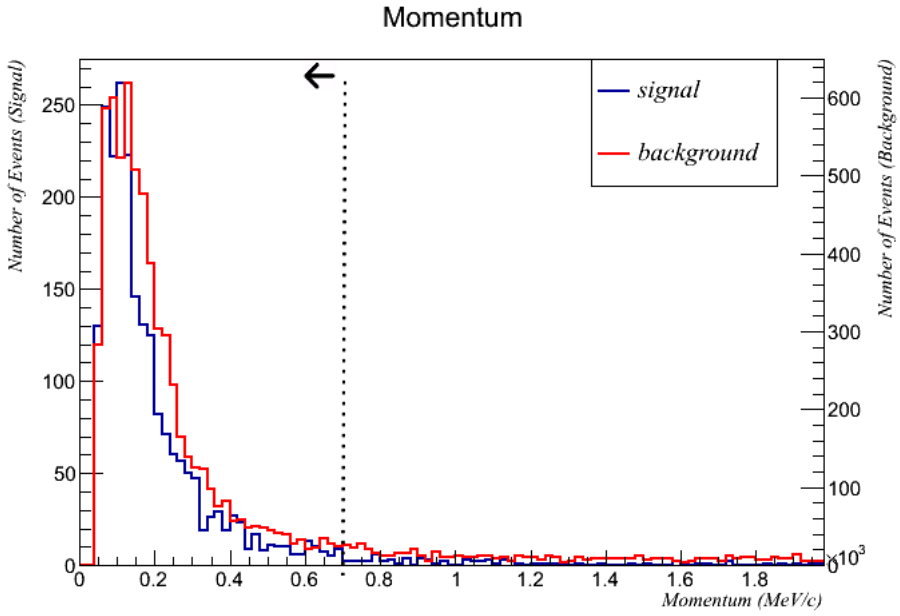

Figure 5.20: Momentum of isolated electron candidates when considering a maximum momentum cut.

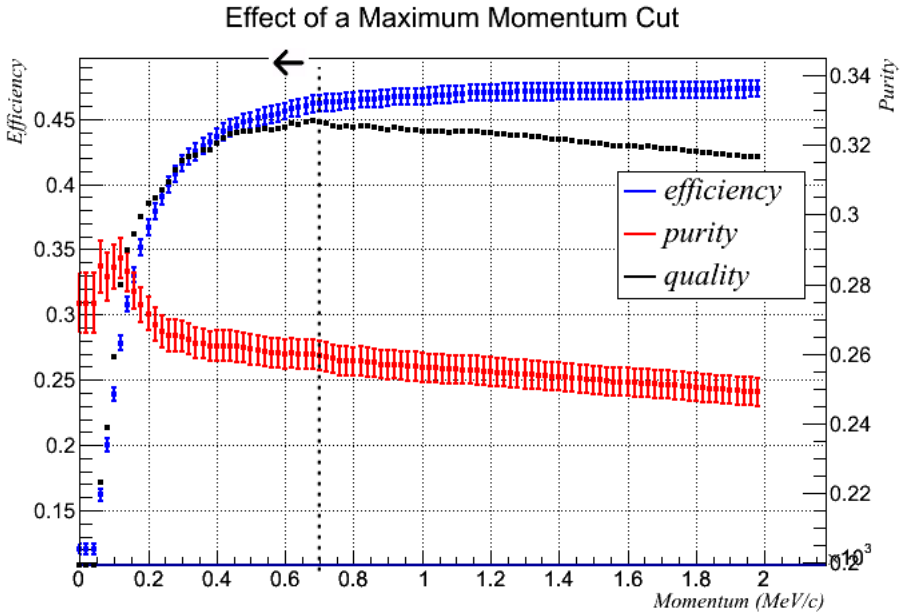

Figure 5.21: Effect of a maximum momentum cut on isolated electron candidates.

pair partners (extrapolated distance is defined in section 5.1.4).

- The invariant mass of the isolated lepton candidate and its pair partners.
- The electron pulls of pair partners.
- The momentum of pair partners.

For each property in the above list, a threshold cut value was selected. If an isolated lepton candidate has a pair partner which passes all these cuts, referred to as exclusion cuts, then the isolated lepton candidate is rejected as it is considered to be part of an electron/positron pair.

However, when investigating these cuts it became clear that rejecting any isolated lepton candidates because they had suitable pair partners drastically lowered the overall efficiency of the analysis as well as the quality. Figure 5.22 shows the minimum invariant mass of isolated electron candidates and their potential pair partners. An exclusion invariant mass cut would reject isolated

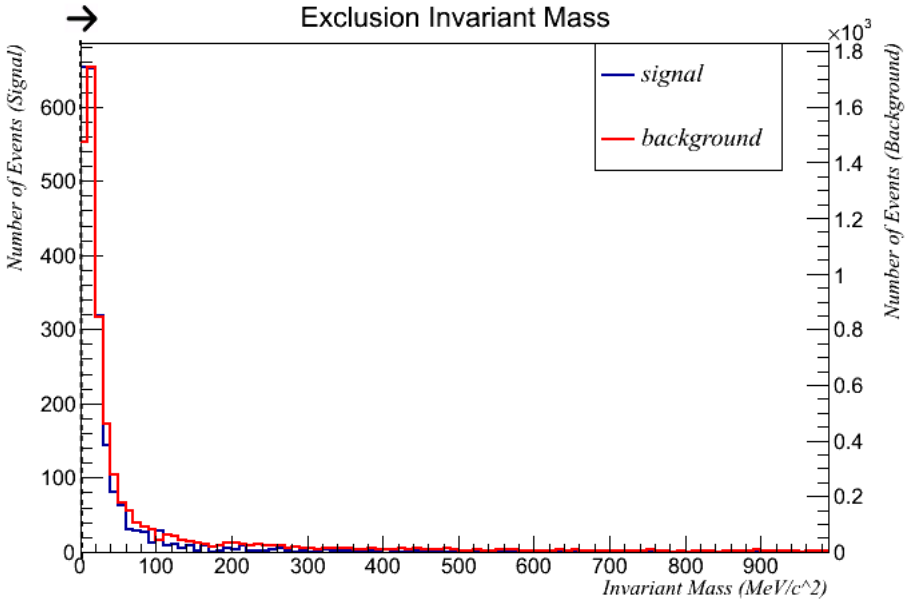

Figure 5.22: Exclusion invariant mass histogram for isolated electron candidates. For each isolated electron candidate, the lowest invariant mass of its pair partners is calculated and plotted.

electron candidates if they had any pair partners where the invariant mass of the candidate and its pair partner was lower than the cut value (the reasoning being that this candidate is actually part of a pair).

In figure 5.22 the signal and background distributions are almost identical and no value was gained from an exclusion invariant mass cut, shown in figure 5.23. This is a trend continued in all of the potential exclusion cuts for both isolated electrons and isolated positrons. The minimum separation distance between isolated electron candidates and potential pair partners is shown in figure 5.24, with the effect of an exclusion separation distance cut shown in figure 5.25. Again, the signal and background distributions are highly similar and no value is gained by adding an exclusion cut.

For this analysis, no exclusion cuts are applied for isolated electron candidates or isolated positron candidates. The effect of this is that the tracks which make up electron/positron pair candidates are also considered isolated lepton can-

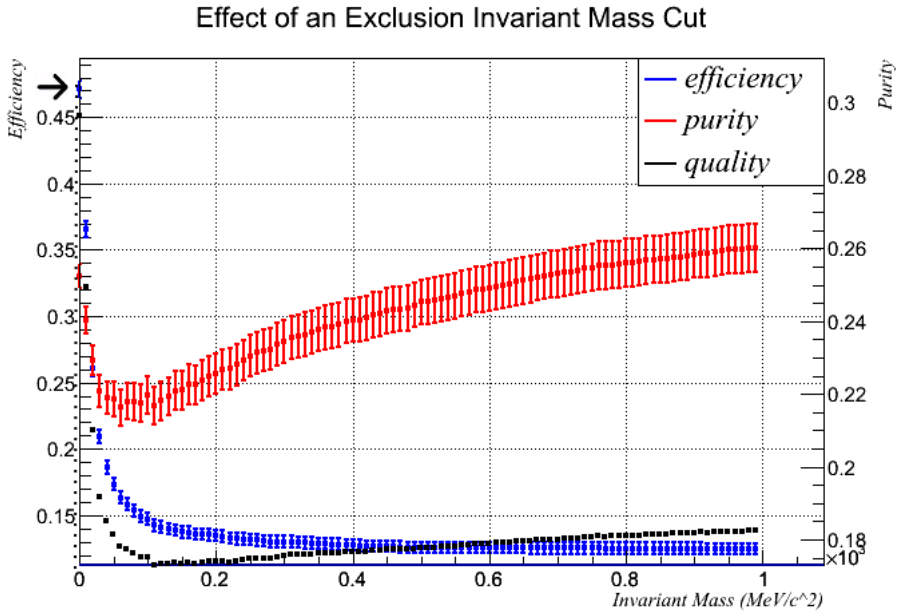

Figure 5.23: Effect of an exclusion invariant mass cut for isolated electron candidates. Rejecting an isolated electron candidate based on the invariant mass it makes with potential pair partners does not improve the quality of the analysis and so no cut was applied.

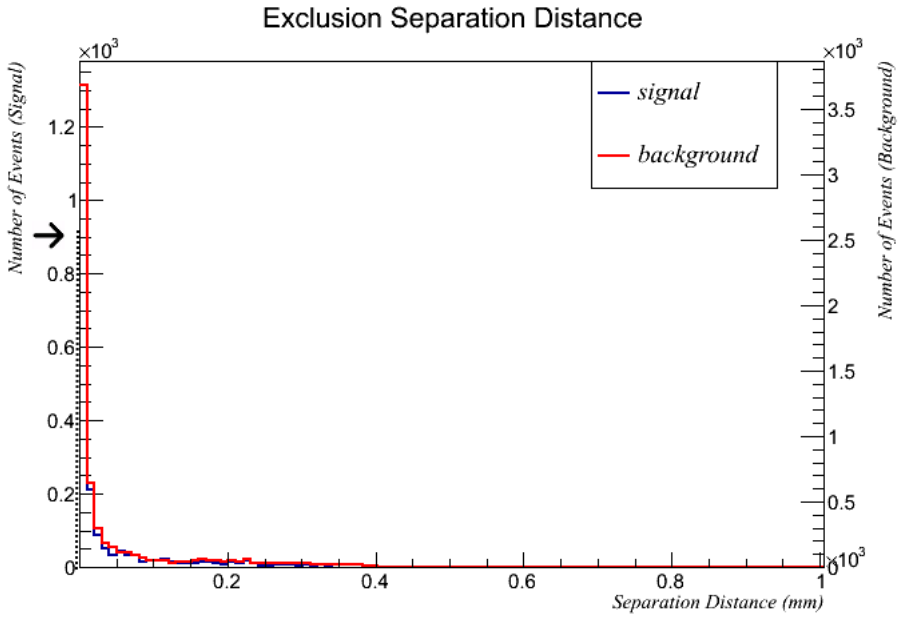

Figure 5.24: Exclusion separation distance histogram for isolated electron candidates. For each isolated electron candidate, the shortest distance between its starting point and that of its pair partners is calculated and plotted.

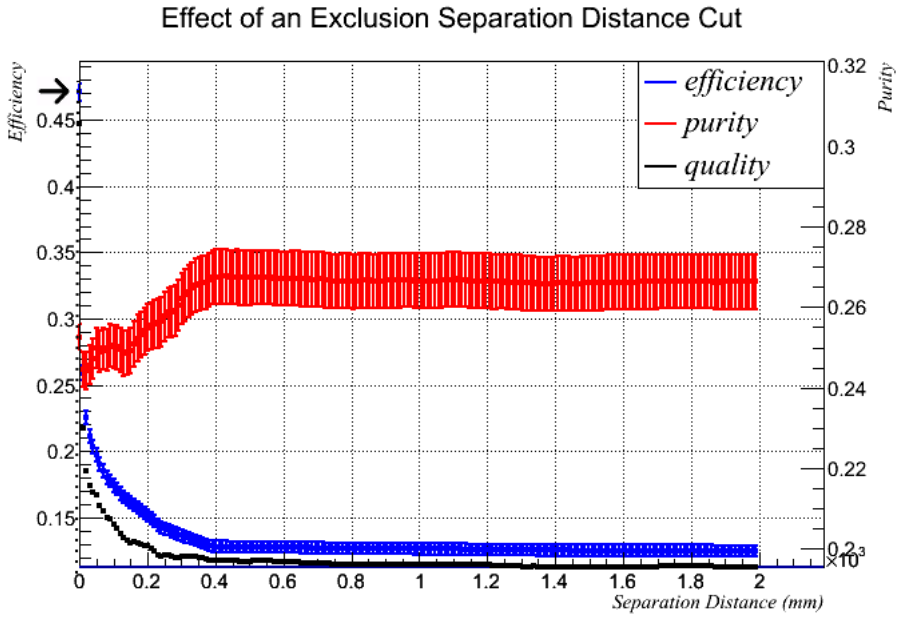

Figure 5.25: Effect of an exclusion separation distance cut for isolated electron candidates.

didates, subject to the tracks also passing the remainder of the isolated lepton selection cuts (momentum and electron pull). This conclusion was initially counter-intuitive, however by considering these tracks as isolated leptons their chances of being eventually reconstructed as π^0 candidates increases. This increase in the number of π^0 candidates adds more signal and relatively small amounts of background, increasing the overall quality of the analysis.

5.2.5 Summary of first Generation Cuts

This section has documented a comprehensive investigation into all possible methods of increasing statistics and reducing background when attempting to reconstruct photons out of electron/positron pairs. A summary of the final cut values used for pairs, isolated electrons and isolated positrons are given in tables 5.1 and 5.2.

Table 5.1: Summary of first generation selection cuts applied to pair candidates.

Cut Type	Pair
Invariant Mass (MeV/c^2)	110
Min Momentum (MeV/c)	50
Max Momentum (MeV/c)	1000
Min Electron Pull (MeV/c)	-4
Max Electron Pull (MeV/c)	+5
Extrapolated Distance (mm)	none
Separation Distance (mm)	1200

Table 5.2: Summary of first generation selection cuts applied to isolated electron candidates.

Cut Type	Isolated Leptons
Min Electron Pull	-2
Max Electron Pull	+2
Min Momentum (MeV/c)	50
Max Momentum (MeV/c)	700
Invariant Mass Exclusion (MeV/c^2	none
Separation Distance Exclusion (mm)	none
Extrapolated Distance (mm)	none
Min Auxiliary Momentum (MeV/c)	none
Max Auxiliary Momentum (MeV/c)	none
Min Auxiliary Electron Pull	none
Max Auxiliary Electron Pull	none

Table 5.3: Summary of first generation selection cuts applied to isolated positron candidates.

Cut Type	Isolated Leptons
Min Electron Pull	-2
Max Electron Pull	+1
Min Momentum (MeV/c)	50
Max Momentum (MeV/c)	700
Invariant Mass Exclusion (MeV/c^2	120
Separation Distance Exclusion (mm)	none
Extrapolated Distance (mm)	none
Min Auxiliary Momentum (MeV/c)	none
Max Auxiliary Momentum (MeV/c)	none
Min Auxiliary Electron Pull	none
Max Auxiliary Electron Pull	none

Chapter 6

Evaluating Second Generation Selection Cuts: π^0 Vertex Reconstruction

To reconstruct π^0 vertices for a given event, the 4-momenta of the progenitor photons of all pair and isolated lepton candidates are reconstructed by adding the 4-momenta of the leptons used to reconstruct the photon candidate. These photon candidates are then projected upstream through the detector in straight lines. The position where the separation distance between the two photon candidates is the smallest is then calculated. If this position is not in either FGD1 or FGD2, the π^0 candidate is rejected.

Selection cuts are then applied for each of the six topologies described in the list in section 4.1.3. The selection cuts applied to each topology are as follows:

1. The tracker multiplicity of the bunch.
2. The ECal multiplicity of the bunch.
3. The isolated ECal multiplicity of the bunch.

4. The P0D multiplicity of the bunch.
5. The TPC1 multiplicity of the bunch.
6. The overall multiplicity of the bunch.
7. The opening angle of the photon candidates.
8. The distance between the points of closest approach of the photon candidates.

The π^0 invariant mass, whilst the ideal discriminator, was not used because it is a key kinematic quantity being measured by this analysis, and using it would introduce a bias. A veto on situations where one photon conversion happens in FGD1 and the second in FGD2 was considered, however after investigation it was found that adding this veto reduces the overall quality of the analysis.
The following subsections analyse the value of these selection cuts for all six topologies. The selection cuts settled upon are documented in table 6.1. The pair and isolated electron selection cuts are evaluated using the base selection due to the very limited amount of statistics available for that topology (the final selection reconstructs sixty pair and isolated electron signal candidates however out of these only seven are not affected by event migration. i.e. for the other fifty three candidates, cuts which would remove them from the pair and isolated electron selection topology would not remove them from the overall analysis because they would be reconstructed as other topology types). All other five topologies use the final selection.

6.1 Event Vetoes: TPC1 Multiplicity and P0D Multiplicity

In this analysis, any event with tracks present in either the P0D subdetector or the TPC1 subdetector are vetoed outright regardless of the topology being investigated. It became apparent early into the analysis that accepting events containing tracks within the TPC1 subdetector or P0D subdetector significantly lowered the overall purity of the analysis. The TPC1 and P0D subdetectors are upstream of both FGD1 and FGD2 (i.e. the neutrino beam enters these subdetectors before entering the FGDs). Therefore it is highly unlikely that any neutral-current π^0-producing neutrino interaction occuring in one of the FGDs will have tracks present in either the P0D or TPC1. Conversely, other types of neutrino interactions originating in the P0D or TPC1 present a source of background that could be easily removed by vetoing any events with tracks present in either of these subdetectors.

Figure 6.3 shows the TPC1 track multiplicities for candidate π^0 events using the isolated electron and isolated positron topology, with figure 6.4 showing the effects of a TPC1 multiplicity cut (the isolated electron and isolated positron topology histograms are used in this section because of the high statistics in this topology, however the conclusions are identical for all six topologies in this case). The quality of the analysis is significantly greater for a cut of zero than any other cut values (i.e. only events with a TPC1 multiplicity of zero are accepted by the analysis). Figure 6.1 shows the P0D track multiplicities for candidate π^0 events using the isolated electron and isolated positron topology, with figure 6.2 showing the effect of a P0D multiplicity cut. Again, quality is greatest at for a cut of zero.

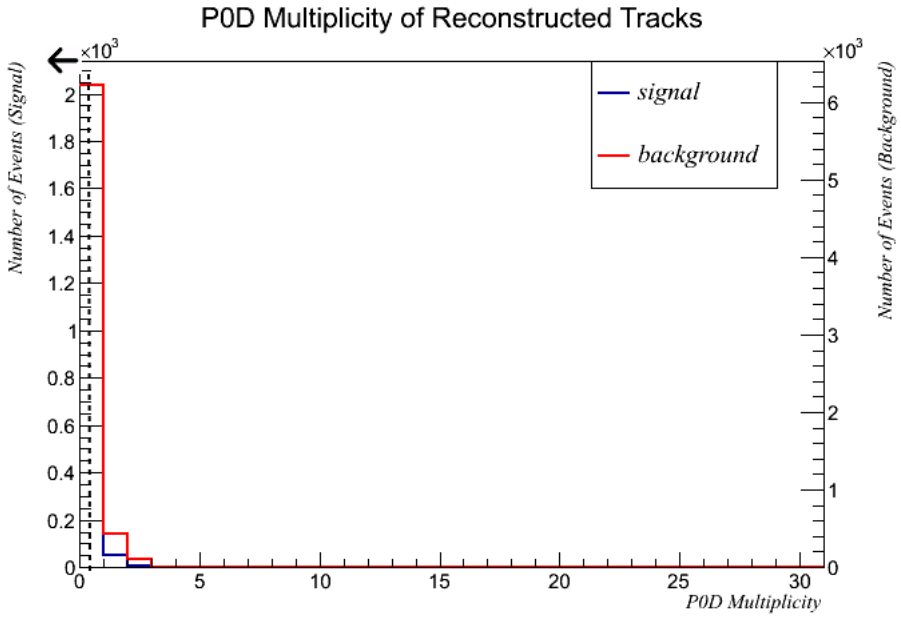

Figure 6.1: P0D multiplicity for candidate π^0 events, isolated electron and isolated positron topology.

For all six topologies, having a TPC1 multiplicity cut of zero and a P0D multiplicity cut of zero maximised quality.

6.2 Twin Pair Reconstruction

A list of all the pair candidates are constructed. The 4-momentum from each pair are added together to reconstruct a candidate photon that produced the pair.

Each of these candidate photons is then combined with the others. For each photon combination, the photons are projected upstream through the detector, where their distance of closest approach is calculated.

The position of the distance of closest approach is then taken as the potential π^0 vertex. If this vertex is within either FGD, the two twin pairs are considered

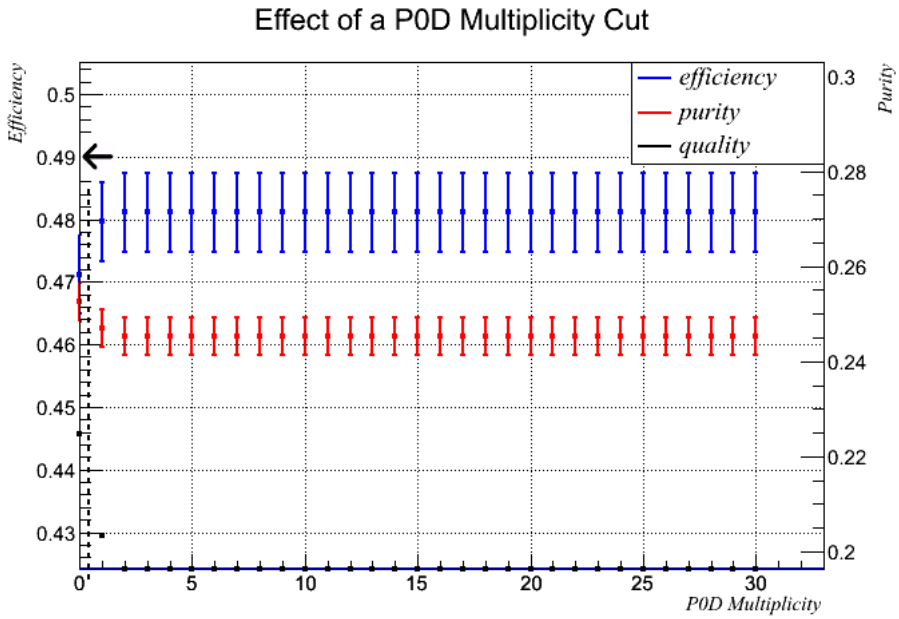

Figure 6.2: Effect of a P0D multiplicity cut, isolated electron and isolated positron topology.

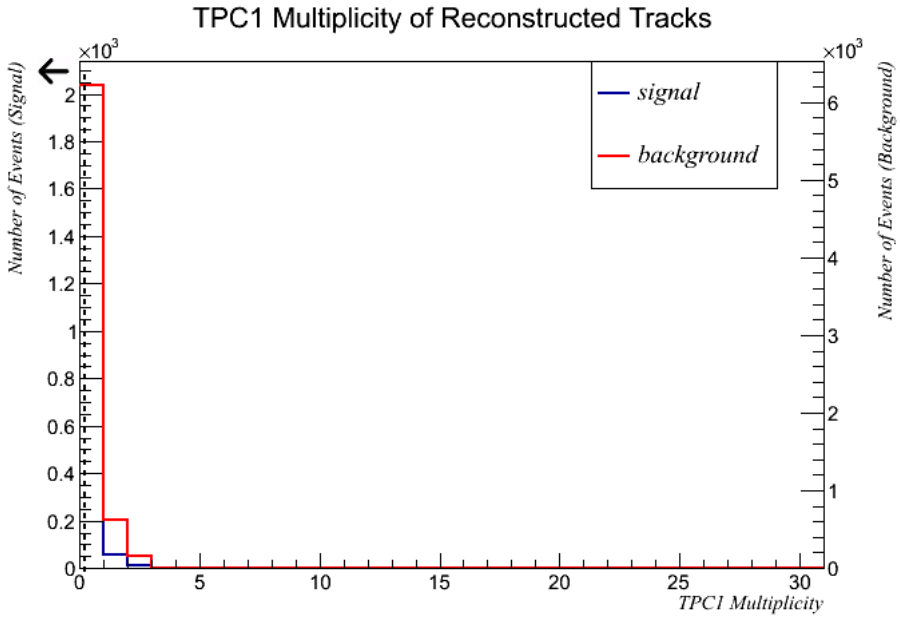

Figure 6.3: TPC1 multiplicity for candidate π^0 events, isolated electron and isolated positron topology.

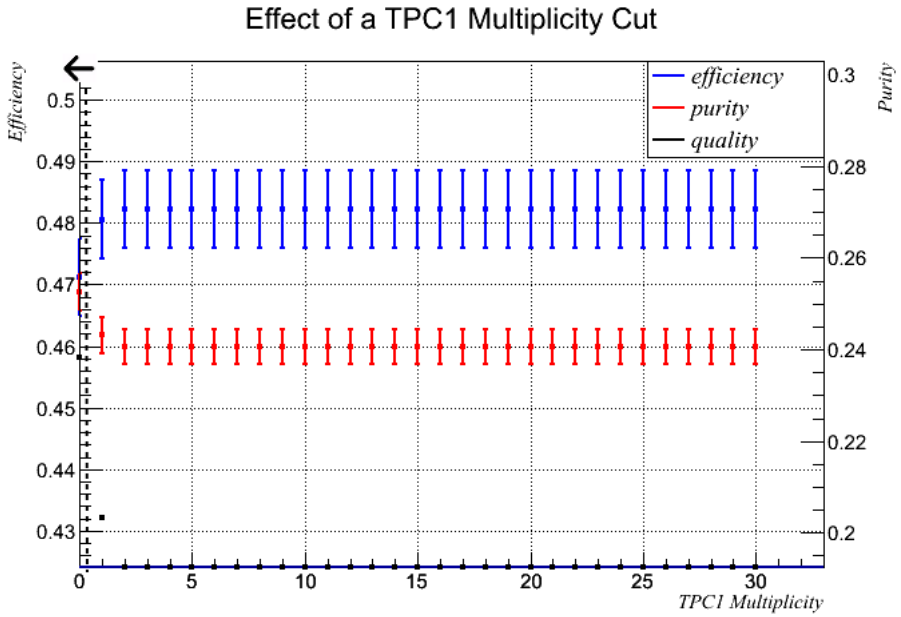

Figure 6.4: Effect of a TPC1 multiplicity cut, isolated electron and isolated positron topology.

a candidate π^0 vertex, dependent on further selection cuts. The final selection was used when evaluating twin pair selection cuts.

6.2.1 Upstream Multiplicity

Upstream multiplicity represents the number of reconstructed tracks with a presence in subdetectors which are upstream of FGD1, the first subdetector within the topology region. Events with reconstructed tracks in either the P0D or TPC1 (two subdetectors upstream of FGD1) are vetoed, therefore any tracks upstream of FGD1 must originate in one of the barrel ECals (the only remaining subdetector upstream of FGD1).

Initially, the upstream multiplicity was considered an event veto: events with any upstream multiplicity would be removed from the selection. However, converting the veto into a selection cut was decided upon as the more rigorous

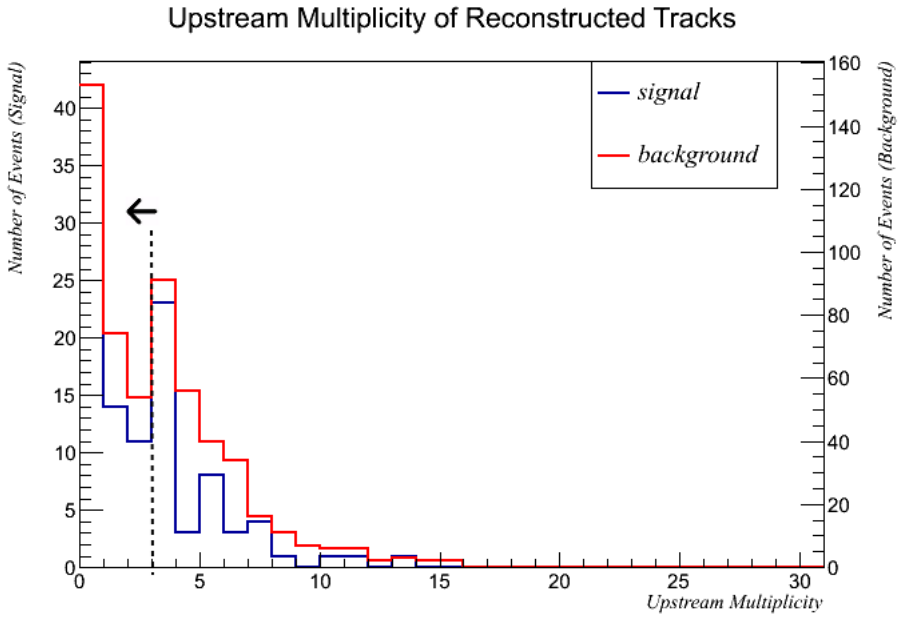

Figure 6.5: Upstream multiplicity of signal and background for twin pair topology.

option; if a veto is required then the quality will be at a maximum with a multiplicity cut value of 0. Figure 6.5 shows the signal and background distributions of the upstream multiplicity on the base selection, and Figure 6.6 shows the effect of an upstream multiplicity cut on the final selection.

Counter-intuitively, having a cut of zero on upstream multiplicities degrades quality and efficiency. A large amount of signal has significant upstream multiplicities. This could be because of backwards scattering tracks into the barrel ECal. A final cut value of three was chosen, as this is the point where the quality distribution plateaus. Having a cut of three instead of no cut slightly improves purity, and will remove outliers that may be present in the data.

6.2.2 Tracker Multiplicity

Tracker multiplicity represents the number of reconstructed tracks with a presence in subdetectors encompassed within the tracker region: FGD1, FGD2,

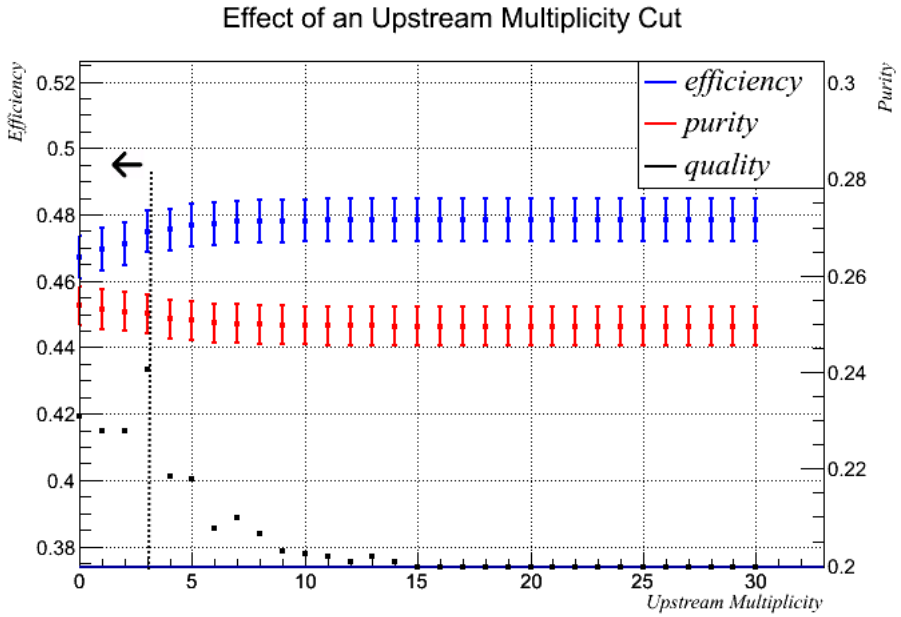

Figure 6.6: Effect of an upstream multiplicity cut, twin pair topology.

TPC1, TPC2, TPC3. This region, with the exception of TPC1, is the area of interest for this analysis. A multiplicity cut could be useful because, when consulting event displays, it appeared that π^0 interactions are quite clean, with relatively low multiplicities. A multiplicity cut could remove messy background events which could potentially pass the reconstruction algorithm, whilst leaving signal relatively intact.

Figure 6.7 shows the tracker multiplicity of the final selection for signal and background, and figure 6.8 shows the efficiency, purity and quality.
For both the signal and background distributions, the tracker multiplicity is spread broadly over a large range of multiplicities (from four to twenty-eight). Given the small maxima at a multiplicity cut of fourteen, the twin pair tracker multiplicity cut is set at this value.

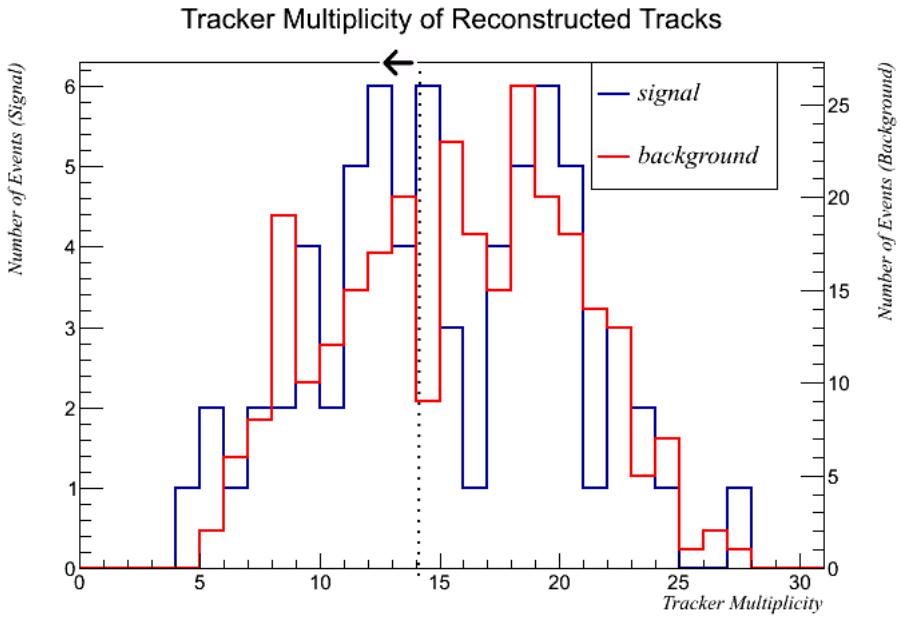

Figure 6.7: Tracker multiplicity of signal and background, twin pair topology.

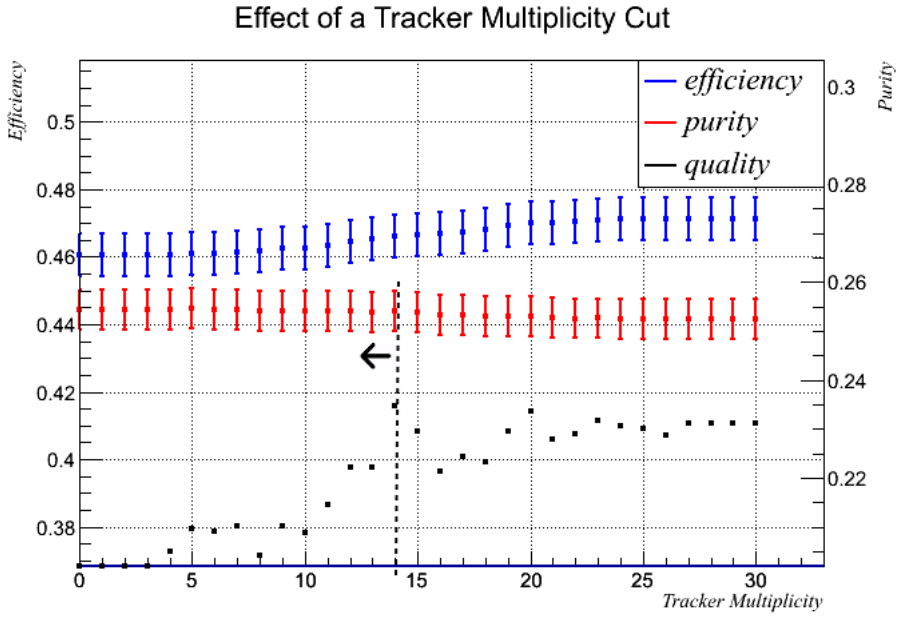

Figure 6.8: Effect of a tracker multiplicity cut. Twin pair topology.

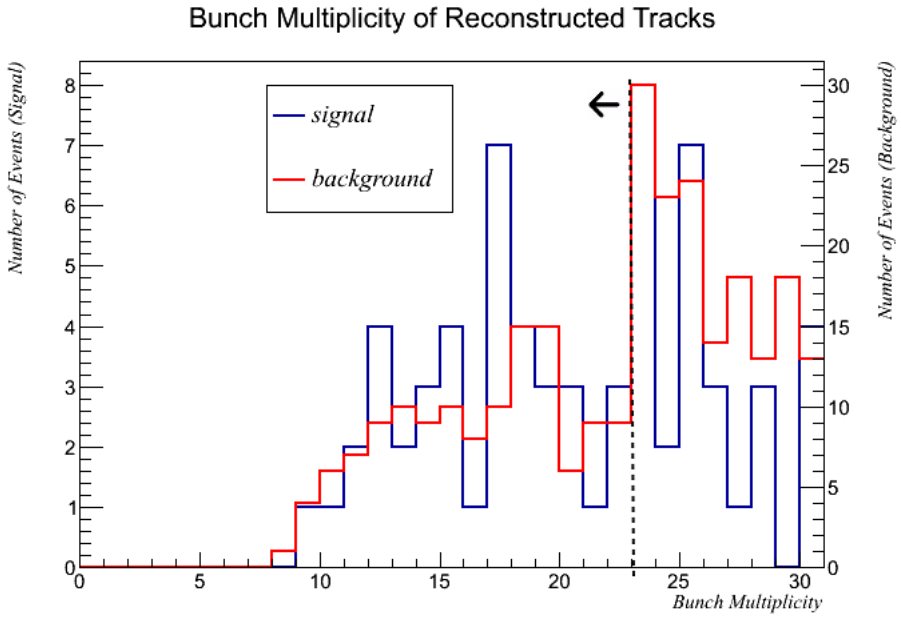

Figure 6.9: Bunch multiplicity of signal and background selections for twin pair topology.

6.2.3 Bunch Multiplicity

Bunch multiplicity represents the number of reconstructed tracks in the entire bunch, across all subdetectors. The reasoning behind such a cut is similar to tracker multiplicity, and covers tracks which escape the tracker region.

Figure 6.9 shows the bunch multiplicity of the base selection for signal and background, and figure 6.10 shows the efficiency, purity and quality.
The quality tails off at high multiplicities and there is a maxima at a multiplicity cut of twenty-three. Therefore the twin pair bunch multiplicity cut was set at twenty-three.

6.2.4 Isolated ECal Multiplicity

The isolated ECal multiplicity refers to the number of tracks which are entirely contained within a barrel or a downstream ECal. The isolated ECal multiplicity

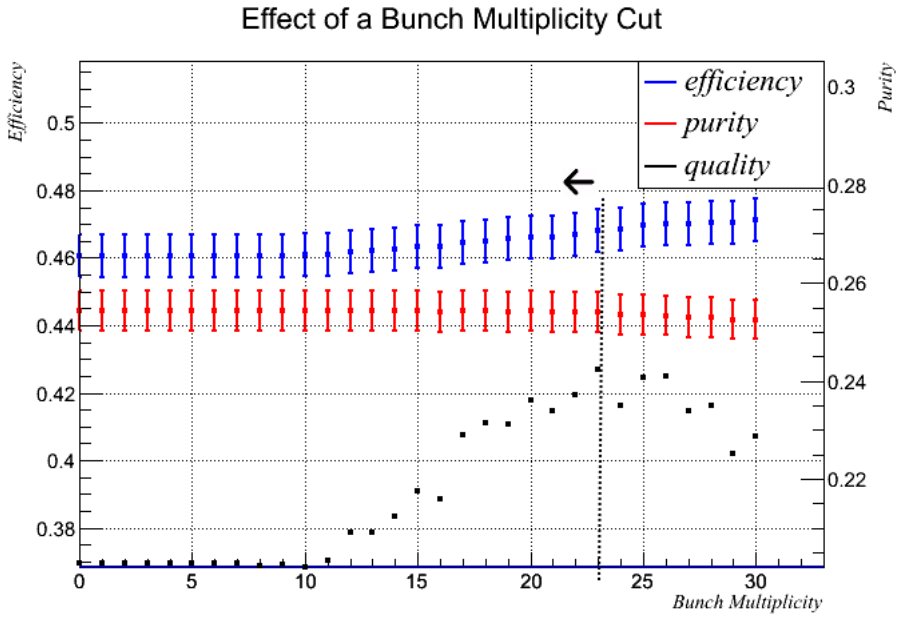

Figure 6.10: Effect of a bunch multiplicity cut. Twin pair topology.

has the potential to be a valuable variable to cut on. Ideally, any NCπ^0 signal interaction should either be contained in the tracker region, or spill into the ECal region; a scenario where a track exists solely in the ECal should be rare, whilst this would not be the case for background. The number of tracks which are isolated in the ECal for a given twin pair π^0 candidate are shown in figure 6.11, with the effect of a cut shown in figure 6.12.

Contrary to what was anticipated, the signal and background distributions are very similar and there is no benefit to applying an isolated ECal multiplicity cut at any level.

6.2.5 ECal Multiplicity

Alongside an isolated ECal multiplicity cut, an ECal multiplicity cut was also investigated. This cuts on the number of tracks present in the ECal for any candidate twin-pair π^0. Ideally, signal events should be mostly constrained to

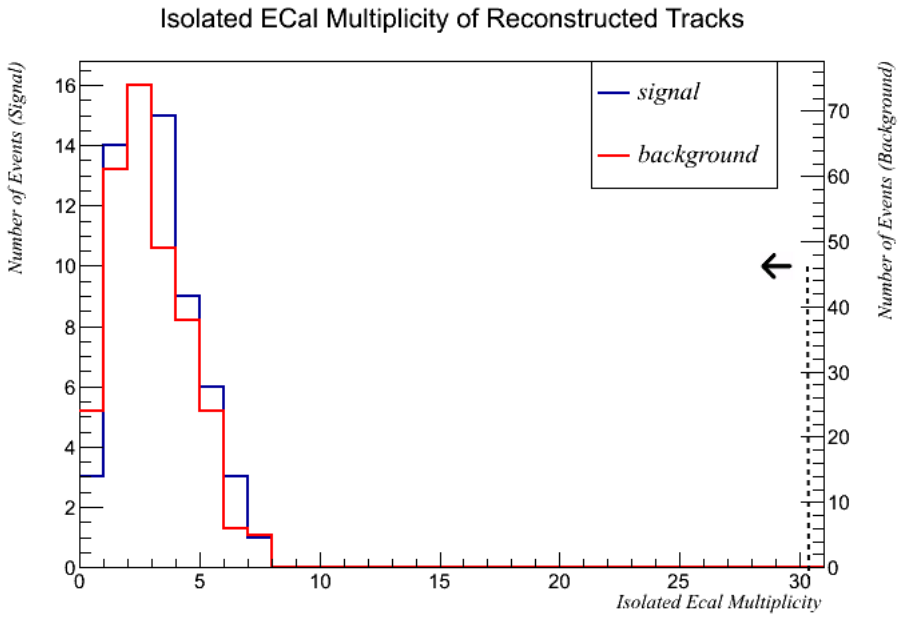

Figure 6.11: Isolated ECal multiplicity of signal and background selections for Twin Pair topology.

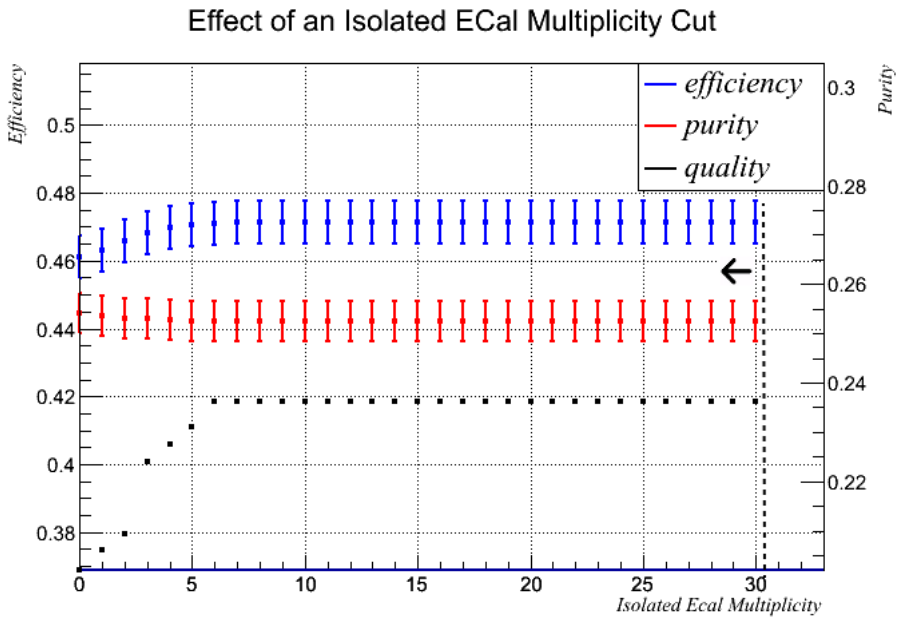

Figure 6.12: Effect of an isolated ECal multiplicity cut.

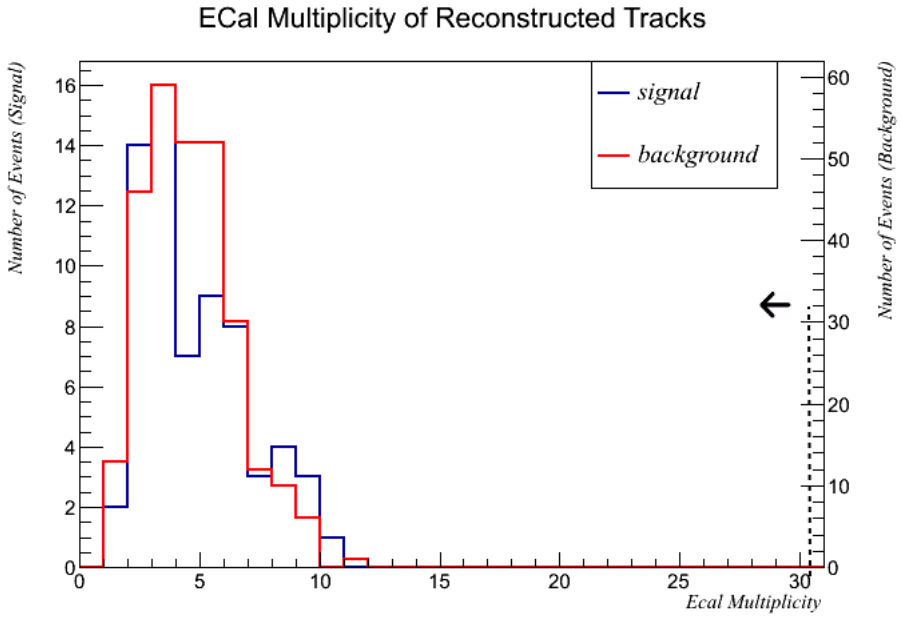

Figure 6.13: ECal multiplicity of signal and background selections for Twin Pair topology, base selection.

the tracker region and should have low ECal multiplicities. Figure 6.13 shows the ECal multiplicity for signal and background for the base selection. These distributions are almost identical, and no value is gained from a cut at any level. The signal does not have particularly low ECal multiplicities, perhaps because of tracks propagating through TPC3 and into the downstream ECal, combined with tracks entering the barrel ECals.

Minimum Opening Angle

Photons produced from a decaying π^0 should be very forward-going, as they can contain significant fractions of the interacting neutrino's kinetic energy, and therefore be highly boosted. A selection cut on the minimum permissible cosine of the opening angle, θ is investigated: a minimum cut means that events with a $\cos\theta$ below the minimum threshold are removed. The opening angle was determined by using the 4-momenta of the projected photons, which were constructed out of the reconstructed lepton pairs.

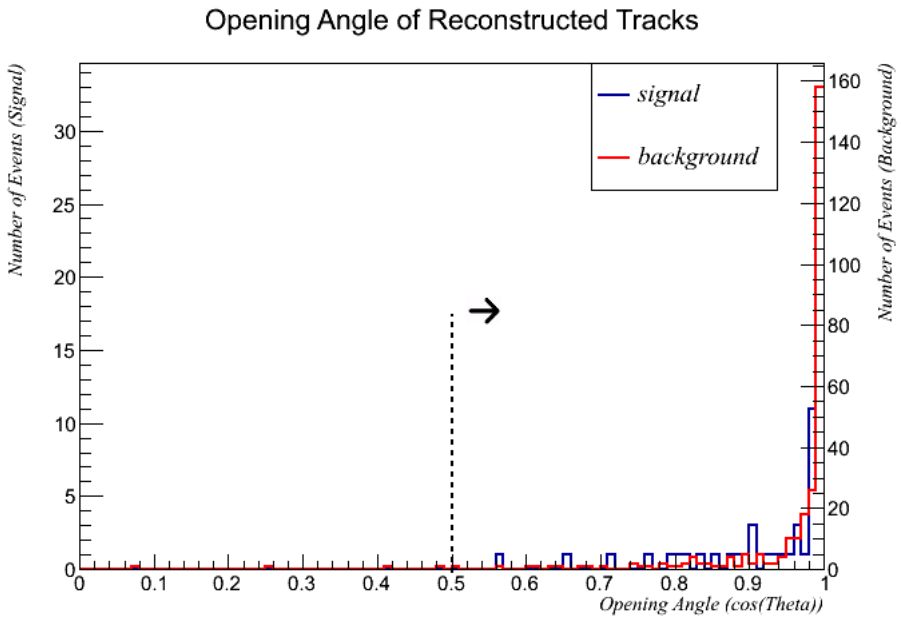

Figure 6.14: Opening angle of projected photons, twin pair topology.

When examining the opening angle for all topologies, the variable being investigated is the cosine of the opening angle.

Therefore, removing potential π^0 candidates with large opening angles could reduce backgrounds whilst preserving signal. Figure 6.14 shows the opening angle of signal and background for the base selection. Whilst the opening angles for both signal and background selections are broad, the gradient of the signal selection is greater than that of the background. Figure 6.15 shows the effect of a minimum opening angle cut on the base selection.

Although the improvements in quality were not as large as anticipated, there is a marginal gain in quality to be had by setting the minimum opening angle at $\cos\theta = 0.5$.

There was no advantage to applying a maximum opening angle cut. Due to the

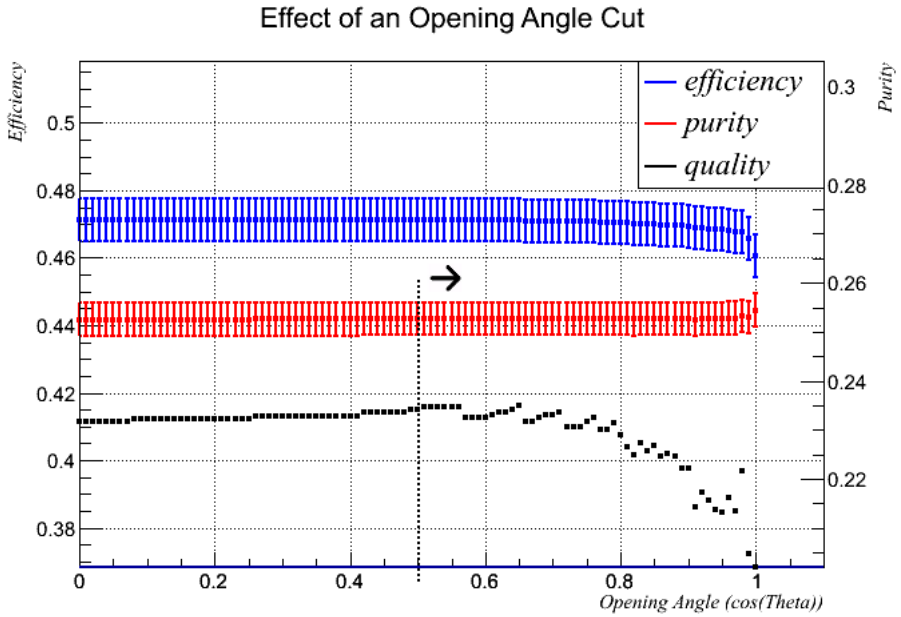

Figure 6.15: Effect of a minimum opening angle cut on the base selection.

forward-going nature of the signal this was not surprising.

6.2.6 Polar Angle

Similar to the opening angle, a polar angle cut was also considered. The 4-momenta of the two photon candidates are added to create the 4-momenta of the π^0 candidate. The angle this π^0 makes with the z-axis is the polar angle. Signal events should have a low polar angle as the π^0 should carry a lot of momentum from the initial neutrino, which has a polar angle of 0. The polar angle for signal and background is shown in figure 6.16, with the effect of a polar angle cut shown in figure 6.17. When examining the polar angle, the cosine of the polar angle is the variable being measured.

There is a clear benefit in a polar angle cut, with a minimum polar angle cut of $\cos\theta = 0.97$ providing greater quality than having no minimum cut.

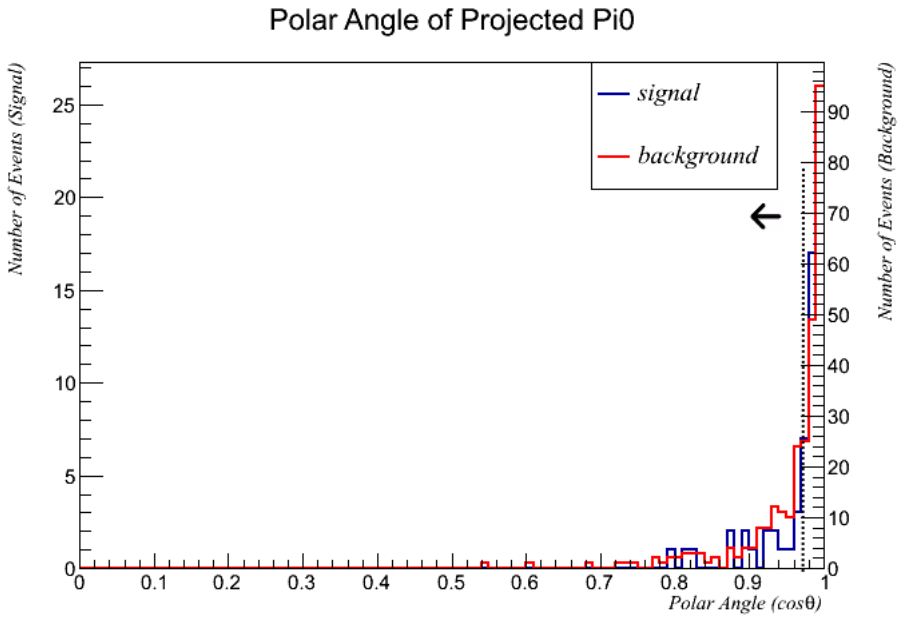

Figure 6.16: Cosine of the polar angle of projected photons, twin pair topology

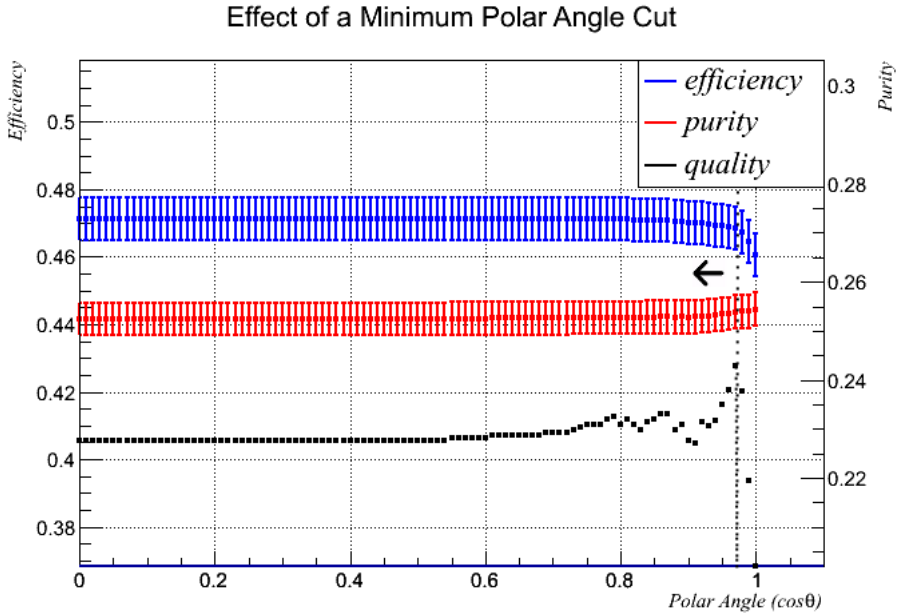

Figure 6.17: Effect of a polar angle cut, twin pair topology

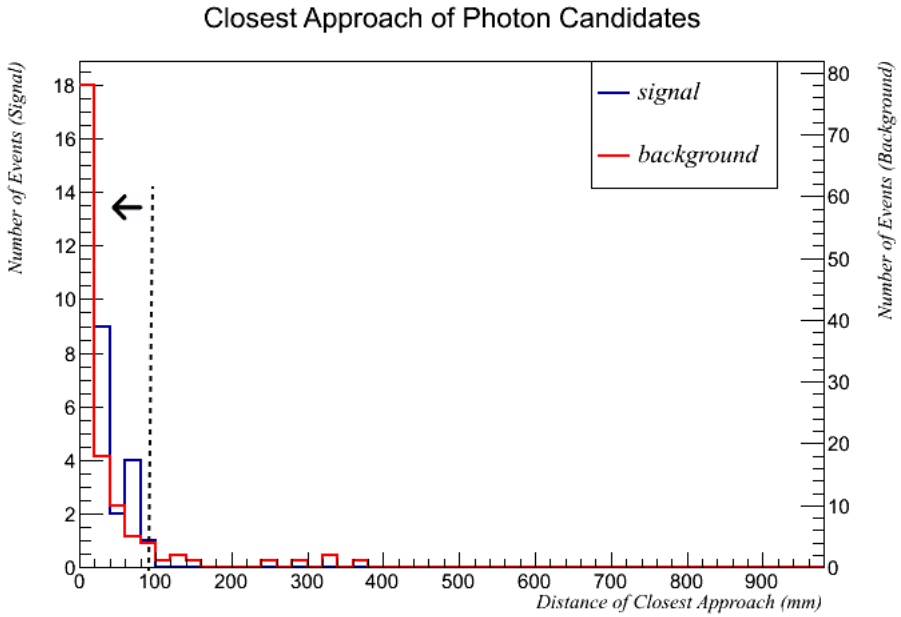

Figure 6.18: Closest approach of projected photons, twin pair topology

6.2.7 Closest Approach

The closest approach is the minimum distance between two photon candidates when their trajectories are projected upstream through the detector. Ideally, if the kinematics of the leptons have been perfectly reconstructed, the distance of closest approach for a true π^0 candidate should be 0. Therefore the value of cutting on this value was investigated.

Figure 6.18 shows the closest approach for twin pair π^0 candidates for both signal and background, with figure 6.19 showing the effectiveness of a cut.

There is some small quality gains to be had by introducing a closest approach, with the quality peaking at a cut of 90mm.

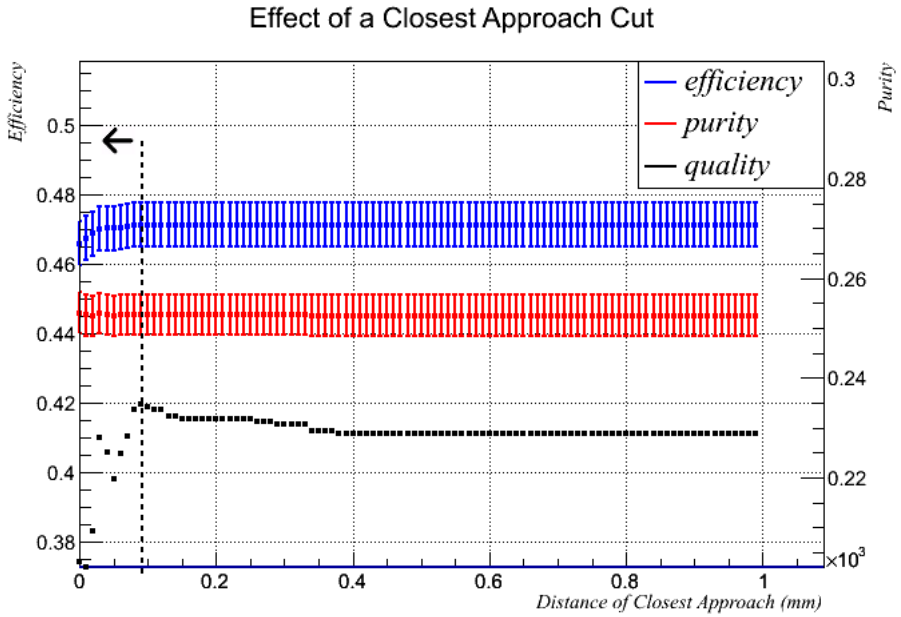

Figure 6.19: Effect of a maximum closest approach cut on the base selection.

6.2.8 Pair Gap Cut

The 'pair gap' is the distance between the average starting positions of the two reconstructed lepton pairs. Ideally this should be low and so could prove a useful variable to cut on. The signal and background pair gap distributions are shown in figure 6.20, with the effect of a cut shown in figure 6.21.

Figure 6.20 shows two distinct clustering of events, one between 0mm and 800mm, with the other between roughly 1100mm and 1800mm. The second cluster is likely due to the two pair candidates originating in separate FGDs. However when examining fig 6.21, cutting out this second cluster does not increase quality. Instead, the quality plateaus at high cut values. Whilst there is a small maxima at 1650mm, this could be due to limited statistics as there are only a handful of reconstructed twin pair candidates with large pair-gaps. As a result no pair gap selection cut was applied to the twin pair topology.

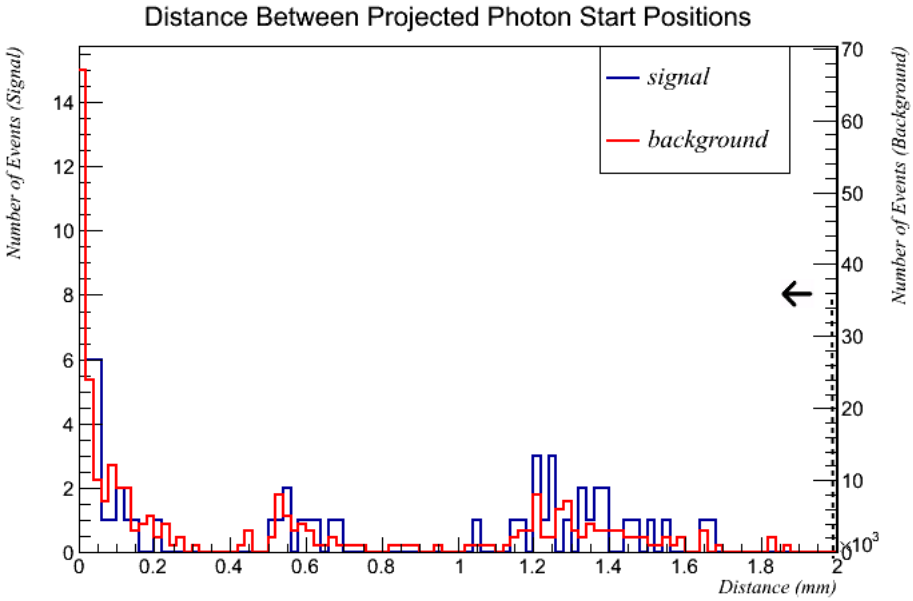

Figure 6.20: Distance between lepton pair start positions, twin pair topology

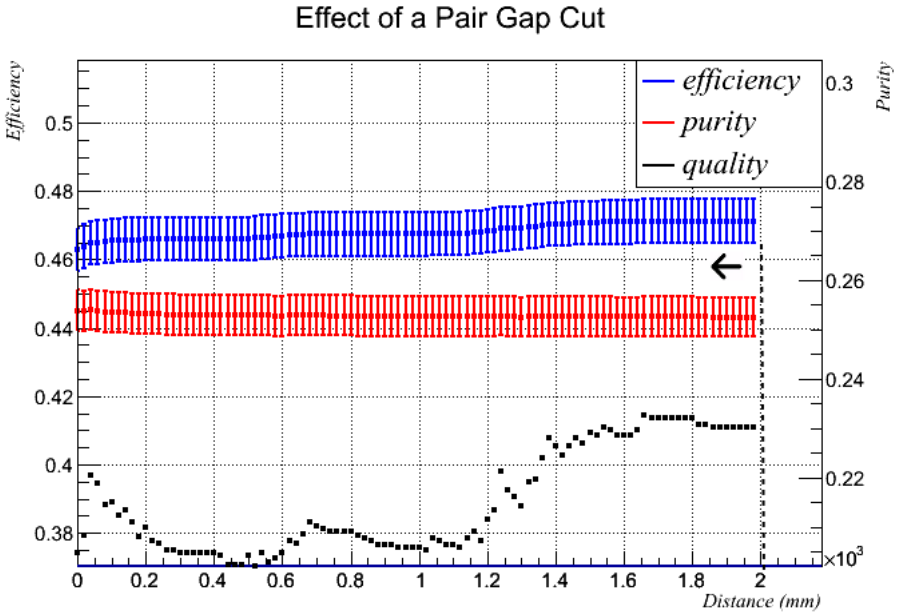

Figure 6.21: Effect of a maximum pair gap cut on the base selection.

6.3 Pair And Isolated Lepton Reconstruction

The pair and isolated electron, and pair and isolated positron topologies were evaluated separately. However, due to their similarities, they are covered in a single section to limit the quantity of histograms being displayed.
The pair and isolated electron sample had very limited statistics available, with only seven reconstructed events in total from all Monte Carlo samples. Because of the low statistics, the base selection was used when developing selection cuts for this topology, outlined in section 4.2.
The pair and isolated positron sample is evaluated using the final set of selection cuts.

6.3.1 Upstream Multiplicity

The upstream multiplicity histograms for the pair and isolated positron topology are shown in figure 6.22, with efficiency, purity and quality shown in figure 6.23.

There is a clear maximum at a multiplicity value of three, therefore a cut value of three was chosen for the pair and isolated positron topology. For the pair and isolated positron topology, the upstream multiplicity distribution is shown in figure 6.24, with the effects of an upstream multiplicity cut shown in figure 6.25. There is a clear maxima in the quality, with a peak of one. Therefore for the pair and isolated electron topology the upstream multiplicity cut was set to one.

6.3.2 Tracker Multiplicity

The tracker multiplicity for the pair and isolated positron topology, final selection, is shown in fig 6.26, with the effect of a potential cut shown in fig 6.27. The

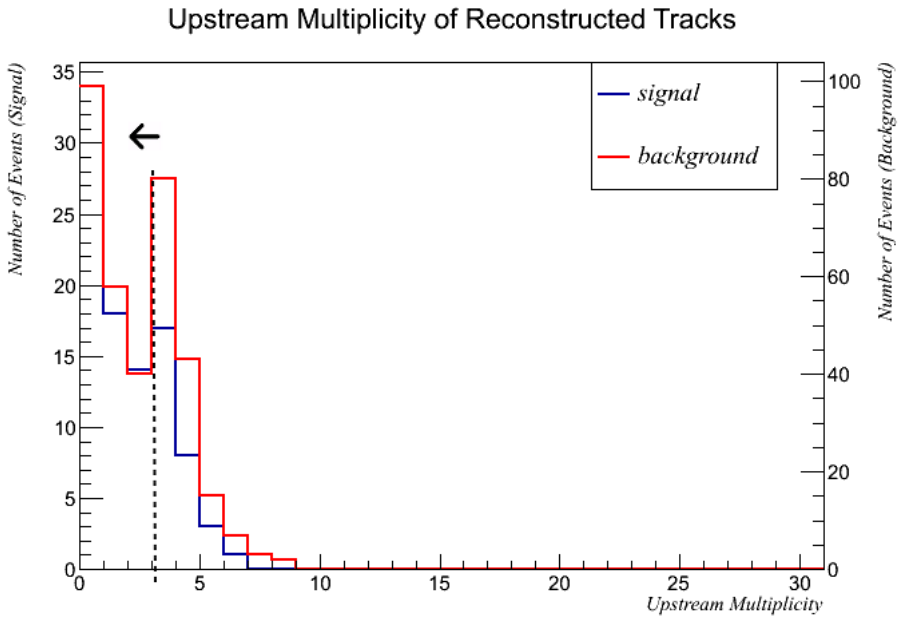

Figure 6.22: Upstream multiplicity of signal and background. Final selection, pair and isolated positron topology.

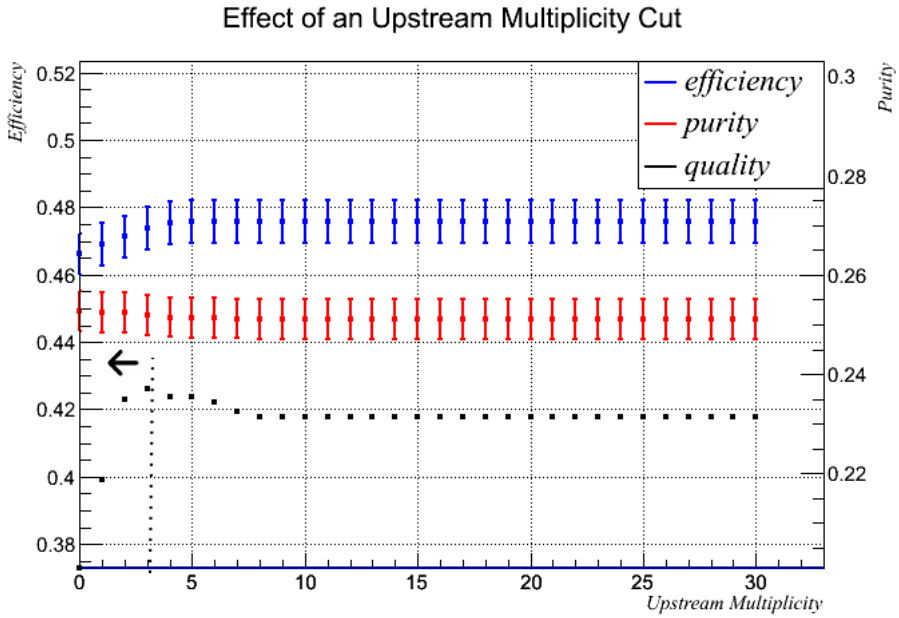

Figure 6.23: Effect of an upstream multiplicity cut on final selection, pair and isolated positron topology.

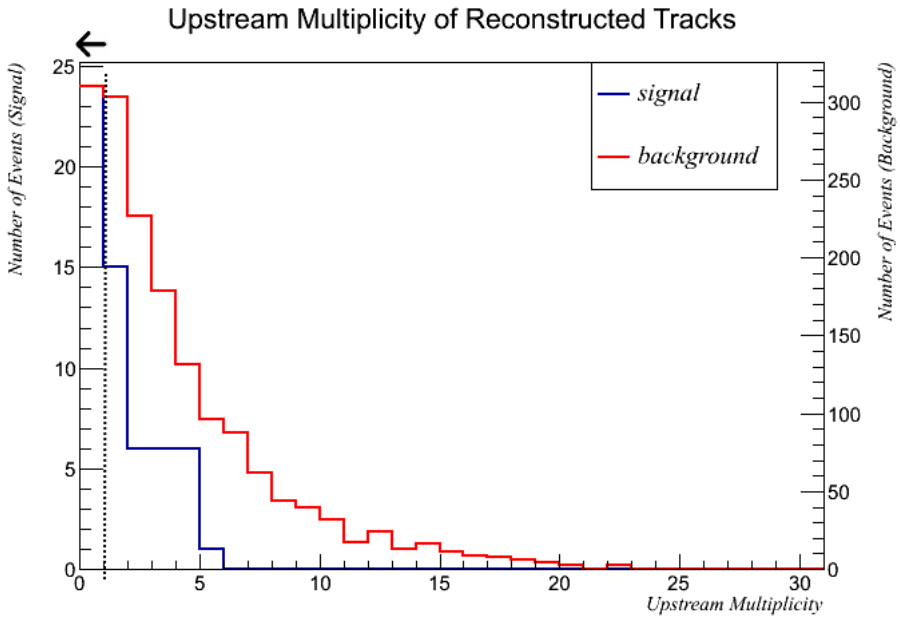

Figure 6.24: Upstream multiplicity of signal and background. Base selection, pair and isolated electron topology.

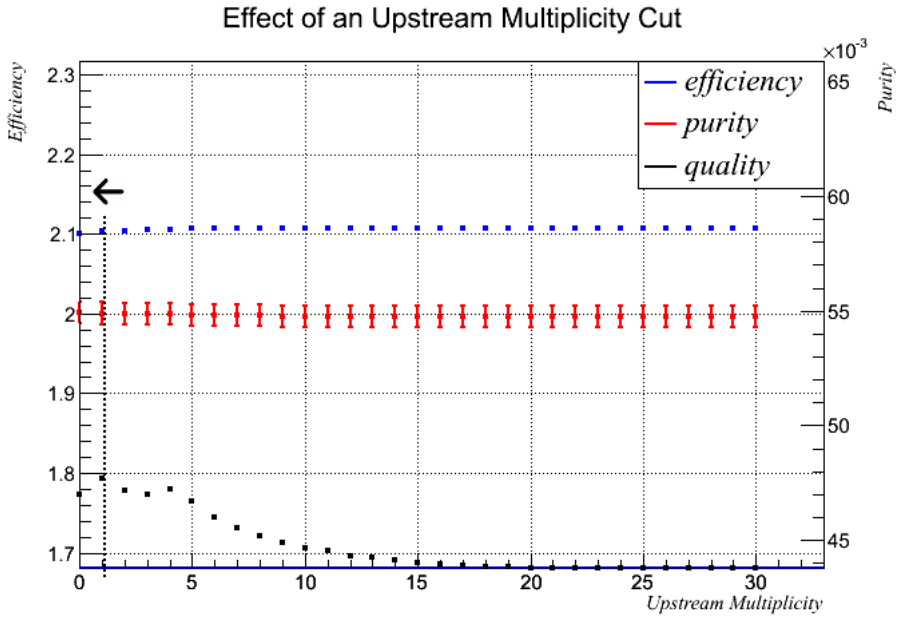

Figure 6.25: Effect of an upstream multiplicity cut on base selection, pair and isolated electron topology.

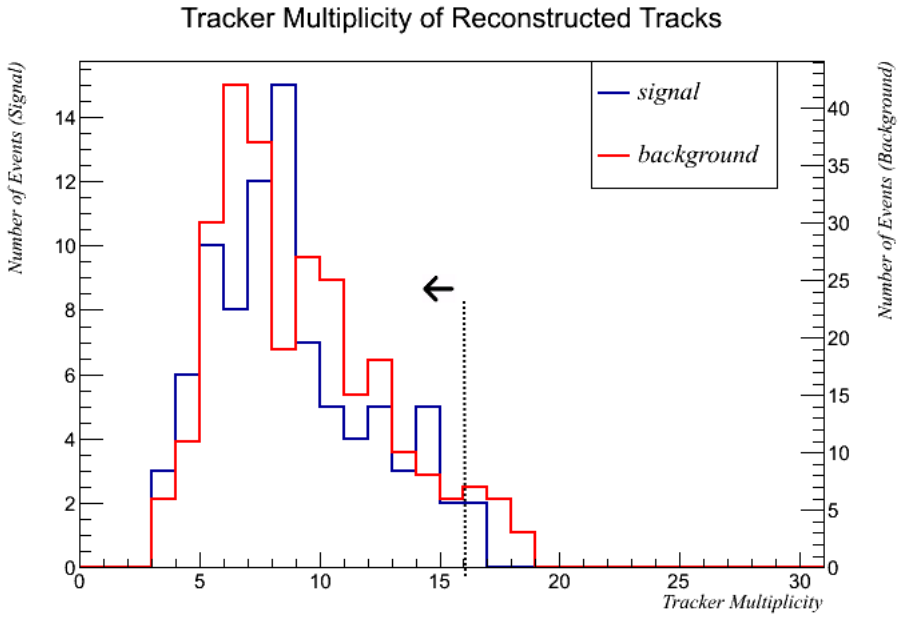

Figure 6.26: Tracker multiplicity of signal and background. Final selection, pair and isolated positron topology.

background distribution appears to have a longer tail than the signal distribution, although the limited statistics make it difficult to say definatively. The quality has a slight maxima at a multiplicity of sixteen, which is where the cut value was set.

For the pair and isolated electron topology, the tracker multiplicity for signal and background is shown in figure 6.28. The effect of a tracker multiplicity cut is shown in figure 6.29. There is a very broad maxima, with its peak at a multiplicity cut of seven, which is where the final cut value was set.

6.3.3 ECal Multiplicity

The ECal multiplicity for the pair and isolated positron topology can be seen in fig 6.30, with the effects of a cut shown in fig 6.31.

There is a peak in quality at a multiplicity cut of four. The effect of an ECal

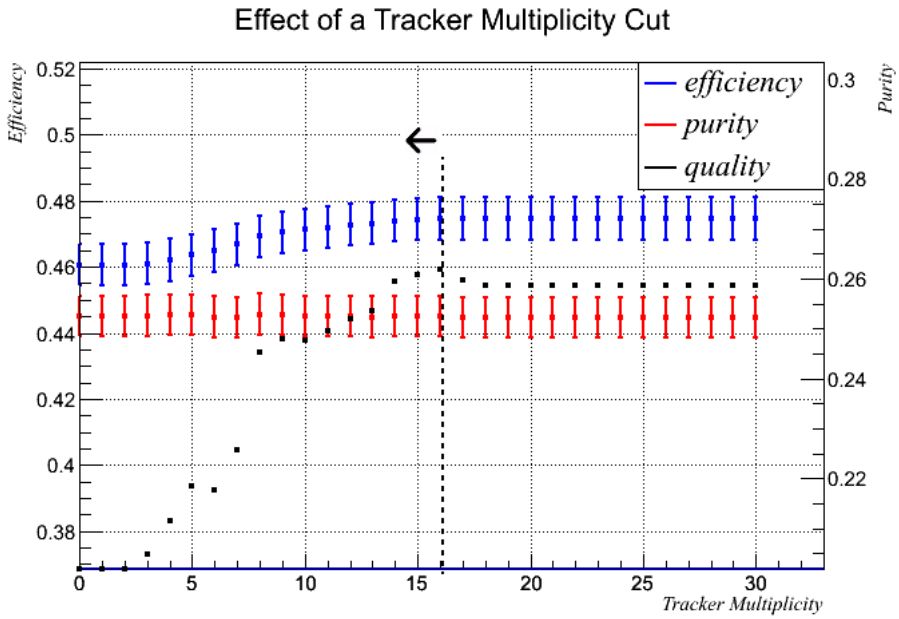

Figure 6.27: Effect of a tracker multiplicity cut on final selection, pair and isolated positron topology.

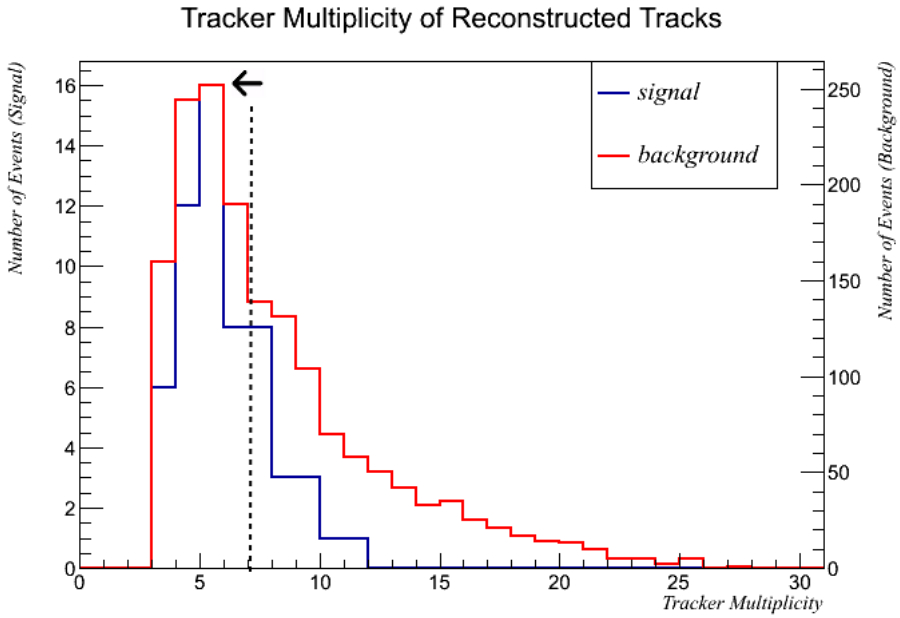

Figure 6.28: Tracker multiplicity of signal and background. Final selection, pair and isolated electron topology.

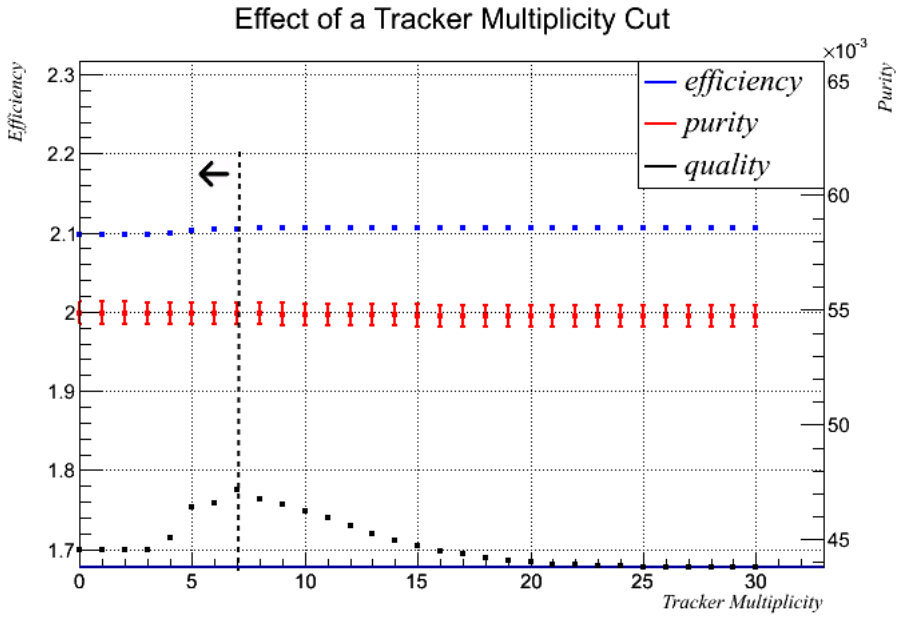

Figure 6.29: Effect of a tracker multiplicity cut on final selection, pair and isolated electron topology.

multiplicity cut for pair and isolated electrons is shown in figure 6.32. There is a slight maxima at a cut value of two, however the gains in quality are similar to a multiplicity cut of zero. Considering the limited statistics available the final cut value was set at two, as this cut permits a greater number of signal events compared to a multiplicity cut of zero.

6.3.4 Isolated ECal Multiplicity

The Isolated ECal multiplicity for the pair and isolated positron topology is shown in figure 6.33, with the effect of a cut shown in figure 6.34.

Similarly to the twin pair scenario, there is no peak in the quality at any value and no isolated ECal multiplicity cut is applied.

The Isolated ECal multiplicity for the pair and isolated electron topology is

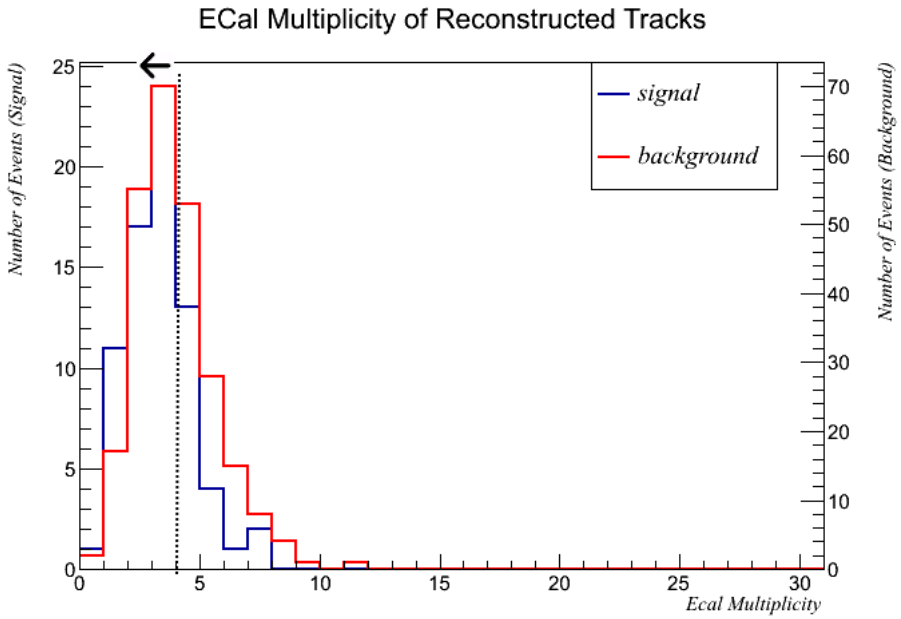

Figure 6.30: ECal multiplicity of signal and background. Final selection, pair and isolated positron topology.

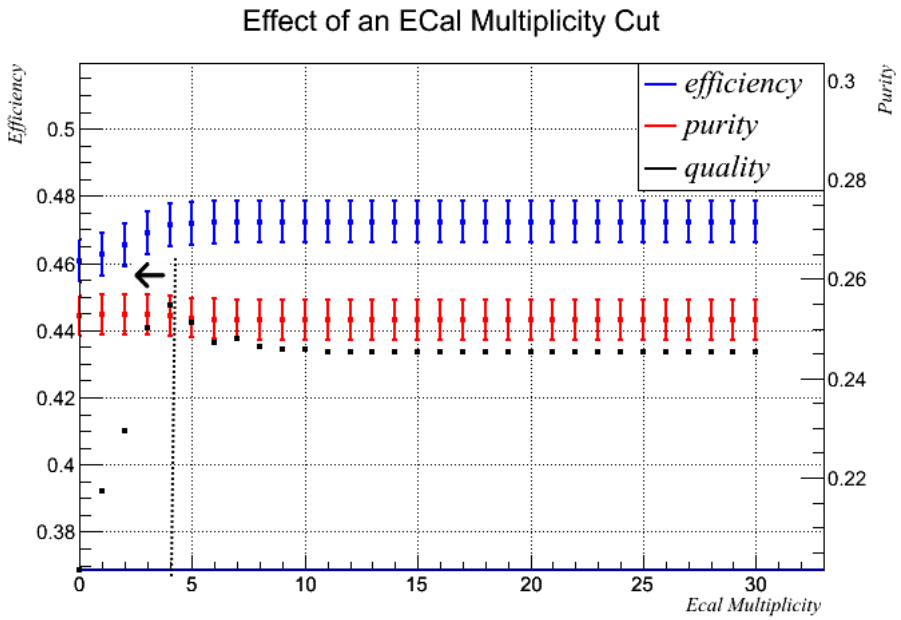

Figure 6.31: Effect of an ECal multiplicity cut on final selection, pair and isolated positron topology.

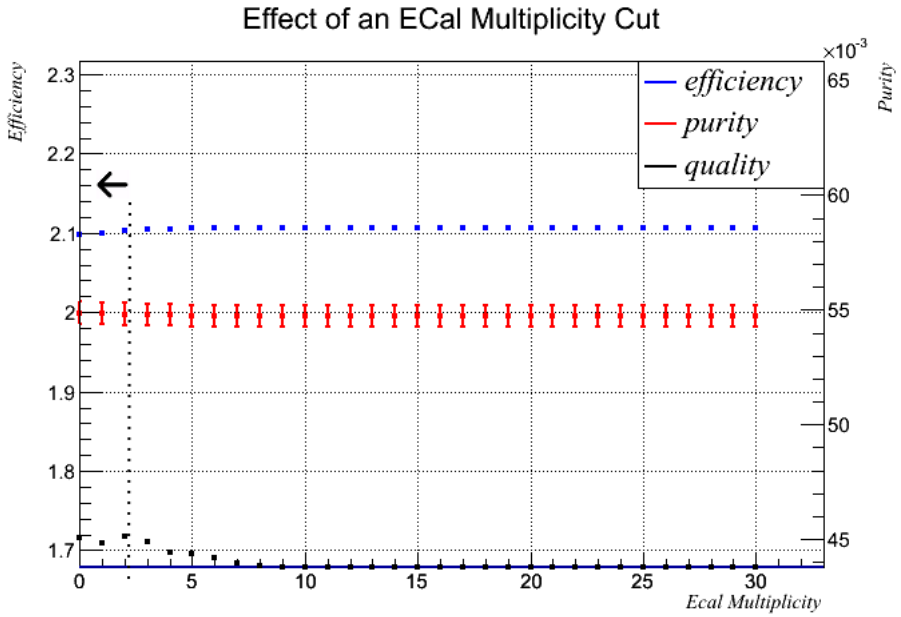

Figure 6.32: Effect of an ECal multiplicity cut on base selection, pair and isolated electron topology.

shown in figure 6.35, with the effect of a cut shown in figure 6.36. Unlike the pair and isolated positron topology there is a peak in quality at a multiplicity of one.

6.3.5 Bunch Multiplicity

The bunch multiplicity for signal and background in the final selection for the pair and isolated positron topology is shown in figure 6.37, with the effect of a selection cut shown in figure 6.38.

The quality plateaus at high multiplicity values, however there is a marginal maxima at a multiplicity cut of nineteen. It is possible that this maxima is a purely statistical effect as there are a limited amount of reconstructed pair and isolated positron signal events. However, a cut at nineteen removes few signal events from the selection, therefore the final cut value was set to nineteen to

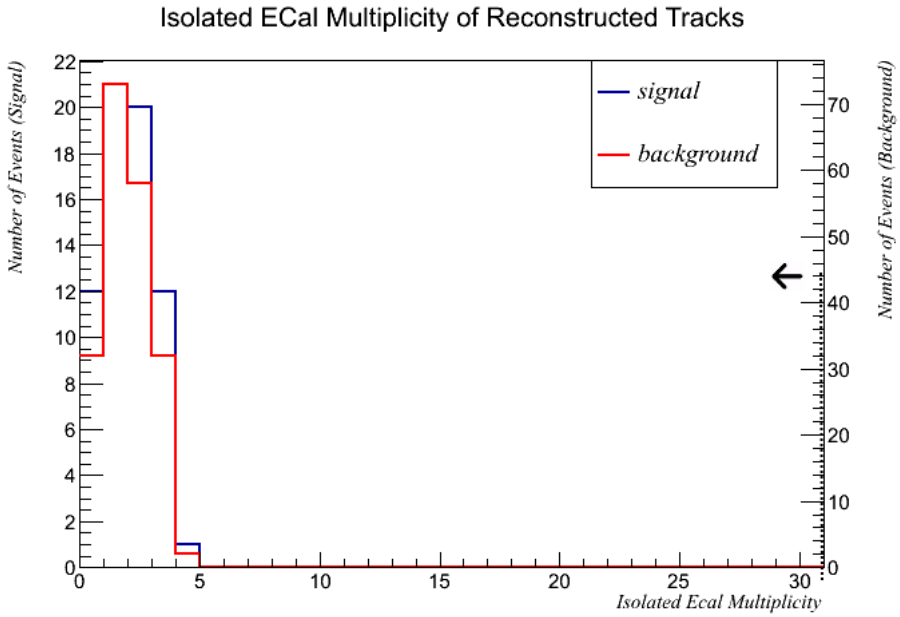

Figure 6.33: Isolated ECal multiplicity of signal and background for the pair and isolated positron topology.

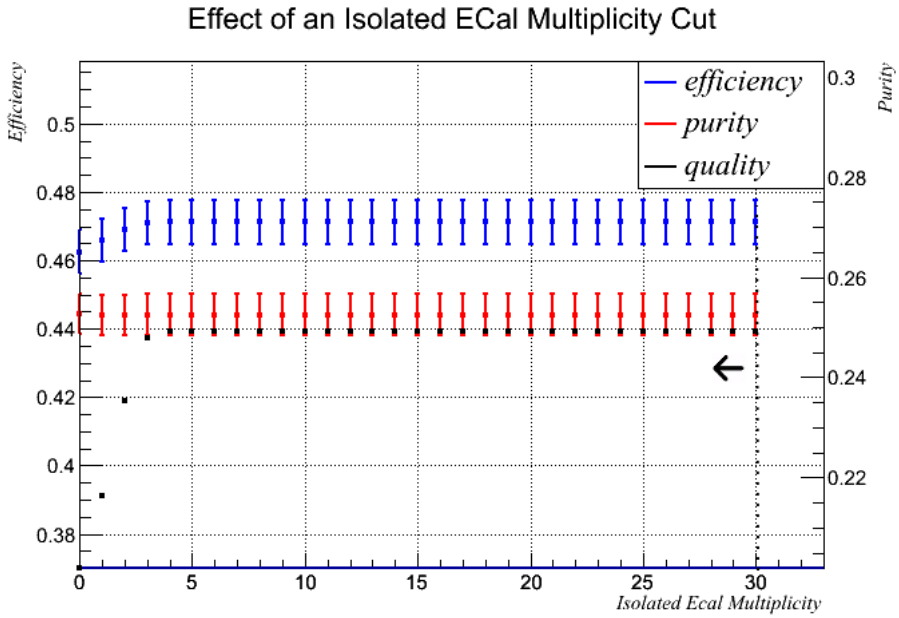

Figure 6.34: Effect of an Isolated ECal multiplicity cut on final selection, pair and isolated positron topology.

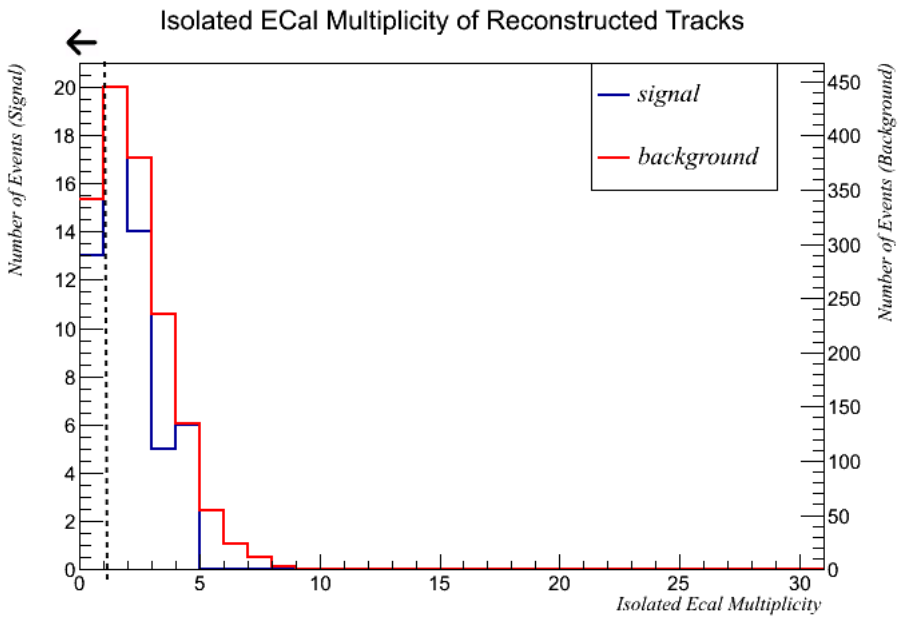

Figure 6.35: Isolated ECal multiplicity of pair and isolated electron candidates, base selection.

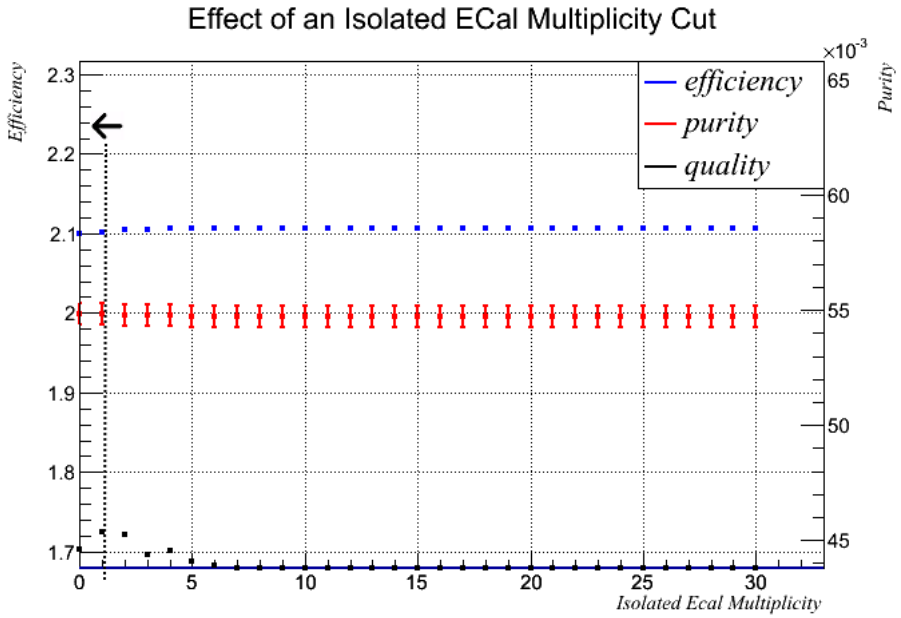

Figure 6.36: Effect of an Isolated ECal multiplicity cut, pair and isolated electron topology, base selection.

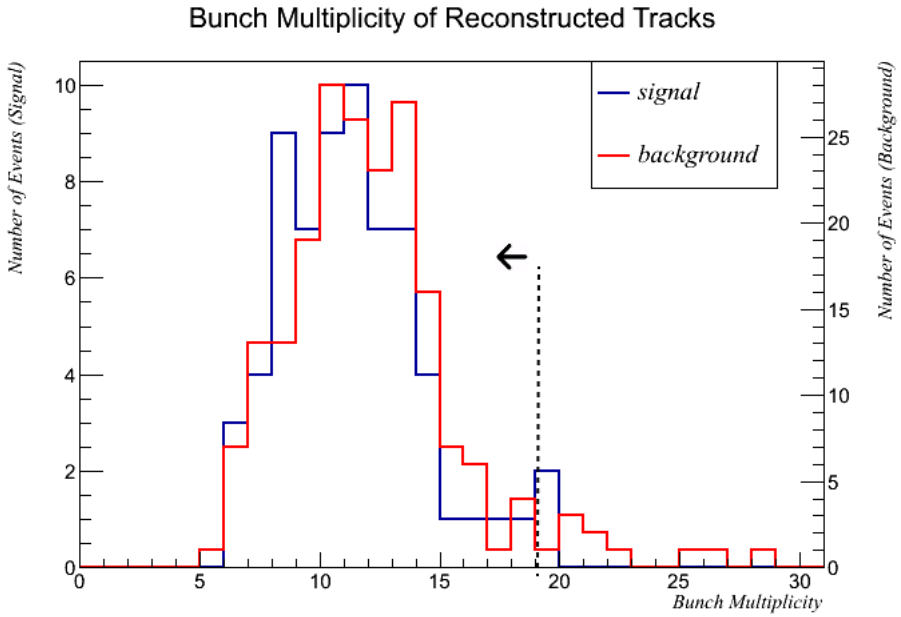

Figure 6.37: Bunch multiplicity cut of final selection, pair and isolated positron topology.

improve the purity of the analysis.

For the pair and isolated electron topology, the effect of a bunch multiplicity cut is shown in figure 6.39. There is a very broad maxima with a peak at twelve, so the final cut value was set at twelve. The large gains shown by this histogram are likely to be exaggerated however. It is plausible that selection of events removed by other selection cuts, namely the tracker and ECal multiplicity cuts, overlap with the selection of events a bunch multiplicity cut would remove.

6.3.6 Opening Angle

The opening angle of the photon candidates for the pair and isolated positron topology is shown, for the final selection, in figure 6.40. The effect of a selection cut is shown in figure 6.41.

The quality histogram plateaus at low opening angle values. Whilst there is a marginal increase in quality for a cut of $\cos\theta = 0.06$ compared with hav-

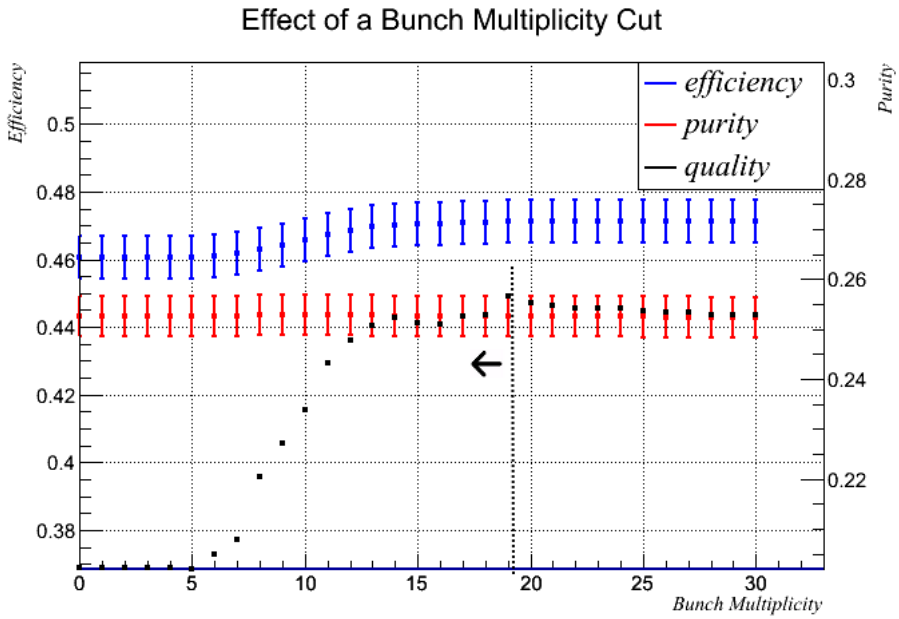

Figure 6.38: Effect of an bunch multiplicity cut on the final selection, pair and isolated positron topology.

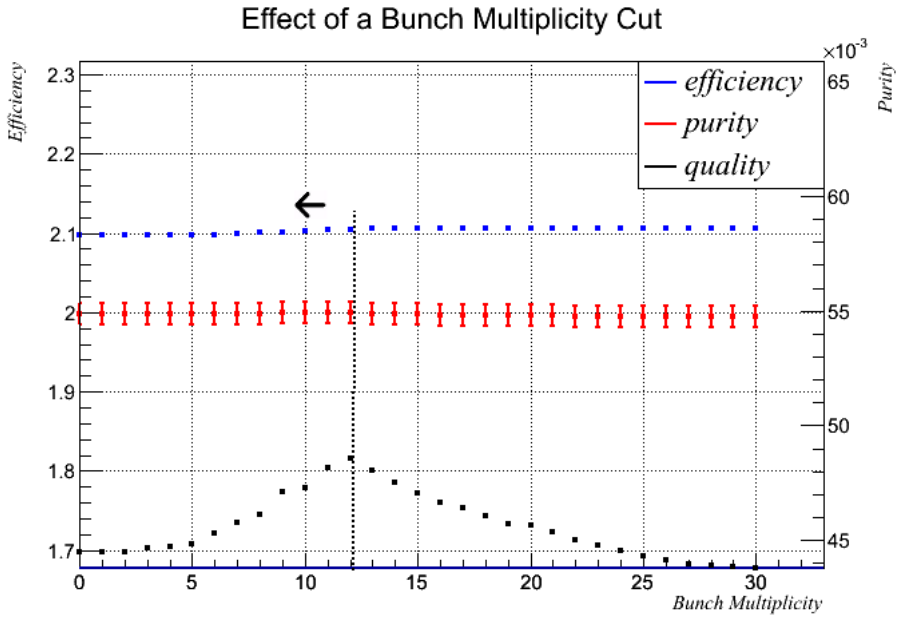

Figure 6.39: Effect of an bunch multiplicity cut on the base selection, pair and isolated electron topology.

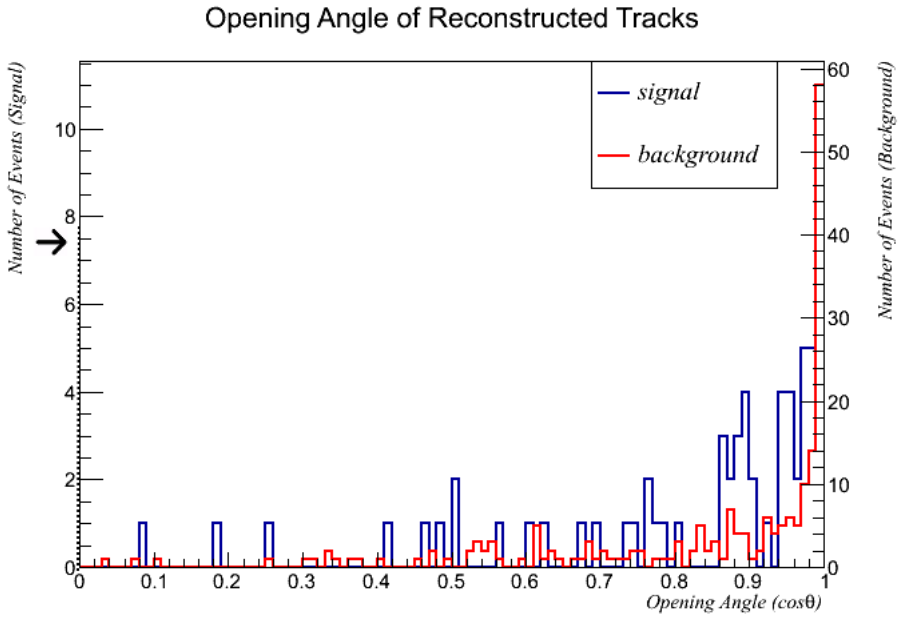

Figure 6.40: Opening angle of final selection, pair and isolated positron topology.

ing no cut, this is likely an artifact of limited statistics. The vast majority of signal and background events are clustereda round $\cos\theta = 1$ and low-opening angle cut thresholds exclude very few events (the 'peak' at 0.06 excludes two background events). Therefore no opening angle cut is applied to the pair and isolated positron topology. This situation is reflected in the pair and isolated electron topology, where there is no benefit to applying an opening angle cut.

6.3.7 Polar Angle

The polar angles for reconstructed pair and isolated positron topology events is shown in figure 6.42. The signal distribution's tail is not as broad as the background distribution, therefore any minimum polar angle cut does not increase the quality. This is shown in figure 6.43, which shows the effect of a polar angle cut on the pair and isolated positron topology. The same conclusions were reached for the pair and isolated electron topology.

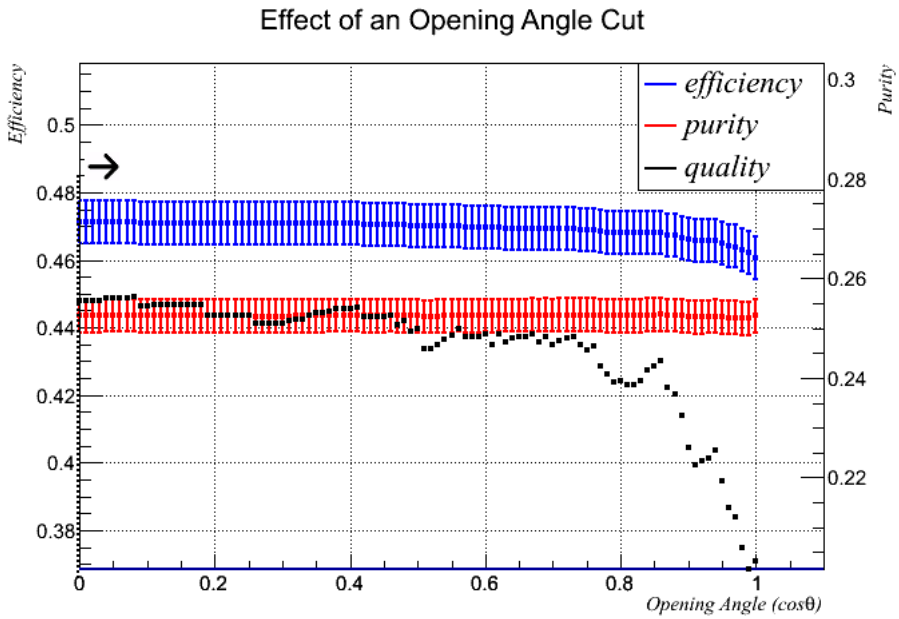

Figure 6.41: Effect of an opening angle cut on final selection, pair and isolated positron topology.

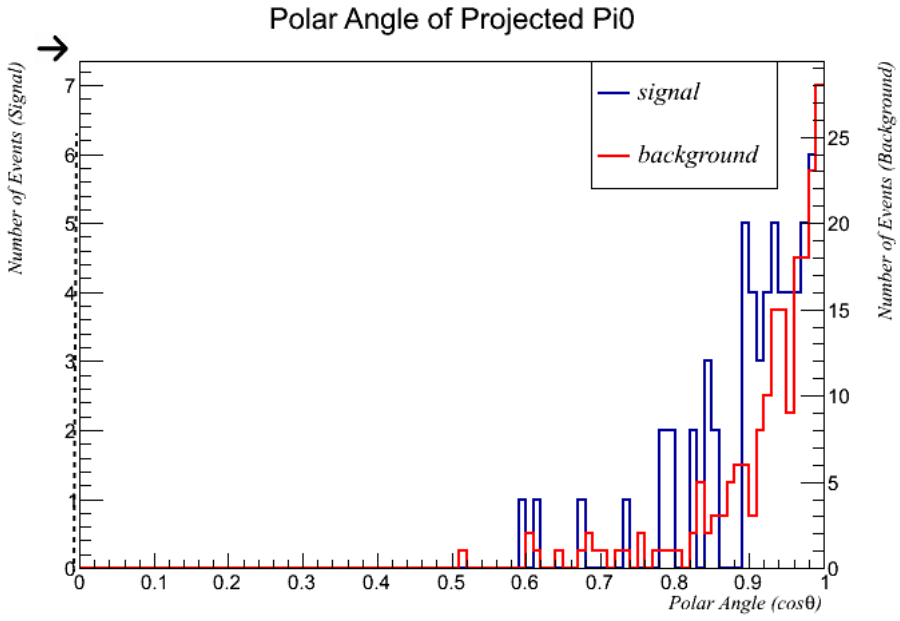

Figure 6.42: Polar angle of final selection, pair and isolated electron topology.

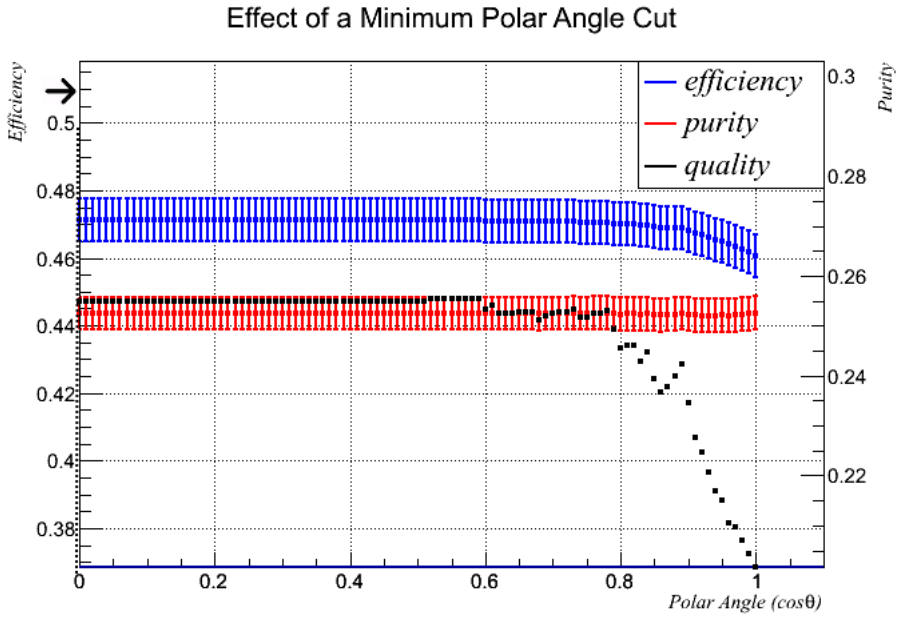

Figure 6.43: Effect of a polar angle cut, pair and isolated electron topology.

6.3.8 Closest Approach

The closest approach of candidate photons for the pair and isolated positron topology is shown in figure 6.44, with the effects of a cut shown in figure 6.45.

The quality appears to plateau at high values, with very little benefit to having a cut at any level. Whilst there is a marginal increase in quality at a cut value of 820mm this is, like the opening angle cut in section 6.3.6, an artifact resulting from very low statistics. This situation is reflected in the pair and isolated electron topology, where there is no benefit to applying a closest approach cut.

6.3.9 Pair Gap

The pair gap for reconstructed pair and isolated positron topology events is shown in figure 6.46. The pair gap for the pair and isolated lepton topology is the difference between the average starting position of the lepton pair and

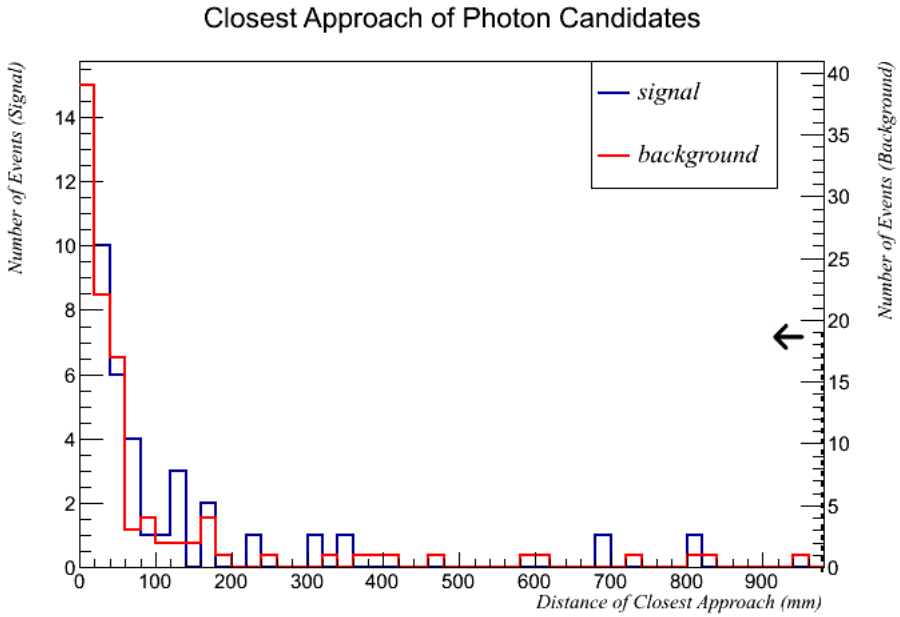

Figure 6.44: Closest approach of candidate photons using final selection, pair and isolated positron topology.

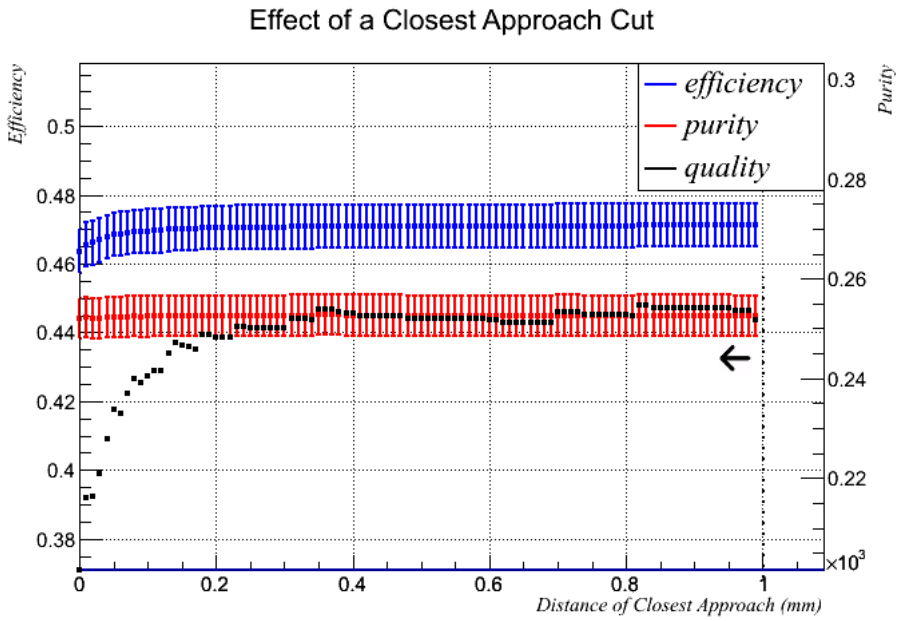

Figure 6.45: Effect of a closest approach cut, pair and isolated positron topology.

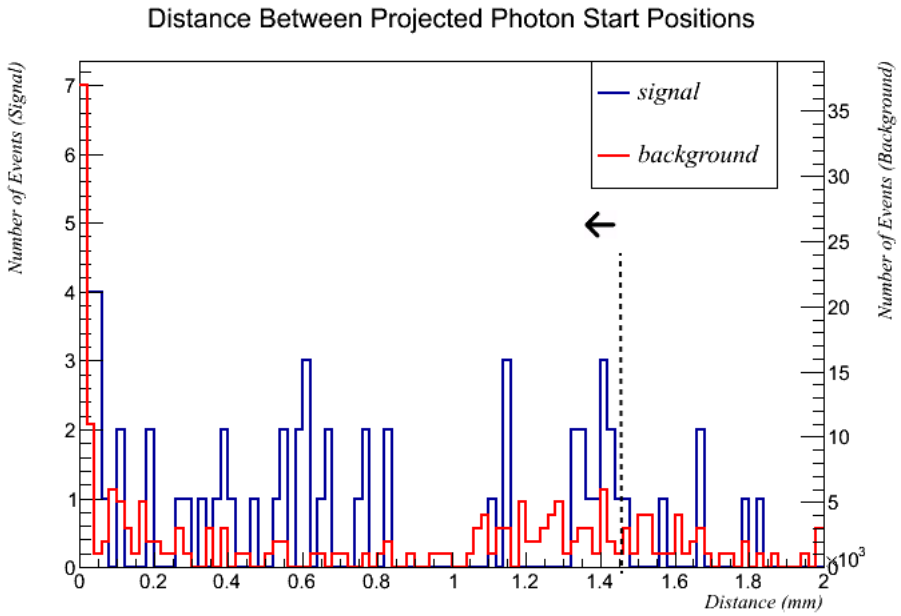

Figure 6.46: Distance between photon end positions of final selection, pair and isolated positron topology.

the starting position of the isolated lepton. The limited statistics are spread over a broad spectrum of values, however one can see the events in fig 6.46 are clustered into two regions: one from 0mm to approximately 800mm, and the other from about 1200mm to 1800mm. This is similar to the structure seen in the histograms for the twin pair topology (section 6.3.9). The second region of events arises when tracks for a given π^0 candidate originate in different FGDs. The quality histogram has a maxima at 1450mm, which excludes the majority of these mis-matched events.

The effect of a pair gap cut for the pair and isolated electron topology is shown in figure 6.48, where there is also a peak in quality, this time at 600mm.

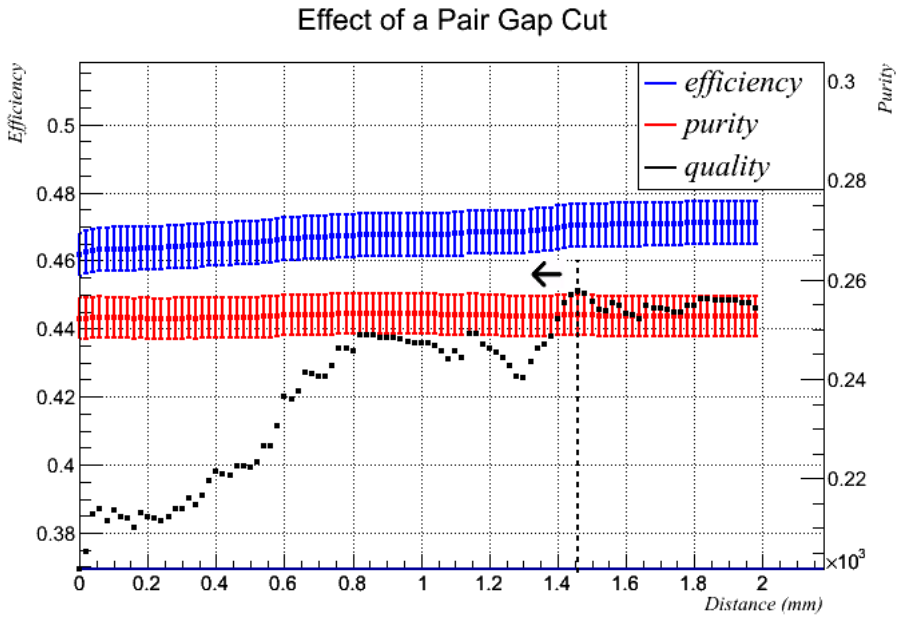

Figure 6.47: Effect of a pair gap selection cut, pair and isolated positron topology.

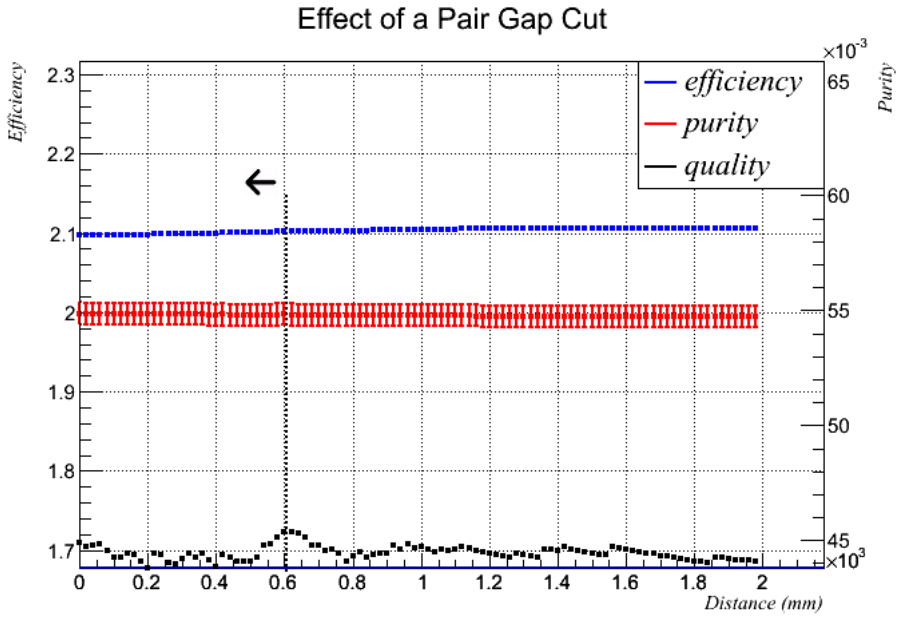

Figure 6.48: Effect of a pair gap selection cut on base selection, pair and isolated electron topology.

6.4 Twin Isolated Lepton Reconstruction

Three event topologies have been combined into this subsection due to their similarities: twin isolated electron, twin isolated positron, and isolated electron and isolated positron.

These three topologies have the least amount of kinematic information available, with only two reconstructed tracks per π^0 vertex. As a result backgrounds are high. It is important to apply selection cuts to try and increase the purity of this set of events.

6.4.1 Bunch Multiplicity

The bunch multiplicity for the twin isolated electron final sample is shown in figure 6.49. The signal follows a very similar distribution to previous topologies, however the background has a slightly different distribution. The two-track topologies are more likely to accept high-multiplicity background events because there is less information that can be used to discriminate between signal and background.

The efficiency, purity and quality histograms are shown in figure 6.50. There is a broad peak with a maximum at a multiplicity cut of 13.

The bunch multiplicity for the isolated electron and isolated positron topology is shown in figure 6.51, with the effect of a cut shown in figure 6.52. In figure 6.51 the background appears to be relatively greater than the signal at high multiplicities, however applying a cut at any level reduces the efficiency of the analysis too severely for there to be any gain in quality. The quality plateaus at high multiplicities and there is no value in applying a bunch multiplicity cut at

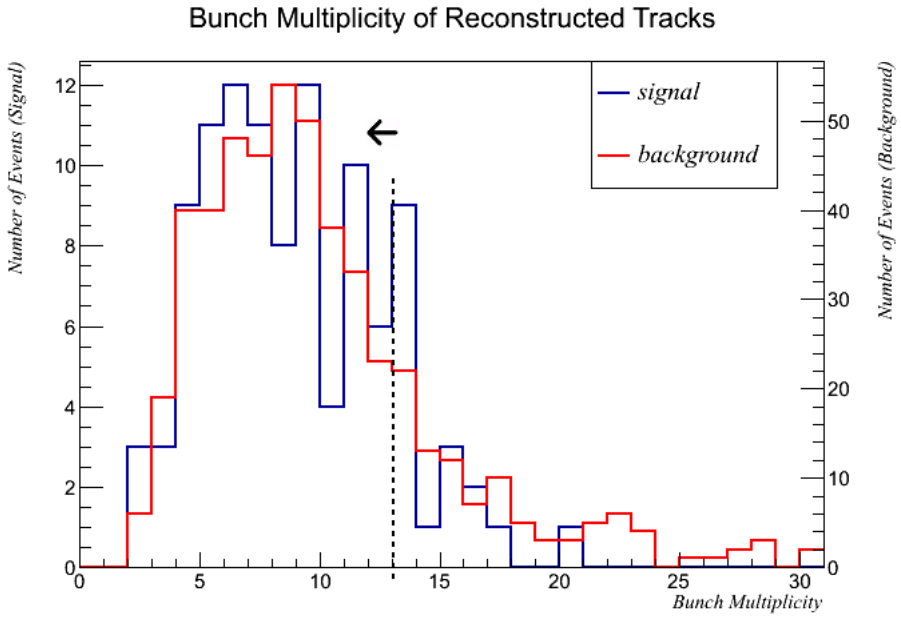

Figure 6.49: Bunch multiplicity of final selection, twin isolated electron topology.

any level. This situation is also reflected in the twin isolated positron topology, where no bunch multiplicity cut is applied.

6.4.2 Tracker Multiplicity

The tracker multiplicity for signal and background,for the isolated electron and isolated positron final selection is shown in figure 6.53. The distributions are highly similar, and the other two topologies also showed similar distributions. The effect of a cut is shown in figure 6.54.

No value is seen in applying a tracker multiplicity cut at any level for any topology. The hypothesis was that the signal events would have relatively low multiplicities compared to the background, however the actual multiplicity distribution is very broad. Reconstructing signal events from only two tracks introduces high-multiplicity, poorly reconstructed events and reduces the ability to

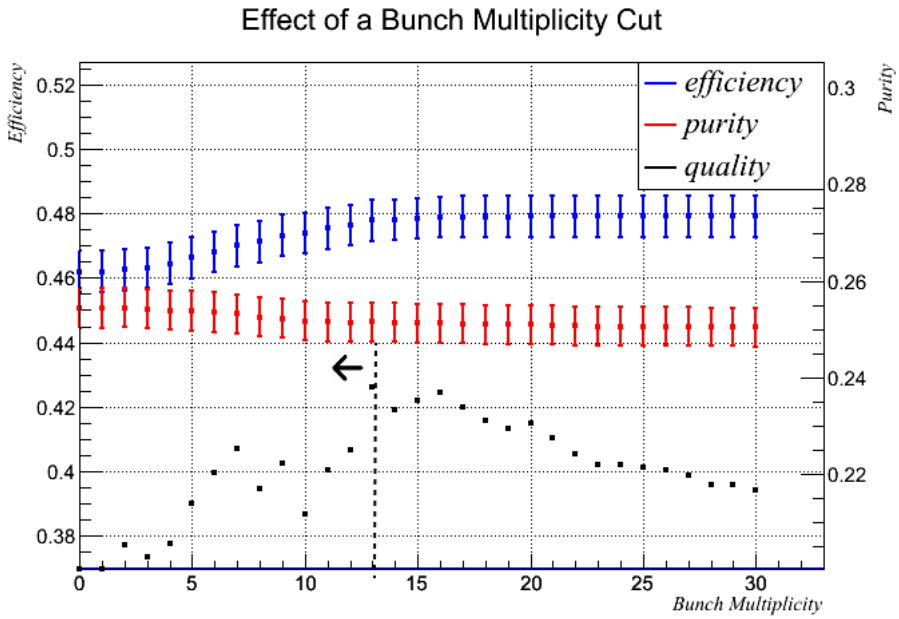

Figure 6.50: Effect of a bunch multiplicity cut on final selection, twin isolated electron topology.

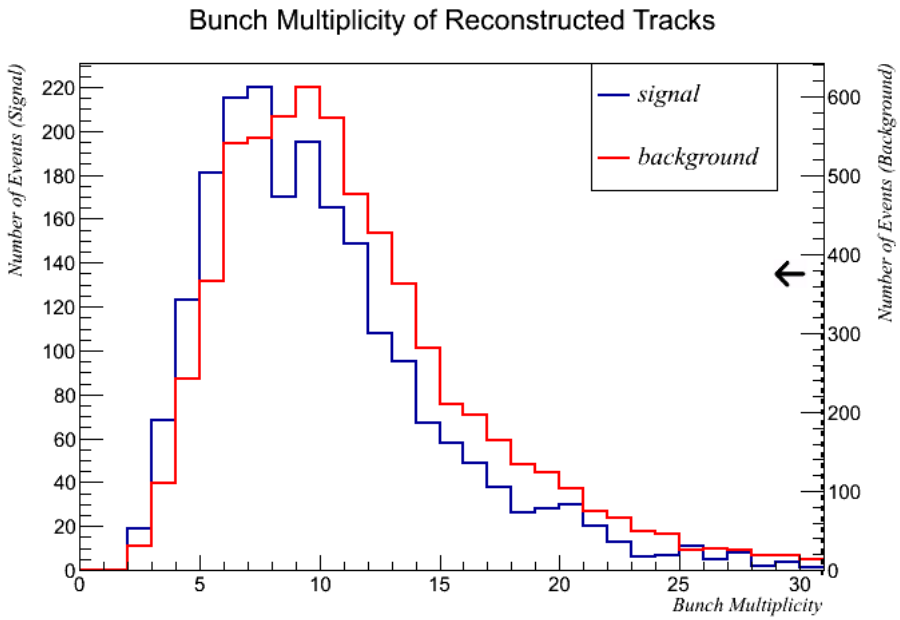

Figure 6.51: Bunch multiplicity of final selection, pair and isolated electron topology.

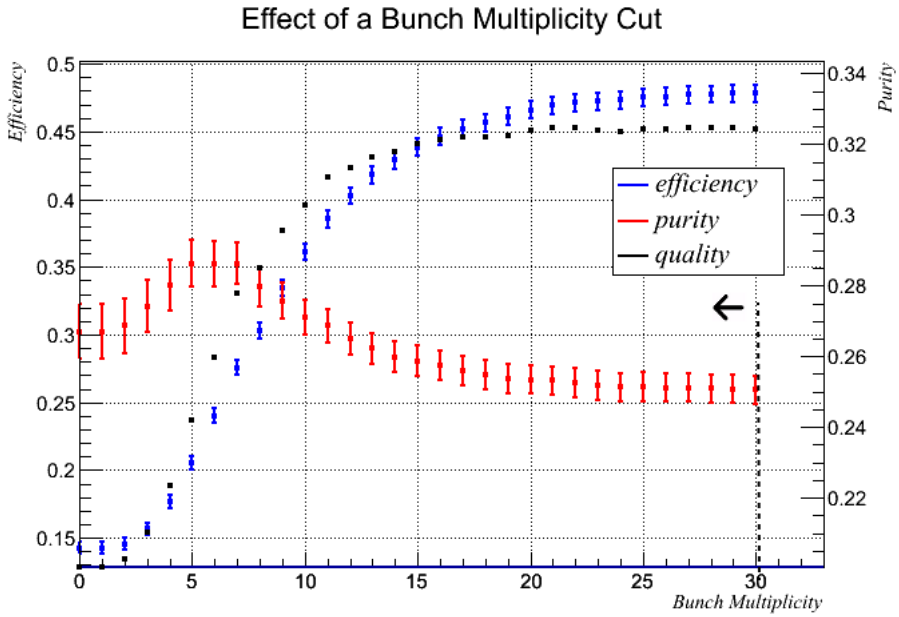

Figure 6.52: Effect of a bunch multiplicity cut on final selection, isolated electron and isolated positron topology.

discriminate between signal and background based on track multiplicity alone.

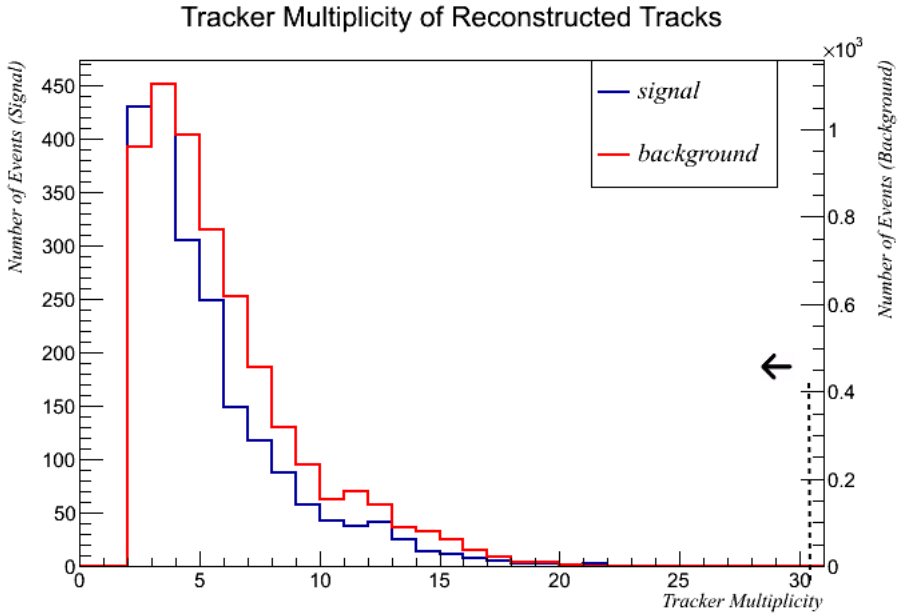

Figure 6.53: Tracker multiplicity of final selection, pair and isolated electron topology.

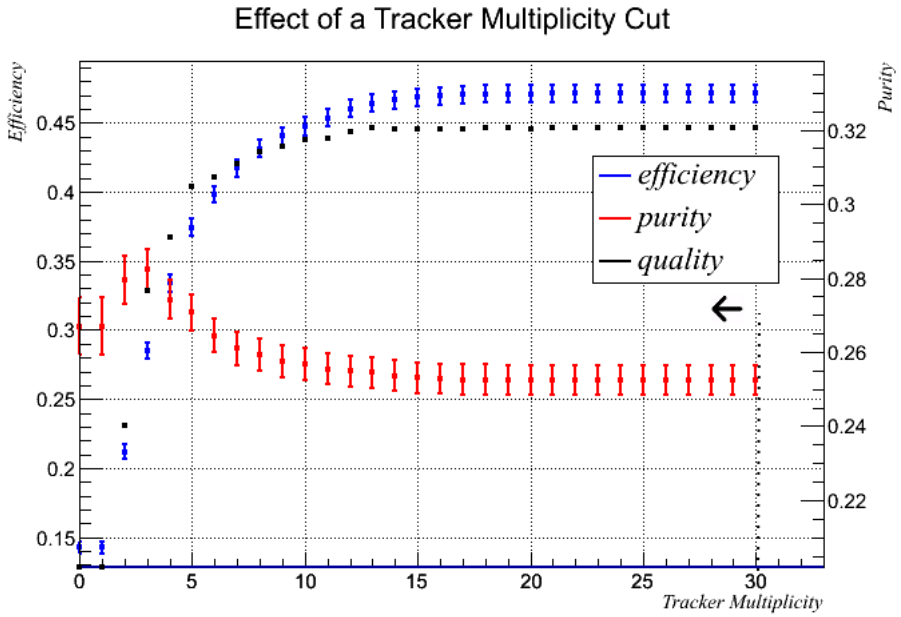

Figure 6.54: Effect of a tracker multiplicity cut on final selection, isolated electron and isolated positron topology.

6.4.3 Upstream Multiplicity

The upstream multiplicity for the twin isolated positron topology is shown in figure 6.55, with the effects of a cut shown in figure 6.56.

There is a clear benefit in terms of purity and quality to having an upstream multiplicity cut of 0. For both the twin isolated electron topology and the isolated electron and isolated positron topology the cut value was set at 0 for identical reasons.

6.4.4 Isolated ECal Multiplicity

The isolated ECal multiplicity is shown for the isolated electron and isolated positron topology, in figure 6.57. The effect of the cut for the tight selection is shown in figure 6.58. There is no maximum, and this pattern is repeated for the other two topologies. Therefore no isolated ECal multiplicity cut was used.

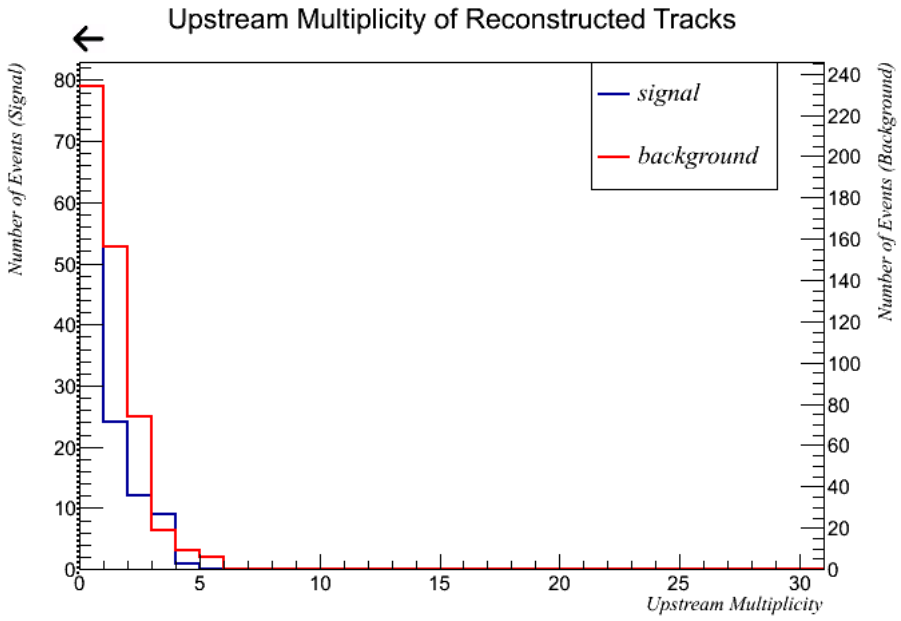

Figure 6.55: Upstream multiplicity of final selection, twin isolated positron topology.

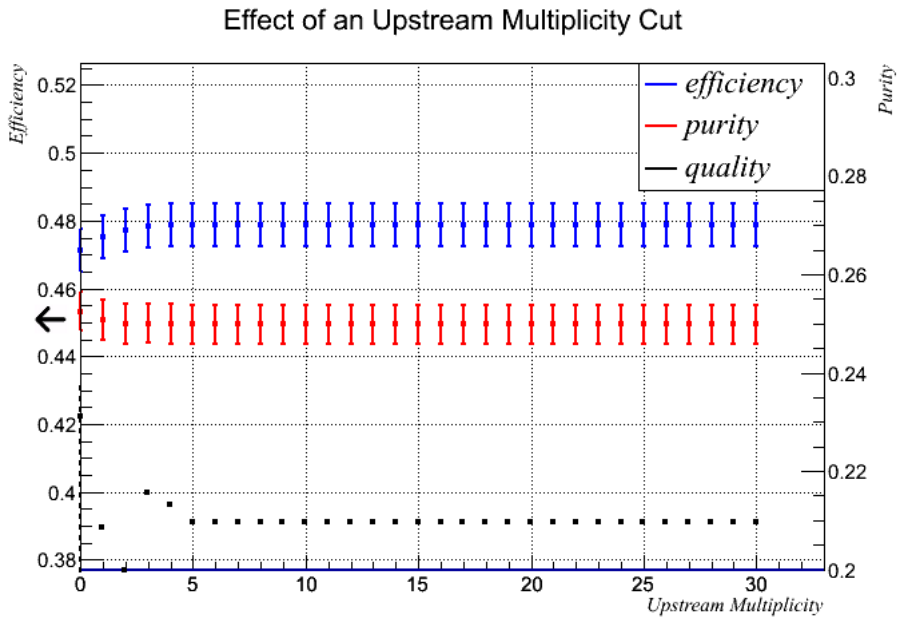

Figure 6.56: Effect of an upstream multiplicity cut on final selection, twin isolated positron topology.

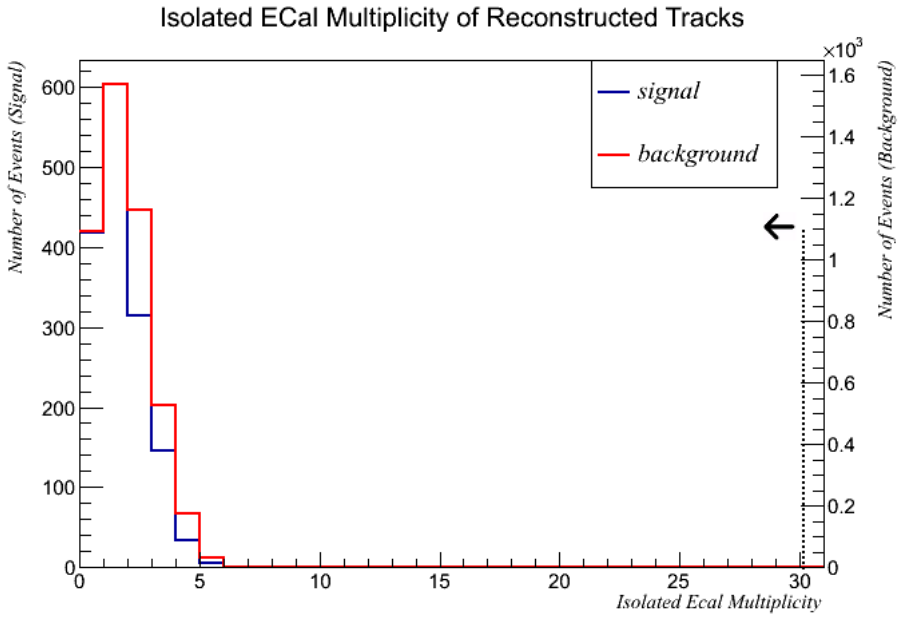

Figure 6.57: Isolated ECal multiplicity of final selection, isolated electron and isolated positron topology.

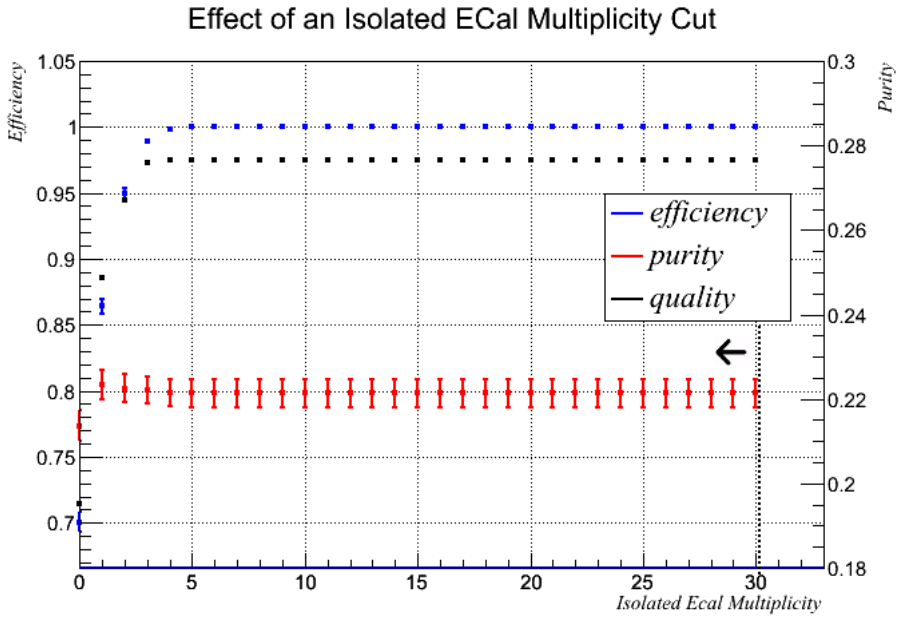

Figure 6.58: Effect of an isolated ECal multiplicity cut on final selection, isolated electron and isolated positron topology.

6.4.5 ECal Multiplicity

The ECal multiplicity for signal and background are displayed, for the twin isolated electron topology final selection in figure 6.59. The effect of a cut is shown in figure 6.60.

When looking at the effect of a cut for the twin isolated electron topology, there is a peak at a multiplicity cut of 0. The ECal multiplicity distribution for the win isolated electron topology is shown in figure 6.59 and the effect of an ECal multiplicity cut is shown in figure 6.60.

For the twin isolated positron topology, the quality peaks at an ECal multiplicity cut value of two. The ECal multiplicity distribution for the win isolated positron topology is shown in figure 6.61 and the effect of an ECal multiplicity cut is shown in figure 6.62.

Figure 6.63 shows the effect of an ECal multiplicity cut for the isolated electron and isolated topology. For this topoology, no benefit is gained from an ECal multiplicity cut as the quality plateaus at high multiplicities.

6.4.6 Opening Angle

The opening angles for the twin isolated lepton topologies are larger than the other topologies because each photon candidate has missing kinematic information, creating differences between the estimated properties of the photons and their true counterparts. However because of the wider acceptance, the background should also have a larger opening angle spread. The opening angle for the isolated electron and isolated positron topology, final selection, is shown in figure 6.64. Both signal and background have a wide spread of an-

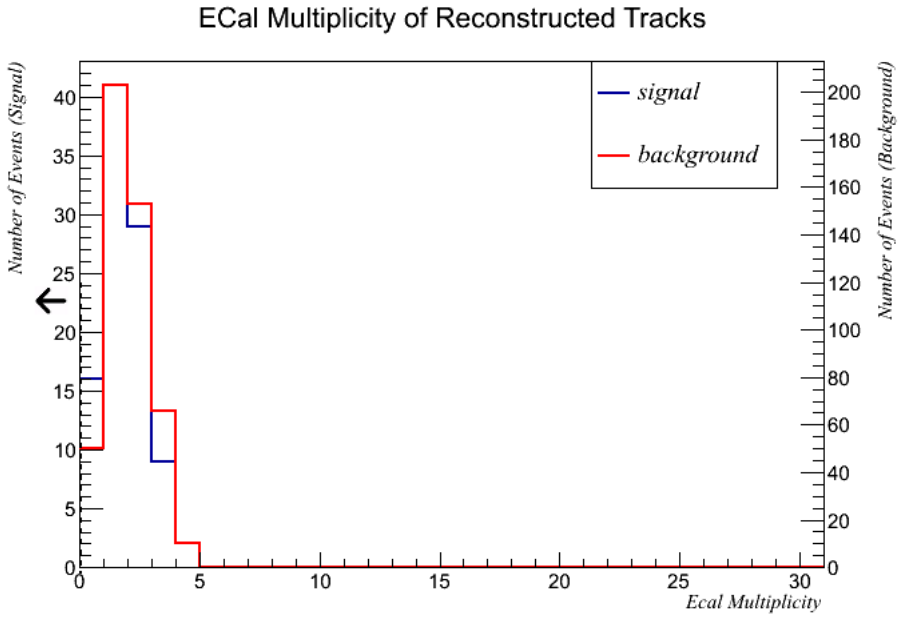

Figure 6.59: ECal multiplicity of final selection, twin isolated electron topology.

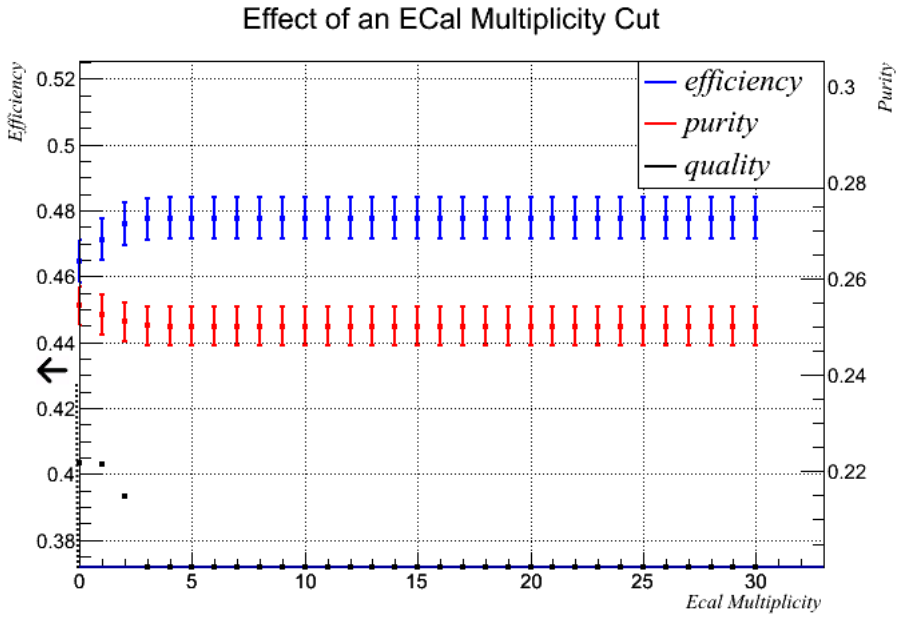

Figure 6.60: Effect of an ECal multiplicity cut on final selection, twin isolated electron topology.

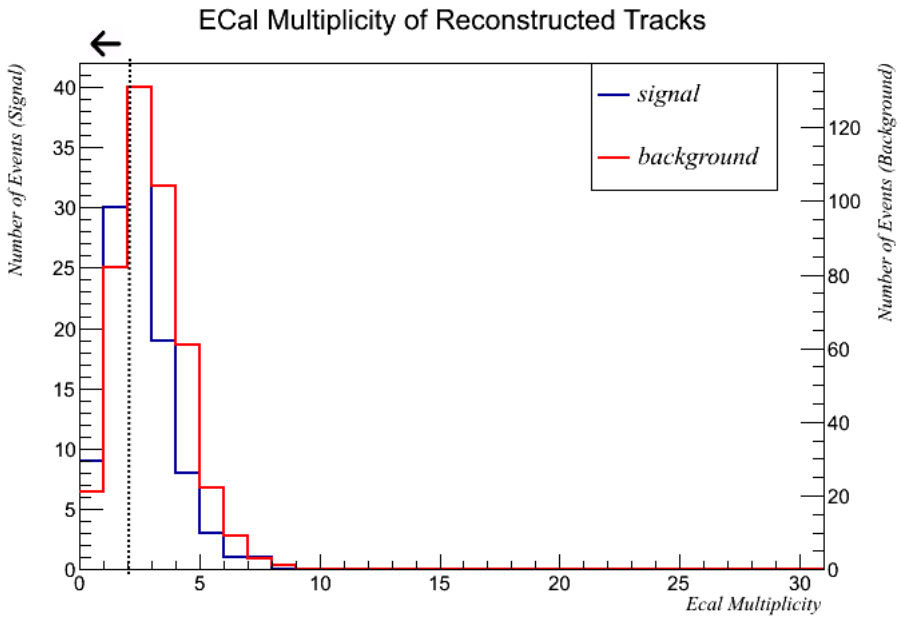

Figure 6.61: ECal multiplicity of final selection, twin isolated positron topology.

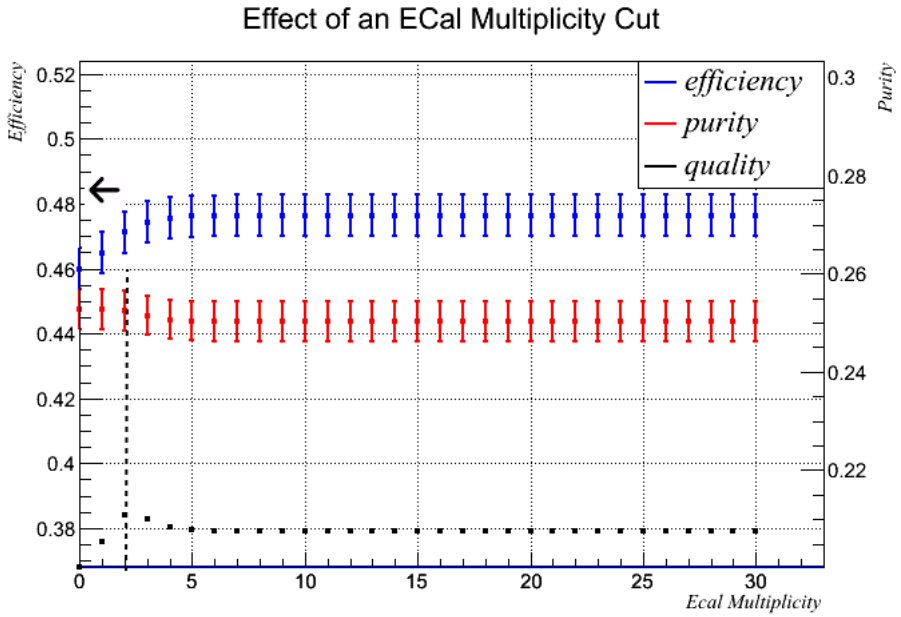

Figure 6.62: Effect of an ECal multiplicity cut on final selection, twin isolated positron topology.

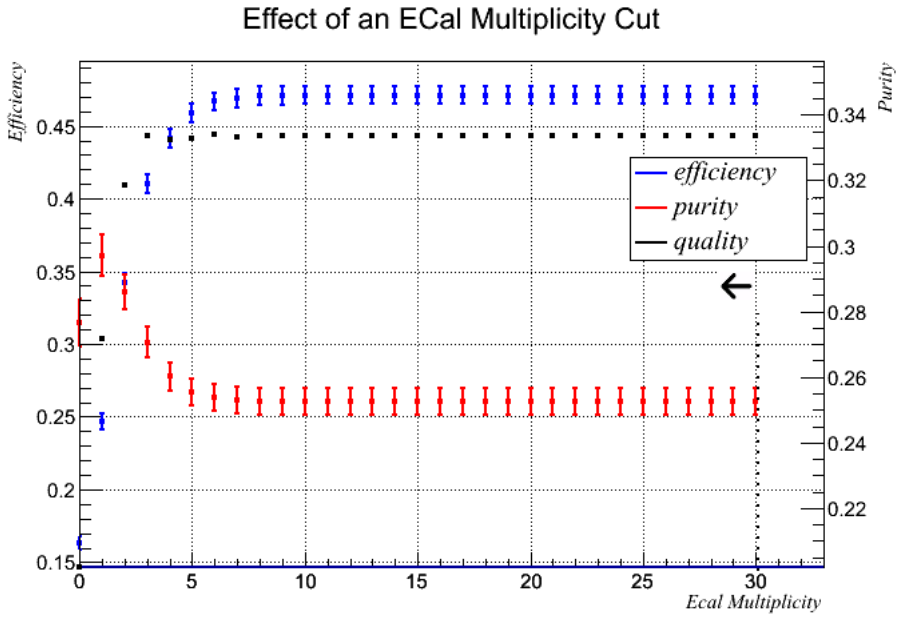

Figure 6.63: Effect of an ECal multiplicity cut on final selection, isolated electron and isolated positron topology.

gles, but the gradient for the background is shallower, suggesting value in a cut.

The effect of such a cut is shown in figure 6.65. Whilst there is clear value in having a cut, the maximum is very broad. The eventual cut value settled upon was $\cos\theta = 0.6$. This cut value was chosen to preserve statistics, as a tight cut would have significantly lowered the efficiency, by up to 5%.

Both the twin isolated electron and twin isolated positron topologies benefit from a looser cut value, with the maximum instead located at roughly 0.5, as shown for twin isolated electrons in figure 6.66.

6.4.7 Polar Angle

For isolated electrons and isolated positrons, the polar angle for signal and background using the final selection is shown in figure 6.67, with the effects

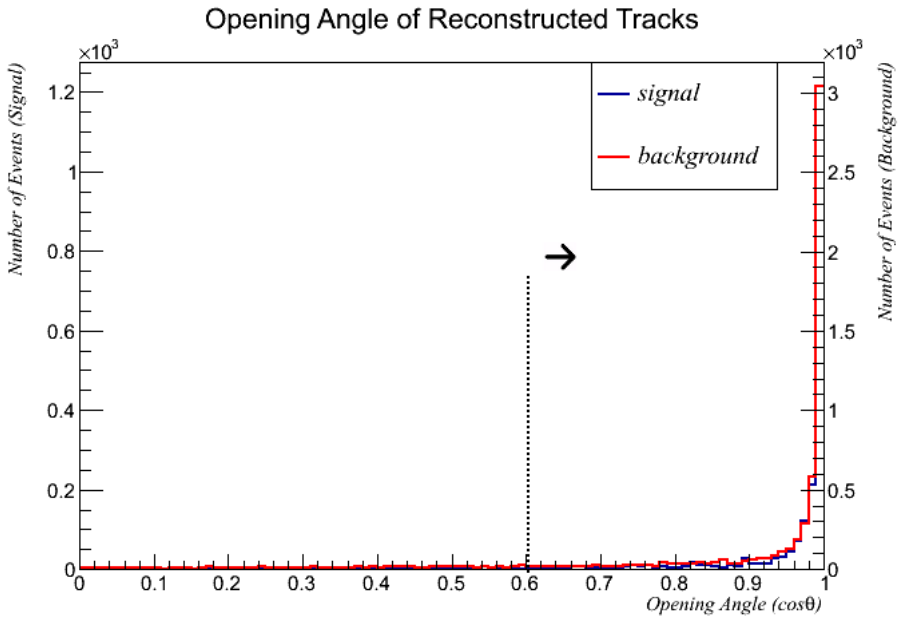

Figure 6.64: Projected opening angle between decay photons of π^0 candidates, isolated electron and isolated positron topology.

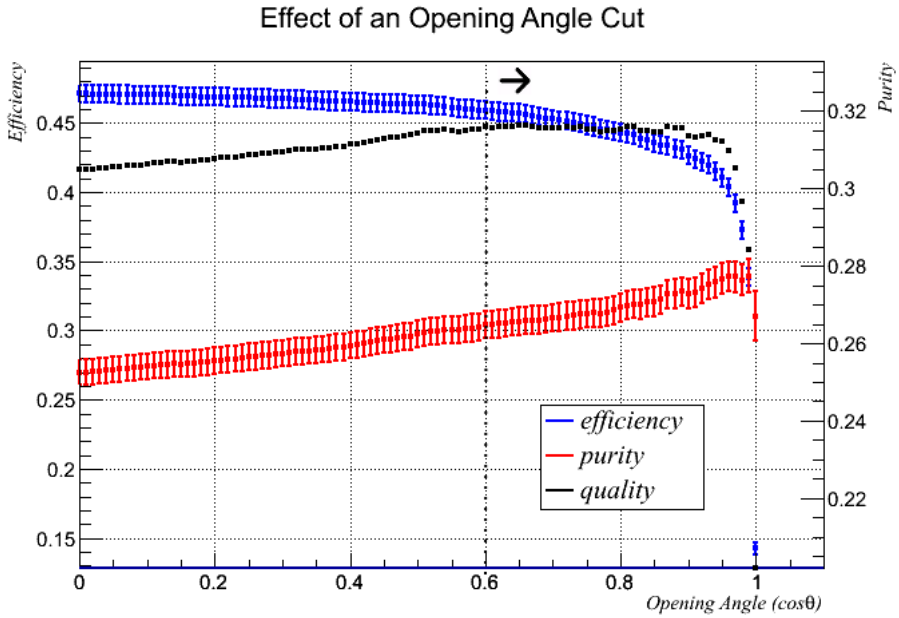

Figure 6.65: Effect of a minimum opening angle cut on final selection, isolated electron and isolated positron topology.

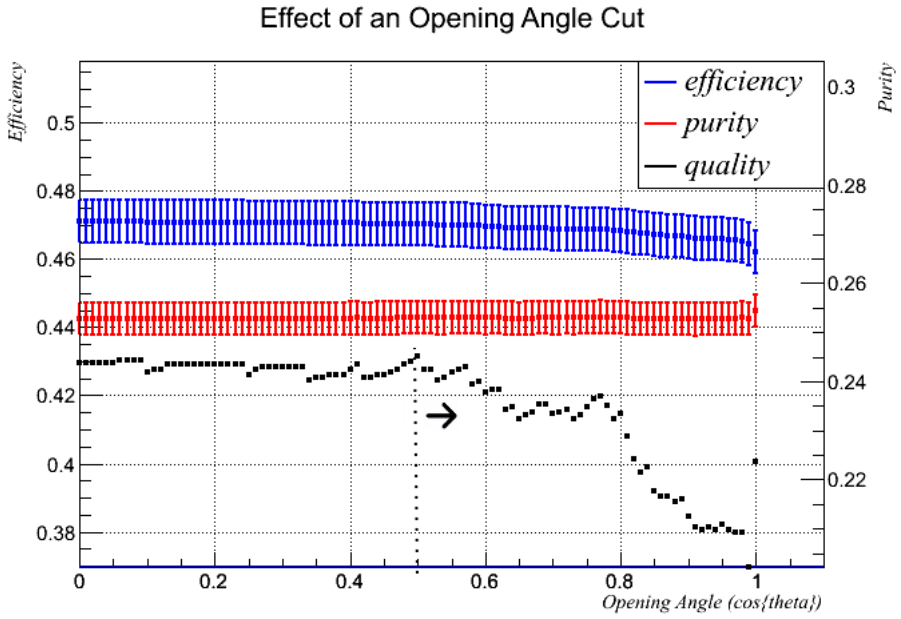

Figure 6.66: Effect of a minimum opening angle cut on final selection, twin isolated electron topology.

of a cut shown in figure 6.68.

The signal and background distributions are very similar. The polar angle cut provides little to no benefit for any of the topologies covered in this subsection. No polar angle cuts were applied to any topology. This is likely due to the fact that created reconstructed π^0 candidates from only two tracks means there is missing kinematic information. Because of this the signal events cover a wider spectrum of polar angles and therefore reduces the effectiveness of a polar angle cut.

6.4.8 Closest Approach

The distance of closest approach for photon candidates constructed from the twin isolated positron topology is shown in figure 6.69, using the final selection. The effect of a closest approach cut for this topology is shown in figure 6.70.

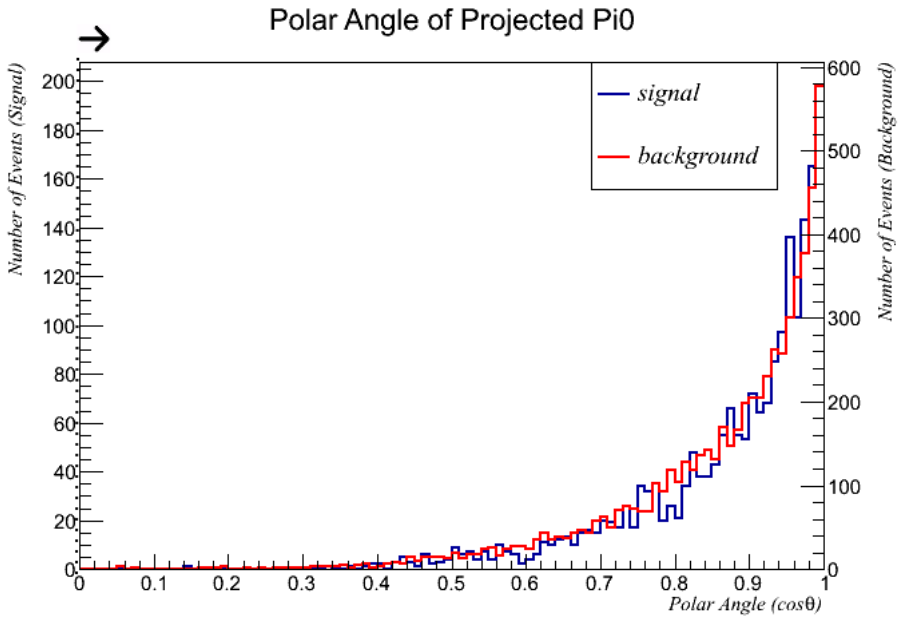

Figure 6.67: Polar angle of reconstructed π^0 candidate, isolated electron and isolated positron topology.

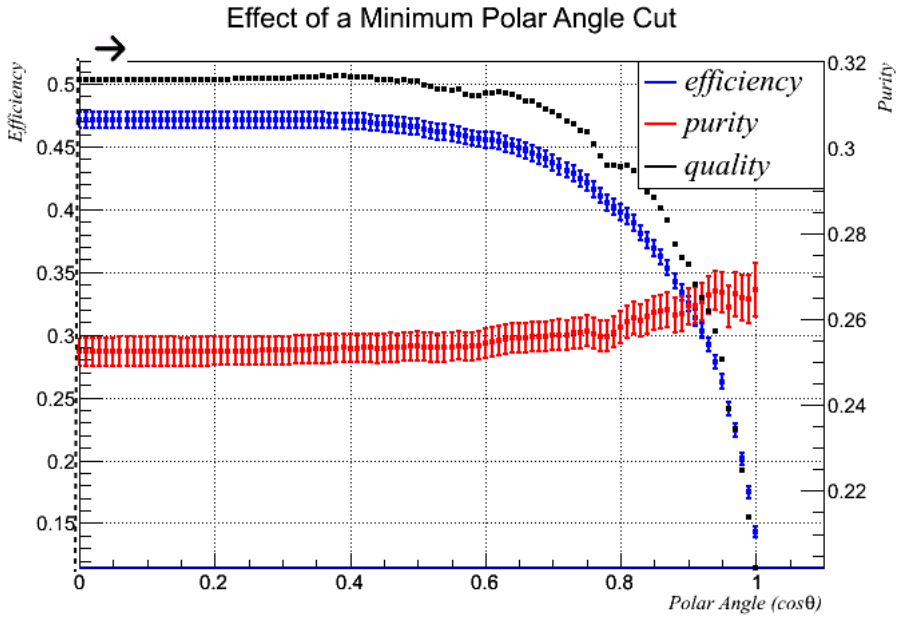

Figure 6.68: Effect of a minimum polar angle cut on final selection, isolated electron and isolated positron topology.

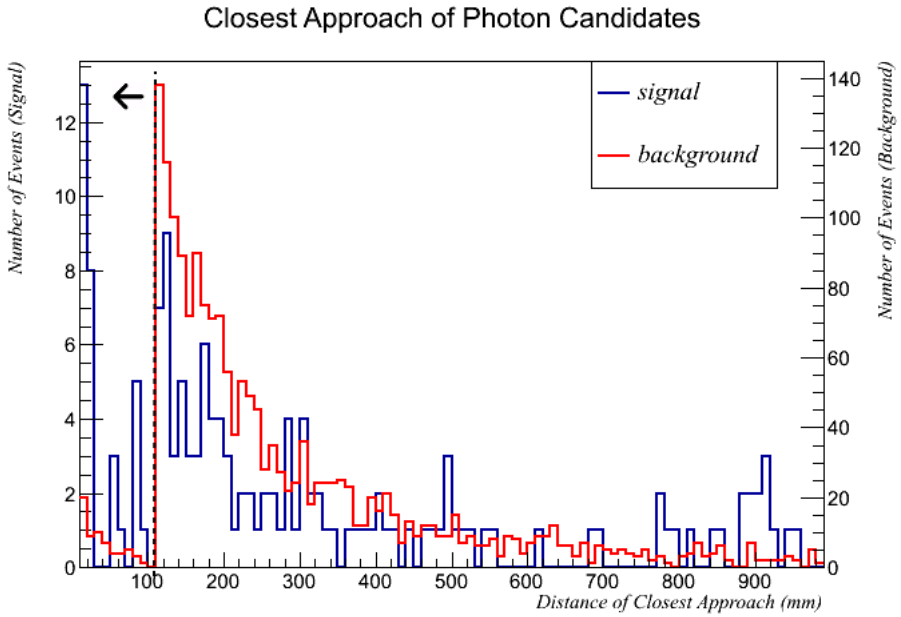

Figure 6.69: Distance of closest approach values for final selection, twin isolated positron topology.

There is a secondary peak in this distribution which is absent from previous distributions, however it is also present for the twin isolated electron and isolated electron and isolated positron positron topologies. This peak is likely due to positrons being matched together from separate FGDs. A closest approach cut of 100mm for this topology effectively vetoes this cross-FGD matching from happening, and greatly improves quality.

The effect of a closest approach cut for twin isolated electrons is shown in figure 6.71. Figure 6.72 shows the effect of the cut for isolated electrons and isolated positrons. The cut values for these two topologies are set at 130mm and 75mm respectively.

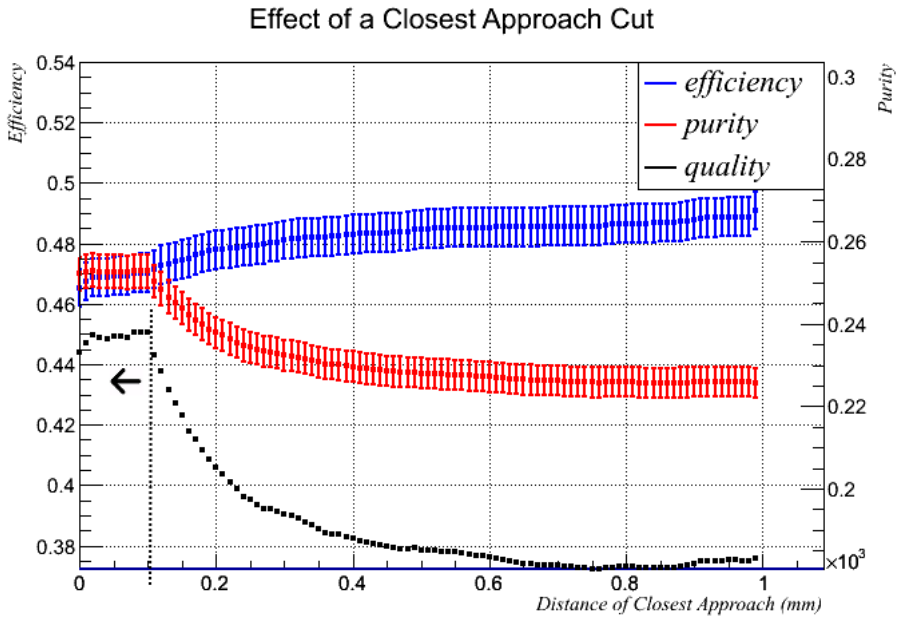

Figure 6.70: Effect of a closest approach cut on final selection, twin isolated positron topology.

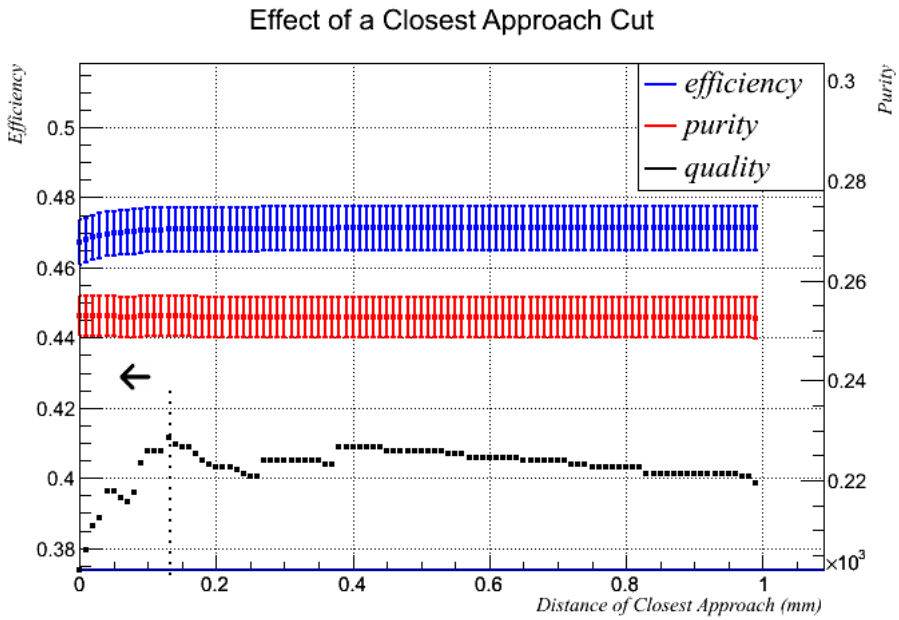

Figure 6.71: Effect of a closest approach cut on final selection, twin isolated electron topology.

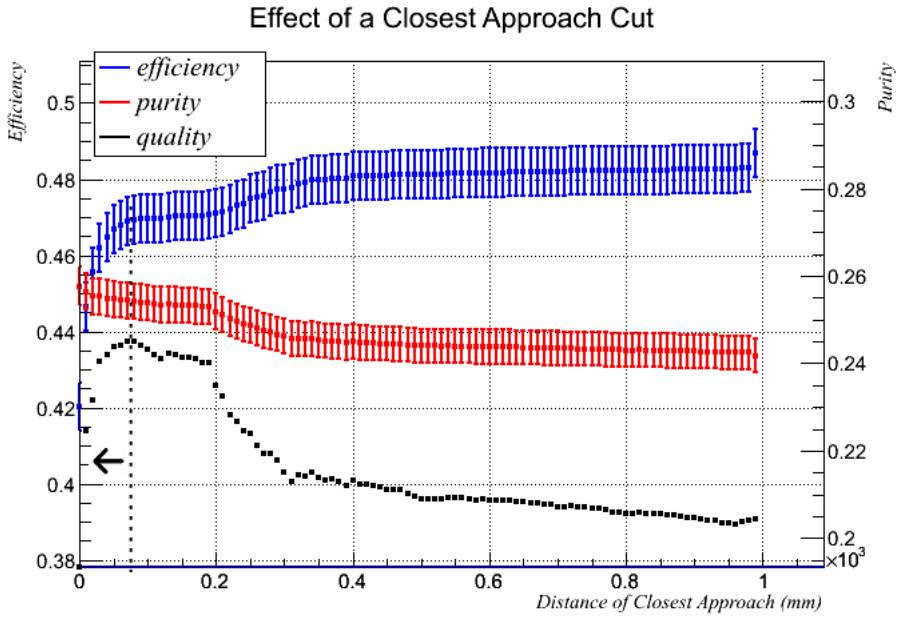

Figure 6.72: Effect of a closest approach cut on final selection, isolated electron and isolated positron topology.

6.4.9 Pair Gap

The distance between the origin points of the two tracks making up isolated electron and isolated positron candidates is shown in figure 6.73. The distributions appear similar at first, however the tail of the distribution is greater for the background than it is for the signal. This is reflected in figure 6.74, which shows the effect of a pair gap cut. There is a distinct peak in the quality by having a pair gap cut at 400mm. A cut at this value helps to eliminate situations where the isolated electron and isolated positronoriginate in different FGDs.

Unlike the isolated electron and isolated positron topology, there is no value in applying a pair gap for the twin isolated electron topology and the twin isolated positron topology. The effect of a pair gap for twin isolated positrons is shown in figure 6.75. The signal and background distributions are extremely similar, and no value can be extracted from a cut.

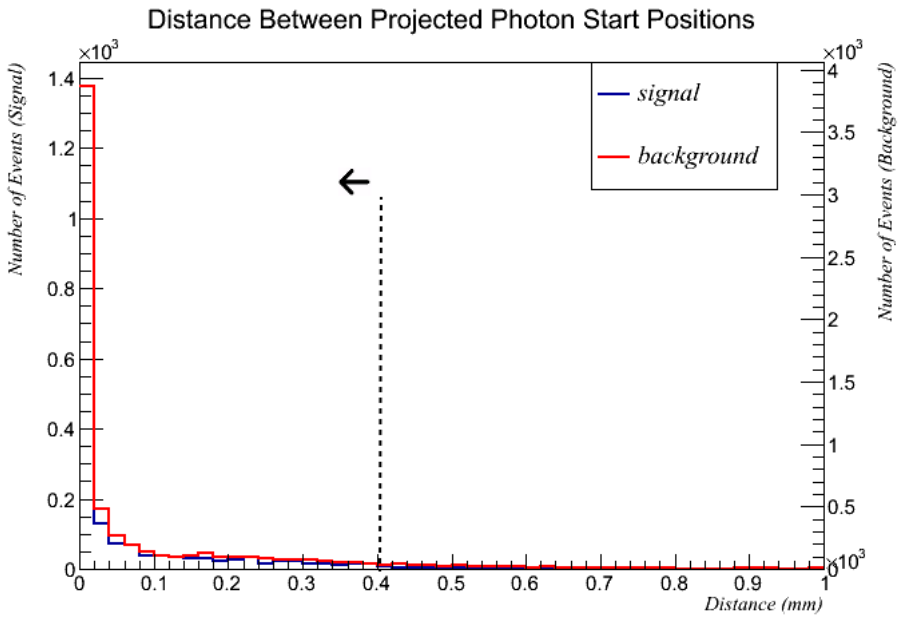

Figure 6.73: Pair gap of final selection, isolated electron and isolated positron topology.

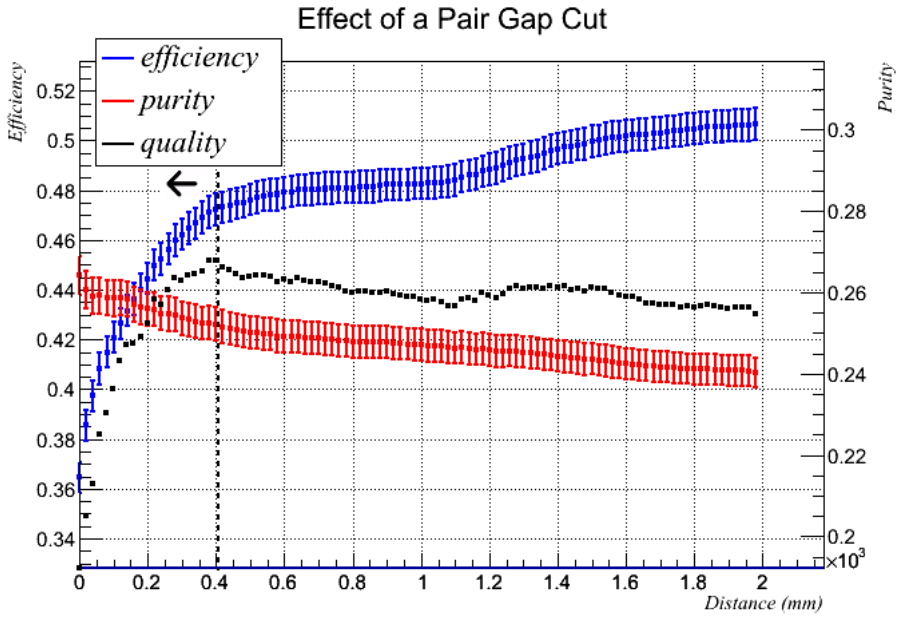

Figure 6.74: Effect of a pair gap cut, isolated electron and isolated positron topology.

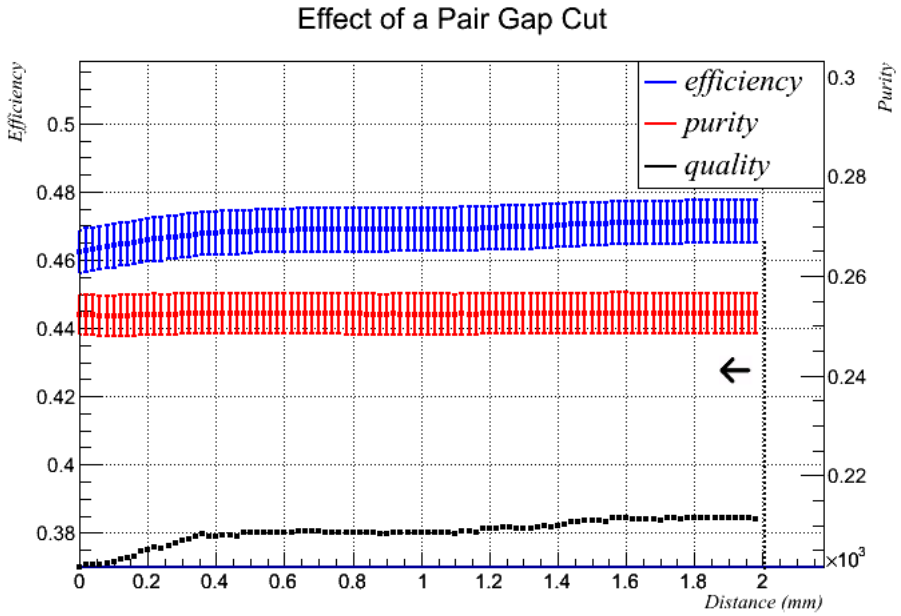

Figure 6.75: Effect of a pair gap cut, twin isolated positron topology.

6.5 Summary of Cuts

This chapter has described the development a selection cuts-based analysis to detect tracker-contained neutral-current π^0 interactions, an analysis that was developed from the ground up without relying on prior work. As a result, a wide range of selection cuts needed to be investigated and validated.

The analysis was split into two tiers, referred to as first generation and second generation, the former to reconstruct π^0 decay photons, the latter to combine those decay photons into π^0 candidates. Initially, only π^0 candidates with two reconstructable electron/positron pairs was considered. However, this approach only yielded a predicted 2 π^0s using data-level statistics.

As a result, the analysis was diversified to reconstruct decay photons where one of the leptons produced from the converting photon was not reconstructed. This approach has yielded significantly increased statistics, however the num-

ber of selection cuts being applied has increased dramatically with the addition of 5 other second generation topologies. The final results of this analysis are presented in chapter 9.

Table 6.1 summarises the second generation cuts.
It is difficult to identify any one single selection cut as having a significant effect. The overall quality of the analysis relies on a large number of selection cuts each adding incremental improvements. However the largest gains were made from applying TPC1 and P0D multiplicity vetoes, as well as constraining the momentum range of pair and isolated lepton candidates.

6.5.1 Sand Muon Analysis

One final source of background that has not yet been considered is that of sand muons. Neutrinos from the beam can interact with the rock and sand surrounding the ND280 detector pit, producing muons via a charged-current interaction. These muons can then penetrate into ND280 and be logged as data. The timing and angular direction cuts which remove cosmic muons do not apply for sand muons. The timing cut is passed because the sand muons are only created during beam spills. The angular direction cut is passed because the muons have roughly the same directionality as the beam.

Because of the muon cut in the analysis, and the fact that a single muon is unlikely to mimic an NCπ^0 event, backgrounds are expected to be small but must nevertheless be accounted for. When scaled amount of statistics available for the data, 38.8 sand muon events are predicted. This is only 5.2% of the total background.

Table 6.1: Summary of 2nd generation final selection cuts.

Cut Type	Twin Pair	Pair and Iso e^-	Pair and Iso e^+	Twin Iso e^-	Twin Iso e^+	Iso e^- + Iso e^+
Bunch Multiplicity	none	12	19	13	none	none
Tracker Multiplicity	14	7	16	none	none	none
Upstream Multiplicity	3	1	3	0	0	0
ECal Multiplicity	none	2	4	2	0	none
Isolated ECal Multiplicity	none	1	0	none	none	none
P0D Multiplicity	0	0	0	0	0	0
TPC1 Multiplicity	0	0	0	0	0	0
Minimum Opening Angle ($\cos\theta$)	0.5	none	none	0.5	0.5	0.6
Minimum Polar Angle ($\cos\theta$)	0.97	none	none	none	none	none
Closest Approach (mm)	90	none	none	130	100	70
Pair Gap (mm)	none	600	1450	none	none	400

Chapter 7

Data vs Monte Carlo Comparisons

So far in this analysis, Monte Carlo information has been used to develop the selections for π^0 candidates. An important step is to validate that the data and the Monte Carlo agree so that any problems with the detector, simulation or reconstruction problems can be recognised and diagnosed. This is particularly apparent because such a data vs Monte Carlo analysis was undertaken for a π^0 analysis involving the tracker and ECal region and significant discrepancies between ECal reconstruction and ECal simulation were observed. This was because the detector simulation software only simulated the magnetic field in the tracker region and not in the magnet flux return. This in turn caused some ECal modules to predict too many Monte Carlo events and some to predict too few Monte Carlo events [75]. Because this analysis is focused around the tracker region, that issue does not greatly effect this analysis and has since been resolved.

Extensive comparisons between data and the simulated Monte Carlo have been carried out by other individuals, either as part of software testing or as part of other analyses [79] [78] [80]. Most of these have shown good agreement and the places where disagreement has been seen have either been understood and corrected, or worked around. As the π^0 event signature studied in this analysis

is very different from most other analyses using ND280 and the selection of tracker events is new, an important further step is to compare the data with Monte Carlo after the cuts developed in this analysis have been applied. In an ideal case, the data and Monte Carlo will agree well with each other, however there are uncertainties in the Monte Carlo associated with poorly understood interaction cross-sections (such as the NCπ^0 neutrino interaction which is the focus of this thesis), so some level of disagreement might be expected.
For this chapter, the same analysis is applied to all available data sets and also to the Monte Carlo. For events where a candidate π^0 has been reconstructed, important variables are taken and plotted on histograms, for both the data run and the Monte Carlo run. When these comparisons were being made, the overall analysis was not yet complete; a normalisation difference between the data and Monte Carlo could indicate a discrepancy in π^0 production cross sections. As this is the measurement this analysis intends to make, prematurely observing the existence of any normalisation difference could introduce a bias. To avoid this, these histograms have been normalised to an area of one.

The data and Monte Carlo have been plotted with their associated statistical errors. The selection cuts used on both data and Monte Carlo samples are the selection cuts listed in secions 5.2.5 and 6.5. The data and Monte Carlo used is outlined in section 4.1.6.

7.1 Data vs MC: Track Momentum

Figure 7.1 shows the momentum of tracks which were used to reconstruct π^0 candidates (see chapter 4). Whilst the errors on the data are moderate due to limited statistics, the data and Monte Carlo distributions are highly similar. Given the uncertainties associated with low-momentum track reconstruction

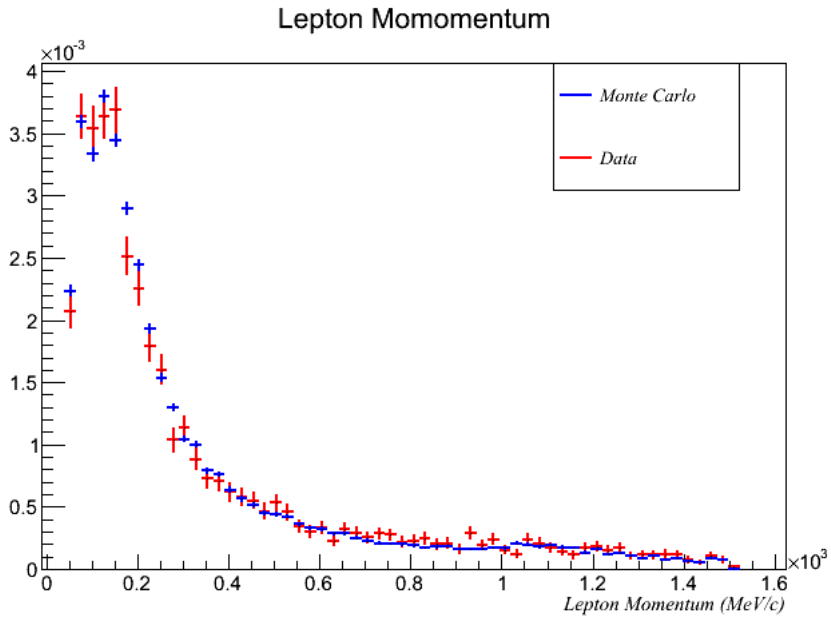

Figure 7.1: Momentum of tracks used to reconstruct candidate π^0 vertices for data and Monte Carlo.

in ND280 5.1.3, this is promising.

7.2 Data vs MC: Track Z-Position

The Z-coordinate of reconstructed tracks is an important variable to compare, because discrepancies would indicate either a misalignment between the simulated FGDs in the Monte Carlo and the actual FGDs, or problems with reconstructing tracks in the edge-layers of the FGDs. However, figure 7.2 shows a good agreement between the data and the MC. The binning on the histogram is coarse due to limited statistics; finer binning would increase the statistical errors on the individual data points such that no inferences could be drawn from the histogram. Whilst more Monte Carlo would have been benificial, generating Monte Carlo is computationally intensive. As a result, all the Monte Carlo information available was generated by a central T2K computing group.

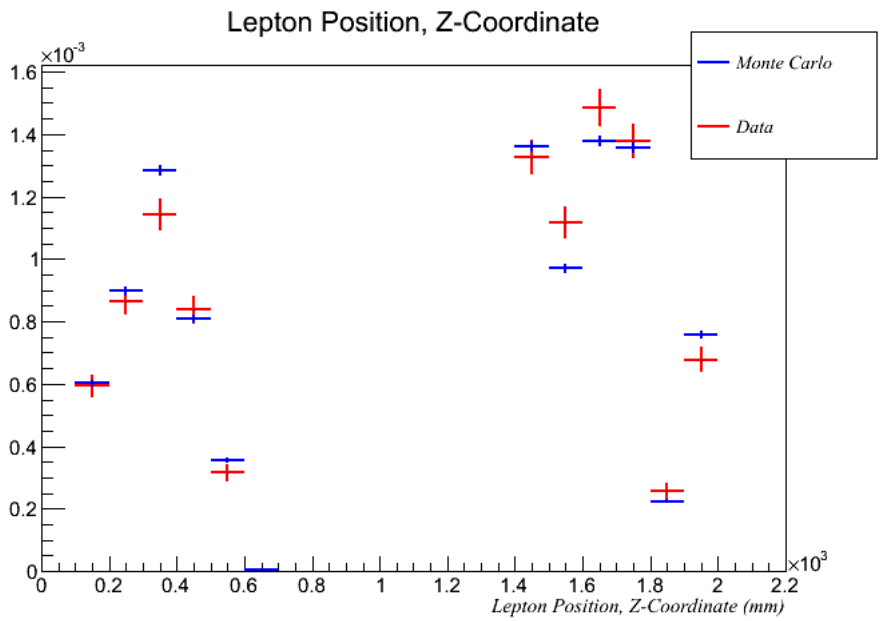

Figure 7.2: Z-position of tracks used to reconstruct candidate π^0 vertices for data and Monte Carlo.

7.3 Data vs MC: Polar Angle

The variable being plotted in figure 7.3 is the polar angle: the angle made between the reconstructed track and an imaginary line drawn along the z-axis. Particles produced from neutrino interactions in ND280 will have a Lorentz boost along the z-axis and so this variable measures the ability of the reconstruction software to accurately reconstruct the directionality of tracks undergoing a strong Lorentz boost.

Figure 7.3 shows the angle tracks make with the z-axis ($\cos\theta$) for data and MC. The histogram shows agreement over most of the range of the histogram, however at high $\cos\theta$ there are differences between the data and the Monte Carlo in two bins, which could be statistical fluctuations.

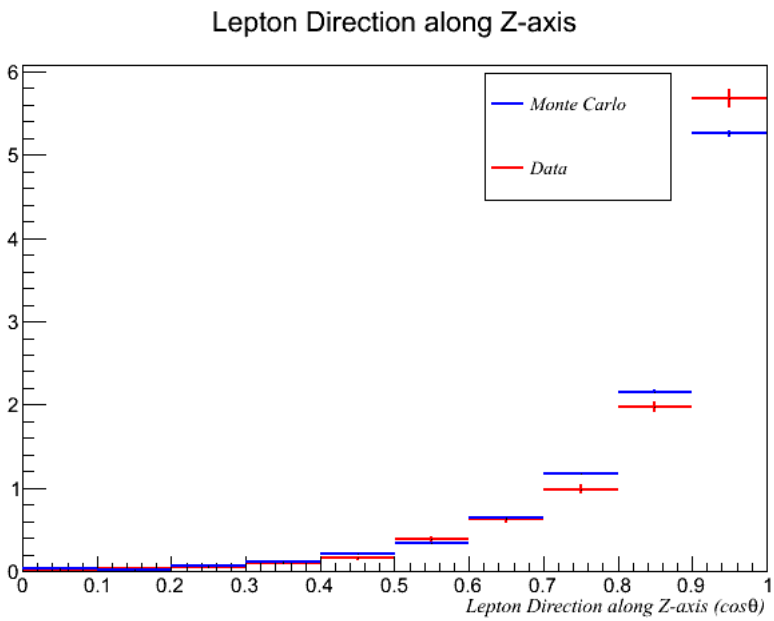

Figure 7.3: Z-direction of tracks used to reconstruct candidate π^0 vertices for data and Monte Carlo.

7.4 Data vs MC: Multiplicity

Comparing the number of reconstructed tracks for data and Monte Carlo is a check that the tracker subdetectors are reconstructing the expected number of tracks. Discrepancies could point to problems with identifying tracks which have entered the tracker region from another subdetector, or problems with modelling of the physics, or problems with reconstructing tracks with limited information (although this would also be identified in the momentum comparisons).

The multiplicity of an event refers to the number of reconstructed tracks. Subdetector multiplicities refer to the number of reconstructed tracks with a presence in the relevant subdetector. A track does not have to be isolated to a specific subdetector for it to contribute to that subdetector's multiplicity count.

Tracker multiplicity refers to the number of reconstructed tracks with a presence in the tracker region (TPCs + FGDs). Figure 7.4 shows the data vs Monte Carlo comparison for tracker multiplicity. There is a good agreement between the two sets of information.

Bunch multiplicity refers to the number of reconstructed tracks throughout the entire ND280 detector. Figure 7.5 shows the data vs Monte Carlo comparison for bunch multiplicity. There is good agreement here.

ECal multiplicity refers to the number of reconstructed tracks with a presence in the downstream and barrel ECal regions. Figure 7.6 shows the data vs Monte Carlo comparison for ECal multiplicity. In this case, there is a difference between the data and Monte Carlo for high multiplicity tracks. This difference is likely due to issues with the ECal simulation software, discussed at the start of this chapter. Nevertheless, the ECal region of ND280 has a limited scope in this analysis, only being used as a multiplicity veto and so the consequences of this issue are not severe.

7.5 Summary

The data and the Monte Carlo have been compared. The data and Monte Carlo distributions match well for important variables which are new to this analysis. There were no surprising features seen which warranted further investigation. Once data vs Monte Carlo comparisons were finished, the analysis had reached a mature stage where a rate measurement could be made. However, an evaluation of the sources of error in the analysis was required. This is discussed in chapter 8.

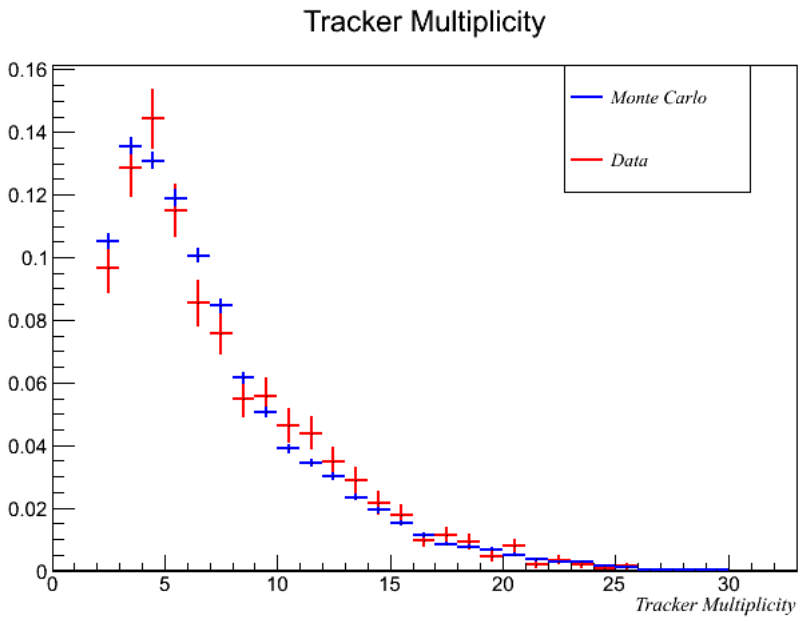

Figure 7.4: Tracker multiplicity of events containing reconstructed π^0 vertices for data and Monte Carlo.

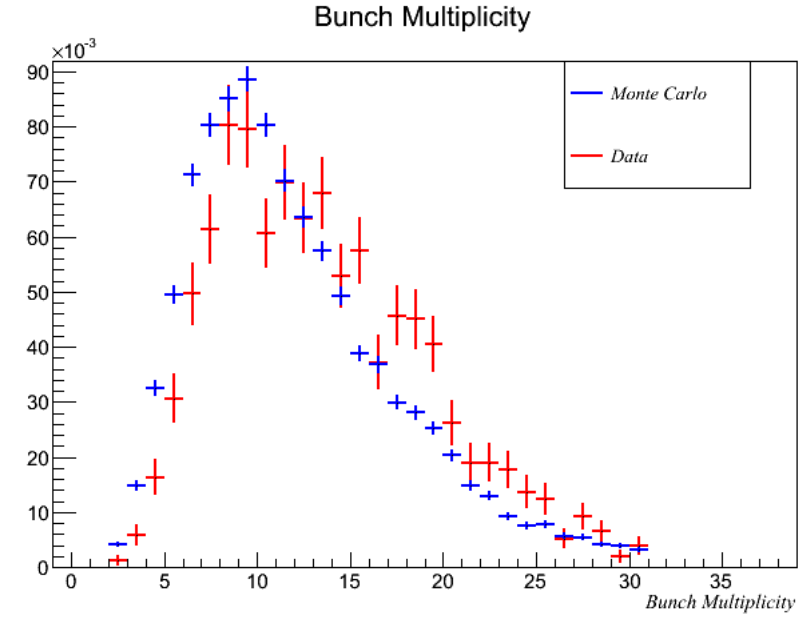

Figure 7.5: Bunch multiplicity of events containing reconstructed π^0 vertices for data and Monte Carlo.

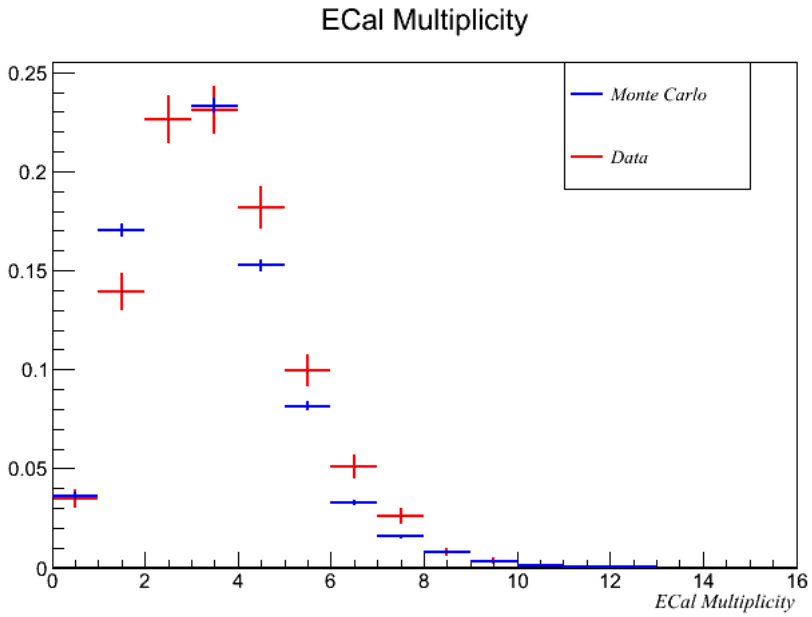

Figure 7.6: ECal multiplicity of events containing reconstructed π^0 vertices for data and Monte Carlo.

Chapter 8

Systematic Error Evaluation

8.1 Overview of Systematics

This chapter gives an overview of the three sources of systematic error in this analysis. These are:

1. The uncertainties in the cross sections of neutrino interactions that form the background of this analysis.

2. Systematic errors associated with uncertainties in the neutrino beam flux.

3. Systematics associated with uncertanties in the ND280 detector.

The cross section uncertainty is measured by parametrising the cross-sections of neutrino interactions in ND280, and then using the uncertainties on these parameters to build up an overall uncertainty on the amount of measured background. This is performed using a piece of T2K-specific software called T2KReWeight. This parametrisation was developed by an internal T2K group called the Beam and ND280 Flux and Cross-section Extrapolation Task Force (BANFF) [14].

Next, systematic errors associated with uncertainties in the neutrino beam flux

are described. The overall result of this analysis is left as a rate measurement (see chapter 9), however eventually the result will be as a cross-section measurement. This cross-section measurement will be made as a ratio to neutrino interaction cross-sections already made in ND280 (specifically, the inclusive charged-current quasi-elastic ν_μ-induced cross-section [15]). This factors the flux uncertainty measurement out of this analysis. Nevertheless, the sources of flux uncertainty and their parametrisation are discussed in this chapter.

Finally, the systematics errors associated with uncertainties in the ND280 detector are discussed. Development of the systematic errors associated with the tracker region is not yet mature for π^0 interactions. At this stage, only a qualitative overview of the errors involved has been made. Previous detector systematic studies from other ND280 analyses indicate this error will be small relative to the flux and cross section errors [15]
Table 8.6 in section 8.6 shows each systematic error associated with this analysis, the uncertainty associated with the paramter and its source.

8.2 Cross-Section Systematic Errors

This chapter describes the procedure used to calculate the systematic error on the background associated with cross-section uncertainties in the Monte Carlo event generator, NEUT. This is primarily done with a piece of T2K-specific software called T2KReweight. A description of the T2KReweight software and the cross-section parameters used to describe neutrino interactions at ND280 follows.

8.2.1 Signal and Background Samples

Since the goal of this analysis is to measure the cross-section for neutral-current π^0 producing neutrino interactions, no cross-section systematic uncertainties have been applied to the signal selection for this analysis, as this is the parameter which is trying to be measured. The systematic cross-section uncertainty applies exlusively to the background of this analysis.

8.3 T2KReweight

The cross-sections of the spectrum of neutrino interactions that can occur in ND280 are described by several parameters, as will be described in 8.3.1. Each of these parameters has an associated uncertainty, constrained by external data. Using the uncertainties on each parameter, it is possible to build up a total expected uncertainty on the measured number of background events in this analysis. The alternative to this approach would be to take the Monte Carlo information this analysis uses and to recreate the Monte Carlo multiple times, each time modifying the cross-sections of every neutrino interaction type observed in this analysis to reflect their known uncertainties. This approach would require an unfeasibly large amount of computational resources.

The T2KReweight software framework was developed by the T2K collaboration [81] and is designed to help calculate systematic errors associated with the uncertainties associated with neutrino interaction cross-section parameters.

The neutrino interactions in ND280 broadly fall under four categories: Charged-current quasi-elastic interactions, neutral-current interactions, deep inelastic scattering, and resonances produced through charged-current and neutral-current

interactions. A detailed breakdown of the cross-section parameters is in the following section.

When the NEUT event generator generates Monte Carlo [82], our best estimte of neutrino cross-section parameters are used to determine the number of neutrino interactions produced. The range of uncertainty on these parameters introduces a systematic error.

T2KReweight employs the concept of 'throws' to evaluate systematic errors: each event has an associated uncertainty based around uncertainty in the cross-section of the neutrino interaction that generated the event. This uncertainty is modeled as a Gaussian. T2KReweight picks a random uncertainty for each of the parameters on the Gaussian spectrum and applies it to the cross-section of the interaction type which created the event, generating a weight, w. For example, a weight of 1.1 means an uncertainty of 10% has been applied, and a weight of 1.0 means no uncertainty has been applied. The weight is calculated as:

$$w = \frac{\sigma_{\text{updated}}}{\sigma_{\text{nominal}}} \tag{8.1}$$

This process is then repeated 1000 times, where each repeat is called a 'throw'. After 1000 throws the distribution of uncertainties for the event gives an estimate of the systematic uncertainty caused by the variations of the parameters. This resembles a normal distribution. The mean of the distribution will be the average weight, and the standard deviation is the uncertainty associated with the event.

All of the background events returned by the analysis are treated this way. After running the output of the analysis over a single throw, an average weight of

the events accepted by the analysis can be calculated by summing the weights applied to each event. After 1000 throws, a distribution of these weights can be created. This should resemble a normal distribution, whose standard deviation is equal to the cross-section uncertainty associated with the entire analysis. This is then used as the cross-section systematic error for the rate measurement. The systematic comes in the cross-section uncertainty on the background measurement.

8.3.1 Cross-Section Parameters

As previously mentioned in section 8.1, parameters used to describe the cross sections of neutrino interactions in ND280 have been developed by an internal T2K subgroup called the Beam and ND280 Flux and Cross-section Extrapolation Task Force (BANFF) [14]. The parametrisation was initially developed for ND280's inclusive charged-current quasi-elastic neutrino interaction cross-section measurement [15]. Each parameter is discussed in the following section. Aside from external data constraints, the cross-section parameters are further constrained by internal data taken by the ND280 experiment. Further descriptions can be found in internal T2K technical note 108 [14].

8.3.2 Charged-Current Quasi-Elastic Parameters

There are four parameters associated with charged-current quasi-elastic interactions: CCQE1, CCQE2, CCQE3 and MaQE. The first three modify the charged-current quasi-elastic cross section spectrum in three separate energy regions. E1 corresponds from 0-1.5 GeV, E2 from 1.5 GeV to 3 GeV and E3 for energies $>$3 GeV. This parametrisation is split into three different energy regions because the cross-section at lower energies (0-1.5 GeV) is much better understood [83]. This lower energy region (0-1.5 GeV) is also much more relevant to interaction

Parameter	E_ν Range	Nominal	Error
M_A^{QE}	all	1.21 GeV/c^2	0.45
CCQE E1	$0 < E_\nu < 1.5$	1.0	0.11
CCQE E2	$1.5 < E_\nu < 3.5$	1.0	0.30
CCQE E3	$E_\nu > 3.5$	1.0	0.30
p_F ^{12}C	all	217 MeV/c	30
E_B ^{12}C *	all	25 MeV	9
SF ^{12}C	all	0 (off)	1 (on)
M_A^{RES}	all	1.16 GeV/c^2	0.11
CC1π E1	$0 < E_\nu < 2.5$	1.63	0.43
NC1π^0	all	1.19	0.43
CC1π E2	$E_\nu > 2.5$	1.0	0.40
CC Coherent	all	1.0	1.0
CC Oth shp	all	0.0	0.40
NC Other	all	1.0	0.30
W Shape	all	87.7 MeV/c^2	45.3

Table 8.1: Cross-section parameters for the BANFF fit. Table shows the relevant energy range for each parameter, their nominal value and associated uncertainty. Table taken from [14].

in ND280 because the T2K beam energy peaks at 600 MeV. The MaQE parameter modifies the overall shape of the CCQE cross-section distribution.

The uncertainties on the CCQE cross-section are largely taken from MiniBooNE data published in 2001 [83]. The uncertainty on the CCQE parameters is calculated by comparing published MiniBooNE data with a Monte Carlo set generated by NEUT and applying a χ^2 fit.

The cross-section uncertainty from the MiniBooNE result falls within the energy range denoted by CCQEE1. The other two energy regimes have little available data available to define an error with. Both parameters are derived from extrapolated MiniBooNE data, and set at 30%. The T2K neutrino beam averages 700 MeV, and the proportion of neutrino interactions above 1.5 GeV is

small, hence these large errors from CCQEE2 and CCQEE3 do not greatly affect this analysis.

The uncertainty on MaQE is also acquired through analysing MiniBooNE data [83].

8.3.3 Single Pion Production

Single pion production cross-section errors are calculated from three different MiniBooNE data sets: CC1π^0 [84], CC1π^+ [85] and NC1π^0 [16].
From these three data sets, nine cross-section parameters are fitted: MaRes, W shape, CC coh, CC1pi, CCother, NCcoh, NC1pi0, NC1piplus, NCother.
The CC-other and NC-other parameters represent charged-current and neutral-current interactions not covered by quasi-elastic or single-pion interactions. They are described in section 8.3.4.
MaRes is a parameter which modifies the overall shape of the cross-section distribution for resonant interactions. W-shape is a data-driven term which is used to modify the shape of NEUT's NC1π^0 cross-section distribution to better represent available data.
Coherent pion production occurs when the interacting neutrino interacts with the nucleus as an unfragmented single entity. The nucleus undergoes a slight recoil, but little energy transfer occurs [86]. At higher energies, a pion can be created. Both charged-current and neutral-current coherent pion production can occur, and are parametrised by the two parameters 'NC-coh' and 'CC-coh' respectively. These parameters act as a scaling function on the overall cross-section.

The CC1π^+, CC1π^0 and NC1π^0 terms scale the cross-sections for charged-current single π^+ production, charged-current single π^- production and neutral-current single π^0 production.

These nine parameters are fitted using a joint fit to the three MiniBooNE data sets, which is possible because the three forms of single pion production are connected by common parameters in NEUT. The fit is performed by using a χ^2 minimisation of the nine parameters.

The fit provides very little insight into the errors associated with NC1π^0 interactions, CC-other and NC-other. For NC1π^0 interactions, this lack of data is the precise reason for this analysis.

For charged-current coherent pion production, data from MiniBooNE suggests there is no coherent pion production from neutrino interactions with energies of order 1GeV. The uncertainty of CC-coherent pion production is therefore set at 100%.

For neutral-current coherent pion production, comparing SciBooNE data with NEUT models suggests an error of 15%.

The fitted values of all nine parameters can be found in table 8.1.

8.3.4 Other CC Interactions

CCOther describes charged-current interactions not described by CCQE or single-pion production. This parameter combines three interaction types: deep inelastic scattering (DIS), charged-current multiple-pion production and charged-current resonant $\eta/K/\gamma$ production. Limited data is available for the energy range of the T2K beam, however figure 8.1 shows the cross-section of these interaction types (plotted in red) as a function of neutrino energy. At the energy range of the T2K beam, the amount of CCOther interactions is minimal, and so a lighter treatment of the associated errors does not undermine the validity of

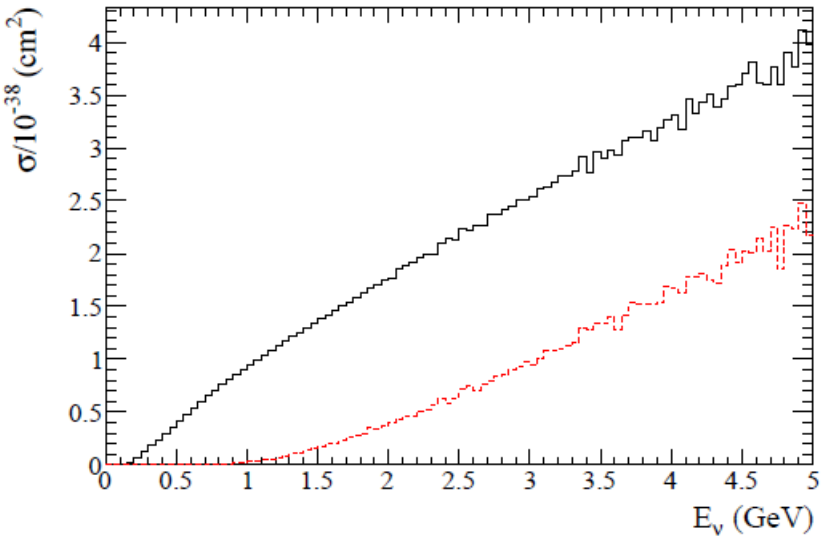

Figure 8.1: Cross section per nucleon of CC inclusive (black) and CC multi-π + DIS (red) as a function of neutrino energy. At the T2K neutrino beam's peak energy (0.6 MeV), the contribution from the latter is negligible.

the analysis.

The uncertainty of CCOther, using external data from MINOS [87], is 10% at 4 GeV. The error for a given energy is then defined using equation 8.2.

$$\sigma_{\text{CCother}} = \frac{0.4}{E_\nu} \tag{8.2}$$

8.3.5 p_F and E_b: Fermi Gas Model parameters

For charged-current quasi-elastic interactions, the scattering nucleus is modelled by NEUT as a Fermi gas [82]. This model is important because it would be insufficient to assume that the nucleon the neutrino interacts with is at rest in the lab frame. The Fermi gas model assumes the nucleus is constructed out of non-interacting fermions in a potential well. The uncertainties for two pa-

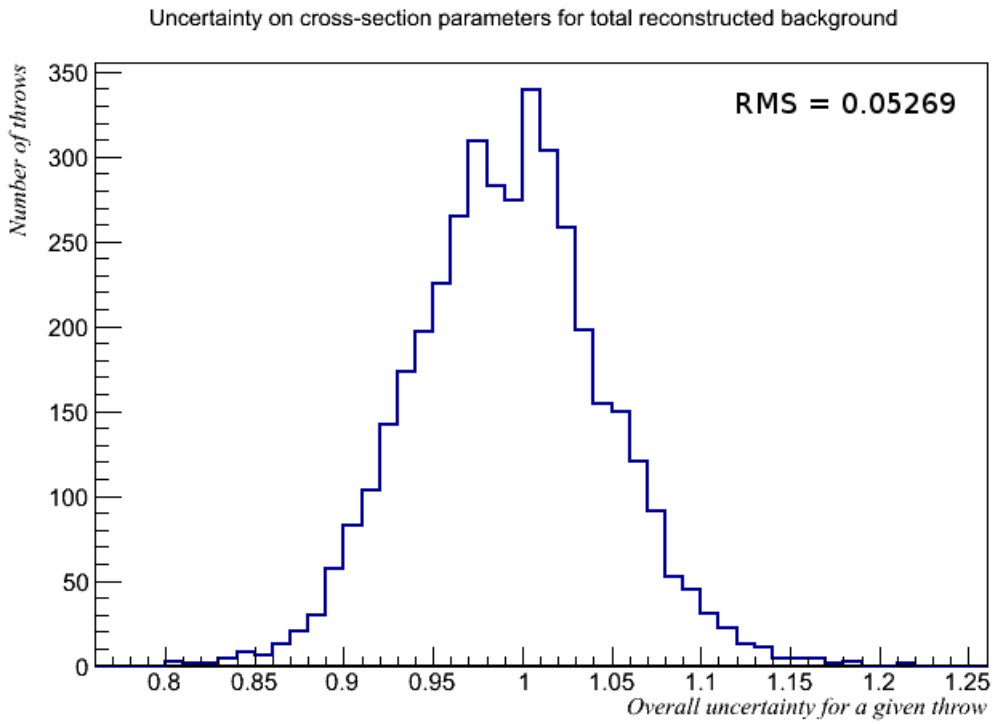

Figure 8.2: Results of reweighting cross-section systematics through T2KReWeight. Each entry represents the output of the entire analysis, where each reconstructed event has been reweighted using distributions drawn from the uncertainties given in table 8.1. This is repeated 4000 times to produce a Gaussian distribution whose RMS is the 1-σ cross-section uncertainty.

rameters used to construct a Fermi gas model: p_F and E_b, are calculated using electron scattering data [88]. p_F is associated with the average momentum of the Fermi gas, and E_b represents the nuclear binding energy.

8.3.6 Cross-Section Systematic Error: Result

When the output of this analysis was run over the T2KReweight 'throw' procedure described earlier, an estimated cross-section systematic of 5.3% was obtained. This is shown in figure 8.2. The histogram shows the overall weight of all reconstructed background for each individual 'throw'; the RMS of these throw weights being the standard error.

However, given the table of errors in table 8.1, this number is suspiciously low. Roughly 70% of the background of this analysis comes from either charged-current pion production and charged-current deep inelastic scattering (described later, see table 9.3). The cross-section parameters in table 8.1 were developed for T2K's ν_μ oscillation analysis [15], an analysis where charged-current pion production and deep inelastic scattering did not significantly feature as a background. Applying T2KReweight in the context of neutral-current interactions has not been done before, and further analysis of the cross-section parameters may be required to deliver accurate results, as well as an improved integration of the T2KReweight software package into the π^0 analysis.

As an interim measure to provide a cross-section error, the built-in NEUT uncertainties on cross-section systematic parameters are used with T2KReweight to get a rough estimate of the overall cross-section uncertainty. This approach does not involve the parametrisation given by the BANFF group. This resulting uncertainty is 17.6%. Although the BANFF group's parametrisation is not used, the approach uses the T2KReWeight software and methodology outlined in section 8.3. This result is likely to be an overestimate due to the fact that the cross-section systematics are not independent, and the BANFF group have considerably reduced the cross-section systematic errors by using external and internal data constraints. More work needs to be performed by the π^0 research group to integrate research from the BANFF group into π^0 analyses, the 17.6% error value is used as the cross-section systematic uncertainty for the results of this analysis.

8.4 Beam Flux Systematics

When making a cross-section measurement, the result can be given as a ratio to the ν_μ inclusive charged-current quasi-elastic interaction cross section. This has the advantage of removing the flux measurement from the analysis, as well as providing a better estimate of the beam flux error. What follows is an overview of the components of the beam flux systematic error. The flux systematic parametrisation was developed by the BANFF group, and further information about the beam flux uncertainties can be found in [6].
The beam flux uncertainty systematic affects both the NCπ^0 rate measurement and the corresponding cross-section interaction measurement, as both measurements are a function of the total number of neutrinos entering the ND280 detector.

8.4.1 Hadron Interaction Uncertainties

Hadron interaction uncertainties introduce systematic errors from several sources. There are three major sources or uncertainty. Experimental uncertainties in cross-section rates in interactions of protons with the target, scaling data from one momentum region to momenta regions where data is missing and extrapolating particle yields to different target materials. This subsection discusses the interaction uncertainties of several relevant interactions which are involved in producing the T2K neutrino beam: pion production, kaon production and secondary nucleon production.

Hadron production measurements from dedicated accelerator experiments were used to develop the hadronic interaction uncertainty systematic error. These are the NA61/SHINE experiment [89], the BNL-E910 experiment [90], Eichten *et al.*[91] and Allaby *et al.*[92]. These experiments fire proton beams onto a car-

bon target and measure the momentum and angles of the particles produced. NA61/SHINE is a dedicated experiment, run partly by T2K members, which exposes parts of our target to beams with the same energy as the T2K beam.

The overall flux uncertainty introduced from hadronic interactions is shown in figure 8.6

Pion Production Uncertainties

The three sources of pion production uncertainty are:

1. Uncertainty associated with NA61/SHINE data used to re-weight pion production uncertainty.
2. Scaling down available NA61/SHINE data to interactions with lower momentum incident nucleons.
3. NA61/SHINE data does not cover the T2K beam's entire phase-space.

Data on pion production is taken from the NA61/SHINE collaboration, the error on the measurement is dominated by systematic errors, described in detail in their own paper [93]. Figure 8.3 shows the total errors with respect to pion momentum and pion opening angle. The total errors for phase-space most relevant for T2K's neutrino beam is between 5% and 10% [14].

The error associated with downscaling the NA61/SHINE pion multiplicity data to lower incident nucleon momenta is studied by studying the BNL-E910 data and comparing it with NA61/SHINE. The BNL-E910 experiment provides proton-on-beryllium data at 12.3 GeV and 17.5 GeV respectively. The parametrisation used to fit the two data sets was developed by Bonesini *et al.*[94].

The NA61/SHINE data covers the majority of the phase space for the T2K beam's secondary pions. To cover the remaining phase space, the NA61/SHINE data is fitted to the parametrisation developed by Bonesini, extrapolating data into the uncovered region, which is very small. Figure 8.4 shows the parametrisation succeeds in fitting to the NA61/SHINE data. To improve the overall fit, π^+ and π^- data is fitted separately.

Kaon Production Uncertainties

The treatment of uncertainties associated with kaon production are largely similar to those of pion production, however with less input data until NA61/SHINE data became available. Data from NA61/SHINE [93], Eichten [91] and Allaby *et al.* [92] are used to develop the kaon production systematics. The overall uncertainty associated with kaon production are documented in table 8.5.

The uncertainties in table 8.5 are split into three categories. The first is the uncertainty for the overall normalisation, σ_N. The second is uncertainties associated for a given kaon momenta and kaon opening angle, $\sigma_{\Delta p \Delta \theta}$. These uncertainties are dominated by statistical errors. The final category is the uncertainties in normalisation for a given angular bin, $\sigma_{\Delta\theta}$.

Secondary Nucleon Production Uncertainties

Interactions of secondary protons and neutrons inside the beam target contribute to 16% and 5% of the overall neutrino flux respectively [6]. This uncertainty is measured by comparing the simulated beam model with data from proton production measurements from Eichten *et al.*[91] and Allaby *et al.*[92]. This uncertainty is shown in figure 8.6.

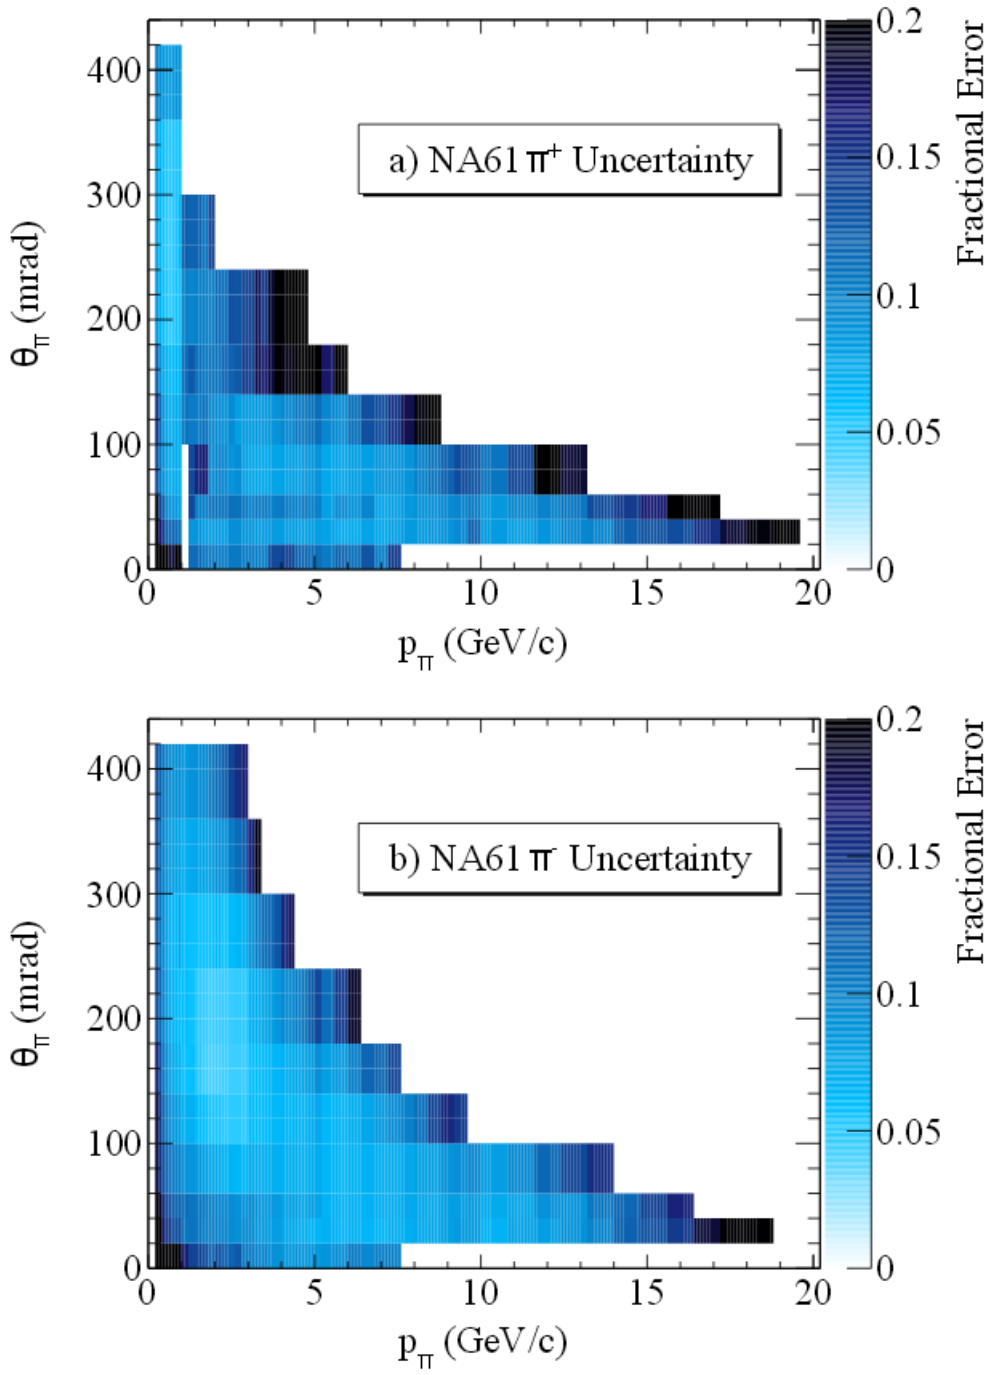

Figure 8.3: Fractional error on NA61/SHINE measurements for momentum/opening-angle bins. Figure from [6]

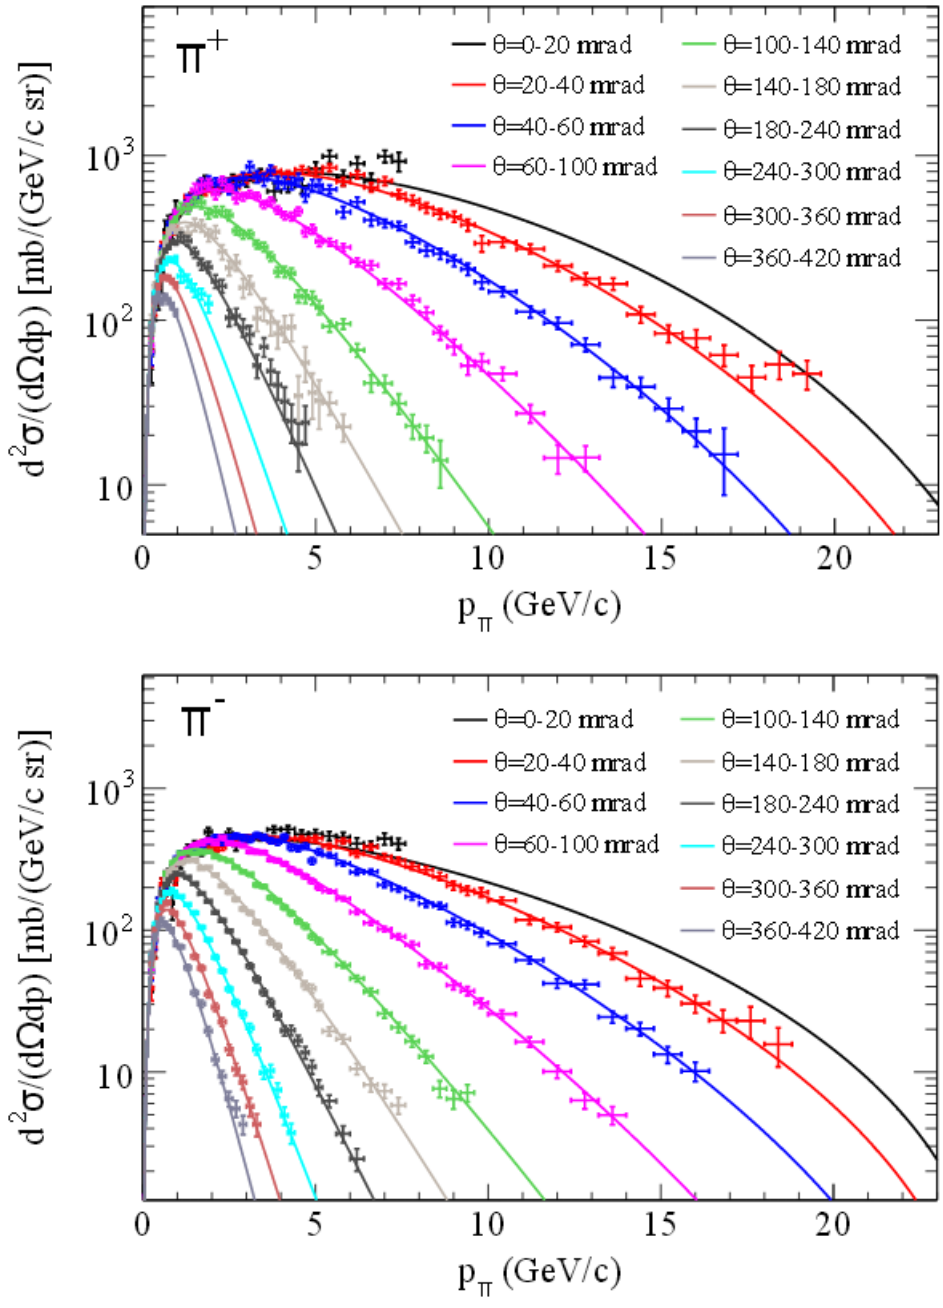

Figure 8.4: Fit developed by Bonesini, applied to NA61/SHINE production data. [6].

	σ_N	$\sigma_{\Delta p \Delta \theta}$	$\sigma_{\Delta \theta}$
NA61/SHINE	2.3%	$11-24\%$	–
Eichten *et al.*	15%	4%	5%
Allaby *et al.*	10%	$2-5\%$	10%

Figure 8.5: Summary of uncertainties related to kaon production. The uncertainty for the overall normalisation is σ_N. [6]

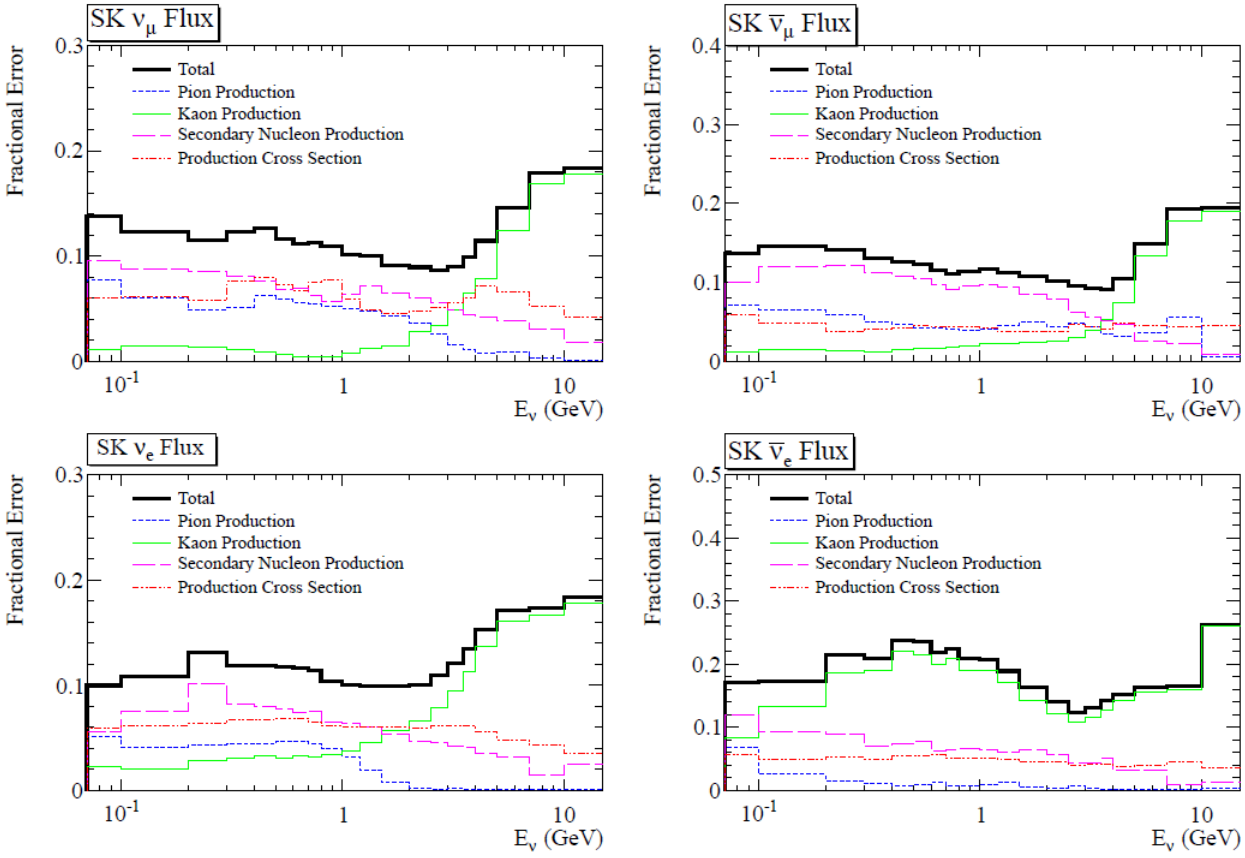

Figure 8.6: Overall neutrino flux uncertainties from hadronic interactions. From the top-left going clockwise: muon neutrino flux, anti-muon neutrino flux, anti-electron neutrino flux, electron-neutrino flux [6]

8.4.2 Proton Beam And Off-Axis Angle Uncertainties

There are three sources of systematic uncertainty relating to the proton beam:

1. The alignment uncertainty of the proton beam monitors.
2. The alignment uncertainty between the primary proton beamline and the target station.
3. Systematic errors in the position measurements by the proton beam monitors.

The proton beam is monitored by five current transformers, twenty one electrostatic monitors and nineteen segmented secondary emission monitors. The overall systematic uncertainty associated with the proton beam is approximately 2%.

The uncertainty in the location of ND280 relative to the neutrino beam is 0.44 mrad [6]. The effect of the off-axis angle uncertainty on the flux uncertainty is determined by moving the ND280 detector by 0.44 mrad in the JNUBEAM simulation (JNUBEAM is a software program that simulates the T2K neutrino beam [6]). However, off-axis angle uncertainties only become significant relative to other systematic errors for measurements made at Super-Kamiokande. At ND280, any off-axis angle uncertainty is negligible due to the small distance between the detector and beam source.

8.4.3 Target and Horn Alignment Uncertainties

The T2K neutrino beam is created by colliding protons onto a graphite target to produce pions. These pions then decay to produce the neutrino beam.

The horns are three magnets which are used to focus pions, that decay to create neutrinos, forming the main component of the neutrino beam. Flux uncertainties are introduced from uncertainties in the positioning of the horns, and positioning of the proton interaction target.

The uncertainties in target positioning were calculated using the JNUBEAM simulation, by rotating the target. Doing this with the actual neutrino beam would be impractical.

The beam target was rotated by 1.3 mrad in the horizontal plane and 0.1 mrad in the vertical plane, which introduces a flux uncertainty on the order of a few percent. 1.3 mrad and 0.1 mrad are the relative horizontal/vertical uncertainties in the alignment between the beam target and the horn [6]. Figure 8.7 gives a complete description of the target alignment uncertainty.

The horn alignment uncertainties come in two forms; positional and angular. For positional uncertainty, movements along all three axes were studied, however only uncertainties along the y-axis introduced noticeable neutrino flux uncertainties.

For angular uncertainties, the effects of an 0.2 mrad rotation of the horns in the horizontal and vertical planes was studied. The effects of the horn target and horn alignment uncertainties on the ν_μ flux are shown in Fig 8.8.

For the horn angular alignment uncertainties, the relevant measurement is the uncertainty in the relative alignment between the primary (proton) and secondary (pion) beamlines and so the angular uncertainties between the horns are treated as fully correlated.

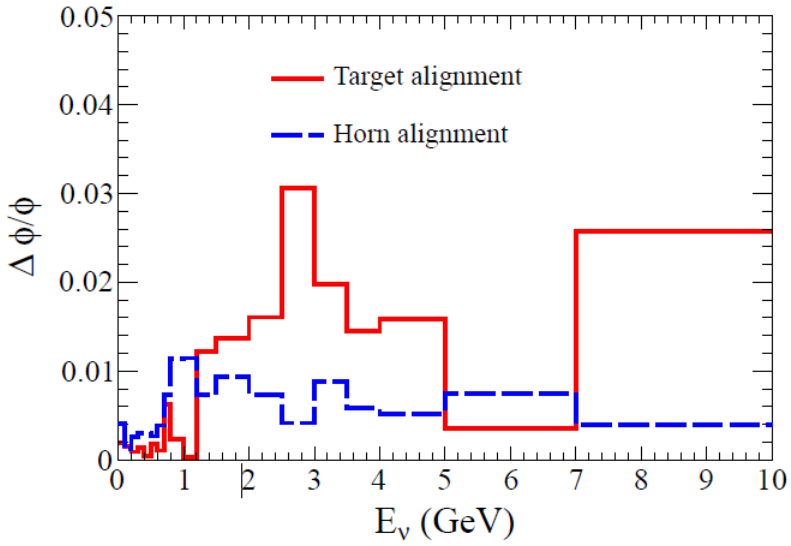

Figure 8.7: Fractional uncertainties due to the target and horn alignment in the ν_μ flux [6].

8.4.4 Horn Current and Magnetic Field Uncertainties

There is a 1.3% uncertainty regarding the current used to generate the horn's magnetic field; the magnetic field itself has a measured 2% uncertainty, contributeing to an overall 2% uncertainty in the neutrino flux [6].

8.4.5 Summary of Flux Uncertainties

The previous sections have described the sources which contribute towards the uncertainty in the T2K neutrino beam flux. Ultimately the results of this π^0 analysis will be in the form of a cross-section measurement, however the measurement in this thesis will be given as a ratio to the inclusive charged-current quasi-elastic ν_μ-induced interaction cross-section. The overall flux uncertainty of this cross-section is 11% [15]. This 11% uncertainty is incorporated into the overall systematic uncertainty of the analysis presented in this thesis.

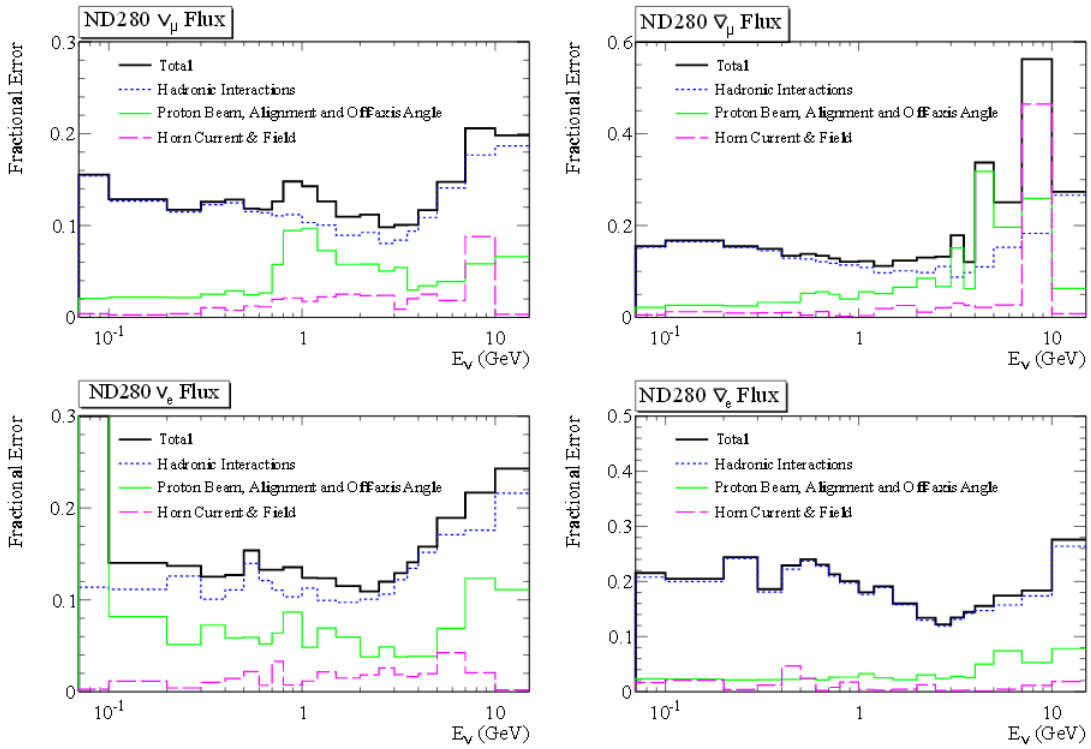

Figure 8.8: Overall beam flux uncertainties. From the top-left going clockwise: muon neutrino flux, anti-muon neutrino flux, anti-electron neutrino flux, electron-neutrino flux [6]

8.5 Detector Systematics

Currently, development of the systematic errors associated with the tracker region is not fully developed for π^0 analyses. At this stage, only a qualitative overview of the errors involved can be made.

8.5.1 TPC Errors

There are several key variables that are measured using the TPC which introduce systematic uncertainties, these are discussed in this section. Information about the TPCs systematics was taken from internal technical note 93 [78] referring to T2K's ν_μ CCQE analysis.

Track Finding Efficiency

The TPC track finding efficiency refers to the probability that a track propagating through the TPC will be identified by the reconstruction software and tagged as a track. This was studied in an internal technical note [78], and no significant difference was found between data and Monte Carlo simulations. However, further work is required to properly extend the analysis to the low momentum regions used in this thesis.

TPC Particle ID

The TPC particle ID (PID) systematic refers to uncertainties in the energy loss measurement, $\frac{dE}{dx}$, in the TPC, and its application towards an electron hypothesis and a muon hypothesis; the electron pull and muon pull values used earlier in the analysis.

The systematic errors on these pull values is studied in [79]. These are studied by comparing data with Monte Carlo for a high-purity muon sample. The average $\frac{dE}{dx}$ pull under the muon hypothesis is compared with the data and Monte Carlo for each TPC and each run period independently. The uncertainty in pull values is of the order of 1%. The same conclusions were reached for the uncertainty on the electron pull.

Momentum Uncertainty

The TPC momentum uncertainty is linked to the uncertainty in the strength of the magnetic field produced by ND280s magnet. If the magnetic field strength is misunderstood, then the momentum attributed to a given bending radius will also be incorrect. The uncertainties in the magnetic field are documented in internal T2K technical note [95]. A summary of this uncertainty measure-

ment follows.

A non-magnetic device was built at CERN to map the magnetic field in ND280, which is movable in two dimensions. It was placed inside the ND280 'basket' region (which covers the tracker region and the P0D subdetector) with an accuracy of within 0.1mm, and has the capability to map the magnetic field of the entire basket region. These measurements were carried out using several different magnetic field strengths corresponding to a coil current of 0 A, 250 A, 500 A, 750 A and 1000 A.

The main source of uncertainty was the nonlinearity of the magnetic field. The combined error in magnetic field along the y-axis (i.e. the bending axis) was 0.5%.

Charge Misidentification

One of the first cuts applied to reconstructed tracks is a charge cut, to separate tracks in to electron candidates and positron candidates. Assigning an incorrect charge to tracks in the TPC affects the quality of this analysis and therefore needs to be investigated.

Internal technical note 48 [96] investigates charge misidentification, by examining charge misidentification within the Monte Carlo and using data vs Monte Carlo comparisons to extrapolate this rate into the data. For tracks between 0-1.3 GeV/c, a charge uncertainty of 0.8% was measured.

8.5.2 FGD Errors

Position Uncertainty

The FGD scintillator bars have dimensions of 9.61 mm x 9.61 mm x 1864.3 mm [7], which limits the position resolution to approximately 1 cm. This is the dominant systematic when measuring the uncertainty in FGD positions.

Momentum Uncertainty

The FGD momentum uncertainty is dominated by two effects; the magnetic field strength and limited positional resolution. The field-strength systematic has been discussed in section 8.5.1, with the 1 cm granularity of track hits adding an additional uncertainty. However, any track used in this analysis must have both a TPC component as well as an FGD component. The superior momentum reconstruction of the TPCs means that the FGD momentum reconstruction is not used. Therefore systematics associated with FGD momentum reconstruction do not feature in this analysis.

8.5.3 ECal Errors

Track/ECal matching efficiency

The most detailed consideration so far of the ECal detector systematics can be found in Dr. Sam Short's thesis [73]. Tracks which have a presence in both the tracker and ECal might not be matched correctly, introducing an uncertainty. To investigate this, a high purity control sample is selected for both data and Monte Carlo. The systematic uncertainty is linked to any differences in efficiency between the data and the Monte Carlo, combined with the statistical error on the data sample. This investigation is still ongoing.

ECal hit reconstruction efficiency

When reconstructing hits in the ECal, there is an uncertainty in the amount of detected light recorded. The ECal will not register a hit unless the charge registered by the ECal is greater than 3.5 photon equivalent units [97]. Uncertainties in the charge resolution of the ECal therefore contribute towards a systematic error.

The uncertainty in ECal component dimensions is also a factor that must be considered. The area of the scintillator bars in the x and y dimensions are 40 mm $\pm$ 0.4 mm $\times$ 10 mm $\pm$ 0.4 mm. The length of each scintillator bar varies, but the average length uncertainty is 0.1 mm. The length, width and thickness uncertainties of the lead sheets in the ECal are 3 mm, 3 mm and 0.1 mm respectively.

8.5.4 Summary of Detector Systematics

This section gives a qualitative overview of the type of detector systematic errors that contribute towards the overall systematic error. Further investigation is required to produce a quantitative result. Considering the large cross-section systematic of this analysis, 17.6%, and an expected flux systematic of approximately 11% [15], detector systematics would have to be significant in order to contribute greatly towards the overall systematic error. Previous tracker-based studies have shown this not to be the case, with detector systematics varying between 0.6% and 7% but mostly focused on the lower end of this range [15]. Further investigation is required to fully develop accurate detector systematics for a neutral-current π^0 analysis.

8.6 Summary of Systematic Errors

Table 8.6 shows each of the individual systematic errors associated with this analysis and the source of the value. The overall neutrino flux uncertainty is 10.9%, determined when developing the charged-current quasi-elastic neutrino cross-section [15]. The cross-section systematic uncertainty requires further development, particularly with regards to developing the cross-section parameters associated with sources of background not found when measuring charged-current quasi elastic neutrino interactions but are present when measuring neutral-current π^0-producing interactions (such as coherent pion production). The current cross-section systematic uncertainty is 17.6% but this can be reduced with further development. The detector systematics require further work in order to accurately calculate, however previous tracker-based analyses have carefully developed their own detector systematics and found them to be small relative to the flux and cross-section systematic uncertainties [15].

Parameter	Uncertainty on Parameter	Source
Cross Section Systematics		
CCQE shape Uncertainty (8.3.2)	45%	Calculated by ND280 [14]
CCQE $(0 < E_\nu < 1.5)$ (8.3.2)	11%	Calculated by ND280 [14]
CCQE $E_\nu > 1.5$ (8.3.2)	30%	Calculated by ND280 [14]
Fermi Gas factor ${\mathrm{p}_F}^{12}C$ (8.3.5)	30%	Calculated by ND280 [14]
Fermi momentum factor ${E_B}^{12}C$ (8.3.5)	9%	Calculated by ND280 [14]
CC coherent shape uncertainty (8.3.3)	11%	Calculated by ND280 [14]
$CC1\pi$ $E_\nu < 2.5$ (8.3.3)	43%	Calculated by ND280 [14]
$CC1\pi$ $E_\nu > 2.5$ (8.3.3)	40%	Calculated by ND280 [14]
CC coherent (8.3.4)	100%	Calculated by ND280 [14]
CC other shape (8.3.4)	40%	Calculated by ND280 [14]
NC other (8.3.3)	30%	Calculated by ND280 [14]
W shape (8.3.3)	45.3%	Calculated by ND280 [14]
Total Cross Section Uncertainty (8.3.6)	17.6%	From NEUT uncertainties
Flux Systematics		
Secondary nucleon production (8.4.1)	6.9%	Calculated by ND280 [15]
Pion multiplicity (8.4.1)	5.0%	Calculated by ND280 [15]
Kaon multiplicity (8.4.1)	0.8%	Calculated by ND280 [15]
Off-axis angle (8.4.2)	1.6%	Calculated by ND280 [15]
Proton beam (8.4.2)	1.1%	Calculated by ND280 [15]
Horn absolute current (8.4.4)	0.9%	Calculated by ND280 [15]
Horn angular alignment (8.4.3)	0.5%	Calculated by ND280 [15]
Horn field asymmetry (8.4.4)	0.3%	Calculated by ND280 [15]
Target alignment (8.4.3)	0.2%	Calculated by ND280 [15]
Total Flux Uncertainty (8.4.5)	10.9%	Calculated by ND280 [15]
Detector Systematics		
TPC Tracking Efficiency (8.5.1)		
TPC Particle ID (8.5.1)		
TPC Momentum Uncertainty (8.5.1)	Not Included	
TPC Charge Misidentification (8.5.1)	(see section 8.5.4)	
FGD Position Uncertainty (8.5.2)		
FGD Momentum Unertainty (8.5.2)		
Track : ECal Matching Efficiency (8.5.3)		
ECal Hit Reconstruction Efficiency (8.5.3)		

Table 8.2: Summary of systematic error sources for this analysis. This table lists the sources, the uncertainty related to each source and whether the error has been calculated yet. The detector systematic errors are to be calculated by future analyses. This analysis was a first-generation study into the feasibility of tracker-based π^0 analyses and developing comprehensive detector systematics was beyond the scope of this analysis. From the previously published tracker-based CCQE analysis [15] the detector systematics are likely to be small. All the systematics in this table are relevant for both the NCπ^0 rate measurement and corresponding cross-section interaction measurement.

Chapter 9

Rate Measurement and Results

9.1 Results of Monte Carlo Analysis

After running the analysis algorithms on the Monte Carlo, using the final set of selection cuts, the number of reconstructed signal and background events for each topology type is summarised in table 9.1. The 'scaled' numbers of signal and background are entries where the number of events has been scaled to reflect the amount of T2K data available. The signal for this analysis is defined as a neutral-current π^0-producing neutrino interaction occurring within either FGD1 or FGD2. π^0s that are produced through other neutrino interaction types (for example, charged-current coherent pion production) are not considered signal. The background is defined as any neutrino interaction that does not meet the signal definition. The purity of each topology, defined in section 4.1.3, is also shown. The efficiency of each topology is not shown, as the distinctions between topology types are defined by limitations in the software; trying to categorise true events as of a particular topology is both difficult and arbitrary. What matters is the *overall* efficiency of the analysis. This overall efficiency is 47.1%. The denominator of this efficiency value is defined as every event which matches the description of signal. i.e. an NCπ^0 neutrino interaction occurring

Topology	# Signal	# Background	# Scaled Signal	# Scaled Background	Purity (%)
Twin Pair	280	844	23.5	68.6	25.5
Pair + Iso e^-	81	207	6.8	16.82	28.8
Pair + Iso e^+	353	771	29.64	62.66	32.1
Twin Iso e^-	60	258	5.1	21.0	19.4
Twin Iso e^+	85	247	7.1	20.1	26.2
Iso e^- + Iso e^+	2063	6323	173.2	513.8	25.2
Sand Muons	0	456	0	38.8	0
Total	2921	9104	245.4	741.7	24.9

Table 9.1: Results of the analysis on all available Monte Carlo. Monte Carlo generated through NEUT event generator. For more information see section 4.1.6

within an FGD where both photons convert in the tracker region, and at least two of the photon's daughter tracks having momenta of over 50 MeV/c.
Another factor to consider is event migration between topologies. A reconstructed event can only be classed as one topology type. When an event can be classed into multiple topologies, the topology with the smallest opening angle is chosen. This is somewhat arbitrary, and little attention was paid to discriminating between the topologies. What matters is the overall rate measurement; whether an event is considered twin pair or twin isolated electron does not matter in this context.

To summarise, neutral-current π^0-producing interactions suffer from the twin problems of limited statistics and high backgrounds. With an overall efficiency of 47.1% and an overall purity of 24.9%, this analysis has attempted to maximise both, with a slight emphasis on purity (the efficiency $\times$ purity^2 is the quantity which is being maximised). The final purity is still rather low, considering the large systematic errors involved from background subtraction, however increasing it significantly would have a seriously adverse effect on statistics. Given a predicted 245 reconstructed signal events, compromising further on statistics would have a serious impact on the total statistical error associated

with the analysis. With the current amount of available statistics, the author feels that the right balance was struck between removing background and retaining the signal.

Table 9.3 shows the breakdown of the reconstructed background from Monte Carlo samples in terms of the interaction type at the vertex of the reconstructed event. A small proportion of the background, 6%, is from π^0-producing neutral-current interactions. These aren't classed as signal, because the interaction vertex was outside of the specific tracker region that is the focus of this analysis. When all the NCπ^0 analyses are combined, this type of background can be subtracted out of the sample. The majority of background is from charged-current and neutral-current deep-inelastic-scattering (DIS). Combined, the two sources of deep-inelastic-scattering produce over 30% of the total reconstructed background. As DIS can produce π^0s, this can be a difficult source of background to isolate and remove. Curiously, 0.9% of the background came from neutral-current elastic scatting. This is likely due to kinetic energy transfer to the nucleons being sufficient to generate high energy nucleon interactions.

Table 9.2 shows the breakdown of the reconstructed signal in terms of the interaction type from Monte Carlo samples. 3% of the total came from what was reported as π^+ and π^--producing neutral current interactions. This is likely due to the fact that the charged pions underwent charge-conversion within the interacting nucleus, and so their true final interaction state is actually π^0. Just over half the signal is from single π^0 neutral-current interactions, with the remainder split between multiple-pion production and neutral-current deep inelastic scattering.

Table 9.2: Breakdown of signal sources for reconstructed events

Interaction Type	Percentage of total background
Neutral-current π^0	57.0%
Neutral-current multiple-π	17.4%
Neutral-current deep-inelastic-scattering	22.6%
Neutral-current π^+ and Neutral-current π^-	3.0%

Table 9.3: Breakdown of background sources for reconstructed events

Interaction Type	Percentage of total background
Charged-current quasi-elastic scattering	8.3%
Charged-current π^+	10.0%
Charged-current π^- (from $\bar{\nu}_\mu$	0.9%
Charged-current π^0	9.3%
Charged-current multiple-π	17.6%
Charged-current Δ resonance	5.2%
Charged-current deep-inelastic-scattering	26.3%
Neutral-current π^0	6.0%
Neutral-current π^-	1.2%
Neutral-current π^+	1.2%
Neutral-current multiple-π	2.9%
Neutral-current Δ resonance	3.4%
Neutral-current deep-inelastic-scattering	6.8%
Neutral-current elastic scattering	0.9%

9.2 Data Analysis: Rate Measurement

The analysis algorithms were run over data taken by the ND280 detector during the period between 2010-2013 to select the candidate events. This data-taking period was split into three runs, mentioned in section 4.1.6.

The number of reconstructed candidate π^0 events per run are shown in table 9.4, alongside the event rate predicted by the Monte Carlo, and the number of protons-on-target in each run (referred to as PoT).

The total number of predicted Monte Carlo events is 987.1. This is split into 245.4 signal events, and 741.7 background events. This gives a predicted pu-

Run #	MC Prediction	Data Results	# PoT
Run 2	231.9	263	0.78377×10^{20}
Run 3	399.0	396	1.34821×10^{20}
Run 4	356.4	377	1.20428×10^{20}
Total	987.1	1036	3.33577×10^{20}

Table 9.4: Results of the analysis on data and Monte Carlo.

rity of 24.9%. Using this background, the 1036 reconstructed data events can be split into an estimated 294.4 signal events and 741.7 background events.

Table 9.4 shows that the data results and the Monte Carlo results are in relatively good agreement; certainly within the margins of error involved.

9.3 Statistical Error

After running the analysis on data, 1036 candidate events were reconstructed. The statistical variation on this result can, approximately, be modeled as a normal distribution with a mean of 1036 and a standard deviation of 32.2. Given the predicted 294.4 signal events, this puts the relative statistical uncertainty at 32.2 events out of 294.4, or 10.9%.

9.4 Final Result

The cross-section systematic of this analysis is 130.5 events, or 17.6% of the predicted background (see section 8.3.6. The flux uncertainty systematic, using the 11% flux error from [15], is 114.0 events).

This gives an overall systematic of 176.2 events. Combining the rate measurement, statistical error and cross-section systematic error, the final rate measurement of neutral-current π^0-producing neutrino interactions in FGD1 and FGD2

of ND280 is $294.4 \pm 32.2_{\text{stat}} \pm 173.3_{\text{syst}}$. This final rate measurement was made by subtracting the predicted background from the measured rate. This is in contrast to using the predicted purity from the Monte Carlo analysis and applying it to the data. Doing so would yield 258 signal events and 778 background events. However, this approach is more model dependent than the background subtraction method, as it implies that the signal is well understood by the Monte Carlo. This is not the case, with π^0 cross section measurements being poorly understood.

The rate measurement, when made as a ratio of the predicted Monte Carlo, yields a ratio of $\frac{R_{\text{Data}}}{R_{\text{MC}}} = 1.20 \pm 0.13_{\text{stat}} \pm 0.72_{\text{syst}}$.

Chapter 10

Conclusions and Further Notes

This thesis represents work performed to measure the neutrino interaction cross-section for neutral-current π^0-producing ν_μ interactions. This analysis is the first to explore this using solely the tracker region to detect π^0s, and is unique in both its original goals and implementation.

Of all the π^0 analyses performed at ND280 which have attempted to use the tracker region, this analysis is the most complete. Initially designed as a 'proof-of-concept' investigation to see if tracker-based π^0 reconstruction was possible, original and innovative solutions were required to obtain a statically-relevant result. The analysis began as a conventional analysis attempting to reconstruct electron/positron pairs produced from converting photons produced by a π^0 decay. This resulted in a predicted number of reconstructed neutral-current π^0 events from data of 2 events. The concept of 'isolated' electron/positron reconstruction and trying to reconstruct π^0s using limited, incomplete kinematic information was created and implemented by the author of this thesis. The resulting analysis is capable of reconstructing a π^0 using only two reconstructed tracks with a reconstruction purity of up to 28%.

The final analysis was found to reconstruct neutral-current π^0-producing neutrino interactions with a high efficiency. When measuring efficiency, the denominator is defined as the number of events which are A: neutral-current π^0 producing neutrino interactions, B: originate in the tracker with decay photons which also convert in the tracker, C: both photons produce at least one lepton with a momentum higher than 50MeV/c. This definition was used because, in an ideal case, the entirety of the denominator would be reconstructed and considered signal. With this definition, the overall reconstruction efficiency is 36%. The overall reconstruction purity was 22%. Whilst low, this limited purity was a neccesary comprimise to acquire a statistically relevant result. The initial prediction of 2 reconstructed signal events has increased to 187 reconstructed signal events, in the data set available to T2K.

Due to problems developing a cross-section systematic error, and time-constrains when developing a detector systematic error, the overall error analysis is not yet complete. Without complete errors it would be disingenuous to convert the rate measurement supplied in chapter 9 into a cross-section measurement. As a result, the output of this analysis is being left as a rate measurement, for the time being.

When run on data, the analysis reconstructs 863 events, against a predicted 820 events from Monte-Carlo samples. With a Monte Carlo prediction of 632.6 background events, this results in a measurement of 230.4 signal events. There is a statistical error of $\pm$29.4 events, and a systematic error of $\pm$179.0 events. The cross-section component of the systematic error is likely an overestimation, which can be potentially be significantly reduced after further work. Using the BANFF-developed systematic errors along with the covariance matrix provided by their research group, the cross-section systematic uncertainty could

be reduced to ± 46 events, however further research is needed to validate this.

Measurement of neutral-current π^0-producing neutrino interactions has always been a difficult measurement to make, particularly at low energies. This is demonstrated by the extremely limited amount of published cross-section results. There are still challenging aspects to the analysis of T2K data before the rate measurement made in this thesis can be converted into a cross-section measurement. Nevertheless, an end-to-end analysis has been completed and presented. π^0 production is an important background to electron neutrino oscillation appearance, and will continue to be so for as long as water-Chereknov detectors remain in use within neutrino physics.

Bibliography

[1] The ALEPH Collaboration, The DELPHI Collaboration, The OPAL Collaboration, The SLD Collaboration, and *et al*. Precision electroweak measurements on the Z resonance. *Physics Reports*, 427:254–454, 2006.

[2] J. Hosaka and *et al*. Solar neutrino measurements in Super–KamiokandeI. *Physical Review D*, D:112001, 2006.

[3] The SNO Collaboration. The Sudbury Neutrino Observatory. *Nuclear Instruments and Methods in Physics Research*, A449:172–207, 2000.

[4] S. Hatakeyama and *et al*. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81:1652, 1998.

[5] Abe. S and *et al*. High sensitivity search for $\bar{\nu}_e$'s from the sun and other sources at KamLAND. *Phys. R.ev Lett.*, 92 (7):071301–071305, 2004.

[6] K. Abe and *et al*. T2K neutrino flux prediction. *Phys. Rev. D*, 87:012001, 2013.

[7] The T2K Collaboration. The T2K Experiment. *J. Phys: Conf. Ser.*, 304:012018, 2011.

[8] D. Beavis and *et al*. Long Baseline Neutrino Oscillation Experiment at the AGS. Physics Design Report. *Proposal E889, BNL*, 52459, 1995.

[9] Yurii G Kudenko. Study of neutrino oscillations in long–baseline accelerator experiments. *Physics Uspekhi*, 54:549, 2011.

[10] J. Beringer and *et al*. Particle Data Group. *Phys. Rev. D*, 86:010001, 2012.

[11]

[12] M. Derrick and *et al*. Study of single–pion production by weak neutral currents in low–energy νd interactions. *Phys. Rev. D*, 23:569, 1981.

[13] P. de Perio and *et al*. NEUT Systematic Studies for 2010a Analysis. Technical report, T2K Collaboration, 2010.

[14] Patrick de Perio and *et al*. Cross section parameters for the 2012a oscillation analysis. Technical report, T2K Collaboration, 2012.

[15] K. Abe and *et al*. Measurement of the inclusive ν_μ charged current cross section on carbon in the near detector of the T2K experiment. *Phys. Rev. D*, 87:092003, 2013.

[16] A. A. Aguilar Arevalo and *el al*. Measurement of ν_μ and $\bar{\nu}_\mu$ induced neutral current single π^0 production cross sections on mineral oil at $e_\nu \sim \mathcal{O}$ (1gev). *Phys. Rev. D.*, 81:013005, 2010.

[17] W. Lee and *et al*. Single–Pion Production in Neutrino and Antineutrino Reactions. *Phys. Rev. Lett.*, 38:202, 1977.

[18] P. Nienaber. *PhD thesis*. PhD thesis, University of Illinois at Urbana–Champaign, 1988.

[19] W. Krenz and *et al*. Experimental study of exclusive one–pion production in all neutrino–induced neutral current channels. *Nuclear Physics B*, 135:45–46, 1978.

[20] J. L. Raaf and *et al*. Neutral Current π^0 Interactions at MiniBooNE. *ArXiv pre-print. arXiv:hep-ex/048015v1*, 2004.

[21] K. Abe et al. Indication of Electron Neutrino Appearance from an Accelerator–produced Off–axis Muon Neutrino Beam. *Phys. Rev. Lett.*, 107:032002, 2011.

[22] K. Abe and *et al*. Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam. *Phys. Rev. Lett*, 112:061802, 2014.

[23] F.P.An and *et al* (Daya Bay Collaboration). Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.

[24] F.P.An and *et al* (Daya Bay Collaboration). Independent measurement of the neutrino mixing angle θ_{13} via neutron capture on hydrogen at Daya Bay. *Phys. Rev. D.*, 90:071101, 2014.

[25] C. L. Cowan Jr., F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the Free Neutrino: a Confirmation. *Science*, 20, 1956.

[26] S. L. Glashow. Partial–symmetries of weak interactions. *Nucl. Physics*, 22:579–588, 1961.

[27] S. Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264, 1967.

[28] Abdus Salam. Renormalizability of Gauge Theories. *Phys. Rev.*, 127:331, 1962.

[29] The Nobel Prize in Physics 1979, 2013.

[30] Enrico Fermi. Tentativo di una teoria dei raggi β. *Ricerca Scientifica*.

[31] F. J. Hastert and *et al*. Search for elastic muon–neutrino electron scattering. *Physics Letters B*, 46:121–124, 1973.

[32] F. J. Hastert and *et al*. Question of Parity Conservation in Weak Interactions. *Phys. Rev.*, 104:254, 1956.

[33] C. S. Wu and *et al*. Experimental Test of Parity Conservation in Beta Decay. *Phys. Rev.*, 105:1413–1415, 1957.

[34] The ALEPH Collaboration, The DELPHI Collaboration, The OPAL Collaboration, The SLD Collaboration, the LEP Electroweak Working Group, and *et al*. Precision electroweak measurements and constraints on the Standard Model. *CERN–PH–EP*, 095, 2010.

[35] M. Goldhaber and A. W. Sunyar L. Grodzins. Helicity of Neutrinos. *Phys. Rev.*, 109:1015, 1958.

[36] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.

[37] Davis R. Jr. Solar Neutrinos II, Experimental. *Phys. Rev. Lett.*, 12 (11):300, 1964.

[38] Davis R. Jr. Search for Neutrinos from the Sun. *BNL–12981*, Meeting on Neutrino Physics, Moscow, 1968.

[39] The SNO Collaboration. Measurement of Charged Current Interactions Produced by ^{8}B Solar Neutrinos at the Sudbury Neutrino Observatory. *Phys. R.ev Lett.*, 87:071301, 2001.

[40] B. Pontecorvo. Mesonium and anti–mesonium. *Sov. Phys. JETP*, 6:429, 1957.

[41] E. Majorana. Teoria simmetrica dell'elettrone e del positrone. *Nuovo Cimento*, 14:171, 1937.

[42] N. Cabibbo. Unitary Symmetry and Leptonic Decays. *Physical Review Letters*, 10 (12):531–533, 1963.

[43] M. Kobayashi and T. Maskawa. CP–Violation in the Renormalizable Theory of Weak Interactions. *Progress of Theoretical Physics*, 49 (2):652–657, 1973.

[44] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28:870–880, 1962.

[45] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.

[46] Y. Fukuda and *et al*. Solar Neutrino Data Covering Solar Cycle 22. *Physical Review Letters*, 77:1683–1686, 1996.

[47] K. S. Hirata and *et al*. Observation in the Kamiokande–II detector of the neutrino burst from supernova SN1987A. *Physical Review D*, 38:448, 1988.

[48] K. S. Hirata and *et al*. Experimental study of the atmospheric neutrino flux. *Physical Review Letters*, B205:416, 1988.

[49] Abe S and *et al*. Searech for neutrino oscillations on a long base–line at the CHOOZ nuclear power station. *Eur. Phys. J*, C27:331–374, 2003.

[50] W. W. Allison *et al*. Neutrino Oscillation Effects in Soudan–2 Upward–stopping Muons. *Physical Review D*, D:052005, 2005.

[51] M. H. Ahn and *et al*. Measurement of Neutrino Oscillation by the K2K Experiment. *Physical Review D*, 74:072003, 2006.

[52] F. Ardellier and *et al*. Double Chooz, A Search for the Neutrino Mixing Angle θ_{13}. *Eur. Phys. J*, C27:331–374, 2006.

[53] Y. Abe and *et al*. Reactor $\bar{\nu}_e$ disappearance in the Double Chooz experiment. *Phys. Rev. D*, 86:052008, 2012.

[54] Y. Abe and *et al*. Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector, journal =.

[55] KamLAND Collaboration. KamLAND: a liquid scintillator anti–neutrino detector at the Kamioka site. *RCNS-98-15*, 1998.

[56] M. H. Ahn and *et al*. Long Baseline neutrino oscillation experiment from KEK to Kamioka (K2K). *High Energy Accelerator Research Organization*, 2002.

[57] A. Peterson and *et al*. The MINOS experiment. *Nuclear Physics A*, 731:C503–C504, 2003.

[58] P. Adamson and *et al*. Measurement of the neutrino mass splitting and flavor mixing by MINOS. *Physical Review Letters*, 106 (18), 2011.

[59] G. Tzanakos and *et al*. A Proposal to FNAL to run MINOS with the medium energy NuMI beam, 2011.

[60] R. B. Patterspm and *et al*. The NOvA Experiment: Status and Outlook. *arXiv*, 1209.0716, 2012.

[61] R. Acquafredda and *et al*. The OPERA experiment in the CERN to Gran Sasso neutrino beam. *Journal of Instrumentation*, 4, 2009.

[62] R. Acquafredda and *et al*. Observation of a first ν_τ candidate event in the OPERA experiment in the CNGS beam. *Phys. Lett. B*, 691:138–145, 2010.

[63] R. Acquafredda and *et al*. New results on $\nu_\mu \rightarrow \nu_\tau$ appearence with the OPERA experiment in the CNGS beam. *Journal of High Energy Physics*, 11:036.

[64] R. Acquafredda and *et al*. Measurement of the neutrino velocity with the OPERA detector in the CNGS beam. *arXiv*, 1109.4897, 2011.

[65] R. Acquafredda and *et al*. Measurement of the neutrino velocity with the OPERA detector in the CNGS beam using the 2012 dedicated data. *Journal of High Energy Physics*, 1210, 2012.

[66] Athanassopoulos and *el al*. The Liquid Scintillator Neutrino Detector and LAMPF Neutrino Source. *Nuclear Instruments and Methods A*, 388:149–172, 1997.

[67] A. Aguilar and *el al*. Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Physical Review D*, 64:112007, 2001.

[68] A. A. Aguilar Arevalo and *et al*. A Search for Electron Neutrino Appearance at the Δm^2 1 ev^2 Scale. *Physical Review Letters*, 98:231801, 2006.

[69] Corey Adams and *et al*. The Long–Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe. *arXiv*, 1307.7335, 2014.

[70] J. Altegoer and *et al*. The NOMAD experiment at the CERN SPS. *NIM A404*, pages 96–128, 1998.

[71] H. Bethe. Zur Theorie des Durchgangs schnelter Korpus–kularstrahlen durch materie. *Annalen der Physik*, 397:325–400, 1930.

[72] C. Giganti and Marco Zito. Particle Identification with the T2K TPC. Technical Report 01, T2K Collaboration, 2009.

[73] Samantha Short. *Study of Neutrino–Induced Neutral Current Neutral Pion Production in the T2K Neare Detector*. PhD thesis, Imperial College London, 2013.

[74] Pawel Guzowski. *Reconstruction of neutrino induced neutral current neutral pion events with the T2K ND280 Tracker and ECAL*. PhD thesis, Imperial College London, 2012.

[75] Personal correspondance with dr. helen o'keeffe, 2013.

[76] T2K Collaboration. FGD Detector Geometry, internal T2K page, 2011.

[77] Benno List. Why and When to Optimize Efficiency Times Purity. *ETH Zurich internal note*, 2002.

[78] C. Bojechko and *et al*. CCQE–like and CC–non–QE–like ν_μ event selections in the ND280 tracker using Run 1+2 data. Technical report, T2K Collaboration, 2012.

[79] A. Cervera Villanueva and *et al*. PID systematics of the Time Projection Chambers for the ν_μ analysis. Technical report, T2K Collaboration, 2011.

[80] D. B. Roberge. Comparison of Hit Association in the FGD Between Data and Monte Carlo. Technical report, T2K Collaboration, 2010.

[81] C. Andrepoulos and *et al*. Handling Neutrino Interaction Uncertainties using Event Reweighting. Technical report, T2K Collaboration, 2009.

[82] Y. Hayato and *et al*. NEUT. *Nucl. Phys. B – Proceedings Supplements*, 112:171–176, 2002.

[83] A. A. Aguilar Arevalo and *et al*. Measurement of muon neutrino quasielastic scattering on carbon. *Phys. Rev. Lett.*, 100:032301, 2008.

[84] A. A. Aguilar Arevalo and *et al*. Measurement of ν_μ and $\bar{\nu}_\mu$ induced charged current neutral pion production cross sections on mineral oil at $e_\nu \sim \mathcal{O}$ 0.5-2.0gev. *Phys. Rev. D*, 83:052009, 2011.

[85] A. A. Aguilar Arevalo and *et al*. Measurement of ν_μ and $\bar{\nu}_\mu$ induced charged current charged pion production cross sections on mineral oil at $e_\nu \sim \mathcal{O}$ 0.5–2.0gev. *Phys. Rev. D*, 83:052007, 2011.

[86] Daniel Scully. *Neutrino Induced Coherent Pion Production*. PhD thesis, University of Warwick, 2013.

[87] P. Adamson and *et al*. Neutrino and antineutrino inclusive charged–current cross section measurements with the minos near detector. *Phys. Rev. D*, 81:072002, 2010.

[88] E. J. Moniz and *et al*. Nuclear fermi momenta from quasielastic electron scattering. *Phys. Rev. Lett.*, 26:445–448, 1971.

[89] N. Abgrall and *et al*. Measurements of cross sections and charged pion spectra in proton–carbon interactions at 31 gev/c. *Phys. Rev. C*, 84:034604, 2011.

[90] I. Chemakin and *et al*. Pion production by photons on a thin beryllium target at 6.1, 12.3 and 17.5 gev/c incident proton momenta. *Phys. Rev. C*, 77:015209, 2008.

[91] T. Eichten and *et al*. Particle production in proton interactions in nuclei at 25 gev/c. *Nucl. Phys. B*, 44, 1972.

[92] J. V. Allaby and *et al*. High–energy particle spectra from proton interactions at 19.2 gev/c. *Tech. Rep.*, 70-12, 1970.

[93] N. Abgrall and *et al*. Measurement of production proterties of positively charged kaons in proton–carbon interactions at 31 gev/c. *Phys. Rev. C*, 85:035210, 2012.

[94] M. Bonesini and *et al*. On particle production for high energy neutrino beams. *Eur. Phys. J. C*, 20:015209, 2001.

[95] E. Frank and *et al*. B–field calibration and systematic errors. Technical report, T2K Collaboration, 2011.

[96] J. Caravaca and *et al*. Charge Misidentification in local and global reconstruction. Technical report, T2K Collaboration, 2011.

[97] D. Allan and *et al*. The Electromagnetic Calorimeter for the T2K Near Detector ND280. *Journal of Instrumentation*, 8, 2013.