

T2K ND280 π^0 Electromagnetic Calorimeter

Benjamin Still



A thesis submitted to the University of Sheffield in partial fulfilment
of the requirements for the degree of Doctor of Philosophy
in the subject of Physics

Department of Physics & Astronomy

April 2009

Abstract

The Tokai to Kamioka (T2K) experiment will begin taking data in late 2009; using the most intense source of man made neutrinos to probe the suppressed oscillation of $\nu_\mu \rightarrow \nu_e$ and therefore the third and final unmeasured θ_{13} lepton mixing angle. T2K achieves this with near (ND280) and far detector (Super Kamiokande) combination characterising the beam content before and after oscillation at distance of 295 km.

The confident measurement of small numbers of ν_e interactions at Super Kamiokande require accurate understanding of all background processes, most notably the neutral current production of π^0 ; a dedicated π^0 detector (P0D), part of the ND280, will measure this and other π^0 cross-sections.

The task of reconstructing a π^0 is challenging, requiring the conversion and containment of the low energy decay product photons which may travel some distance or escape detection in the P0D entirely. To aid in the task a veto detector, the π^0 electromagnetic calorimeter (P0DECal), is placed around the outer edges of the P0D to catch any missing energy or particles.

This thesis focuses upon the final design and development of reconstruction algorithms for the P0DECal. The performance of these methods is evaluated using Monte Carlo simulated neutrino data.

MC data of all types of expected neutrino interaction types are then used to estimate the impact of P0DECal information upon the selection of NC1 π^0 events using the P0D alone. It is shown that, especially during early running and water out running modes, the P0DECal is an essential addition improving both statistics and purity of NC1 π^0 samples.

Acknowledgements

I would firstly like to thank my girlfriend, Leila, for all of her selfless and unfaltering love, patience and support without which this could not be possible.

Thank you to my parents, grandparents and family for your loving support and encouragement in all I have chosen to do, and every advantage in life to be able to pursue them.

I owe a great debt of thanks to many colleagues who have grown to be friends.

Thanks to Susan Cartwright for her supervision and especially keen and criticising literary eye when reading drafts of the thesis; Neil McCauley for continued advice and help on reconstruction issues and for his comments upon this thesis; Yoshi Uchida for his continuing support and especially for his time and advice on the design study.

Thank you to my colleagues at Sheffield: Matt Robinson and Paul Hodgson for the fantastic computing cluster, their patience with me and their computing and programming advice; Rob Duxfield and Ed Daw for many discussions on data analysis and statistics common sense; Jon Perkin for being the font of ROOT knowledge; Mark Ward for sharing his hardware expertise and the odd late night beer.

I would also like to extend my gratitude to my fellow T2K students: Ian Taylor for discussions of clustering algorithms; Le Phoc Trung for discussions on the P0D reconstruction and π^0 analyses; Antony Carver for discussions on multivariate techniques.

Finally I would like to thank STFC and the University of Sheffield for the funding and opportunity to attend the many essential and exciting, conferences and summer schools worldwide.

Contents

1	Introduction	1
2	Neutrinos	3
2.1	From β -Decay to Neutrinos	3
2.2	Neutrino Interactions	5
2.3	Mass and Flavour Mixing	8
2.4	Neutrino Oscillations	14
2.5	Neutrino Oscillation Physics	19
2.6	Three Neutrino Mixing and The Future	27
3	The Tokai to Kamioka T2K Experiment	35
3.1	Overview of the experiment	35
3.2	The J-PARC neutrino beam	36
3.3	Super-Kamiokande	38
3.3.1	The Detector	38
3.3.2	PID at Super-Kamiokande	39
3.3.3	Neutrino CCQE backgrounds at Super-K	42
3.4	Near Detector Suite	43
3.4.1	INGrid	43
3.5	The ND280 Near Detector	44
3.5.1	Overview and Magnet	45
3.5.2	Tracker	45

3.5.3	Side Muon Range Detector, SMRD	50
3.5.4	The π^0 Detector, P0D	51
3.5.5	Electromagnetic Calorimeter, ECal	53
3.5.6	Electronics	55
3.6	Time Line	57
4	Light Yield Measurements of ECal Scintillator Bars	61
4.1	Multi-Pixel Photon Counters, MPPCs	61
4.2	Scintillator and Wavelength Shifting (WLS) Fibre	63
4.3	Experimental Setup	65
4.4	Analysis	67
4.5	Results	69
4.5.1	First Batch Prototype ECal Bar	69
4.5.2	Second Batch ECal Prototype	70
4.5.3	SciBar	71
4.6	Conclusions and Discussions	73
5	ND280 Offline Software	75
5.1	Neutrino Interaction Monte Carlo Generation	75
5.2	Software Framework	77
5.2.1	The ND280 Data Framework: oaEvent	77
5.2.2	Offline Database	80
5.2.3	User Interface: nd280Control	80
5.3	Monte Carlo (MC) Simulation of the ND280	80
5.3.1	Detector Monte Carlo Simulation: nd280mc	80
5.3.2	Electronics Simulation: elecSim	81
5.4	Event Reconstruction	82
5.4.1	Global Reconstruction: oaRecon	82
5.4.2	Tracker Reconstruction	82

5.4.3	P0D Reconstruction: p0dRecon	84
5.4.4	Tracker and Downstream ECal Reconstruction: ecalRecon	85
5.4.5	P0DECal Reconstruction: ecalRecon	86
5.5	Data Analysis	87
5.5.1	Event Display	87
5.5.2	Analysis Framework, oaAnalysis	87
6	Design of the P0DECal	89
6.1	P0DECal Overview	89
6.1.1	Motivation and Physics Goals	89
6.1.2	Detector	90
6.2	Photon Conversion and Containment	92
6.2.1	Photon Conversion	92
6.2.2	Photon Containment	94
6.2.3	Discussion	97
6.3	Energy Resolution	98
6.3.1	Discussion	101
6.4	Discriminating Variables	101
6.4.1	Discussion	103
6.5	Conclusions	105
7	MC Analysis of P0DECal Reconstruction	107
7.1	Hit Preparation: TECALPrepareHits	107
7.2	TP0DECALBasicClustering	109
7.3	TP0DECALCreatePid	110
7.3.1	Identify Clusters as Track-like or Shower-like	110
7.3.2	Determine Cluster Position	110
7.3.3	Determine Direction	111
7.3.4	Output Results	111

7.4	True Clustering	111
7.5	Monte Carlo Test Samples	112
7.5.1	Particle Gun Monte Carlo	112
7.5.2	Background	114
7.6	Clustering Efficiency and Purity	115
7.6.1	Cluster Finding	115
7.6.2	Clustering Efficiency	120
7.6.3	Clustering Purity	126
7.6.4	Conclusion	126
7.7	Position and Direction Reconstruction	126
7.7.1	Position Resolution	126
7.7.2	Direction Resolution	128
7.7.3	Conclusion	132
7.8	Energy Reconstruction	132
7.8.1	Reconstruction Algorithm	132
7.8.2	Calibration	133
7.8.3	Conclusion	136
7.9	Conclusion	136
8	P0DECal Particle Identification	137
8.1	Training Samples	137
8.1.1	Tracks	138
8.1.2	Showers	138
8.2	Discriminating Variables	139
8.2.1	Track Length	140
8.2.2	Axis Max Ratio, AMR	141
8.2.3	Path Charge Ratio	142
8.2.4	Max Charge Ratio	142
8.2.5	Charge Standard Deviation	143

8.2.6	Charge Mean	143
8.2.7	Normalised Charged Standard Deviation	143
8.2.8	Max Width	144
8.2.9	Charge Weighted Track Width	144
8.2.10	Variable Correlations	144
8.2.11	Variable Choice	147
8.3	Analysing MVA Output	149
8.3.1	Performance and The Receiver Operation Characteristic (ROC) Curve	149
8.3.2	Robustness: Overtraining Checks	151
8.4	Fisher Linear Discriminant	152
8.4.1	Description	152
8.4.2	Training	153
8.5	Maximum Likelihood	154
8.5.1	Description	154
8.6	Boosted Decision Trees	157
8.6.1	Theory	157
8.7	Multilayer Perceptron	160
8.7.1	Theory	161
8.7.2	Conclusion	167
8.8	Comparison and Conclusion	168
9	P0D + P0DECal NC1π^0 Event Selection	169
9.1	Monte Carlo Generation	170
9.2	Analysis	171
9.3	Results	174
9.4	Conclusion and Discussion	177
10	Summary and Conclusions	183

List of Tables

2.1	Limits on the neutrino mass from various sources.	10
2.2	Best fit values with 1σ errors and 2σ and 3σ confidence level intervals. These results are from a fit of global data from solar, atmospheric, reactor and accelerator experiments from [1].	31
3.1	Breakdown of ECal modules.	54
8.1	Separation of discriminating variables with and without decorrelation. .	148
9.1	Breakdown of NC1 π^0 selection from GENIE water in P0D data. . . .	180
9.2	Breakdown of NC1 π^0 selection from GENIE water out P0D data. . . .	181

List of Figures

2.1	Neutrino interactions.	8
2.2	KamLAND L/E.	18
2.3	Super-K atmospheric oscillation.	20
2.4	Solar and atmospheric oscillation data.	22
2.5	Neutrino mass hierarchies.	26
2.6	T2K Sensitivity to θ_{13} and δ_{CP} .	33
3.1	ν_μ and ν_e spectrum from the J-PARC neutrino beam at 280 m plus an overview of the beamline.	37
3.2	Schematic of Super-K plus example muon and electron event displays.	41
3.3	The INGrid on-axis detector.	44
3.4	The ND280 off axis detector.	46
3.5	ND280 TPC and MICROMEGAS.	49
3.6	ND280 SMRD Scintillator	50
3.7	ND280 P0D detector.	52
4.1	Hamamatsu MPPC	62
4.2	MPPC pedestal and first p.e. peak.	64
4.3	Experimental setup for light yield study.	66
4.4	MPPC waveform persistency plot.	67
4.5	Truncated MPPC waveform.	68
4.6	End on view of ND280 ECal scintillator.	70
4.7	First prototype ND280 ECal scintillator light yield.	71

4.8	Second prototype ND280 ECal and SciBar scintillator light yield.	72
5.1	Data objects in the nd280 offline software.	77
5.2	Data flow and dependency of all nd280 offline software.	79
5.3	Example event displays.	87
6.1	Photon interaction cross-section in lead.	92
6.2	Fraction missing photon events as a function of energy.	94
6.3	The layer of photon interaction.	95
6.4	Fraction of initial photon energy lost from the P0DECal.	96
6.5	Relative energy deposition as a function of initial incident photon energy.	99
6.6	Idealised energy resolution as a function of initial photon energy. . . .	100
6.7	TrackLength of photons and muons.	102
6.8	MeanChargeDeposited of photons and muons.	103
7.1	Flow of event data in P0DECal Reconstruction.	108
7.2	Angle convention in P0DECal.	113
7.3	Venn diagram outlining clustering efficiency and purity.	116
7.4	Photon event cluster finding efficiencies as a function of incident energy and angle.	118
7.5	Muon event cluster finding efficiencies as a function of incident energy and angle.	119
7.6	P0DECalRecon charge clustering efficiency.	122
7.7	P0DECalRecon hit clustering efficiency.	123
7.8	P0DECalRecon charge clustering purity.	124
7.9	P0DECalRecon hit clustering purity.	125
7.10	P0DECalRecon position resolution.	127
7.11	P0DECalRecon angular resolution as a function of energy.	129
7.12	P0DECalRecon reconstructed direction as a function of true incident direction.	130
7.13	P0DECalRecon direction resolution.	131

7.14	Energy reconstruction Likelihood variables.	134
7.15	Reconstructed energies of true 400 MeV photons.	135
8.1	Schematic of track and shower events in the P0DECal	140
8.2	Correlation matices of PID Variables with each other and true incident energy and angle on incidence.	145
8.3	Distributions of PID discriminating variables.	146
8.4	Venn diagram of ROC curve construction.	150
8.5	Example output and ROC curve of a MVA technique.	151
8.6	Example of a Fisher discriminant output.	153
8.7	Example of a likelihood discriminant output.	155
8.8	Schematic of a decision tree.	157
8.9	BDT overtraining check.	159
8.10	Example of a BDT discriminant output.	160
8.11	Example MLP network.	162
8.12	MLP ROC curve performance as a function of hidden neurons.	163
8.13	MLP overtraining check.	165
8.14	MLP Performance as a function of number of epochs and training events.	165
8.15	MLP Performance as a function of number of epochas and training events.	165
8.16	MLP Performance as a function of number of epochas and training events.	166
8.17	Example of a MLP discriminant output.	167
8.18	Comparison ROC curves for different MVA techniques.	168
9.1	Flow of P0D + P0DECal NC1 π^0 event selection.	172
9.2	Theta comparison cut, P0DECal track and P0D vertex position.	173
9.3	P0D and P0D+P0DECal NC1 π^0 event selection efficiency and purity. .	175
9.4	P0D+P0DECal NC1 π^0 event selection efficiency multiplied by purity as a function of timing and theta analysis cut.	176
9.5	The efficiency and purity of NC1 π^0 event selection with cuts.	178

Chapter 1

Introduction

The Standard Model of particle physics has withstood many experimental challenges since its development in the 1970s [2]. To date, the only area in which physics beyond the Standard Model can be studied in detail is the neutrino sector, where the Standard Model picture of massless, left-handed, flavour-conserving neutrinos has been thoroughly overturned in the last 20 years in favour of a complex pattern of non-zero masses, mixing angles and, possibly, \mathcal{CP} -violating phases.

The Tokai to Kamioka (T2K) experiment [3] will be the first in a new generation of neutrino superbeam experiments, using a neutrino beam of unprecedented power observed over a 295 km baseline. The main physics aim of T2K is to measure the third mixing angle in the lepton mixing matrix, while also improving the precision of the so-called “atmospheric” mixing angle and mass squared difference. These goals will be achieved using the well known Super-Kamiokande water Čerenkov experiment [4] as the far detector, with a new purpose built near detector (the ND280)[5] to characterise the unoscillated beam and provide improved measurements of key neutrino cross-sections.

Many of these key cross-sections involve the production of neutral pions, which are critical to the success of T2K because they present a background to electron identification in Super-K. The ND280 therefore includes a dedicated Pi-0 Detector or P0D, which is designed specifically to detect and reconstruct photons from π^0 decays. However, the low energy of the T2K beam (peaking at ~ 600 MeV) results in low energy photons which will often travel considerable distances through material before initiating an electromagnetic shower. To reduce the acceptance losses and consequent systematic errors caused by escaping photons, the P0D is surrounded by an electromagnetic calorimeter, the P0DECal, whose task is to contain and identify photons and muons which might otherwise have gone undetected.

Construction of the electromagnetic calorimeter, including the P0DECal, is a UK responsibility. This thesis describes the final design study (Chapter 6), development and performance of reconstruction software, and physics impact of the P0DECal. In the UK group I have been solely responsible for developing the P0DECal software packages (Chapters 7 and 8) and investigating its impact on physics (Chapter 9). I also carried out a brief study of light yield from ECal scintillator (Chapter 4), to ensure that it met the design goals.

The structure of the thesis is as follows. In Chapter 2, we begin by charting the colourful history of the neutrino from the postulation of its existence to modern day massive neutrinos. At each stage we review the experiments which measured the characteristics of this elusive particle and finish with a discussion of what the immediate, near and not so near future holds.

In Chapter 3 the Tokai to Kamioka (T2K) experiment is introduced, with a discussion of the physics goals and set-up of the experiment. I then concentrate upon the ND280 which is designed to characterise the neutrino beam before oscillation and measure relevant cross-sections; including Super Kamiokande (Super-K) far detector backgrounds such as the neutral current production of single neutral pions ($NC1\pi^0$), described in section 3.3.3.

The next Chapter takes a slight detour, presenting the results of my light yield measurements of ND280 electromagnetic calorimeter (ECal) scintillator bars.

Chapter 5 gives an overview of the ND280 offline software; Monte Carlo data simulation, reconstruction and physics analysis packages.

I present a study in Chapter 6 from which the design of the P0DECal sub-detector was finalised. The performances of different detector configurations are investigated using Monte Carlo simulation data at the truth level.

The chain of reconstruction algorithms which I developed for the P0DECal is described in Chapter 7, along with their associated performance on Monte Carlo data.

The essential ability to discriminate between showering particles and those leaving a track is covered in Chapter 8. Discriminating variables are discussed and chosen before a number of multivariate analysis (MVA) techniques, linear and non, are trained and tested with different reconstructed Monte Carlo data.

An MVA technique is selected and the fully reconstructed P0DECal data is used in conjunction with reconstructed P0D data to perform a feasibility study of $NC1\pi^0$ vertex selection which is presented in Chapter 9. We then conclude with an overview of the P0DECal's impact upon the measurement of this important ν_e background.

Chapter 2

Neutrinos

2.1 From β -Decay to Neutrinos

The Neutrino

The neutrino was born in 1930 from the mind of Wolfgang Pauli, in a letter written to the participants of a radioactivity conference in Tübingen. He postulated the existence of a light neutral particle as a ‘desperate remedy’ to the apparent violation of energy-conservation in β -decay. β -decay was believed to be a two body process in which a nuclear transition resulted in the emission of an electron. In such a two body decay one expects mono-energetic electrons with energy equal to the nuclear transition energy. Instead continuous spectra of electron energies were observed up to each expected energetic maximum. No other particles were detected in the vicinity of these decays, which led Pauli to the conclusion that, assuming energy conservation was not being violated in this decay, there is a second light neutral spin 1/2 particle, which he called a neutron, created alongside the electron.

In 1932 James Chadwick discovered [6] neutral particles with a mass similar to that of the proton. The name neutron was immediately adopted, but it was not the particle that Pauli postulated. In 1934 Enrico Fermi developed a theoretical framework for β -decay incorporating Pauli’s neutral particle [7]. From this he deduced that the neutron that James Chadwick had detected was far too heavy to be Pauli’s neutral particle. Realising that the neutral particle must be at least as light as the electron Fermi dubbed the new neutral particle the neutrino (“little neutral one”). Indirect evidence of the neutrino continued to build as Fermi’s theory was seen to accurately describe the electron spectra seen in an array of β -decay experiments. Later that year Bethe and Peierls took Fermi’s theory of neutrino interaction and calculated the cross-section of

neutrinos [8] with ordinary matter. They deduced a seemingly unobservable small cross-section, concluding “there is no practically possible way of observing the neutrino”. Their calculation was remarkably close to modern experimental and theoretical values, although it took over two decades for experimental techniques to advance to the point where their pessimism was proven unfounded.

The field of experimental neutrino physics began 22 years later with Reines and Cowan. In 1953 the pair made the first tentative observation [9] of (anti)neutrinos from the Hanford reactor, through scintillation observation of positrons from the reaction $\bar{\nu} + p \rightarrow n + e^+$. The large background of this experiment kept scepticism alive, prompting others to try to detect the neutrino through various methods. In 1955 a radiochemical experiment devised by Ray Davis [10] was situated at the same reactor. The experiment was sensitive to the reaction $\nu + {}^{37}\text{Cl} \rightarrow e + {}^{37}\text{Ar}$ but it measured no signal. Definitive direct evidence of antineutrino observation came in 1956 [11] when Reines and Cowan performed a more sensitive version of their earlier experiment at another reactor at Savannah River.

In these intervening years there was a series of developments that would eventually impact upon neutrino physics. In the mid-forties it was suggested by Sakata and Inoue[12] and independently by Marshak [13] that Yukawa’s meson [14] was in fact two distinct particles; the pion and the muon. Their suggestions were confirmed by the observations of Lattes et al. [15]. It was soon noted that the decay of pions to muons was a similar process to β -decay in that the muons were accompanied by a neutrino. Was this the same neutrino as that seen in β -decay?

Neutrinos

In 1959, whilst talking with colleagues, Melvin Schwartz realised that one could create an intense beam of neutrinos by first generating a beam of charged pions and allowing them to decay; charged pions are the main product of colliding a proton beam with a fixed target. While Schwartz developed his ideas, Lee and Yang were discussing the lack of $\mu \rightarrow e^- + \gamma$ events [16], with the γ coming from $\nu\bar{\nu} \rightarrow \gamma$ annihilation. This would lead one to assume that there are two distinct types of neutrino associated with the electron and muon respectively. The Summer of 1962 saw Schwartz and colleagues publish results on the first ever accelerator generated neutrinos from their experiment at the Brookhaven Laboratory. They showed that neutrinos produced from a beam of charged pions interact to produce muons only [17]. Here for the first time was irrefutable evidence that there existed more than one generation of neutrino. The neutrino in β -decay was indeed different from that seen by Lattes et al.

By 1974 experiments had discovered two complete generations of both quarks and leptons. In 1975 Martin Perl and colleagues in the SLAC-LBL collaboration added a third lepton to the mix: the tau, τ . Two years later Leon Lederman and colleagues, at the Fermi Laboratory, brought the bottom, b , quark into the picture. So with three generations of charged lepton and quark, neutrino physicists once again had a challenge on their hands to directly detect the third generation of neutrino. Because of the tau's short lifetime and the ever present phantom nature of the neutrino it took until 2000 before the tau neutrino, ν_τ , was discovered by the DONUT experiment at the Fermi Laboratory [18].

2.2 Neutrino Interactions

V-A

We now take our story of the neutrino down a different route and focus on the nature of its interactions. For β -decay one is able to write five classes of Lorentz invariant interaction: Scalar (S), Vector (V), Tensor (T), Axial vector (A) and Pseudo-scalar(P). Each of these interaction types has an individual coupling, C_j , which must be experimentally determined. The matrix element for allowed beta decay can be written as a linear combination of these terms,

$$\mathcal{M} = \frac{G_F}{\sqrt{2}} \sum_{nucleons} \sum_{j=SVTAP} \int C_j \bar{\psi}_p(\mathbf{x}) O_j \psi_n(\mathbf{x}) \bar{\psi}_e(\mathbf{x}) O_j \psi_{\bar{e}}(\mathbf{x}) d^3x, \quad (2.1)$$

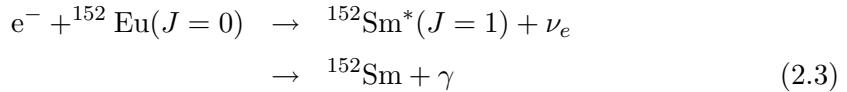
where the ψ are single-particle Dirac wavefunctions. The operators, $O_j = 1, \gamma_\lambda, \sigma_{\lambda\nu}, \gamma_\lambda \gamma_5, \gamma_5$ mediate the interactions for $j = S, V, T, A, P$ respectively.

In the flurry of meson discoveries, the same that prompted the pion-muon discovery, data began to show that two particles, τ and θ , were identical in every way apart from their parity as measured from their decay products. Lee and Yang published a paper [19], around the same time as Reines and Cowan's first observations, noting that parity conservation (\mathcal{P} , symmetry under spatial inversion; $\mathbf{x} \rightarrow -\mathbf{x}$) had only been assumed and was not a pre-requisite of Fermi's theory. In the paper they suggested direct experimental tests of parity violation in β -decay (as well as muon decay and others). Their hypothesis was quickly confirmed by Wu et al. [20] and subsequently Garwin et al. [21]. It was also noted, from these experimental results and theory [22], that if parity is violated then so must charge conjugation symmetry (\mathcal{C}) be. When one accepts that parity violation is possible the extra degree of freedom allows pseudo-scalar

versions of each term with an additional set of respective coupling constants,

$$\mathcal{M} = \frac{G_F}{\sqrt{2}} \sum_{\text{nucleons}} \sum_{j=SVTAP} [\bar{\psi}_p O_j \psi_n] [\bar{\psi}_e O_j (C_j - C'_j \gamma^5) \psi_{\bar{\nu}}] d^3 x. \quad (2.2)$$

All coupling constants were constrained and subsequently measured by just a handful of different experiments. Using Mott-Møller scattering a number of experiments directly observed the parity of electrons and positrons from β decay. These experiments concluded that all electrons were emitted left-handed and positrons right-handed. The experiment of Goldhaber et al. [23] determined the parity of the neutrino just a year after its first detection. After electron capture by a Europium nucleus the excited daughter Samarium nucleus de-excites through emission of a single photon. Keeping a record of the parity of the nuclei one finds that helicity of the de-excitation photon directly correlates with the parity of the neutrino emitted after electron capture. They concluded that all neutrinos emitted in this process were left-handed.



With only (right)left-handed (anti)leptons seen interacting in nature, this highly constrains the relationship between each class of scalar and pseudo-scalar coupling constants. For the S, T and P the scalar and pseudo-scalar coupling are equivalent, $C_j = C'_j$. The V and A interaction modes are constrained by the scalar and pseudo-scalar coupling having opposite phase, $C_j = -C'_j$. These results led to the suggestion by Sudarshan and Marshak [24], also by Feynman and Gell-Mann [25], that weak interactions have a V-A structure. This form results in interactions that maximally violate \mathcal{P} and \mathcal{C} but preserve \mathcal{CP} and time reversal symmetry (\mathcal{T} ; $t \rightarrow -t$). The fact that \mathcal{P} is maximally violated also reopened the question of whether the neutrino and antineutrino are really distinct. With hindsight we can understand Davis' 1955 failure to observe the neutrino [10] as the weak interaction does not couple the particles involved. By the time the V-A model was developed it was realised that the point interaction of Fermi's theory must be a low energy approximation, and that the short range of the weak interaction was probably due to the exchange of a massive gauge boson.

Simply putting in a mass for the boson leads to divergences in the theory. The solution was developed over the next decade, into what is now known as the Glashow-Weinberg-Salam (GWS) model [26, 27, 28, 29]. Using local SU(2) gauge symmetry breaking the theory provided a mass for the mediating bosons. This process also unified the new

massive weak bosons with electromagnetism into a single electroweak force.

Alongside the required W^\pm bosons and the photon the theory also predicted the existence of a massive neutral weak boson, the Z^0 . The neutral current weak interactions of the Z^0 boson were predicted to be similar to those of a short ranged (massive) photon. Unlike the electromagnetic interactions of the photon, however, the Z^0 could interact with the neutrino. This made a beam of neutrinos, such as the one used by Melvin Schwartz et al. [17], an obvious place to search for such neutral interactions. In 1973 the Gargamelle collaboration published the first evidence of a neutral current both in the reaction $\nu_\mu + e \rightarrow \nu_\mu + e$ [30] and in the reaction $\nu_\mu + N \rightarrow \nu_\mu + X$ [31]. Combined with the model of strong interactions, the GWS model of electroweak interactions forms the basis of what is now known as the Standard Model of particle physics.

Neutrino Nucleus Interactions

For neutrinos the Standard Model is unchanged from the GWS model. Neutrinos are therefore (right)left-handed (anti)particles which only interact with the W^\pm and Z^0 bosons. From the the earliest experiments the charged current interactions, mediated by the W^\pm , have been the main channel in which to measure neutrinos. This is because it is often easy to characterise the type and energy of the charged lepton daughter.

After their creation at neutrino interaction vertices the weak bosons then proceed to interact with a nucleus or nucleons. The amount of momentum transferred by the boson determines the final weak hadronic state. At low energies the momentum transfers are classified by the observed final hadronic state as either coherent, low momentum transfer, or incoherent, high momentum transfer.

At the scale of the nucleus, interactions between nucleons can be said to be mediated by the exchange of off shell (virtual) pions [14]. When the momentum transferred by the mediating weak boson is low, then the boson's long wavelength leaves it insensitive to the internal structure of the nucleus. Low momentum transfer weak bosons can interact with the entire nucleus, acting as a coherent whole, which results in one or more of these virtual pions being brought on shell. This process, called coherent production, results in a forward peaked pion momentum, an indicator of the small momentum transferred. Although small, the coherent contribution to the total neutrino cross-section is non-negligible.

Interactions at the nucleon level entail a larger momentum transfer, resulting in a much broader spectrum of final state momenta, and are conversely called incoherent interactions. These incoherent processes are by far the dominant mechanism of interaction.

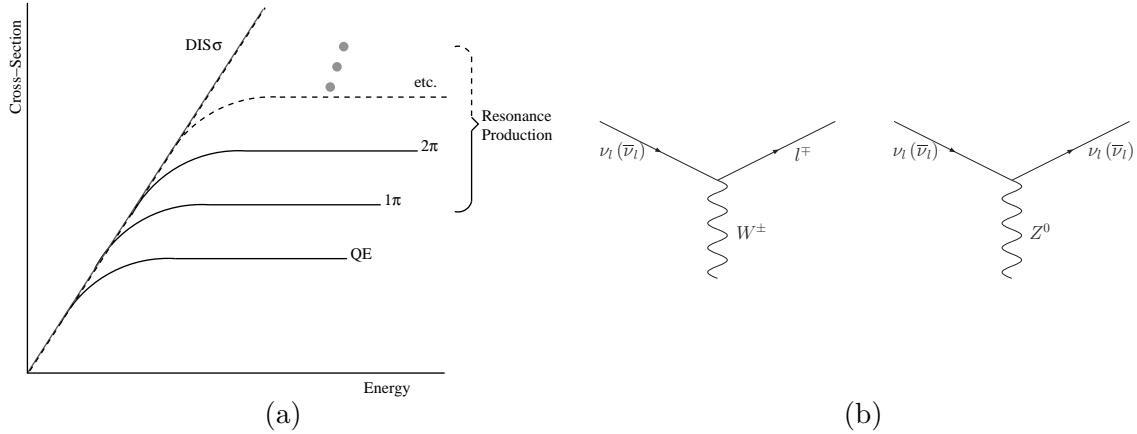


Figure 2.1: (a) Schematic of incoherent resonance cross-section as a function of energy. The total cross-section rises linearly with energy and consists of consecutive pionic resonances after the tree level quasi-elastic interactions. (b) Feynman diagrams of the standard model weak interactions of the neutrino.

The magnitude of momentum transferred, which is dependent on incident neutrino energy, determines the possible spectrum of final states. The tree level diagrams are quasi-elastic scattering which result in a recoiling nucleon; the same target nucleon in neutral current and a new nucleon in charged current interactions. Higher level diagrams produce pions of appropriate charge with multiplicity dependent upon the momentum transferred by the boson [Fig 2.1(L)].

The final state particles from both classes of interaction can interact with the sea of off-shell pions or decay on their exit from a nucleus. These final state interactions (FSI) can lead to a different pion multiplicity from expected. Charged pions can produce additional charged and/or neutral pions while neutral pions may decay before exiting the nucleus.

2.3 Mass and Flavour Mixing

Dirac mass terms enter the Standard Model in the coupling of left and right hand components of a particle field. In the GWS electroweak model only left-handed leptons and quarks exist. Leptons in this model are therefore massless as such cross terms cannot exist, yet the charged leptons are known to be massive. To make the theory consistent with this fact one must add by hand a right-handed singlet field for each generation of charged lepton. At the time of construction there was no evidence to suggest that the neutrinos possessed a mass and so no analogous right-handed singlets

exist for neutrinos in the SM.

Extending the Standard Model to include a right-handed neutrino, we can follow the same path taken for other massive particles in the Standard Model. This Dirac neutrino acquires a mass from the left-right coupling to the Higgs field:

$$\mathcal{L} = -m_D (\bar{\nu}_R \nu_L + \bar{\nu}_L \nu_R). \quad (2.4)$$

The ν_L and ν_R are the left and right-handed (chirality) neutrinos respectively. The Dirac mass m_D is the coupling constant from the neutrino's coupling to the Higgs field and has gauge quantum numbers from the electroweak coupling of the left-handed neutrino.

As the right-handed neutrino couples to no Standard Model force it has zero gauge quantum numbers. Not long after Fermi compiled his theory of β -decay Ettore Majorana published a paper suggesting the possibility of a symmetry between particles and antiparticles [32]. This symmetry is possible only if the particle is void of gauge couplings, as the right-handed neutrino would be. So as soon as one introduces a right-handed neutrino there is also the possibility of including a Majorana term of the form

$$\mathcal{L} = -\frac{1}{2}M \left(\bar{N}_R^C N_R^C + \bar{N}_R^C N_R \right), \quad (2.5)$$

where $N_R^C = i\gamma^0\gamma^2 \left(\bar{N}_R^T \right)$ is the charge conjugate of a right-handed neutrino and M is the Majorana mass. These two formalisms can be combined into a single term

$$\mathcal{L}_{D+M} = -\frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \bar{N}_R^C \end{pmatrix} \begin{pmatrix} 0 & m_D \\ m_D & M \end{pmatrix} \begin{pmatrix} \nu_L^C \\ N_R \end{pmatrix} + \text{h.c..} \quad (2.6)$$

In many theories the natural size of the Majorana mass is of the order of 10^{15} GeV, much larger than Dirac masses generated by the Higgs mechanism. The eigenstates of this Lagrangian, obtained by diagonalising the matrix, are a (nearly) right-handed Majorana neutrino with mass of approximately M , and a (nearly) left-handed Majorana neutrino with a mass of approximately $\frac{m_D^2}{M}$. If m_D is taken to be around 100 GeV, the left-handed neutrinos will be given masses in the region of 10^{-2} eV. This 'see-saw' mechanism is the most popular method of explaining the current evidence for the neutrino's tiny mass; a million times smaller than that of the electron. One of its most important predictions is that neutrinos are Majorana particles. This is a question that has yet to be settled experimentally. If neutrinos are indeed Majorana particles then phenomena such as neutrinoless double β -decay ${}_Z^A X \rightarrow {}_{Z+2}^A X' + 2e^-$ are possible. Several experiments [33, 34, 35, 36, 37] are now looking for evidence of this process

which, in addition to the discovery of the Majorana nature of neutrinos, would provide a measurement of the absolute scale of neutrino masses. Most of neutrino physics, and neutrino oscillations in particular, is however insensitive to the Majorana/Dirac question.

The absolute mass scale of the neutrino can also be probed by regular β -decay as well as other particle, astrophysical and cosmological probes. Current global data on these searches are summarised in Table 2.1.

Measurement	Origin of Limit	Limit	Reference
Tritium Decay	$m_{\nu_e}^2$	< 2.0 eV	[38]
π^+ Decay	$m_{\nu_\mu}^2$	< 0.17 MeV	[39]
τ Decay	$m_{\nu_\tau}^2$	< 18.2 MeV	[40]
SN1987A Time-of-Flight	$m_{(\nu_e \rightarrow \nu)}^2$	< 5.7 eV	[41]
Terrestrial Time-of-Flight	$m_{(\nu_e \rightarrow \nu)}^2$	< 50 MeV	[42]
Cosmology	$\sum_i m_{\nu_i}$	< 0.66 eV	[43]

Table 2.1: Limits on the neutrino mass from various sources.

The first hint that the neutrino may not be as it was represented in the Standard Model came in the late 1960's courtesy of Ray Davis Jr. [44] from the results of a large scale version of his earlier detector at Savannah River. A massive tank of dry-cleaning fluid, tetrachlorethylene C_2Cl_4 , was buried deep underground in the Homestake mine to reduce background from cosmic rays. For the first time in history it measured the ν_e flux emanating from the Sun. Davis worked closely with the theorist John Bahcall, whose predicted flux of neutrinos was three times that measured by Davis. At the time it was thought there was something wrong with the theory, the experiment, or both. Bahcall's solar model agreed very well with other solar data such as helioseismology, and further refinements of both theory and experiment improved the precision to which the discrepancy was known, but did not make it go away. This became known as the solar neutrino problem, and would not be resolved for some decades.

Independent confirmation of the Homestake result came in 1989 from the Kamioka Nucleon Decay Experiment (Kamiokande)[45], a 0.68 kt fiducial volume water Čerenkov detector. Further corroboration came from the Gallium radiochemical experiments SAGE [46] and GALLEX (later GNO) [47]. Each of these three experiment detected far less neutrinos than predicted by Bahcall's standard solar model (SSM) theoretical predictions [48]. These varied detector measurements largely ruled out experimental error as a cause of the solar neutrino problem. It wasn't until a decade later, 30 years since Davis first brought the solar neutrino problem to light, that Bahcall's predictions would be confirmed.

A similar discrepancy between theory and experiment was being seen at around the same time in neutrinos produced in cosmic ray showers. The Kamiokande [49] and IMB [50, 51] proton decay water Čerenkov experiments were both measuring their background of neutrinos produced in cosmic ray interactions with the atmosphere. Both experiments saw about a 60% difference between the measured and predicted interaction rates of ν_μ . The rate of ν_e interactions, however, agreed with their predicted flux, so an error in these calculations could be ruled out. In the late 1980's, two more experiments, Fréjus and NUSEX, both ionisation chamber detectors, reported results consistent with the theory; however, their statistical errors were too large to refute the earlier results convincingly. Many felt this discrepancy, dubbed the atmospheric neutrino problem, was simply an artefact of Čerenkov detectors. If a real effect, however, the observed depletion of specific flavours of neutrino from two independent sources lent support to the idea of flavour change in the lepton sector.

Flavour change was first postulated in 1963 when Cabibbo proposed an explanation [52] for the fact that strangeness is not always conserved in weak interactions (e.g. in $K^+ \rightarrow \pi^+ + \pi^0$). Cabibbo suggested that the weak interaction coupled to a mixture of strangeness conserving and strangeness violating hadronic currents. In the modern Standard Model this is now understood in terms of mixing between quark states; the u couples not to the d, but to a mixture of d, s and b.

The previous year a paper by Maki, Nakagawa and Sakata [53] attempted to address the same problem using the Nagoya model of baryons. In this model baryons are composed of some kind of baryonic matter plus a lepton. If one interprets the paper in modern Standard Model terms the mixing introduced in the paper was essentially between ν_e and ν_μ .

In 1970 Glashow, Iliopoulos and Maiani showed that by adding a fourth quark to Cabibbo's model, flavour changing neutral currents could be suppressed [54]. Three years later Kobayashi and Maskawa noted [55] that the CP violation observed in the neutral kaon system by Cronin, Fitch et al. [56] could be explained by including a third generation of quarks. The unitary matrix that describes the resultant mixing, which is actually between the strong and weak eigenstates of the quarks, is hence known as the CKM matrix (for Cabibbo, Kobayashi and Maskawa).

An equivalent matrix for the lepton sector was thought not to be necessary because a massless neutrino is defined only by its weak eigenstate. Allowing massive neutrinos introduces an analogous mixing between neutrino weak and mass eigenstates¹. Flavour

¹The matrix describing such mixing is called the PMNS matrix after Pontecorvo, for his 1968 paper on $\nu_e \leftrightarrow \nu_\mu$ oscillation [57], and also Maki, Nakagawa and Sakata who essentially introduced the mixing in explaining strangeness violation in the quark sector [53].

mixing is a consequence of having multiple generations of massive particle. With just one generation (one dimensional space), gauge and mass eigenstates must be aligned. As one increases the number of generations, this condition no longer holds. In general the gauge and mass eigenstates need not be aligned and can be rotated with respect to each other in the particle's state space. More generations lead to a higher number of possible degrees of freedom. These degrees of freedom arise in the Standard Model when diagonalising coupling matrices to obtain real coupling constants that can be measured experimentally. This diagonalisation can be viewed as rotating eigenstate axes until they align and yield single real values.

Until now we have written the Lagrangian terms for a single generation of lepton. For N generations of Dirac neutrino we can write the generalised mass Lagrangian term as

$$\mathcal{L} = \sum_{i,j}^N Y_{\alpha\beta}^\nu \epsilon_{ab} L_{ab}^\dagger \Phi_a^\dagger \nu_{R\beta} + \text{h.c.} \quad (2.7)$$

Here $Y_{\alpha\beta}$ is the matrix of Yukawa couplings of the neutrino of flavours $\alpha/\beta = e, \mu, \tau$ to the Higgs field through interaction with the Higgs doublet $\Phi = \begin{pmatrix} 0 \\ \phi + h(x) \end{pmatrix}$, and $\epsilon_{ab} = 1(a \neq b), 0(a \equiv b)$ ensures coupling to the Higgs field.

The Yukawa matrix is a complex $N \times N$ matrix which contains the couplings of each generation of weak neutrino eigenstate to the Higgs field. An arbitrary complex $N \times N$ matrix M can be transformed into a real diagonal matrix by two $N \times N$ unitary matrices, U , such that $U_1 M U_2^\dagger = M_D$ where M_D is real diagonal. This is done in the case of the Yukawa matrix and the real diagonal form, m_{ij} , is interpreted as the representative mass of each generation of neutrino:

$$Y_{\alpha\beta}^\nu = V_{\alpha i}^L m_{ij}^\nu V_{\beta j}^R. \quad (2.8)$$

The unitary matrices $V_{\alpha i}^\nu$ are uniquely defined except that one can multiply each on the left by the same diagonal phase factor matrix without altering the physics. This is the same as a phase change of the particle field, which must be invariant under such a global symmetry change. One must also diagonalise the Yukawa couplings of the charged leptons, which introduce an additional set of analogous unitary matrices $V_{\alpha i}^l$ and phase factors. These unitary matrices then provide transformations from weak to mass eigenstates for each lepton. To describe weak interaction in terms of mass eigenstates, i , rather than the weak flavour eigenstates, α , we make the substitution

of each left and right-handed, charged lepton and neutrino field:

$$|l_\alpha\rangle = \sum_i^N V_{\alpha i}^l |l_i\rangle \quad |\nu_\alpha\rangle = \sum_i^N V_{\alpha i}^\nu |\nu_i\rangle. \quad (2.9)$$

The Yukawa couplings of left and right-handed leptons are equivalent and so this substitution leaves all but W^\pm boson interaction terms unchanged. The charge carrying W^\pm bosons mediate interactions between charged leptons and neutrinos. When substituting for the weak fields one is left with the product of charged lepton and neutrino unitary matrices, $V_{\alpha i}^{l_L} V_{\beta i}^{\dagger \nu_L}$, which itself is a unitary matrix, U . This new matrix in essence provides the transformation from charged lepton flavour to neutrino mass eigenstate, much as the CKM matrix transforms from strong to weak eigenstates in the quark sector. This single unitary matrix now provides the combined diagonalisation of the Yukawa matrix $U^\dagger Y U = m$. With this knowledge one can choose a different but equivalent substitution to the one above;

$$|l_\alpha\rangle = |l_i\rangle \quad |\nu_\alpha\rangle = \sum_i^N U_{\alpha i} |\nu_i\rangle, \quad (2.10)$$

where $U_{\alpha i} = V_{\alpha j}^{l_L} V_{\beta i}^{\dagger \nu_L}$, l denotes the charged lepton fields, ν the neutrino fields, α the charged lepton flavour and i the neutrino mass eigenstate.

The matrix U is an $N \times N$ unitary matrix, the parameters of which are not determined within the theory. The $N \times N$ matrix elements may be complex, which at first glance leaves $2N^2$ parameters to be measured. Each column of the matrix U describes eigenvectors of neutrino mass eigenstates, which dictates the unitary nature, imposing constraints upon the matrix. Each eigenvector must be a unit vector of length $\sum_\alpha |U_{\alpha i}|^2 = 1; i = 1, N$. This imposes N constraints upon the matrix. The eigenvectors must also be orthonormal, $\sum_\alpha U_{\alpha i}^* U_{\alpha j} = \delta_{ij}$ where $i, j = 1, N$. With the $N(N - 1)/2$ orthogonality conditions having both a real and complex part, this provides a further $N(N - 1)$ constraints, leaving N^2 parameters to account for.

The phases associated with the arbitrary phase transformation of the charged lepton fields provide a further N constraints on the matrix U . One could also try to continue with this idea to remove a further $N - 1$ degrees of freedom from the phase change of the neutrino fields. One phase remains, as only differences of the lepton and neutrino phases appear in the final U matrix, e.g. $U'_{e1} = \exp(i\beta_e - \alpha_1) U_{e1}$. If the neutrino is allowed to be its own antiparticle, then the arbitrary phase change of the neutrino field does not leave the mass Lagrangian invariant. Therefore if one assumes the neutrino to be Majorana then the re-phasing of the field cannot occur and so one cannot impose any

further constraints. For completeness we continue to include these Majorana phases, ξ_1 and ξ_2 , which lie along the diagonal and so do not affect the mixing between lepton and neutrino eigenstates seen in charged current weak interactions.

This leaves the matrix U with $N^2 - N$ real parameters. Half are mixing angles; parameters that would be present even if U were real. Half are complex phases. The popular parameterisation used in [58] is that of three real rotation matrices with an additional fourth complex phase matrix containing the Majorana phases. The final overall phase that cannot be removed through rotation must be assigned to one of the real mixing angles. This phase, which does have an impact upon neutrino oscillation, is analogous to the \mathcal{CP} violating phase in the CKM matrix. It is associated with the third and final unmeasured mixing angle θ_{13} .

$$U_{li} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} \exp(i\delta) \\ 0 & 1 & 0 \\ -s_{13} \exp(i\delta) & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} \exp(i\xi_1/2) & 0 & 0 \\ 0 & \exp(i\xi_2/2) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (2.11)$$

Here $c_{ij} = \cos \theta_{ij}$, $s_{ij} = \sin \theta_{ij}$ where θ_{ij} are the three mixing angles. The δ , ξ_1 and ξ_2 quantities are the three complex phases of the mass eigenstates with respect to the weak eigenstates. The three mixing matrices describe mixing between specific mass eigenstates. The $e^{i\delta}$ term in the second matrix is the Dirac CP violation phase.

2.4 Neutrino Oscillations

Flavour oscillation in the lepton sector developed quite separately from the electroweak model. A change of particle is a result of the mixing of different particle fields. Mixing was first predicted in the particle/antiparticle, K^0 and $\overline{K^0}$, neutral kaon meson system [59, 60]. Bruno Pontecorvo expanded on this idea to investigate the possibility of particle/antiparticle transitions in other particle systems. He first considered the bound $(\mu^\pm e^\pm)$ system [61], then neutrino/antineutrino transitions [62] and concluded that mixing would be inevitable if not prohibited by some conserved quantity associated with the neutrino. With the discovery of the muon neutrino, ν_μ , by Schwartz et al., lepton flavour seemed to provide the conserved quantity required to forbid these type of

mixings. In 1968 Pontecorvo then suggested the possibility of $\nu_e \leftrightarrow \nu_\mu$ flavour-change mixing [57].

The oscillation mechanism suggested by Pontecorvo was a leading candidate to explain the flavour change mechanism seen in both the solar and atmospheric neutrino problems. As with the kaon oscillations, it could be parameterised by an amplitude and a length scale; the parameters of oscillation are directly related to the mixing of neutrino states.

The amplitude for the neutrino flavour change $\nu_\alpha \rightarrow \nu_\beta$ between source and detector is given by

$$Amp(\nu_\alpha \rightarrow \nu_\beta) = \sum_i U_{\alpha i}^* Prop(\nu_i) U_{\beta i}. \quad (2.12)$$

$U_{\alpha i}^*$ is the amplitude for neutrino of flavour ν_α to be the mass eigenstate ν_i (see section 2.3). The factor $Prop(\nu_i)$ is the kinematic amplitude for the mass eigenstate ν_i to propagate from source to detector. The final factor $U_{\beta i}$ is the amplitude of the mass eigenstate ν_i being the weak eigenstate ν_β .

Neutrino Propagation

If one makes the assumption that the neutrino propagates as a plane wave, then the propagation amplitude $Prop(\nu_i)$ is $\exp[-m_i \tau_i]$ where m_i is the mass of ν_i and τ_i the proper time elapsed in the ν_i rest frame. From Lorentz invariance $m_i \tau_i = E_i t - p_i L$ where L is the lab-frame distance from source to detector, t is the lab-frame time it takes the ν_i to travel this distance, and E_i and p_i are the lab-frame energy and momentum of ν_i respectively. In realistic detectors one cannot measure this amplitude directly but instead measures a bulk probability $P(\nu_\alpha \rightarrow \nu_\beta) = |Amp(\nu_\alpha \rightarrow \nu_\beta)|^2$, in which only relative phase differences will have physical significance. The relative phase of $Prop(\nu_i)$ and $Prop(\nu_j)$, $\delta\phi_{ij}$ is given by

$$\delta\phi_{ij} = (p_i - p_j) L - (E_i - E_j) t. \quad (2.13)$$

Experiments do not in practice measure the transit time, t , but this may be substituted by, to an excellent approximation [63], L/\bar{v} where

$$\bar{v} = \frac{p_i + p_j}{E_i + E_j} \quad (2.14)$$

is an approximation of the average group velocity of the ν_i and ν_j components. Then

$$\delta\phi_{ij} \cong \frac{p_i^2 - p_j^2}{p_i + p_j} L - \frac{E_i^2 - E_j^2}{p_i + p_j} L \cong (m_i^2 - m_j^2) \frac{L}{2E}, \quad (2.15)$$

where in the final stage we have taken the highly relativistic limit that the momenta of the two components are approximately equal to the beam energy, E . From this one can make the substitution $Prop(\nu_i) \cong \exp(-im_i^2 L/2E)$ as this yields the desired probability. So the amplitude of $\nu_\alpha \rightarrow \nu_\beta$ propagation is

$$Amp(\nu_\alpha \rightarrow \nu_\beta) = \sum_i U_{\alpha i}^* \exp(-im_i^2 L/2E) U_{\beta i}. \quad (2.16)$$

The probability can then be written as

$$P(\nu_\alpha \rightarrow \nu_\beta) = |Amp(\nu_\alpha \rightarrow \nu_\beta)|^2 = \sum_{ij} U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j} \exp(-i(m_i^2 - m_j^2) L/2E). \quad (2.17)$$

Rewriting $m_i^2 - m_j^2 = \Delta m_{ij}^2$ and expanding the sum we get

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_i U_{\beta i} U_{\alpha i}^* \sum_j U_{\beta j}^* U_{\alpha j} + \sum_{ij} U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j} (\exp(-i\Delta m_{ij}^2 L/2E) - 1). \quad (2.18)$$

The first term is $\delta_{\alpha\beta}$ because U is unitary. The terms with $j > i$ in the subsequent terms are just complex conjugates of the terms with $j < i$; also when $j = i$ the phase difference is zero. Hence

$$P(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} + 2 \sum_{i>j} \Re [U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j} (\exp(-i\Delta m_{ij}^2 L/2E) - 1)]. \quad (2.19)$$

Finally we split the last term into real and imaginary components and, using a little trigonometry, get the form

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= \delta_{\alpha\beta} \\ &- 4 \sum_{i>j} \Re [U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j}] \sin^2 \left(1.27 \Delta m_{ij}^2 \frac{L}{E} \right) \\ &+ 2 \sum_{i>j} \Im [U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j}] \sin \left(2.54 \Delta m_{ij}^2 \frac{L}{E} \right). \end{aligned} \quad (2.20)$$

Here we have included the previously omitted factors of \hbar and c , so that the variable $\Delta m_{ij}^2 = m_i^2 - m_j^2$ is in units of eV^2 , L is in km and E in GeV. So when one considers travelling neutrinos we gain a further two free parameters that must be measured

experimentally, Δm_{12}^2 and Δm_{23}^2 , which measure the mass differences between the 1 and 2 and the 2 and 3 neutrino eigenstates respectively.

The first term is the result if the neutrino were massless and its mass and weak eigenstates aligned; this can also be interpreted as an infinite oscillation length. The second term explains the standard oscillation that occurs from the three real mixing angles due to the misalignment of the weak and mass eigenstates. The third term contains all of the complex phase information, coupled with mixing angle terms. It disappears if U can be expressed as a real matrix and also if $\alpha \equiv \beta$:

$$\Im \mathfrak{m} [U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j}] = \Im \mathfrak{m} [U_{\alpha i} U_{\alpha i}^* U_{\alpha j}^* U_{\alpha j}] = \Im \mathfrak{m} [|U_{\alpha i}|^2 |U_{\alpha j}|^2] = 0. \quad (2.21)$$

Solar and atmospheric experiments, measuring the disappearance of particular neutrino flavours, quantify this survival probability. In these experiments, the imaginary term in equation 2.20 does not exist;

$$P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - 4 \sum_{i>j} |U_{\alpha i}|^2 |U_{\alpha j}|^2 \sin^2 \left(1.27 \Delta m_{ij}^2 \frac{L}{E} \right). \quad (2.22)$$

The analogous CKM quark mixing matrix sector was measured to have very small values for each of the four mixing parameters. To account for the experimental observations in solar and cosmic-ray neutrino experiments, the corresponding values in the PMNS matrix would have to be very different. Other theories for the depletion of neutrino number were springing up alongside data.

Experimental Evidence for Oscillation

In 2002 the decades old solar neutrino problem was laid to rest by the Sudbury Neutrino Observatory (SNO) experiment. As well as measuring the charged current interactions of the neutrino, from Čerenkov light, it was the first solar neutrino experiment to measure neutral current interactions which are sensitive to all flavours of active neutrino. It measured the neutral current flux of ${}^8\text{B}$ solar neutrinos using the reaction $\nu + d \rightarrow p + n + \nu$ [64]. The flux measurement of this process, which is equally sensitive to all flavours of active neutrino, agreed with Bahcall's SSM. This firmly established flavour change between active neutrinos as the most plausible explanation of the solar neutrino problem. We now believe that electron neutrinos coming from the sun are depleted by the MSW effect (section 2.5), a form of neutrino oscillations that occurs for neutrinos passing through matter.

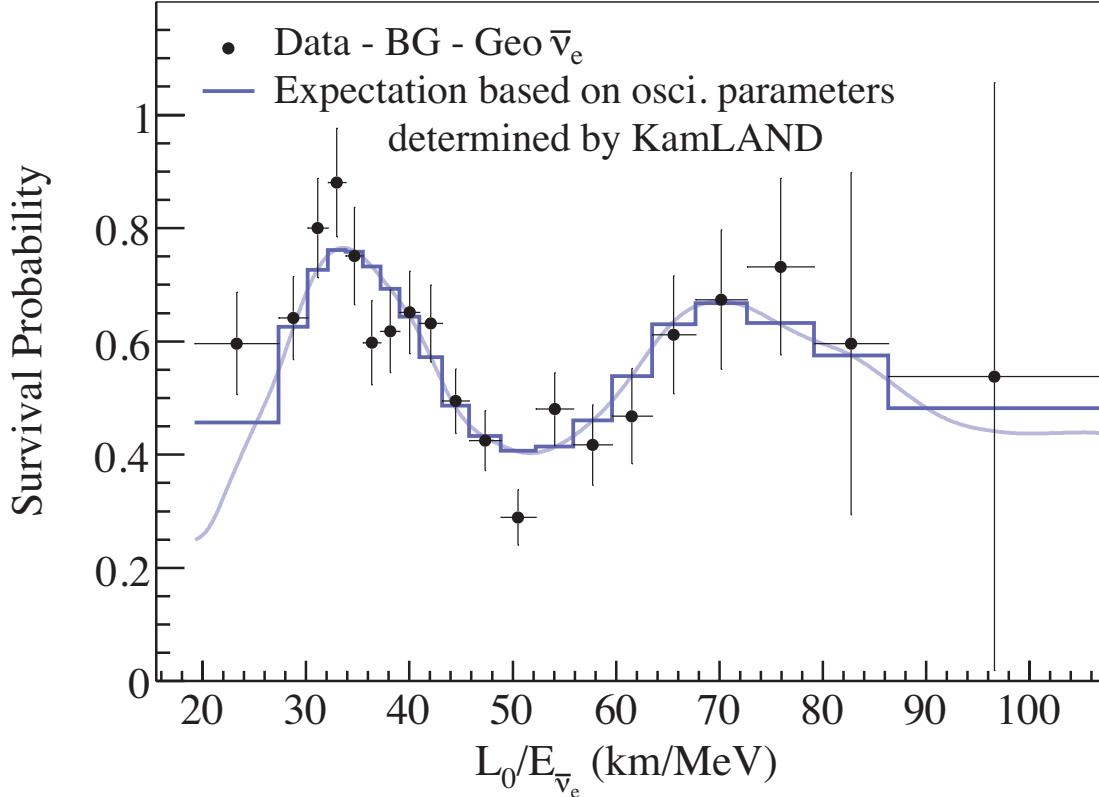


Figure 2.2: Event distribution in L_0/E , where L_0 is taken to be 180 km. The blue histogram are expected data from best fit oscillation which is displayed by the underlying curve [68].

Another way to study the mixing responsible for the solar neutrino problem, assuming CPT invariance, is to look at the survival of electron antineutrinos produced in nuclear fission. Several experiments were sited near (~ 1 km from) nuclear reactors in order to measure the probability of electron neutrino survival. None of these experiments saw any depletion of the $\bar{\nu}_e$ flux. Although flavour change was not seen by these experiments their null results, in particular CHOOZ [65] and Palo Verde [66], set important constraints upon all other neutrino parameters.

The reason that none of these reactor experiments saw any evidence of mixing was that they were simply too close to their neutrino sources: the key L/E oscillation parameter was too small for significant oscillation. The Kamioka Liquid AntiNeutrino Detector (KamLAND) experiment was placed much further (~ 180 km) from the reactors that were its source of antineutrinos. In 2005, KamLAND reported not only a depletion of the ν_e flux but also an energy dependence, the latest version of this analysis is shown in Fig 2.2, that is characteristic of neutrino oscillations [67].

By this time neutrino flavour change was also well accepted as the explanation of the atmospheric neutrino problem. The iron tracking calorimeter experiment Soudan-2 published results in 1996 [69] which were in agreement with the IMB and Kamiokande experiments. This result ruled out systematic problems affecting just water-Čerenkov experiments. The most convincing evidence for oscillation came from the Super-Kamiokande experiment in 1998 [70]. This giant successor to the Kamiokande experiment was around 30 times larger at 22.5 kt, providing more statistics and higher energy containment (for more on the Super-Kamiokande experiment see section 3.3.1). Super-Kamiokande’s most recent physics results [71], before closing for its stage IV upgrade, are shown in Fig 2.3: the interaction rates for electron and muon type neutrinos, separated into two ranges of energy and binned as a function of the cosine of the zenith angle ϕ_z . Neutrinos coming from above have $\cos \phi_z = 1$, travelling around 30 km before passing through the detector. Those arriving from below, $\cos \phi_z = -1$, travel around 13,000 km before interacting in the detector. Any zenith angle dependence on neutrino flux can therefore be directly related to the distance travelled by the neutrino. The results are well fitted by the L/E dependence of oscillations. An analysis by the Super-Kamiokande collaboration [72], that uses only events which have a good resolution in L/E , strongly favours the oscillation model.

2.5 Neutrino Oscillation Physics

With oscillations accepted as the most plausible mechanism for lepton flavour change, the task was now to measure the mixing parameters.

Electron (anti)neutrinos undergoing oscillation in both solar and reactor experiments are not energetic enough to produce muons or taus if oscillation occurs. This means that such experiments can only measure the depletion of electron neutrinos into some second generic unknown flavour x which is a weighted mean of the other neutrino flavours. Atmospheric experiments view a depletion of muon-like neutrinos but do not observe an increase in electron flavour neutrinos. The analysis of data from these experiments is therefore often performed using the simplification of just two neutrino flavours and two neutrino mass eigenstates.

The two-neutrino version of the vacuum oscillation survival probability, equation 2.22, employed in reactor and atmospheric data analysis,

$$P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - \sin^2 2\theta \sin^2 \left(1.27 \Delta m^2 \frac{L}{E} \right), \quad (2.23)$$

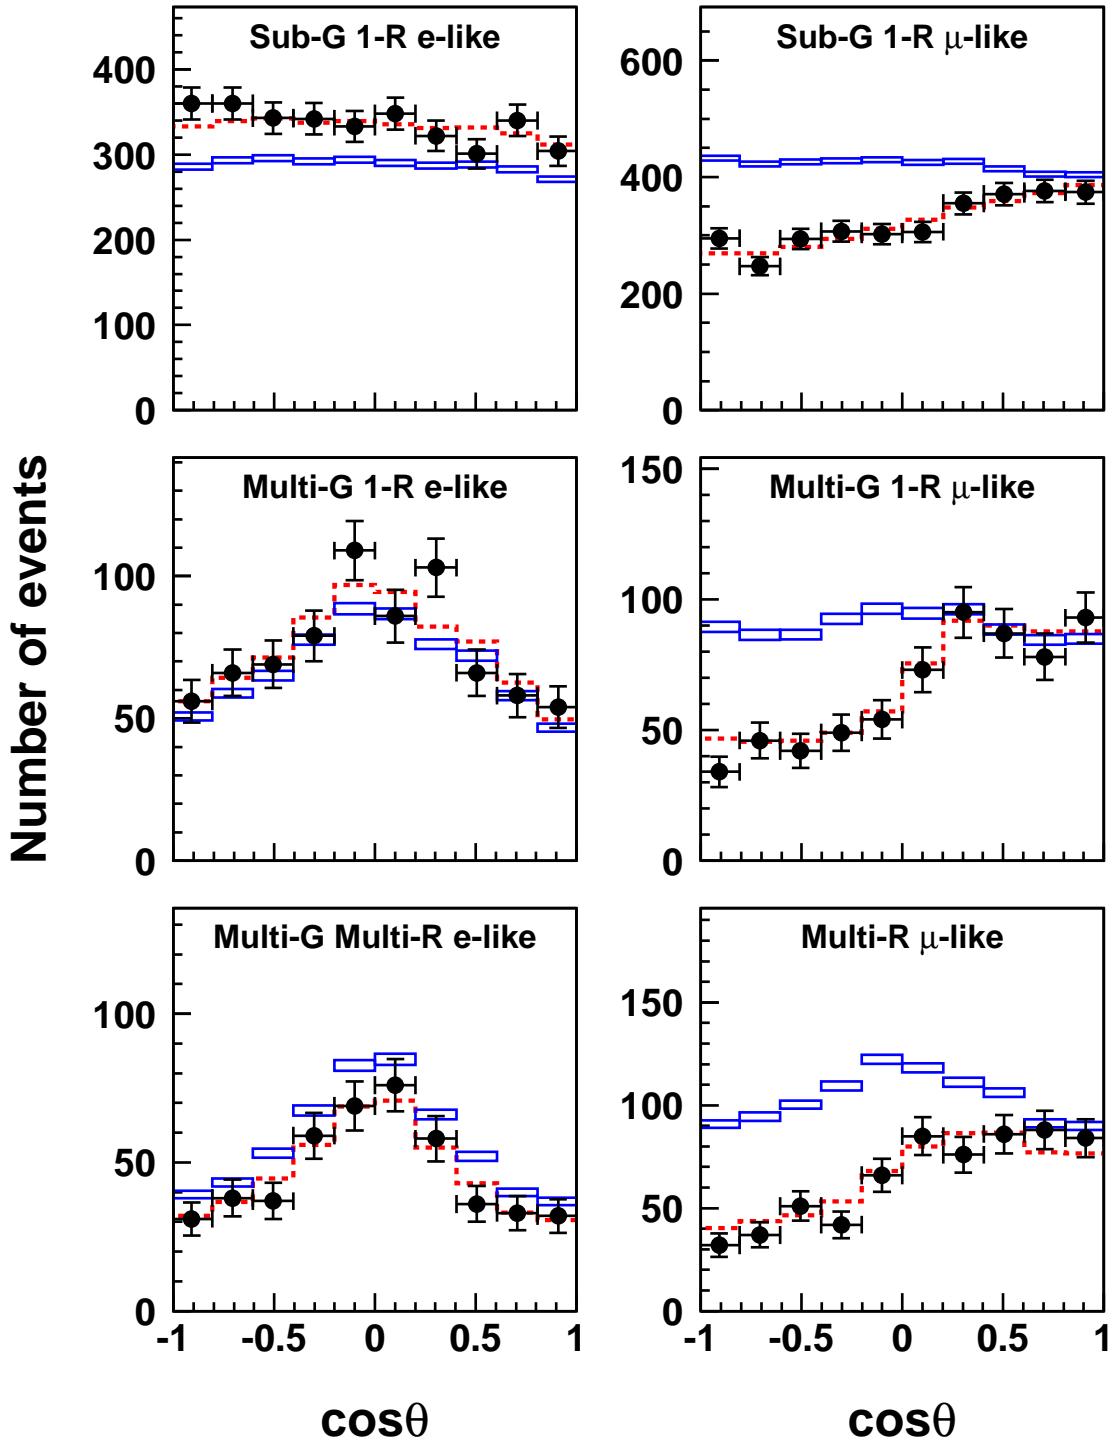


Figure 2.3: Number of single ring sub-GeV, single ring multi-GeV and multi-ring multi-GeV (from top to bottom) for both e-like (left) and μ -like (right) events, reproduced from [71]. The black points are the data, blue boxes show the expected result from no oscillation and the red dashed line the best fit oscillation scenario.

arises from the much simplified 2D rotation matrix

$$U_2 = \begin{pmatrix} \nu_i & \nu_j \\ \nu_\alpha & \cos \theta & \sin \theta \\ \nu_\beta & -\sin \theta & \cos \theta \end{pmatrix}. \quad (2.24)$$

Solar Neutrinos and Matter Effects

In the solar regime measurement of the parameters $(\theta_\odot, \Delta m_\odot^2)$ come from several solar neutrino experiments: gallium and chlorine radiochemical, early proton decay experiments, Super-Kamiokande and SNO. The same parameters are also measured, assuming \mathcal{CPT} invariance, by the long baseline reactor experiment KamLAND. All of these contributions can be seen on Fig 2.4(a).

The analysis of data from oscillation of solar neutrinos has an extra level of complexity from the neutrinos' passage through the solar matter. As neutrinos pass through matter they can coherently forward scatter off particles they encounter by the exchange of weak bosons. These interactions can significantly modify their propagation [74] which can result in a rather different probability from that seen in the vacuum oscillations [75]. Flavour change that occurs in matter arises from the interplay between flavour-nonchanging neutrino-matter interactions and neutrino mass and mixing which we have been discussing. This phenomenon is known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect, after pioneers of the idea [74, 75], which I will now develop in the two-neutrino framework.

As mentioned earlier in deriving the approximate form of the propagation factor one can describe the propagation of an ultrarelativistic neutrino via a Schrödinger-like equation.

$$\mathcal{H} \begin{pmatrix} \nu_e \\ \nu_x \end{pmatrix} = i \frac{\partial}{\partial t} \begin{pmatrix} \nu_e \\ \nu_x \end{pmatrix}. \quad (2.25)$$

The effective Hamiltonian in the equation, a matrix \mathcal{H} in neutrino flavour space, differs from its vacuum counterpart by the addition of interaction potentials arising from the coherent forward neutrino-scattering. Both ν_e and ν_x can interact via the exchange of neutral Z^0 bosons with all types of baryonic matter. When one is considering only the possibility of transitions to active neutrino flavours, i.e. that sterile neutrinos do not exist, then this common interaction potential merely adds a multiple of the identity matrix to the Hamiltonian. Such an addition has no effect on flavour transitions and so we will not include this in the following discussion. The ν_e , however, has additional

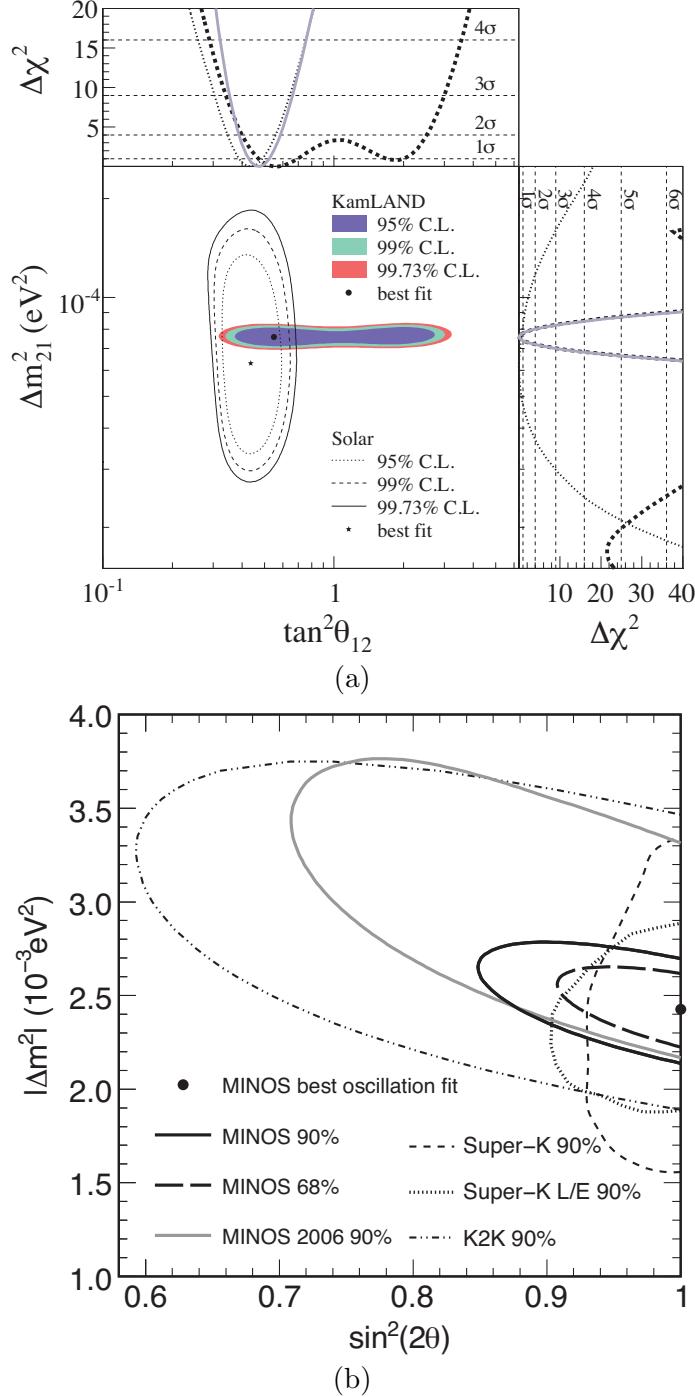


Figure 2.4: These plots show the region in which the two-neutrino solar (a) and atmospheric (b) mixing parameters lie. (a) is reproduced from the latest KamLAND data release [68] and shows the sensitivity of solar experiments to the mixing angle and KamLAND's sensitivity to the mass splitting. (b) is reproduced from the latest MINOS data release [73] and also shows the constraints from the Super-K atmospheric and K2K accelerator neutrino experiments.

W^\pm charged current interactions possible with electrons with a potential

$$V = \sqrt{2}G_F N_e. \quad (2.26)$$

Here G_F is the Fermi constant, still used in these low energy interactions, and N_e is the number density of electrons (per unit volume). We can then write down a Hamiltonian which describes the interplay of neutrino vacuum oscillation and matter effects. We do this by adding this matter Hamiltonian, which we write here as a function of radial distance from the centre of the Sun, to the vacuum oscillation Hamiltonian.

$$\begin{aligned} \mathcal{H} &= \mathcal{H}_V + \mathcal{H}_M \\ &= \frac{\Delta m_\odot^2}{2E} \begin{bmatrix} \sin^2 \theta_\odot & -\sin \theta_\odot \cos \theta_\odot \\ -\sin \theta_\odot \cos \theta_\odot & \cos^2 \theta_\odot \end{bmatrix} + \begin{bmatrix} V(r) & 0 \\ 0 & 0 \end{bmatrix} \\ &= \frac{\Delta m_\odot^2}{2E} \begin{bmatrix} \epsilon + \sin^2 \theta_\odot & -\sin \theta_\odot \cos \theta_\odot \\ -\sin \theta_\odot \cos \theta_\odot & \cos^2 \theta_\odot \end{bmatrix}. \end{aligned} \quad (2.27)$$

Here we have written the interaction potential as a function of radius from the centre of the Sun as $V(r) = \sqrt{2}G_F N_e(r)$ and in the last step I introduce the substitution $\epsilon = 4EV(r)/\Delta m_\odot^2$ for ease of notation.

To obtain analogous matter oscillation parameters Δm_m^2 and θ_m we must rearrange the above Hamiltonian to fit the original form

$$\mathcal{H} = \frac{\Delta m_m^2}{4E} \begin{bmatrix} \sin^2 \theta_m & -\sin \theta_m \cos \theta_m \\ -\sin \theta_m \cos \theta_m & \cos^2 \theta_m \end{bmatrix}. \quad (2.28)$$

We first diagonalise the matrix in equation 2.27 and solve the resulting quadratic equation to get the eigenvalues $\lambda_\pm = \lambda_0 \pm \delta\lambda/2$, where

$$\lambda_0 = \frac{1+\epsilon}{2} \quad \delta\lambda = \sqrt{[\epsilon - \cos 2\theta_\odot]^2 + \sin^2 2\theta_\odot}. \quad (2.29)$$

As in the vacuum oscillation case, the mass squared splitting arises from the difference in eigenvalues and so $\Delta m_m^2 = \Delta m_\odot^2 \delta\lambda$. To get an expression for θ_m we equate the off-diagonal terms of the matrices in equations 2.27 and 2.28 to get $\delta\lambda \sin 2\theta_m = \sin 2\theta_\odot$. Substituting in for $\delta\lambda$ in these expressions gives

$$\Delta m_m^2 = \Delta m_\odot^2 \sqrt{[\epsilon - \cos 2\theta_\odot]^2 + \sin^2 2\theta_\odot} \quad (2.30)$$

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta_\odot}{[\epsilon - \cos 2\theta_\odot]^2 + \sin^2 2\theta_\odot}. \quad (2.31)$$

Two-neutrino oscillation in a vacuum cannot distinguish between a mixing angle θ_\odot and an angle $\theta'_\odot = \pi/2 - \theta_\odot$. These two solutions for the mixing angle represent physically different situations: looking at the 2D mixing matrix, one can see that if $\theta_\odot < \pi/4$ then the lightest mass eigenstate, denoted ν_1 , is predominantly ν_e while the second θ'_\odot solution is exactly the converse with the predominant flavour in ν_1 being ν_x . The addition of matter effects destroys this 2 solution degeneracy. For electron neutrinos the interaction potential is of definite positive sign. This means that the $\nu_e - \nu_e$ element of the MSW Hamiltonian has two different sizes when the mixing angle is θ_\odot or $\theta'_\odot = \pi/2 - \theta_\odot$, which results in different flavour contents emanating from the Sun.

Let us think of the oscillation of an electron neutrino from production in the Sun at $r \simeq 0$, where the electron number density is of the order of $N_e \simeq 6 \times 10^{25}/\text{cm}^3$, which yields $V(r \simeq 0) = 0.75 \times 10^{-5} \text{ eV}^2/\text{MeV}$. We consider ν_e production in this inner solar region from ${}^8\text{B}$ decay: these neutrinos are the most dominant fraction at the high energy end of the solar spectrum $E \sim 6 - 7 \text{ MeV}$ [48] and are therefore seen by all classes of solar neutrino experiment. In this region \mathcal{H}_M dominates over \mathcal{H}_V meaning that to a first approximation $\mathcal{H}(r \simeq 0)$ is diagonal. This means that a ${}^8\text{B}$ neutrino is born not only in a ν_e flavour eigenstate but also in an MSW mass eigenstate (an eigenstate of the Hamiltonian $\mathcal{H}(r \simeq 0)$). As $V > 0$, the neutrino is born into the heavier of the two eigenstates. As the neutrino journeys out of the Sun the electron density changes slowly enough that we may solve Schrödinger's equation in a piecewise fashion and patch the solutions. This adiabatic propagation leaves the neutrino in the same MSW mass eigenstate throughout its solar journey until it emerges from the Sun as the heavier ν_2 \mathcal{H}_V eigenstate,

$$\nu_2 = \nu_e \sin \theta_\odot + \nu_x \cos \theta_\odot. \quad (2.32)$$

As this is an eigenstate of the vacuum Hamiltonian, the neutrino remains as ν_2 until detection on Earth. The probability therefore that a ${}^8\text{B}$ ν_e from the Sun be detected as an electron neutrino is simply

$$P(\nu_e \rightarrow \nu_e) = \sin^2 \theta_\odot. \quad (2.33)$$

This is true only for those neutrinos which are produced sufficiently deep in the Sun that the electron density is above the critical density $N_{e,crit}$.

Reactor neutrino experiments have relatively short baselines through the relatively transparent earth and are therefore void of any such matter effects. Reactor experiments use the vacuum oscillation framework to extract their parameters.

The determination of Δm_{\odot}^2 is dominated by the KamLAND reactor neutrino experiment, which at the beginning of 2008 published results [68] that increased its volume of data almost four-fold, due to a longer live time coupled with an increased fiducial volume. SNO is the main contributor to the determination of θ_{\odot} due to its unique neutral current channel. SNO also recently published the results of its last phase [76] in which the neutral current flux was determined, not only from the CC correlated 6.25 MeV delayed photon from neutron capture, but also from the uncorrelated measurements of the proportional counting ${}^3\text{He}$ neutral current detectors. Super-Kamiokande and SNO can measure the directionality and energy of neutrino events, which allows them to probe the effect of different baselines and energies upon the survival probability. With the ability to measure the shape of the oscillated energy spectrum in the absence of the MSW effect, KamLAND can precisely measure the mass splitting Δm_{\odot}^2 . In their latest data release [76] the SNO collaboration also perform a two neutrino analysis of all solar data and calculate world average values of²

$$\Delta m_{\odot}^2 = (7.59^{+0.19}_{-0.21}) \times 10^{-5} \text{ eV}^2, \quad \theta_{\odot} = (34.4^{+1.3}_{-1.2})^{\circ} \quad (2.34)$$

Atmospheric Neutrinos and Mass Hierarchies

Atmospherically produced neutrinos have been measured by the early proton decay experiments, the previously mentioned Super-Kamiokande and Soudan-2 experiments and the MACRO experiment [77]. These experiments are primarily sensitive to the mixing angle θ_{atm} . A complementary approach has been taken by the K2K [78] and MINOS [73] experiments, which use beams of accelerator-produced muon neutrinos from in-flight pion decay. In each experiment, far detectors are positioned at a moderate/long baseline and the energy of the beam selected to maximise the flavour oscillation. The range of energies in the beam and the fixed distance that the neutrinos travel allow these experiments to build up an accurate dependence of the oscillation on L/E , and hence they are more sensitive to Δm_{atm}^2 . The results from the latest MINOS data release [73] are

$$|\Delta m_{atm}^2| = (2.43 \pm 0.13) \times 10^{-3} \text{ eV}^2, \quad \sin^2(2\theta_{atm}) > 0.9 \text{ (90\% C.L.)} \quad (2.35)$$

Their measurement of the mixing angle θ_{atm} is consistent with that measured by all previous experiments, Fig 2.4(b). The mass squared difference quoted here is unsigned: as mentioned previously, the sign of the mass squared splitting can only be obtained from neutrinos that have been affected by matter effects. Extra long baselines are in

²unless otherwise stated the errors on all values are $\pm 1\sigma$ or 68% C.L.

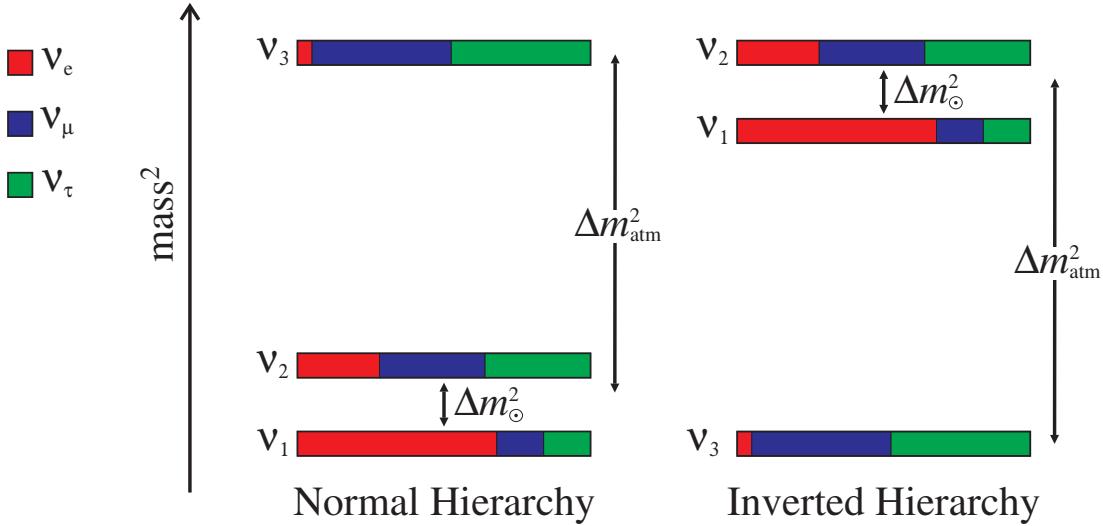


Figure 2.5: Normal and inverted neutrino mass hierarchies. The flavour content of each eigenstate is shown, approximately reflecting current data, with the addition that U_{e3} has been set to $0.05 \exp^{i\pi/2}$

design for future neutrino, β -beam or neutrino factory, experiments.

The separate two-neutrino analyses of these regimes give two mass squared splittings, $|\Delta m_{atm}^2| \gg |\Delta m_{\odot}^2|$. As well as the magnitudes of these splittings, the MSW-mediated oscillation of solar neutrinos has provided the sign of the solar mass squared splitting from the ability to probe the exact flavour content of a specific mass eigenstate. The sign of the atmospheric mass squared difference has yet to be measured. Within a three neutrino framework this leaves only two possible hierarchies of masses. If we associate the solar mass squared splitting to that between mass eigenstates 1 and 2, $\Delta m_{12}^2 \simeq \Delta m_{\odot}^2$, and the atmospheric splitting to that between eigenstates 2 and 3 we can write the two possible hierarchies as $m_3 \gg m_2 > m_1$, which we call the normal hierarchy, and $m_2 > m_1 \gg m_3$ which we call the inverted hierarchy. The normal hierarchy is so named as it follows the same pattern seen in the other Standard Model particles of increasing mass differences with each generation. The inverted is named so as the lightest mass eigenstate consists primarily of the heavier charged lepton flavour eigenstates μ and τ .

Another neutrino oscillation regime has also been suggested in results from the Liquid Scintillator Neutrino Detector (LSND) experiment. This regime does not fit with those mentioned previously, within a three neutrino framework. It suggests a much larger $|\Delta m_{LSND}^2| \simeq 1 \text{ eV}^2$ splitting between mass eigenstates, which would require at least one extra sterile neutrino as $|\Delta m_{12}^2| + |\Delta m_{23}^2| \ll 1 \text{ eV}^2$. LSND reported seeing $\bar{\nu}_e$

appearance in a $\bar{\nu}_\mu$ beam. Analysis of data from several experiments in the same regime as LSND, notably KARMEN2 [79], and NOMAD [80], failed to confirm this result, and most recently MiniBooNE also saw no excess in the predicted region [81]. A recent analysis of the principal relevant experiments [82] finds only a 3.94% probability that the results are consistent with two-neutrino oscillations. It is therefore standard practice to neglect the LSND result in neutrino oscillation analyses, and this is the position taken for the remainder of this thesis.

2.6 Three Neutrino Mixing and The Future

To date the treatment of solar and atmospheric neutrino oscillations within separate two neutrino frameworks has proven successful. The significant difference in values of the mass squared splitting meant that to first approximation if the L/E of an experiment was optimised to one regime then the effect of oscillation from the other would be negligible. The maximal nature of θ_{atm} associated with $\nu_\mu \rightarrow \nu_\tau$ oscillation constrains the value of other possible mixings to be small. These approximations allow the deconvolution of the complicated three neutrino formalism into the two entirely separate two neutrino frameworks described in the previous section.

The solar and atmospheric oscillation parameters have both been measured and the regions in which their true values can lie are ever shrinking, with each new experiment yielding increasingly accurate and complementary measurements. In the near future this accuracy will approach the scale of effects previously neglected. We must therefore begin to analyse solar and atmospheric sector data within a full three neutrino framework. It is also clear that such a treatment is required if one hopes to determine the final small unmeasured mixing angle and begin to probe \mathcal{CP} violation in the lepton sector.

The full three neutrino appearance probability, equation 2.20, is not useful for practical purposes because of the complexity of many of the resulting expressions. Fortunately approximations can still be facilitated by knowledge of oscillation parameter properties presented to us by two-neutrino analyses.

Adopting the notation used above for mass hierarchies, the solar mass squared difference is between the ν_1 and ν_2 mass eigenstates, $\Delta m_{12}^2 = \Delta m_\odot^2$, and the atmospheric mass squared difference is a weighted mean of the two other mass squared differences, $\Delta m_{31}^2 \simeq \Delta m_{23}^2 \simeq \Delta m_{atm}^2$. This assumption rests upon the ability to neglect the spread of the possible mass squared splittings, which are similar to within $\alpha = \Delta m_\odot / \Delta m_{atm}^2 \simeq 10^{-2}$.

Therefore, the two-neutrino oscillation formalism introduced in section 2.5 is a very good approximation for current experiments with errors greater than this scale factor α . When considering neutrino oscillations within a three neutrino framework it is inevitable that any experiment with L/E optimised for the 23 mixing will also be optimised for 13 mixing. The 23 oscillation is known to involve the ν_μ and ν_τ flavours primarily, so it is not a problem for short baseline reactor experiments, e.g. CHOOZ and Palo Verde, as they are looking at the disappearance of $\bar{\nu}_e$. For atmospheric or accelerator produced muon neutrinos, however, the 13 oscillation will only be seen as a second order effect with a magnitude the order of $\alpha \simeq 10^{-2}$. This is the case for T2K, the experiment considered in this thesis.

The atmospheric mixing is found to be almost entirely associated with $\nu_\mu \rightarrow \nu_\tau$, which means that it represents just a single mixing angle, so we approximate $\theta_{23} \simeq \theta_{atm}$. To obtain the two-neutrino framework one would neglect oscillation effects from the other mass squared difference as either below the resolution of the experiment, Δm_{12}^2 in an atmospheric experiment, or constant over the range of the experiment and therefore having no effect on oscillation probability, Δm_{atm}^2 in solar experiments.

However, the increasing precision of 12 and 23 oscillation measurements implies that the small inaccuracy introduced by using the two-neutrino formalism will become significant compared to the experimental error. Therefore, most θ_{13} experiments require the full three-neutrino formalism of equation 2.20, or a suitable approximation to this, as will any oscillation experiment aiming for very high precision. Global analyses, e.g. Schwetz 2008 [1], also use the three-neutrino formalism to determine values for θ_{12} and θ_{23} , for overall consistency.

When one interprets solar, atmospheric, reactor and accelerator neutrino data in a three neutrino framework, a number of constraints are placed upon the unknowns. Starting from the full three neutrino oscillation, equation 2.20, one can distil the dominant oscillation terms present for each class of detector. There are many three neutrino oscillation analyses performed each with just small nuances in the methods.

Reactor 3ν Mixing

We first discuss the full three neutrino consideration of reactor neutrino experiments, as they provide essential limits that will be used when considering the meaning of data from other classes of experiment. Reactor experiments are looking for a disappearance

of $\bar{\nu}_e$, for which the full probability in the three neutrino oscillation picture is

$$\begin{aligned}
 P(\bar{\nu}_e \rightarrow \bar{\nu}_e) &= 1 - 4 \left[\begin{aligned} &|U_{e1}^*|^2 |U_{e2}^*|^2 \sin^2 \left(1.27 \Delta m_{12}^2 \frac{L}{E} \right) \\ &+ |U_{e1}^*|^2 |U_{e3}^*|^2 \sin^2 \left(1.27 \Delta m_{13}^2 \frac{L}{E} \right) \\ &+ |U_{e2}^*|^2 |U_{e3}^*|^2 \sin^2 \left(1.27 \Delta m_{23}^2 \frac{L}{E} \right) \end{aligned} \right] \\
 &= 1 - \cos^4 \theta_{13} \sin^2 2\theta_{12} \sin^2 \left(1.27 \Delta m_{12}^2 \frac{L}{E} \right) \\
 &\quad - \sin^2 2\theta_{13} \sin^2 \theta_{12} \sin^2 \left(1.27 \Delta m_{23}^2 \frac{L}{E} \right) \\
 &\quad - \sin^2 2\theta_{13} \cos^2 \theta_{12} \sin^2 \left(1.27 \Delta m_{31}^2 \frac{L}{E} \right). \tag{2.36}
 \end{aligned}$$

The short baseline experiments of CHOOZ and Palo Verde lie in a regime where $\Delta m_{atm}^2 L/E \sim 1$, which from two neutrino data implies $\Delta m_{12}^2 L/E \ll 1$ allowing us to neglect the first term. The resulting oscillation is then a combination of the terms in Δm_{13}^2 and Δm_{23}^2 . Replacing these mass squared differences with the weighted mean $\Delta m_{13}^2 \simeq \Delta m_{23}^2 \simeq \Delta m_{atm}^2$ the expression reduces to

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) \simeq 1 - \sin^2(2\theta_{13}) \sin^2 \left(1.27 \Delta m_{atm}^2 \frac{L}{E} \right). \tag{2.37}$$

The null oscillation results from these experiments provide the best constraints upon θ_{13} . The sensitive nature of this extremely small mixing angle and the comparatively large errors of current data means that small nuances in analysis techniques yield different limits. An analysis of data released at Neutrino 2008, by Fogli et al [83], reports 90% confidence level values of θ_{13} that for the first time exclude zero. A non-zero value of θ_{13} is desirable if neutrino oscillations are to determine the \mathcal{CP} violation in the lepton sector, as the violating phase δ is always coupled with the sine of this angle. Most other analyses however do not exclude a zero value and show best fit values that are consistent with zero, $\sin^2 \theta_{13} = 0.010^{+0.016}_{-0.011}$ [1].

Superbeams and Current Estimates of θ_{13}

The accuracy to which we measure the mixing parameters will continue with the next generation of superbeam experiments, T2K (see chapter 3) and NO ν A [84], and reactor neutrino experiments. Superbeams will use the same principle of neutrino production as K2K and MINOS but with much higher power proton beams and pion focusing horns. The high power beams will also allow experiments to take an off-axis configuration whilst still retaining an appreciable neutrino flux. In these configurations the axis

of the neutrino beam is directed at a small angle to that in which the detectors lie. From the decay kinematics of the pion, this results in a much more monochromatic beam, which gives a better defined L/E (see Fig 3.1) and therefore oscillation signal. The increased luminosity of these experiments will see the volume of neutrino event data explode, offering around an order of magnitude improvement in measurements of atmospheric mixing parameters over the next five years. The increased statistics also bring hope of directly measuring subdominant oscillation appearance of ν_e from a beam of $> 99\% \nu_\mu$, the probability of which is

$$P(\nu_\mu \rightarrow \nu_e) = 4 \left[\begin{array}{l} U_{e1} U_{\mu 1}^* U_{e2}^* U_{\mu 2} \sin^2(1.27 \Delta m_{12}^2 \frac{L}{E}) \\ + U_{e1} U_{\mu 1}^* U_{e3}^* U_{\mu 3} \sin^2(1.27 \Delta m_{13}^2 \frac{L}{E}) \\ + U_{e2} U_{\mu 2}^* U_{e3}^* U_{\mu 3} \sin^2(1.27 \Delta m_{23}^2 \frac{L}{E}) \end{array} \right] \quad (2.38)$$

Superbeam experiments have been designed to optimise their sensitivity to the atmospheric parameters by measuring the flux at the first atmospheric oscillation maximum. Here $\Delta m_{23}^2 L/E \sim 1$ and $\Delta m_{12}^2 L/E \ll 1$. Also, assuming that the value of θ_{13} is sufficiently large, we can neglect the first Δm_{12}^2 term and use orthogonality, $\sum_i U_{ei} U_{\mu i}^* = 0$, to reduce the above equation to

$$\begin{aligned} P(\nu_\mu \rightarrow \nu_e) &\simeq 4|U_{e3}|^2 |U_{\mu 3}|^2 \sin^2\left(1.27 \Delta m_{atm}^2 \frac{L}{E}\right) \\ &= \sin^2(\theta_{23}) \sin^2(\theta_{13}) \sin^2\left(1.27 \Delta m_{atm}^2 \frac{L}{E}\right) \end{aligned} \quad (2.39)$$

One can see from this probability that superbeams are going to provide complementary measurements that will improve the limit on θ_{13} two fold: first from a more accurate measurement of the atmospheric parameters, but also from direct detection of this rare process. On their own, measurements of $P(\nu_\mu \rightarrow \nu_e)$ from these experiments cannot uniquely determine the value of θ_{13} or any of the other parameters. This is because there are degenerate combinations of parameters that produce the same oscillation probability in a given experiment. One potential degeneracy can be seen in the equation: for a non-maximal value of $\sin^2(\theta_{23})$ there are two possible values for θ_{23} and hence for a measured value of $P(\nu_\mu \rightarrow \nu_e)$ there are (usually) two possible values of θ_{13} from $\sin^2(\theta_{13})$. Two-fold degeneracies also arise due to the unknown mass hierarchy (sign of Δm_{31}^2) and between different combinations of θ_{13} and δ . Therefore any one experiment in isolation may have an eightfold degeneracy when trying to extract parameter values. The combination of multiple experiments (T2K & NO ν A & Reactor) is very important as it can help to eliminate these degeneracies.

Determinations of θ_{13} in a three-neutrino framework have been performed with recent data by a handful of groups, of which I will focus on [1] and [85] to demonstrate the extremes. All groups split the analysis into solar and atmospheric sector data before combining results to obtain a global best fit for θ_{13} . Although the most strict constraints arise from atmospheric sector short baseline reactor experiments (CHOOZ and Palo Verde) analyses of solar neutrino data yield a non-negligible constraint upon the value of θ_{13} . This is because of the difference in correlation of the matter-affected solar neutrino oscillation and the vacuum oscillation of the KamLAND experiment. Both analyses mentioned show similar best fit values, of $\sin^2 \theta_{13} = 0.021$ and $\sin^2 \theta_{13} = 0.03$, consistent with $\theta_{13} > 0$ at the 1.2σ and 1.5σ level for [85] and [1] when analysing solar and recent KamLAND data, .

The major discrepancies between these analyses arise when considering the effect of atmospheric sector data constraints upon θ_{13} . The analysis of Schwetz et al. [1] neglects the effects of Δm_{12}^2 terms when determining a limit on θ_{13} . They choose to do this as these sub-leading effects are still significantly smaller than current experimental statistical and systematic errors. Their minimum χ^2 fit, from the inclusion of both their solar and atmospheric analyses, gives $\sin^2 \theta_{13} = 0.01$, but with $\theta_{13} = 0$ allowed at 0.9σ : thus there is no significant evidence of a non-zero θ_{13} .

The analysis performed by Fogli et al.[83] includes the effects of the sub-leading Δm_{12}^2 terms [85]. They obtain a slightly higher $\sin^2 \theta_{13} = 0.012 \pm 0.013 (1\sigma)$ in combination of atmospheric, K2K, MINOS and CHOOZ data, which is an insignificant $\sim 0.9\sigma$ hint of $\theta_{13} > 0$. But when combined with their solar data constraint of $\sin^2 \theta_{13} = 0.021 \pm 0.017 (1\sigma)$, this gives a global best fit of $\sin^2 \theta_{13} = 0.016 \pm 0.010 (1\sigma)$ which gives a hint of $\theta_{13} > 0$ at the 1.6σ level.

The results of analysing global data in a full three neutrino framework from [1] are shown in the table below.

Parameter	Best Fit	2σ	3σ
$\Delta m_{12}^2 [\times 10^{-5} eV^2]$	$7.65^{+0.23}_{-0.2}$	$7.25 - 8.11$	$7.05 - 8.34$
$ \Delta m_{13}^2 [\times 10^{-3} eV^2]$	$2.40^{+0.12}_{-0.11}$	$2.18 - 2.64$	$2.07 - 2.75$
$\sin^2 \theta_{12}$	$0.304^{+0.022}_{-0.016}$	$0.27 - 0.35$	$0.25 - 0.37$
$\sin^2 \theta_{23}$	$0.50^{+0.07}_{-0.06}$	$0.39 - 0.63$	$0.36 - 0.67$
$\sin^2 \theta_{13}$	$0.01^{+0.016}_{-0.011}$	≤ 0.040	≤ 0.056

Table 2.2: Best fit values with 1σ errors and 2σ and 3σ confidence level intervals. These results are from a fit of global data from solar, atmospheric, reactor and accelerator experiments from [1].

T2K

The relatively short baseline (travel time) of T2K, through the relatively electron rarified crust of the Earth means that vacuum oscillation terms in the Hamiltonian dominate and matter effects can be neglected when considering $P(\nu_\mu \rightarrow \nu_e)$. If one includes sub-leading effects of the Δm_{21}^2 terms the vacuum probability equation 2.38 becomes

$$\begin{aligned}
 P(\nu_\mu \rightarrow \nu_e) \simeq & \sin^2 2\theta_{13} \sin^2 \theta_{23} \sin^2 \Delta \\
 & \pm \alpha \sin 2\theta_{13} \cos \theta_{13} \sin \delta \sin 2\theta_{12} \sin 2\theta_{23} \sin^3 \Delta \\
 & - \alpha \sin 2\theta_{13} \cos \theta_{13} \cos \delta \sin 2\theta_{12} \sin 2\theta_{23} \cos \Delta \sin 2\Delta \\
 & + \alpha^2 \cos^2 2\theta_{23} \sin^2 2\theta_{12} \sin^2 \Delta
 \end{aligned} \tag{2.40}$$

where $\alpha \equiv \Delta m_{12}^2 / \Delta m_{31}^2$ and $\Delta = \Delta m_{31}^2 L / 4E$. The existing solar and atmospheric sector data show that (section 2.5) $\alpha \simeq 3 \times 10^{-2}$. Choosing to dismiss the θ_{13} limits from the atmospheric sector data, because of current disagreement, CHOOZ + solar neutrino data indicate that $\sin^2 \theta_{13} \sim 0.02$, and of course the value of δ is completely unknown. Aside from α and $\sin 2\theta_{13}$ factors, the L , E and current measurements of other oscillation parameters indicate that all other factors in this expansion are ~ 1 . Looking at this approximate formula, the only term that doesn't depend on θ_{13} is the last one (which is in fact the approximate form for the oscillations that KamLAND sees), but this depends on α^2 and is therefore practically unobservable in T2K. For values of θ_{13} within an order of magnitude of its current upper limit, the first term dominates. This would describe a relatively simple oscillatory appearance of ν_e in ν_μ beam.

This then sets the targets for the next generation of neutrino oscillation experiments: firstly, an improvement in the sensitivity to θ_{13} by at least an order of magnitude in hope of a measurement, Fig 2.6; secondly, the measurement of θ_{23} and Δm_{23}^2 to a greater precision, which is interesting in its own right, not only for extracting the values of unknown parameters but also to test whether θ_{23} is indeed maximal mixing, which is of great interest to model builders.

Neutrino Factories and β -Beams

With a measurements of θ_{13} there still remain the open questions in neutrino physics: What is the magnitude of δ , the lepton \mathcal{CP} violating phase? What hierarchy do the neutrino mass eigenstates obey, what is the sign of Δm_{23}^2 ? Both of these questions can

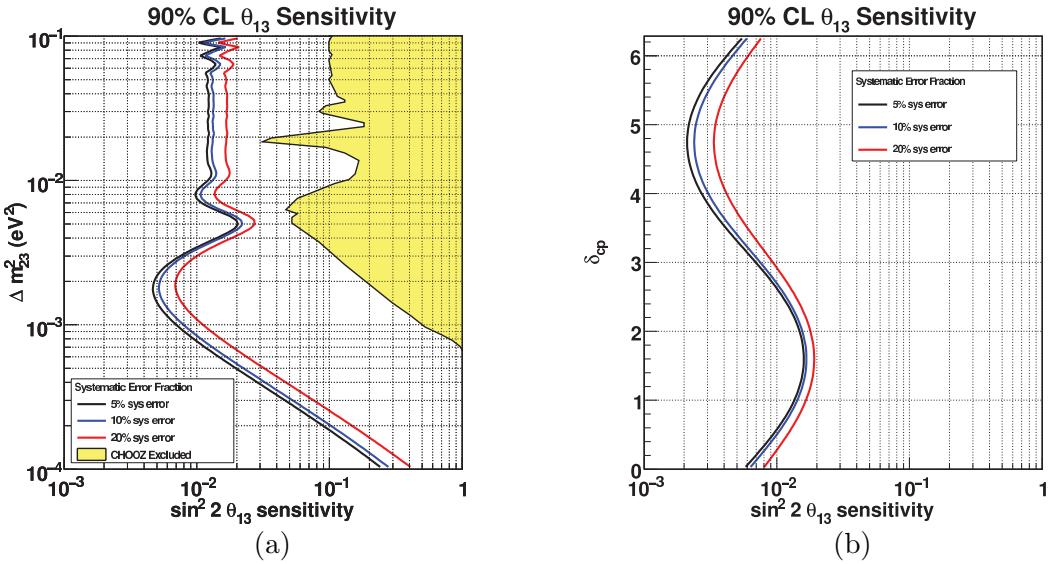


Figure 2.6: (a) Sensitivity of T2K to θ_{13} as a function of Δm_{23}^2 . The yellow region has been excluded to 90% CL by the CHOOZ reactor experiment. (b) Sensitivity to θ_{13} as a function of δ_{CP} . In both plots the beam is assumed to be running at 750 kW for 5 years and each display the 5% (black), 10% (blue) and 20% (red) systematic error fractions. Reproduced from [86]

only be probed with long baseline experiments in which matter effects are no longer negligible.

As shown in the two-neutrino case, matter effects modify the neutrino oscillation probability in a non-trivial way. Including matter effects, again including the sub-leading Δm_{21}^2 effects, it has been shown [87] that the appearance probability can be written as a perturbative expansion of the ratio $\alpha = \Delta m_{21}/\Delta m_{31}^2 \simeq 3 \times 10^{-2}$:

$$P(\nu_\mu \rightarrow \nu_e) \simeq \sin^2 2\theta_{13} T_1 - \alpha \sin 2\theta_{13} T_2 + \alpha \sin 2\theta_{13} T_3 + \alpha^2 T_4 \quad (2.41)$$

$$T_1 = \sin^2 \theta_{23} \frac{\sin^2 [(1-x)\Delta]}{(1-x)^2}, \quad (2.42)$$

$$T_2 = \sin \delta \sin 2\theta_{12} \sin 2\theta_{23} \sin \Delta \frac{\sin (x\Delta)}{x} \frac{\sin [(1-x)\Delta]}{(1-x)}, \quad (2.43)$$

$$T_3 = \cos \delta \sin 2\theta_{12} \sin 2\theta_{23} \cos \Delta \frac{\sin (x\Delta)}{x} \frac{\sin [(1-x)\Delta]}{(1-x)}, \quad (2.44)$$

$$T_4 = \cos^2 2\theta_{23} \sin^2 2\theta_{12} \frac{\sin^2 (x\Delta)}{x^2}. \quad (2.45)$$

Here $\Delta \equiv \Delta m_{31}^2 L / 4E$ is the kinematical phase of the oscillation, and the quantity $x \equiv 2\sqrt{2}G_F N_e E / \Delta m_{31}^2$ is a measure of the importance of MSW matter effects from ambient electrons along a beam's journey from source to detector through the earth.

The T_1 term represents the oscillation due to the atmospheric mass splitting scale. The T_4 term represents the oscillation due to the solar mass splitting scale. The T_2 and T_3 terms are the \mathcal{CP} -violating and \mathcal{CP} -conserving interference terms, respectively. The probability for the corresponding antineutrino oscillation, $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$, is the same as the probability $P(\nu_\mu \rightarrow \nu_e)$ given by the Eqns. above, but with the signs in front of both x and $\sin \delta$ reversed: both the matter effect and \mathcal{CP} violation lead to a difference between the neutrino and anti-neutrino oscillation probabilities.

In view of the dependence of x on Δm_{31}^2 , and in particular the sign dependence, the matter effect can reveal whether the neutrino mass spectrum is of normal or inverted hierarchy. However, to determine the nature of the spectrum, and to establish the presence of CP violation, it is necessary to disentangle the matter effect from CP violation in the neutrino-antineutrino oscillation probability difference that is actually observed. This will require complementary measurements from short baseline, MSW negligible, experiments, reactor experiments Double CHOOZ [88] and Daya Bay [89]; as well as the relatively short baseline T2K superbeam experiment; and long baseline experiments, the long baseline super beam NO ν A [84], upgrades of existing superbeams and future neutrino factories. These complementary measurements will be able to take advantage of the differing dependences on the matter effect and on CP violation.

Chapter 3

The Tokai to Kamioka T2K Experiment

3.1 Overview of the experiment

The Tokai-mura to Kamioka (T2K) experiment is designed to observe the phenomenon of neutrino flavour change, through the disappearance and appearance of specific neutrino type, and measure the parameters that characterise this oscillation. It is the first in the next generation of neutrino ‘superbeam’ experiments using high luminosity neutrino beams to reach new measurement accuracies. The neutrino beam used for T2K is generated at the newly built Japan Proton Accelerator Research Complex (J-PARC) in Tokai-mura on the east coast of Japan.

A suite of near detectors at J-PARC will characterise the neutrino beam, produced from the decay of pions in flight. Coarse detectors determine the beam topology (section 3.4.1) and a single highly granulated magnetised detector, the ND280 (section 3.5), will measure the spectra of electron and muon neutrinos and respective interaction cross-sections.

The characterised neutrino beam then proceeds on a 295 km journey to be sampled again in the upgraded Super-Kamiokande (Super-K) detector deep underground in the Kamioka mine in the west of Japan. With mean neutrino energy in the J-PARC beam of 0.7 GeV (section 3.2) the ratio L/E places T2K in a region extremely sensitive to atmospheric sector mixing, hence the oscillation probability $P(\nu_\mu \rightarrow \nu_\mu)$. The excellent particle identification capability of Super-K, coupled with the high intensity, relatively monochromatic beam from J-PARC, also makes T2K ideally suited to search for rare ν_e appearance events.

It is from these that the T2K experiment draws its two main physics goals. The first is an accurate determination of the mixing parameters Δm_{23}^2 and θ_{23} by observing a depletion of ν_μ events between near and far detectors. The second goal of the T2K experiment is the measurement of the subdominant θ_{13} oscillation parameter from the appearance of ν_e at the Super-K far detector from the predominantly ν_μ beam at the near detectors.

3.2 The J-PARC neutrino beam

The J-PARC complex is a newly built facility situated on the East coast of Japan, at Tokai-mura in the Ibaraki prefecture. It has a number of accelerators which produce proton beams of various energies for use in particle physics, nuclear, material, solid state and structural biology sciences. The linear accelerator (Linac) starts with ionised hydrogen and accelerates it to 400 MeV, injecting, at a rate of 50 Hz, into the Rapid-cycling Synchrotron (RCS), which increases the energy to 3 GeV with a repetition rate of 25 Hz. These protons finally enter the main ring Proton Synchrotron (PS) for a final jump in energy to 50 GeV¹.

The primary neutrino beam line at J-PARC, Fig 3.1c, will fast extract protons from the PS using a set of 5 kicker magnets. The extraction is split into spills lasting a few milliseconds which each contain multiple nanosecond bunches of protons. The proton beam is directed and focused onto a helium cooled, 30 mm diameter 100 mm length, cylindrical graphite target where interactions result in a multitude of charged pions, and a small fraction of kaons.

The target, which is the start point for the secondary beam line, is placed inside the first of 3 focusing horns which produce a magnetic field map designed to collimate the bulk movement of the produced particles. The focusing horns, run with a peak current of 320 kA, can be run in both positive and negative polarity, allowing the selection of π^\pm which then proceed into a decay volume. This volume is filled with inert He gas at 1 atm to minimise the capture upon nuclei of the in-flight pions, which subsequently decay $> 99\%$ of the time by $\pi^\pm \rightarrow \mu^\pm \nu_\mu (\bar{\nu}_\mu)$. The 110 m length decay volume was chosen to maximise the decay of pions whilst minimising the decay of muons. At the end of the decay volume there is a beam dump, of water cooled graphite blocks, which absorbs almost all remaining hadrons. Some high energy muons however penetrate to be detected by the muon monitor (MuMon), positioned just behind the beam dump, and some small distance into the rock beyond. The MuMon will be used to provide

¹The values quoted here are the design values which may take a number of years to achieve.

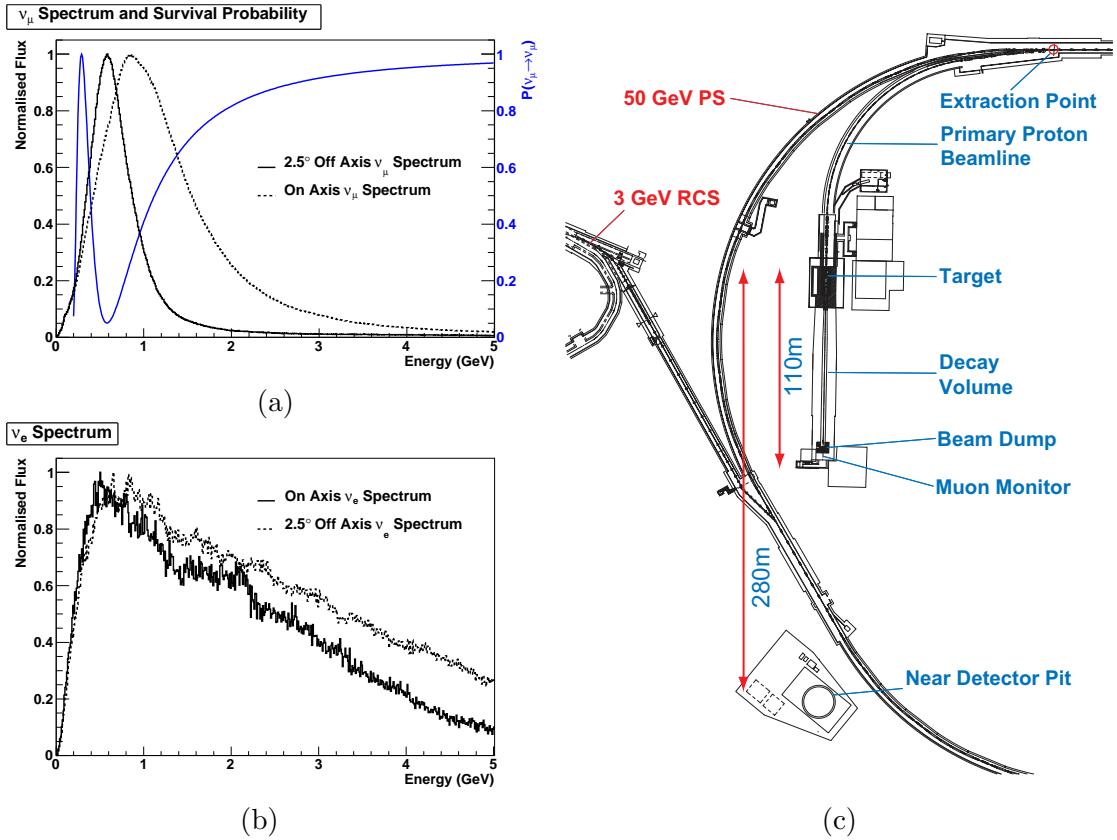


Figure 3.1: Peak normalised ν_μ (a) and ν_e (b) flux, (a) also plots the two-neutrino survival probability of ν_μ at 295 km. The JPARC neutrino beamline is shown in (c) along with sections of the 3 GeV RCS and 50 GeV PS. (c) is modified from [3]

bunch by bunch information on the targeting condition of the proton beam and the pion collection efficiency of the targeting horns.

The entire beam line produces a beam of $> 99\% \nu_\mu$, with a small ν_e contamination from the decay of kaons and muons. The beam centre is actually aimed 2.5° off the axis joining J-PARC and the Super-K far detector. Although this considerably reduces the luminosity, this off-axis configuration of the beam greatly improves the quality of the neutrino energy spectrum by sampling a different region of the pion decay space. The mean energy of the ν_μ beam, Fig 3.1a, reduces as one goes off axis, and at 2.5° the mean energy produces an L/E fraction that coincides with the first minimum in the $P(\nu_\mu \rightarrow \nu_\mu)$ oscillation probability. The spread in neutrino energy at off-axis angles is also reduced, most notably suppressing the high energy tail of neutrinos, which is a large contributor to the production of neutral pions, a major background to electron neutrino signatures in Super-K (section 3.3.2). These two effects give T2K the best

opportunity to not only view a large ν_μ disappearance but also a much more significant ν_e appearance signal.

3.3 Super-Kamiokande

The Kamiokande experiment was originally designed to test grand unified theories, by hoping to measure the proton decay lifetime, but instead found frontier physics in the background to this process; neutrino interactions. Kamiokande’s successor, Super-Kamiokande, was built with neutrino physics as its primary physics target. Since it began data taking on 1st April 1996 it has detected solar, atmospheric and accelerator produced neutrinos.

3.3.1 The Detector

The Super-Kamiokande detector (Fig 3.2 Left) is a water Čerenkov detector located under Mount Ikenoyama in the Gifu prefecture in western Japan. The 50 kt of high purity water is contained within a 42 m high, 36 m diameter, steel vessel. The tank is under 1 km of rock, 2700 m water equivalent (m.w.e.), which attenuates cosmic muons of up to 1.3 TeV in energy.

After an accident in November 2001 which caused the implosion of >50% of the Photo Multiplier Tubes (PMTs) in the inner detector, Super-K redistributed the remaining PMTs and ran at a limited 47% of design efficiency until July 2005 when enough PMTs had been produced to refurbish the detector. This work was completed in July of 2006 returning the inner detector to its previous glory. The inner detector has a total of 11,146 50 cm Hamamatsu PMTs, now covered by fibre-reinforced acrylic windows, providing a 40% area coverage of the 36.2 m high by 33.8 m diameter inner detector cylinder.

The outer detector is optically separated from the inner by 2 Tyvek sheets, which are highly reflective on the outer detector side to improve light collection. It consists of 1,885 20 cm PMTs, facing both inwards and outwards, optically connected to 60 × 60 cm WLS perspex plates to increase light collection efficiency [90]. The outer detector wall is 2.05 m from the sides and 2.2 m from the top and bottom of the steel vessel.

The PMTs detect light produced by charged particle products of neutrino interactions. Čerenkov light is produced when a charged particle travels faster than the speed of

light within a medium. The photons produce a cone of light with opening angle,

$$\cos \theta_c = \frac{c}{vn(\lambda)} = \frac{1}{\beta n(\lambda)}, \quad (3.1)$$

dependent upon the speed of the particle, v , in the medium and the wavelength dependent refractive index, $n(\lambda)$. These particles characteristically produce rings of Čerenkov photons with the inner circle terminating when the charged particle loses energy and its velocity falls below the speed of light in water, $\sim 0.75 c$.

The requirement that a particle is travelling faster than the local speed of light provides a simple minimum energy for each particle type to produce Čerenkov radiation. The number of photons (N) emitted per unit wavelength (λ) per unit distance (x) travelled by a particle may be written as

$$\frac{d^2N}{dxd\lambda} = \frac{2\pi}{\lambda^2} \left(1 - \frac{1}{(\beta n(\lambda))^2} \right) = \frac{2\pi}{\lambda^2} \sin^2 \theta_c, \quad (3.2)$$

where α is the fine structure constant. The actual condition for a particle to be detected is that it must travel a long enough distance with $\beta \simeq 1$ to produce enough photons for the SK PMTs to pick a ring out of electronics noise.

The detector is sensitive to neutrino energies in the range 4 MeV → 1 TeV split into three regimes. The super low energy, SLE, regime from 4.5–6 MeV, which we will not discuss here, and the low energy (LE) regime from 6–30 MeV are primarily interactions from solar neutrinos. The energy of the charged product in the LE regime is calculated from the number of PMTs hit². High energy events >30 MeV, from atmospheric and beam produced neutrinos, reconstruct energy from the net charge deposited in the PMTs hit. Event vertices are constructed from TOF of the Čerenkov photons. The characteristic Čerenkov rings are fitted using a Hough transform image processing algorithm.

3.3.2 PID at Super-Kamiokande

Particle identification [91] at Super Kamiokande takes reconstructed Čerenkov rings and determines whether they were produced by a showering electromagnetic (EM), e^\pm or γ like, particle (Fig 3.2(c)) or a non-showering minimally ionising (MIP), μ^\pm or π^\pm like, particle (Fig 3.2(b)).

Firstly a likelihood estimator is calculated from expected photoelectron (p.e.) distributions of the two types of particle. Expected EM showering particle p.e. distributions are tabulated as a function of Čerenkov opening angle. The distributions were cal-

²the 6 MeV threshold accounting for 29 hits in the reduced coverage of SK II.

culated at three different momenta at a single vertex distance. These data are then extrapolated and interpolated to give data for a wide range of momenta. Data are then corrected for PMT acceptance angle and PMT angle seen by the vertex, and distance effects, such as attenuation and diffusion of the Čerenkov light.

The expected p.e. distributions for non-showering type particles are calculated analytically from the direction Čerenkov photons are emitted from a muon, the intensity of Čerenkov light as a function of this opening angle and a term based on MC simulations that accounts for knock on electrons. Scattered light is finally added to the p.e. distributions of both particle types.

$$L(l) = \prod_{\theta_i < (1.5\theta_c)} \text{prob}(q_i, q_i^{\text{exp}}(e) \text{ or } q_i, q_i^{\text{exp}}(\mu)) \quad (3.3)$$

Here q_i is the p.e. data from the i th PMT and the q^{exp} are the corresponding data points in the two particle types, $l = e$ or μ , expected p.e. distributions. The product is over all PMTs that lie within 1.5 times the Čerenkov opening angle θ_c . The two likelihoods are then maximised through optimising the $q_i^{\text{exp}}(e)$ and $q_i^{\text{exp}}(\mu)$ individually by altering the direction and opening angle of the ring.

These maximised likelihood terms are then transformed into a chi-square (χ^2) distribution.

$$\chi^2(l) = \frac{1}{\log_{10} e} (-\log_{10} L(l)) - \text{Const.} \quad (3.4)$$

From the χ^2 distribution one can write the probability of the pattern being either of the two types of particle:

$$P^{\text{pattern}}(l) = \exp \left(-\frac{1}{2} \left(\frac{\chi^2(l) - \chi^2_{\text{min}}}{\sigma_{\chi^2}} \right)^2 \right) \quad (3.5)$$

Here $\chi^2_{\text{min}} = \min[\chi^2(e), \chi^2(\mu)]$ and σ_{χ^2} is calculated from the number of degrees of freedom, N_D , as $\sigma_{\chi^2} = \sqrt{2N_D}$, here N_D is the number of PMTs hit by the ring.

A second probability is also defined using just the opening angle of the ring. This probability, however, is not used when attempting to reconstruct vertices with π^0 produced. Photons from the subsequent decay of π^0 's sometimes penetrate a non-negligible distance in the water before showering, with a mean free path of ~ 55 cm. This penetration reduces the reconstructed opening angle and consequently reduces the probability that the particle is identified as of showering type. In cases of π^0 vertices one uses only the P^{pattern} to determine particle type.

So, for π^0 vertices, if $P^{\text{pattern}}(e) > P^{\text{pattern}}(\mu)$ then the particle is classified as a

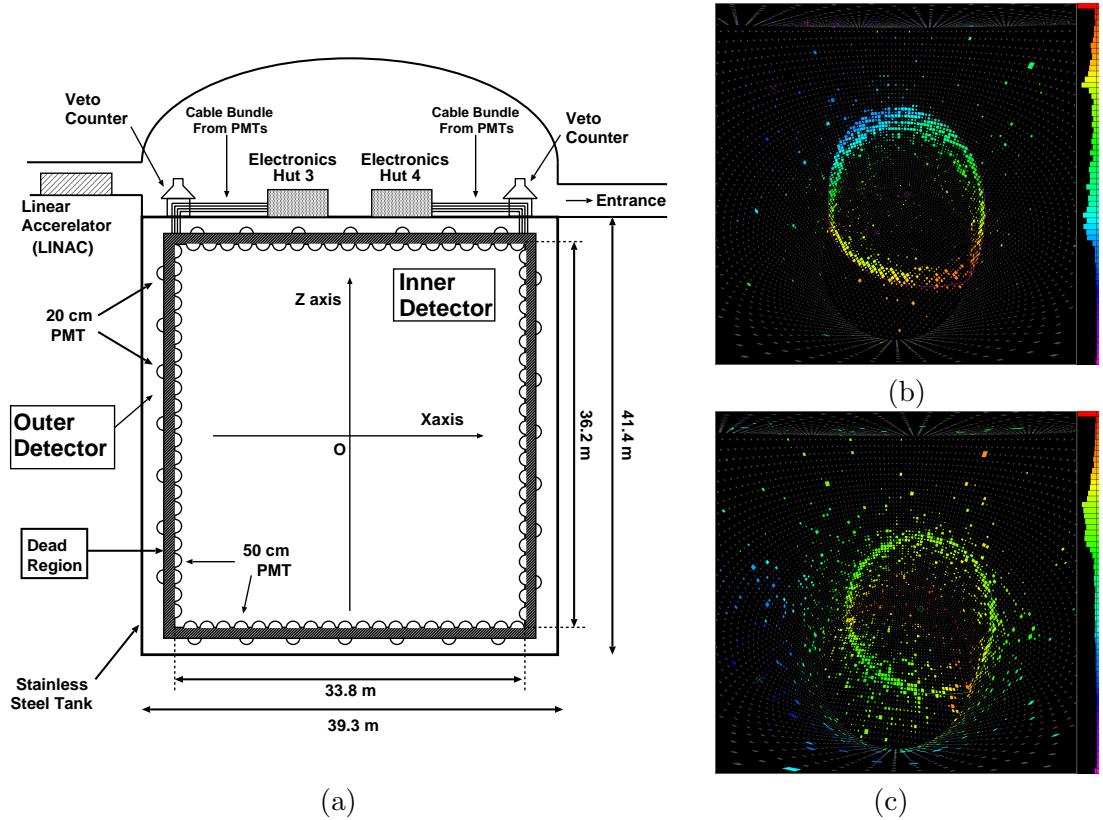


Figure 3.2: (a) A schematic view of the Super Kamiokande detector. The shaded region shows the dead space in the detector formed by the Tyvek optical shielding. Example single ring events in the Super Kamiokande inner detector. (b) A clean ring signal from a stopping muon with a reconstructed energy of 603 MeV (c) A fuzzy ring signal from a showering electron event with a reconstructed energy of 492 MeV. Images reproduced from http://www.ps.uci.edu/~tomba/sk/tscan/compare_mu_e/.

showering or electron like particle and if $P^{\text{pattern}}(\mu) > P^{\text{pattern}}(e)$ then it is classified as a non-showering or muon like particle.

Neutral current events can be used as a measure of the total active flavour content of a neutrino sample. Reconstructed multi-ring π^0 events, where the high energy ring is e-like, have been used in prior Super-K atmospheric analyses [92] to confirm $\nu_\mu \rightarrow \nu_\tau$ as the dominant oscillation mechanism over $\nu_\mu \rightarrow \nu_{\text{sterile}}$. The increased statistics that will come from the T2K beam will improve the statistical and systematic uncertainties of such events dramatically and provide increased sensitivity in the investigation of sterile neutrino models.

3.3.3 Neutrino CCQE backgrounds at Super-K

Super-Kamiokande seeks to measure neutrino interactions from charged current quasi elastic (CCQE) interactions. This base level interaction has a final state with a charged lepton and a recoil nucleon. The single lepton leaves a clear single ring signal within the detector whilst the nucleon remains unseen by Super-K. These single ring events are selected to reconstruct neutrino flux in the detector.

Single pion events produce the largest background to both ν_μ disappearance and ν_e appearance measured from CCQE events. These background cross-sections are to be measured by the ND280 (section 3.5) magnetised near detector.

Both charged current and neutral current production of single charged pions provide a background to ν_μ disappearance. NC single charged pion production generates energetic pions which, from their MIP nature, can be identified as muon candidates, adding false events to the CCQE ν_μ selection. CC single pion production may produce pions with energy below the Čerenkov threshold which are therefore not detected. This has the effect of smearing the momentum of the visible muon and consequently reconstructed neutrino energy, as the 2 body CCQE becomes, unbeknown to us, a more complicated 3 body CC1 π reconstruction. This background can be reduced by looking for Michel electrons, decay electrons from the decay chain $\pi \rightarrow \mu\nu_\mu \rightarrow e\nu_e\nu_\mu$, within a few muon lifetimes in the vicinity of a CCQE tagged event. After such cuts this background still persists and must therefore be understood. The tracker region of the ND280 (section 3.5.2) is designed to measure the differential cross-sections of these processes on both carbon and oxygen rich targets.

To measure the expected small number of subdominant ν_e appearance events one must constrain the systematic error of electron event tagging at Super-K. Aside from intrinsic beam ν_e , the predominant systematic uncertainty comes from neutral current production of single π^0 . Neutral pions decay predominantly into two photons, $98.798 \pm 0.032\%$, with just a small fraction, $1.198 \pm 0.032\%$, producing electron positron pairs (values from the Particle Data Group [38]). The decay is isotropic in the pion rest frame. If just a single decay photon is detected, or the two photons cannot be resolved as separate, then the single e-like ring event can be classed as a CCQE ν_e event.

In the energy regime occupied by the T2K beam, NC1 π^0 neutrino interactions produce relativistic π^0 . Towards higher energies the opening angle of the decay photons can be of the order of Super-K's angular resolution or below. In this case Super-K may not be able to resolve two separate Čerenkov rings and class the event as containing just a single e-like ring. At all energies a fraction of π^0 decays, in which the photons lie close to the π^0 axis of momentum, result in a hard (blue shifted) and a soft (red shifted)

photon in the lab frame. The soft photon will have a large interaction length and may exit the detector before pair producing, or it may produce charged particles with energy insufficient to produce enough Čerenkov light to deposit a ‘signal’ within Super-K. The hard photon will then shower and register a single shower-like ring event. As discussed above (section 3.3.2) PID at Super-K can distinguish between showering, e^\pm or γ like, particles or non-showering, μ^\pm or π^\pm like, particles. Therefore these signals would be misidentified as being a single ring, i.e. CCQE, ν_e event. For this reason it is imperative to understand the cross-section of neutral current single π^0 , NC1 π^0 , events to reduce the systematics of the ν_e appearance background. This background will also be studied at the ND280 near detector, primarily in a dedicated sub-detector aptly named the π^0 detector (P0D, section 3.5.4).

3.4 Near Detector Suite

The suite of near detectors will be located at 280 m from the graphite target, the largest distance possible within the J-PARC site. The ND280 pit will house multiple coarse detectors on the beam axis (hereafter on-axis), the “Interactive Neutrino Grid” or INGrid, designed to measure the beam topology. Alongside these will be a single fine-grained detector, the Near Detector 280m or ND280, positioned along the 2.5° off beam axis (hereafter off-axis), between beam source and Super-K. This detector will measure the flux, energy spectrum and essential cross sections of the ν beam before oscillation.

3.4.1 INGrid

The chosen off-axis angle of the beam can fluctuate during operation due to changes in beam emittance, proton beam optics, horn alignment, and other factors. Proton beam monitors and muon monitors will measure the proton and hadron beam axes. The on-axis INGrid detector provides direct measurement of the neutrino beam direction and relative flux using neutrino interactions. The requirement in the uncertainty of the beam direction is better than 1 mrad, which corresponds to a shift of 14 MeV or 2% in the peak neutrino energy.

The INGrid detector is comprised of fourteen identical modules which sample a $10 \times 10 \text{ m}^2$ area of the beam ($\pm 5 \text{ m } x$, $\pm 5 \text{ m } y$). Each module is composed of nine³

³Original plans were for ten planes but the most downstream plane was removed after modules were found to exceed the ND280 pit crane weight limit. This limit also affected the design of the tracker ECals, section 3.5.5.

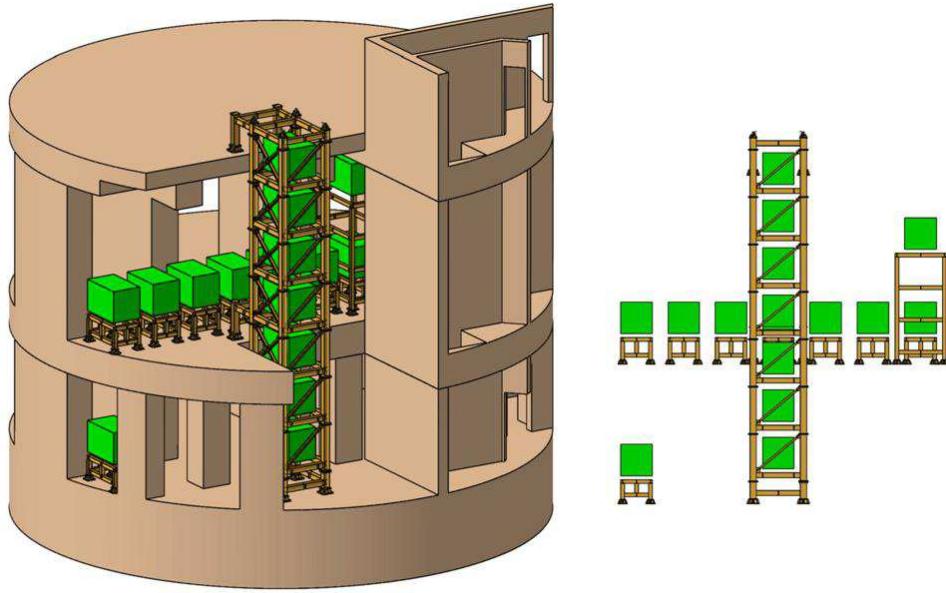


Figure 3.3: The INGrid on-axis detector consists of 16 modules. 14 for a cross sampling a 10×10 m area of the beam while 2 module sample the beam on the diagonal to confirm the off-diagonal flux measurements.

$1 \times 1 \times 0.1$ m³ iron plates which are sandwiched between ten scintillator tracking planes. Each tracking plane consists of two layers of scintillator bars, one horizontal and one vertical in orientation. Single layer veto scintillator layers are instrumented on the other open faces of the module: these are used in event selection. The scintillator strips used by INGrid have dimensions $100 \times 5 \times 1.3$ cm³. These modules will each provide around 7.8 tons of iron as neutrino target. The group's studies have shown that with the expected event statistics, even in the early beam commissioning phase, INGrid will be capable of measuring the beam direction to better than 1 mrad and observing a 1 mm shift at the target.

3.5 The ND280 Near Detector

In addition to beam position monitoring, it is also the goal of the near detectors to characterise the beam and measure relevant neutrino interaction cross-sections. This is achieved with a single multipurpose magnetised detector, the ND280.

3.5.1 Overview and Magnet

The ND280 detector, shown in Fig 3.4, is composed of a collection of finely grained sub-detectors all housed within a magnet. The magnet is recycled from that used in the UA1 experiment at CERN, where a current of 10 kA (power 6 MW) provided a 0.67 T uniform magnetic field perpendicular to the magnet's major axis. In the ND280 it is planned that the coils will be supplied with 3 kA (power 0.6 MW) which will provide a horizontal uniform operational field of 0.2 T. The magnet has inner dimensions of $3.5 \times 3.6 \times 7.0$ m and is split symmetrically in half down its centre, each half comprising 8 C-shaped iron flux return yokes. There are a total of 4 coils with a central hole that originally accommodated the beam pipe of the UA1 experiment. There are insufficient gaps in the magnet for each optical channel in the detector to have a wire or optical fibre fed through for readout, and so all front-end electronics must fit alongside detectors within the inner volume.

An aluminium framed basket of dimensions $2.5 \times 2.6 \times 6.5$ m will house all of the inner detectors. From upstream to down these are the π^0 detector (P0D), the tracker and the downstream ECal (DS ECal). This basket of inner detectors is surrounded on all four sides by electromagnetic calorimeters (ECals) which are fixed to the inner surface of the magnet yokes. The air gaps between the return yokes will also be instrumented with plastic scintillation detectors, the side muon range detector (SMRD), to measure long range muon trajectories originating in the inner detector.

The ND280 detector has been designed to optimise the small space by maximising the amount of active elements and simultaneously provide enough target mass for neutrino interactions. This has been made possible through careful design optimisation in conjunction with development and use of new readout devices and electronics, most notably the magnetic field insensitive minute MPPC solid state photosensors which are described in section 4.1. I will now describe the design of each element of the ND280 off axis detector and mention the physics goals of each.

3.5.2 Tracker

At the heart of the detector is the Tracker region, comprising three Time Projection Chambers (TPCs), interleaved with two Fine Grained Detectors (FGDs).

The ND280 tracker is primarily intended to measure the ν_μ and ν_e fluxes and spectra, as well as various charged current cross sections. Neutrino cross-section measurements and flux determination are coupled. The two are deconvolved using the theoretically well understood charged current quasi-elastic (CCQE) cross-section as a reference; all

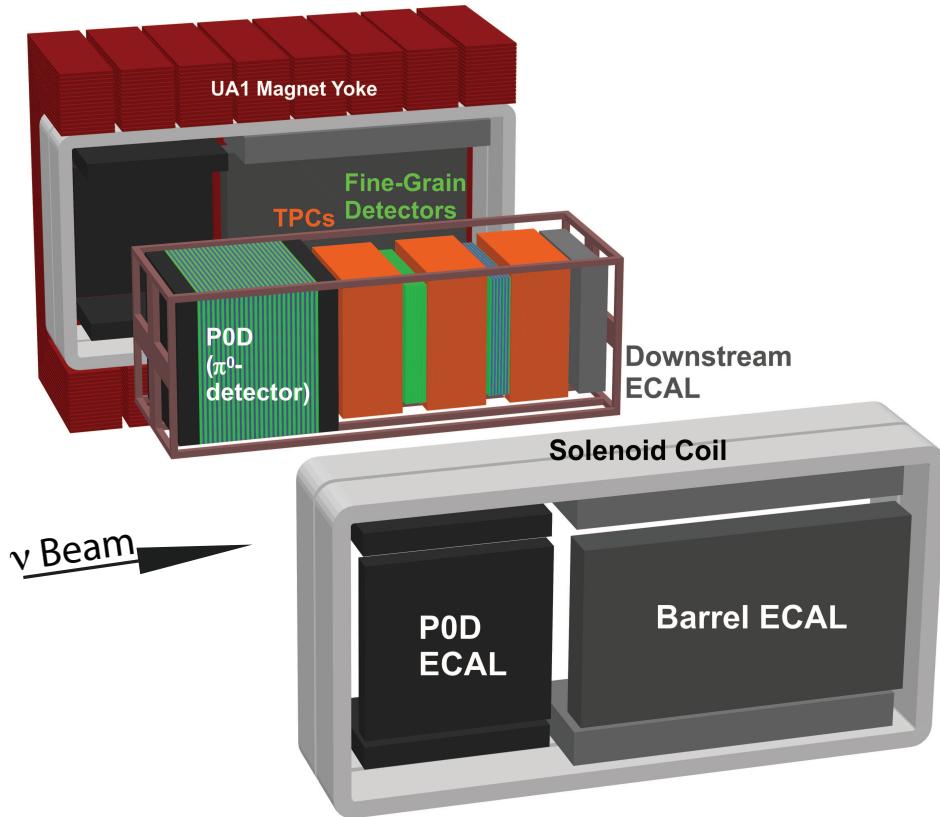


Figure 3.4: The ND280 off axis detector, a composite of the Pi-Zero Detector (P0D), three Time Projection Chambers, two Fine Grained Detectors (FGD) and three distinct Electromagnetic Calorimeter (ECal) regions housed in the old UA1 Magnet. Image courtesy of Yoshi Uchida.

higher order charged current cross sections are measured as a fraction of CCQE.

CCQE interactions, $\nu_e + n \rightarrow e^- + p$, are the most common interaction, $\sim 40\%$, at T2K's beam energy. The energy of an incident neutrino is calculable from just the energy and direction of the final lepton, which makes the two body CCQE reaction ideal for measuring the neutrino energy spectrum.

Some higher order processes form a background to CCQE event selection. The dominant CCQE background process is CC single pion (CC-1 π) production ($\nu_l + N \rightarrow l^- + N' + \pi$), where only the final state charged lepton may be detected. The three body CC-1 π events will smear any energy spectrum measurement, and must be excluded from the energy spectrum analyses in both the near and far detector. At Super-K the majority of such events are removed by selecting just single ring events, with a single

charged lepton in the final state. However CC-1 π events with the pion below Čerenkov threshold form an irreducible background; the tracker is used to measure the size of this and other backgrounds to CCQE interactions at Super-K.

There are a number of cuts that can be used to reject non-QE events. The most obvious is the rejection of events in which there are more than the expected two tracks leaving the interaction vertex. After this cut there will still be those events where very short-ranged pions or protons may not have been identified as separate tracks. Worse still, a pion from a CC-1 π event can be reabsorbed inside the production nucleus in final state interactions. One can fight against these classes of event using kinematic cuts. The direction of the recoil proton can be calculated just from the direction and energy of the charged lepton in CCQE interactions. The angle between the observed proton momentum and the expected proton momentum will be distributed around zero for true CCQE events, but will have a much flatter distribution for non-QE events. Rejection of non-QE events can also arise from the tagging of Michel electrons, produced from the decay of short-ranged pions.

To achieve all of these goals and obtain clean CCQE samples, the recoil proton, possible pions and the charged lepton all need to be reconstructed as coming from a common vertex. This is achieved in the ND280 tracker region by employing two sets of sub-detector, each optimised to perform one or the other reconstruction. There are two massive plastic scintillator-lead sampling fine-grained detectors (FGDs) which provide the target mass for neutrino interactions as well as high resolution tracking of charged particles in the vicinity of the interaction vertex. The FGDs interleave time projection chambers (TPCs) whose goal is to measure the momenta of particles, from curvature in the 0.2 T B-field, and accurately identify particle type, from measurement of dE/dx , of long ranging particles emanating from FGD vertices.

Determining the neutrino spectrum from reconstructed muon momentum requires only a moderate resolution, as the neutrino energy reconstruction is affected by a relatively large unavoidable smearing from Fermi motion of the nucleon target. This smearing is of the order of 10% for neutrinos around the peak beam energy. The muon momentum scale does however play a crucial role in determining the scale of the neutrino energy, and hence in the ν_μ disappearance measurement. The energy scale is one of the main systematics in this measurement. To achieve the T2K physics goals, the muon momentum scale is required to be understood at the 2% level.

The TPCs will also play the key role in determining the flavour composition of the neutrino beam, i.e. the intrinsic ν_e component in the beam. This will be achieved by reconstructing a sample of CCQE events where the high momentum particle is

identified as an electron, using the TPC ionization measurement coupled with ECal (section 3.5.5) and side muon range detector (SMRD, section 3.5.3) information. The use of independent information from different detectors will help to cross-check the selection efficiency and contamination of this analysis. The main difficulty will be to identify electrons reliably in a sample heavily dominated by muons: ν_e are expected to be 0.2% of the ν_μ in the 0.4–0.8 GeV energy range. A further handle in this measurement will be given by the measurement of the neutrino spectrum in the high energy (1 to 3 GeV) region. These neutrinos are predominantly from kaon decays, which are responsible for a large fraction of the ν_e in the beam from the K_{e3} decay channel $K \rightarrow \pi^0 e \nu_e$ ($5.08 \pm 0.05\%$). Measurement of high momentum muons could help understand the hadron production model in the beam Monte Carlo simulations and help predict the ν_e flux more accurately.

A large fraction of low energy neutrino-nucleus interactions involve the production of very low energy hadrons. Monte Carlo simulations of these processes make use of empirical models which have not been well tested, owing to the fact that it is difficult to detect and measure the low energy particles within dense target material. The gas of the TPC will provide additional, low-density, target nuclei. Neutrino interactions within the gas could provide interesting samples of a few thousand events per year which will improve the understanding of neutrino-nucleus interactions, since all of the charged particles produced can be tracked in the TPC.

Fine Grained Detectors

The FGDs are constructed primarily from 0.96 cm \times 0.96 cm \times 184.3 cm bars of extruded polystyrene scintillator, each of which is read out by a wavelength-shifting fibre going down its centre. One end of each fibre is mirrored by vacuum deposition of aluminium, while the other end is attached to a silicon multi-pixel photon counter (section 4.1) and associated electronics (section 3.5.6), which digitize the light signal produced by scintillation inside the bar.

Each FGD has outer dimensions of 230 cm (width) \times 240 cm (height) \times 36.5 cm (depth in beam direction), and contains 1.0 tonnes of target material. The first FGD, FGD1, consists of 5760 scintillator bars, arranged into 30 layers of 192 bars each, with each layer oriented alternatingly in the x and y directions perpendicular to the neutrino beam to allow 3D reconstruction of particle tracks. FGD2 is a water-rich detector which allows separate determination of cross-sections on carbon and water from comparison of interaction rates with FGD1. FGD2 consists of 7 x - y sandwiches of plastic scintillator layers, alternating with six 2.5-cm thick layers of water (for a total

of 2688 active scintillator bars and 15 cm total thickness of water).

Time Projection Chambers, TPCs

Each of the three identical $2.5 \times 2.5 \times 1$ m TPCs has an active volume filled with Ar(95%)CF₄(3%)iC₄H₁₀ gas [93] and provides high resolution measurements of charged particle tracks with minimal interference. Each active volume is split in two by a central cathode plane in y - z . The two ends parallel to this plane are tiled in MICROMEGAS micropattern charge readout modules [94].

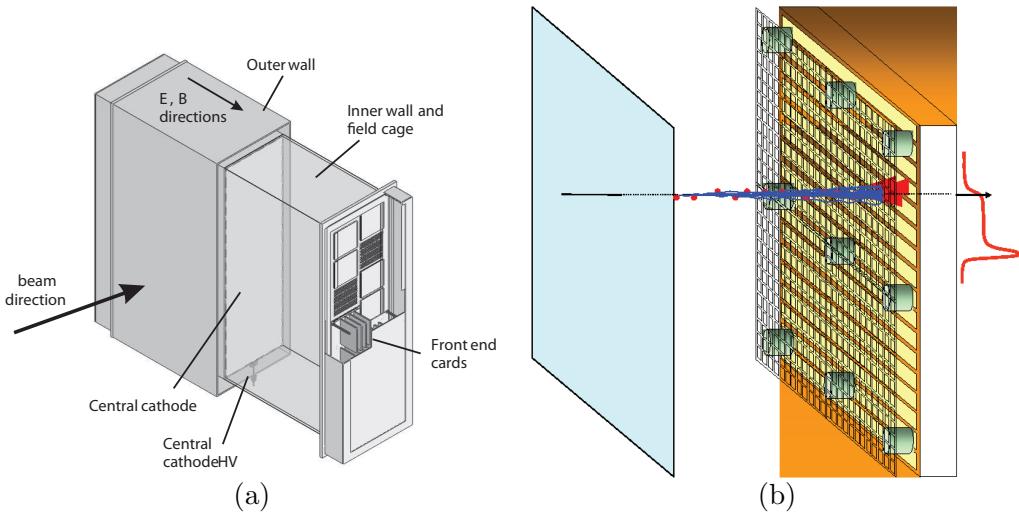


Figure 3.5: (a) Diagram of a TPC, clearly displaying the inner and outer regions and (b) a diagram of one of the MICROMEGAS readout modules. Images reproduced from [93].

Each MICROMEGAS has an array of 48 columns by 36 lines of 6.8×9.7 mm² pads, each of which is read out by a front-end electronic card (FEC). The electrons produced when an ionising particle passes through the gas drift towards the anode of the MICROMEGAS modules. When an electron reaches the micro-mesh, supported 50–100 μ m above the anode, it initiates an avalanche of electrons amplifying the original ionisation.

This structure provides an accurate position for the track in y - z . Relative x positions along the track are calculated from the timing of the signals received. For this calculation one must not only account for the drift speed of the charges but also possess a zero time, t_0 , measurement. This is obtained either from a track crossing the central cathode or from matching the track to objects in the FGD or ECal.

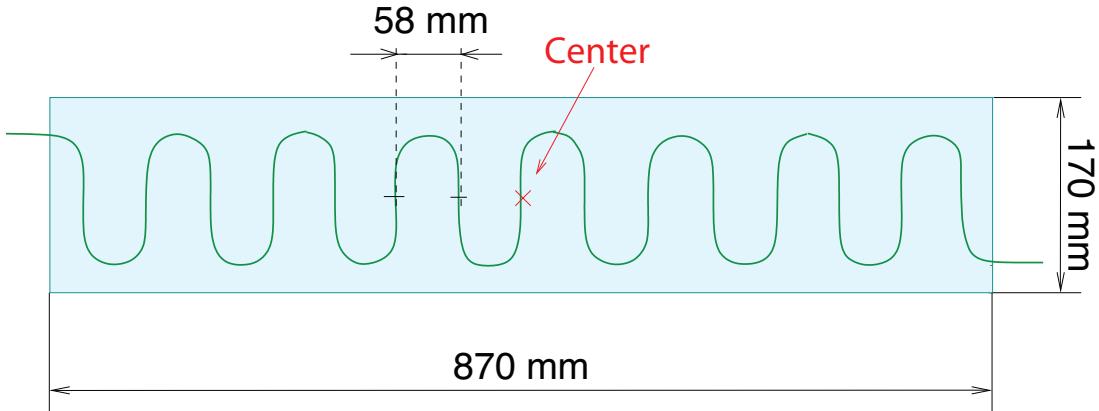


Figure 3.6: A single SMRD scintillation counter. Modified from [5].

3.5.3 Side Muon Range Detector, SMRD

A full spectral measurement of the neutrino beam must include muons with large angles relative to the neutrino beam direction. Muons which escape at large angles with respect to the beam cannot be measured by the TPC. However, these muons will intersect the iron yoke surrounding the entire detector, permitting inference of muon momenta from their range. This is obtained in the ND280 by instrumenting the iron at various depths with plastic scintillator modules, the full set of which is known as the SMRD.

The SMRD will also serve to identify backgrounds entering from outside the detector and secondary particles from beam neutrino interactions in the iron of the magnet yokes. In the same vein, the SMRD will provide a trigger signal in response to cosmic ray muons. Cosmic muons are an invaluable calibration source for all inner detectors, as they provide sample tracks that are very similar in nature to muons that are created in beam neutrino interactions.

The iron yoke consists of 16 C-shaped elements composed of sixteen 48 mm thick iron plates, with 17 mm air gaps in between, and is segmented into 12 azimuthal sections. To build the SMRD, the air gaps will be equipped with SMRD plastic scintillator detectors. The horizontally and vertically oriented gaps have slightly different dimensions, $876 \times 700 \times 17 \text{ mm}^3$ and $876 \times 910 \times 17 \text{ mm}^3$ for horizontal and vertical gaps respectively.

Funding would not allow for all of these air gaps to be instrumented. The decision of which regions to instrument was based on a compromise between two aims. The magnet coils introduce an increased mass before the top and bottom region which

would decrease the number of outgoing particles that would be detected; this favoured the placement of side modules. Downstream sections are also favoured due to the forward going bias of muon production.

For the detection of cosmic ray muons, however, the top and bottom regions are most important and should be instrumented evenly so that all inner sub-detectors receive useful cosmic triggers. To meet this requirement all C-shaped elements of the magnet yoke will have 3 scintillator layers in both the top and bottom. The side regions will vary from 6 layers, for the two most downstream C-shaped elements, to 2 layers for the three most upstream.

The dimensions of each scintillator slab are $870 \times 170 \times 10 \text{ mm}^3$ and $870 \times 180 \times 10 \text{ mm}^3$ for the horizontal and vertical gaps respectively. A 4 mm deep S-shape groove will be machined into each scintillator slab to allow the insertion of a wavelength shifting (WLS) fibre of 1 mm diameter. Fig 3.6 shows a scintillator slab with a S-shaped WLS fibre. Each WLS fibre will be coupled to and read out by two Hamamatsu multi-pixel photon counters (MPPCs) by means of optical connectors at both ends. Photosensors from each azimuthal section of a C-shaped element will go to the same front-end board, so that towers are created which can be used to produce cosmic triggers.

3.5.4 The π^0 Detector, P0D

The detector furthest upstream in the ND280 basket is the Pi-Zero Detector (P0D). As the name suggests, the primary physics goal of the P0D is to measure neutral current π^0 events, most notably those produced in a water target within the expected momentum range of the T2K ν_e appearance signal. Single π^0 NC interactions on water have been measured only at $\sim E_\nu = 1.3 \text{ GeV}$. The measured rate was found to be $6.4 \pm 0.1(\text{stat.}) \pm 0.7(\text{syst.})\%$ [95] of the CC neutrino rate. Although this rate was in agreement with expectation, it was noted that there was disagreement between the observed and theoretically expected π^0 momentum distributions. Since this interaction is poorly studied at T2K neutrino energies, it must be measured at the ND280 so that an accurate background subtraction can be performed at Super-K. The P0D will measure π^0 production through the exclusive channels

$$\begin{aligned} \nu_\mu n &\rightarrow \nu_\mu n \pi^0 \\ \nu_\mu p &\rightarrow \nu_\mu p \pi^0 \\ \nu_\mu n &\rightarrow \mu^- p \pi^0 \end{aligned}$$

The P0D will also make inclusive neutral current and charged current π^0 production

measurements. It is also hoped that the large mass and therefore increased statistics in the P0D will provide a measurement of the intrinsic ν_e flux using quasi-elastic events to complement that made by the tracker.

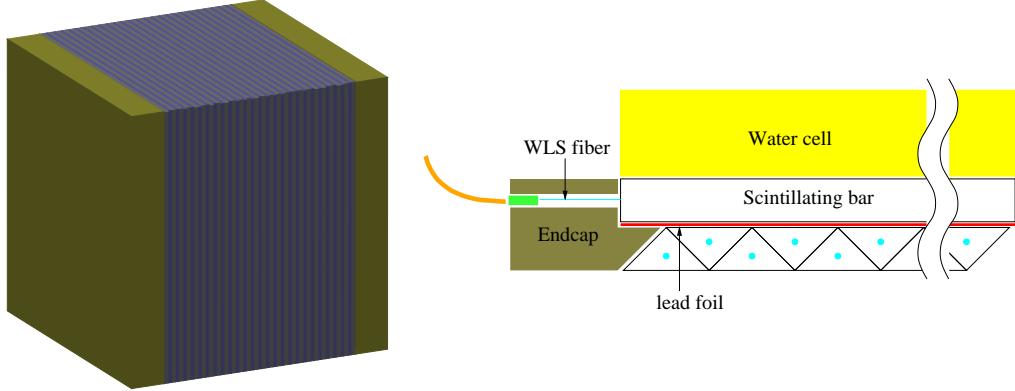


Figure 3.7: (Left) The P0D detector viewed from the upstream end. The central region of the P0D is constructed of alternating water target and scintillator tracking layers. The upstream and downstream regions are composed of lead radiator and scintillator tracking layers. (Right) A schematic of a P0D scintillator layer showing the light-tight end caps and MPPC housing.

The P0D is configured into three separate analysis regions called Super-P0Dules (cf Fig 3.7 Left). Super-P0Dules are constructed from active plastic scintillator layers called P0Dules interleaved with inactive absorber. The inactive layers are split into two classes, radiator and target. Inactive target layers are water layers constructed from semi-flexible pillow bladders, with nominal dimensions $0.03 \times 1.8 \times 2.1 \text{ m}^3$, holding approximately 100 kg of water each. Radiator layers are high Z material introduced to convert photons and enhance the progression of showers.

The central region of the P0D provides the fiducial mass for π^0 measurement. It comprises 25 repeating units of P0Dule, 1.6 mm brass radiator and water target completed with a final P0Dule. Oxygen cross-section measurements will be made by comparing the interaction rate of events collected with the water target regions filled (target in) versus similar running periods with water removed from the target region (target out). The P0D design has a total mass of approximately 15 tons with a fiducial mass of 2 tons of water, and 3 tons of other material. The upstream and downstream regions are configured as EM calorimeters, the upstream (USECal) and central (CECal) ECals respectively. These layers do not have any water target but are instead composed of P0Dule and lead radiator material of thickness 4 mm (USECal) and 4.4 mm (CECal). These regions have seven layers each and provide energy containment of showers and

active shielding.

All scintillator planes are made light-tight with layers of polystyrene on the large faces and end caps. The end caps serve (see Fig 3.7 Right) the dual purposes of providing a manifold to bring WLS fibres out of the scintillating bars, and keeping the ends of the scintillator light-tight. Read-out of individual scintillator bars is performed using a single MPPC at the end of a WLS fibre. For increased spatial resolution and pointing, the P0D design calls for triangular scintillator bars that are 17 mm in height and 32.5 mm in base co-extruded with a TiO_2 surface treatment. Each scintillator module is constructed from 128 (136) bars oriented in the x (y) direction.

3.5.5 Electromagnetic Calorimeter, ECal

The ECal is realised as three separate regions: the downstream ECal (DsECal) occupying the most downstream position of the central basketed detectors, the barrel ECal (BrECal) which surrounds the four outer faces of the tracker region and finally the P0D ECal which surrounds the four outer faces of the P0D region.

Downstream (DsECal) and Barrel (BrECal) ECals

It is the aim of the BrECal and DsECal to measure the energy of particles leaving the tracker region, from vertices in either FGD, with a good resolution. It is also an important task of these ECals to convert, characterise and measure the energy of photons, which would not be seen by the TPC. In making these measurements, the BrECal and DsECal hope to not only improve the systematics of FGD events but also measure a fraction of events producing π^0 which will go largely unnoticed in the tracker.

These ECals are sampling tracker calorimeters composed of active plastic scintillator layers interleaved with lead absorber. All modules have a primary direction that points in towards the tracker so that all particles exiting the region traverse the high Z lead absorber. This is particularly important to convert photons and initiate a shower of electrons and positrons which can be subsequently detected. The scintillator layers are alternately ordered in opposing axes so that the modules can perform 3D reconstruction of showers and tracks.

The BrECal is split into a number of modules, for ease of construction and installation, which are each mounted on the inside of the magnet and coils. The DsECal is a single module which occupies the very last downstream 50 cm of the basketed inner detectors.

The ECal scintillator layers are constructed from extruded plastic scintillator 4 cm wide and 1 cm thick⁴ instrumented with Kuraray Y11(200) WLS fibre which is threaded through a central hole. The length of bar used depends upon the orientation of the layer and the module in question. Table 3.1 gives a concise breakdown of all ECal modules.

Region	Modules	Size (m ³)	Layers	Bars/Layer	Lead (mm)
Downstream	1	$2 \times 2 \times 0.5$	34	50	1.75
Barrel					
Top/Bottom	4	$3.84 \times 1.52 \times 0.5$	30	96/38	1.75
Side	4	$3.84 \times 2.28 \times 0.5$	30	96/58	1.75
P0D					
Top/Bottom	4	$2.34 \times 1.4 \times 0.5$	6	38	4
Side	2	$2.34 \times 2.6 \times 0.5$	6	69	4

Table 3.1: Breakdown of ECal modules.

The DsECal has active layers with bar length of 2 m in both x - z and y - z orientation⁵, each of which is read out at each end with the WLS feeding light into an MPPC photosensor. Because of the high multiplicity of events in the DsECal it was decided to instrument each DsECal bar with MPPC photosensor readout at both ends to maximise the information available.

The lengths of bar in the BrECal range from 1.5 m shortest to 4.2 m longest. The long attenuation length of the WLS fibre, ~ 350 cm, determines that all but the longer bars are read out by a single MPPC with the alternative end Al mirrored. The longer bars may lose too much signal and so are instrumented with MPPC at both ends. All MPPC signals are sent to front end electronics (section 3.5.6) where they are digitised and passed to back end electronics for readout and storage.

The P0D ECal

The P0D ECal is a simpler device which is essentially a descoped version of its barrel cousin. Its aim is not to achieve high resolution energy measurements but just to positively identify photons and muons escaping the P0D. The P0D ECal does not have the energy or spatial resolution required for full π^0 reconstruction. However its presence is necessary to detect photons that do not convert in the active P0D volume or showers only partially contained therein. It can also confirm the passage of charged

⁴For more information on this scintillator please see section 4.2.

⁵Here the first informational axis is given horizontal and the second perpendicular to the face of the ECal stack in question. This notation will be used through this thesis.

tracks, identify MIPs, and so actively discriminate CC and NC events as well as acting as veto for incoming backgrounds.

The six P0D ECal modules consist of 6 active scintillator layers, comprised of 1 cm thick by 4 cm wide plastic scintillator bars, separated by 5 layers of 4 mm thick lead absorber. Each module has an effective depth of $3.6 X_0$. All scintillator layers are made of the same scintillator used in the barrel and downstream ECal. They will be read out at a single end by a Hamamatsu MPPC photosensor with the opposite end mirrored for improved light collection. All bars in all modules and layers are oriented with their long axis in the direction of the beam. This is chosen to minimise the number of photosensor channels and also simplify the furnishing of front end electronics to a single face of the detector. As with the BrECal, the P0D ECal modules are fixed to the inside of the iron yoke of the magnet.

The P0D radiator layers act to convert photons and contain the resulting showers, allowing the characteristic two-shower signature of π^0 decays to be reconstructed. These layers are however perpendicular to the beam, to increase photon detection efficiency in the forward direction. This leads to a decreased photon containment and conversion for photons and reduced efficiency of track measurement of MIPs at high angles to the beam. P0D event reconstruction and analyses compensate for these effects by reducing their fiducial volume perpendicular to the beam. It is the aim of the P0DECal to increase statistics and systematics of such events by measuring such oblique angle particles.

The P0D ECal is the focus of this thesis with the physics goals and final design study discussed in chapter 6 and the performance of reconstruction algorithms explained in chapters 7 and 8. Finally the effects of the P0DECal on NC1 π^0 event selection, in the P0D region of the ND280, are considered in chapter 9.

3.5.6 Electronics

There are two types of front end electronics used in the ND280, using different application specific chips (ASICs). The ECals, P0D, SMRD and INGRID all use the TriP-T (Trigger Pipeline + Timing) chip [96]. The FGD and TPC use the AFTER (ASIC for TPC electronics readout) chip [97] which was developed for TPC Electronics Readout [93, 98] but had time resolution characteristics attractive to the FGD.

Front End Cards, FEC

The central component of the FEC is the AFTER chip, a 72 channel ASIC specifically designed for T2K. In the TPC, each MICROMEGAS pad will connect to a channel and then be continuously sampled into analog memory. The chip allows for 511 samples, with a frequency of 1-50 MHz, and when triggered these are fed into an external ADC.

The resulting output of the electronics is a waveform of charge vs time for each pad. Each peak in the waveform is fitted to a Gaussian and made into a hit in y , z and t , which are then passed to the reconstruction software.

The FGD opted to use the AFTER chip so that the time of a neutrino interaction can be known to higher accuracy than a single timestamp would allow. This is especially important when matching tracks between the FGD and TPC. Another advantage of the better time resolution is the ability to identify a charged pion stopping in an FGD bar. The pion will decay to an almost stationary μ^\pm , which will itself produce a Michel e^\pm , detected by the photosensor significantly later than the initial signal. As the muon decay constant is 2.2 μ s long and the beam spill is 5 μ s, then a 10 μ s live period for the electronics covers the spill width and two muon lifetimes for all particles. The AFTER chip has a maximum sampling frequency of 50 MHz and 511 time slots, which gives 10.220 μ s and matches nicely.

The AFTER cannot provide the full dynamic range needed by the MPPC. The maximum input charge of the AFTER chip is 600 fC, but the maximum output of the MPPC, \sim 100 pC (667 pixels and a gain of 10^6), is significantly higher. Therefore the signal must be attenuated before being passed into the AFTER chip. However, applying the constraint that the electronics noise must be lower than 0.2 times the single p.e. signal and taking the noise at 0.3–0.5 fC, the signal can only be attenuated by 30–50. To satisfy the conflict, each MPPC is fed into two channels, with two different attenuation factors. As the gain and noise levels are still being precisely determined, the exact factors are not yet known, but the FGD group hope to achieve 0.1 to 100 p.e. for the low attenuation channel and 15 to 667 for the high.

TriP-T Front End Boards, TFB

TriP-T Front End Boards (TFBs) read out 64 MPPCs each; 48 TFBs are fed into a single Readout Merger Module (RMM) where data is collated before being sent to the data acquisition system (DAQ).

MPPC photosensor signals are fed into TFBs via miniature coaxial (mini-coax) cable connections. There are 4 TriP-T ASIC chips on each TFB. Each TriP-T chip has

32 input channels consisting of an integrating preamplifier and a second adjustable amplifier which then feeds into an analog pipeline. The pipeline is 48 cells deep, which is perfectly adequate for the operational design of 15 bunches per spill from the beam. So in normal running the preamplifier integration will be triggered by each bunch, with a minimum 50 ns reset time between bunches. After each spill, the pipeline is read out and digitised by an ADC, before being transmitted off the board and the pipeline reset.

The gain of the TriP-T is set globally, which means that one is not able to simultaneously cover the entire range of MPPC signal and accurately distinguish the first p.e. peak, for dark current discrimination. To preserve dynamic range, the MPPC signal is capacitively split and fed into two TriP-T channels. The high capacitance channel, C_{hi} , receives the bulk of the signal and can therefore resolve low p.e. signals. A globally programmable discrimination threshold can be applied to TriP-T channels individually but can only be applied to 16 at any one time. The discrimination is applied to these high gain channels because the threshold is assumed to be a low 3-5 p.e. These discriminated outputs also define the time stamp of detector hits, when passed to the Field Programmable Gate Array (FPGA). The FPGA is a chip which handles not only the timestamping but also all the control requirements and slow controls of the TFB.

Therefore, for each MPPC in the ECal we would expect to receive three values every bunch; ADC values from the high gain and low gain channels and a timestamp. If the charge from an MPPC is not above the discriminator threshold set then the time stamp will essentially be void, and any information possibly useful to the reconstruction will have been lost. If the threshold is lowered to attempt to collect any useful signal lost in this way then noise in the MPPC occurring before a true deposit may trigger the discriminator. This will yield a timestamp with value significantly before the real time. As the reconstruction will rely heavily upon timing information to associate the correct objects, this incorrect timestamp will have a more significant effect than the missing timestamps for low charge signals.

3.6 Time Line

The stage IV upgrade of Super-Kamiokande is complete, with calibration of the new electronics currently underway; the detector will soon be fully operational.

The first protons were extracted into the neutrino beamline, and beam was detected by the MuMON, on the 23rd of April (2009). The limited 0.1% intensity and much shorter bunch time of 10 ns compared with 58 ns in full beam has just one horn installed; modifications have to be made to horns 2 and 3, which gives a reduced 3/4 muon yield

with a much broader profile $\sigma_y = 125$ cm $\rightarrow 140$ cm. Horns 2 and 3 will be installed in the Summer of 2009 while the ND280 and INGrid detectors are being commissioned.

The first INGrid module is undergoing commissioning as I write and hopes to complete after ~ 5 weeks of electrical testing and performance verification and calibration with the DAQ. The complete INGrid array is expected to have been installed and commissioned by mid-July 2009.

The UA1 magnet is installed in the ND280 pit and will be closed, after installation of electronics and commissioning, for magnetic field mapping which is scheduled to take 2 weeks at the start of June 2009. The magnet will be operating initially at just 0.18 T because of an oversight in calculation of power usage when changing cooling devices to active chillers.

The SMRD is the first of the ND280 subdetectors to be installed, starting with the bottom horizontal modules, which began at the end of February 2009, and culminating with the South/right side vertical modules completing installation in September 2009. During this time each installed portion of the SMRD will be commissioned before installation of the next section commences, with commissioning complete at the end of September 2009.

The P0D detector arrived at J-PARC on the 24th April and will be fitted with electronics before cosmic muon and test beam runs. The P0D will be installed in the ND280 pit in four sections over a six-day period beginning early August. The first to be placed in the basket is the P0D downstream ECal, followed by the central fiducial target which is split into two halves, with installation culminating in the P0D upstream ECal. Commissioning of the P0D is then expected to last until mid-September.

After the cosmic run in the UK which finished on the 1st April 2009 the downstream ECal arrived at CERN on the 17th and is being prepared for a test beam run beginning early May, the downstream ECal will then be airfreighted to J-PARC to begin installation at the end of July. The commissioning of this detector is expected to be swift, finishing in the third week of August.

The TPC modules will be assembled in June 2009 at J-PARC, including population of the MICROMEGAS and connection of front end electronics. The FGDs and TPCs are scheduled to be installed in the basket within a five-day period in mid-August, FGDs first followed by TPCs all in downstream order. Commissioning will then commence and is planned to finish at the very end of September 2009.

The first yoke 1 side, bottom and top P0D ECal modules will arrive in mid-July, August and September respectively. After electrical testing these first modules will be installed

over a three-week period beginning 24th August. The final three modules will arrive early 2010 to be installed during an agreed shutdown of ND280 around April 2010.

The only two barrel modules to be installed in 2009 will be the yoke 1 top and side, which will be installed during the first 2 weeks of December 2009. The remaining four modules should be installed during an ND280 shutdown in August or September of 2010.

DAQ commissioning will be ongoing during September and October as detectors are installed, with the TPC commissioning being the last to finish in mid November.

It is hoped that this schedule will see the ND280 begin data taking in earnest with the new year 2010. T2K hopes to present first data on θ_{23} and Δm_{23}^2 at the summer conferences in 2010.

Chapter 4

Light Yield Measurements of ECal Scintillator Bars

In this chapter I present some quality assurance work conducted to test prototype extrusions of ND280 ECal plastic scintillator bars from FNAL. Two prototype extrusions were conducted before full scale production. The light yield of these bars was compared to that of a sample of scintillator, also produced at FNAL, for the K2K SciBar detector [99]. The light readout used the same method as will be employed in the final detector: wavelength shifting (WLS) fibres guiding scintillation light to solid state multi-pixel photon counter (MPPC) photosensors.

4.1 Multi-Pixel Photon Counters, MPPCs

Physical Structure

The MPPC photosensors developed by Hamamatsu are a class of Geiger-mode avalanche multi-pixel photodiode (GAMPD). The GAMPD is an array of identical single avalanche photodiodes (APDs), connected by a common readout chain. Each APD is connected in series with a quenching resistor. This quenching resistor allows each pixel to discharge and reset independently of the rest of the device. The pixels are isolated physically either by guard rings of highly doped material or by trenches etched onto the substrate, Fig 4.1(b).

This multi-pixel structure quantises the readout from such devices, with each pixel contributing a fixed amount of charge to the output. The total charge from a device is therefore generally proportional to the number of pixels fired. This proportionality

decreases as one reaches the limit of the device's dynamic range, which is bounded by the number of pixels on a device. The MPPCs used by the ND280 collaboration have 667 pixels in a square arrangement constructed upon a 1.3×1.3 mm active area, Fig 4.1(a).

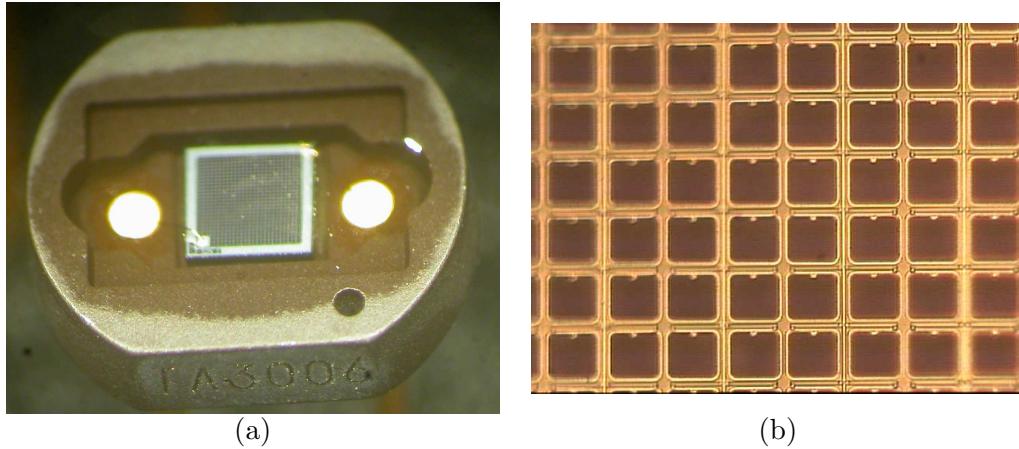


Figure 4.1: (a) Hammamatsu 667 pixel MPPC and (b) a close up of the pixel structure showing the multiple individual APDs separated by trenches etched onto the surface of the silicon substrate. Images courtesy of Mark Ward (a) and Roberto Sacco (b).

Signal and Noise

The pixels themselves are p-n junctions operated in reverse bias in Geiger mode. To operate in Geiger mode, the bias voltage applied (V_b) must be greater than the breakdown voltage (V_{bd}) of the device. The breakdown voltage is the point at which current flow increases exponentially whilst under reverse bias. The quantity $V_b - V_{bd}$ is known as the over-bias and is typically 1-2 V. The charge builds in the depletion region much in the same way as the build up on a capacitor $Q = C(V_b - V_{bd})$.

A pixel will then retain this charge until either photo (signal) or thermal excitation (noise), which produces an electron-hole pair within the depletion region of the junction. This then triggers an avalanche of electron-hole pairs which are swept out of the depletion region resulting in current flow across the junction. This drop in resistance places the bias across the quenching resistor which impedes the otherwise continual charge flow across the junction. Once discharged the resistance of the junction grows once more, as all charge carriers leave the depletion region, and the pixel recharges.

The potential waveform output by the MPPC is integrated to obtain the charge deposited. This is done in situ by analogue to digital converters (ADC) or can be per-

formed by the ADC within an oscilloscope. The pedestal is the integration of electrostatic noise, in the absence of a pixel fire, and thus represents the zero point for integrated charge.

As well as this pedestal there will also be pulses in the absence of photo-excitation, from continuous random thermal excitation causing pixel fires. If one triggers the integration of charge randomly the majority of the time one will see just pedestal events. Occasionally, however, the integration window will fall on one or more of these thermal firings. So if one looks at a histogram of the dark noise from a GAMPD one will see a high statistics pedestal peak, with a Gaussian spread proportional to the electrostatic noise of the electronics system, and low statistics distributions related to a single or more pixel fires, Fig 4.2. The magnitude of these higher charge distributions above the pedestal in conjunction with the length of the integration gate gives a measure of the noise rate of such a device.

Gain

The gain of GAMPD devices is defined by the charge held by each pixel, directly related to the capacitance of each pixel APD. The gain can be calculated by considering the charge stored by each pixel and dividing by the charge on the electron to give the number of electrons in an avalanche, $G = C/e(v_b - V_{bd})$. The gain of a GAMPD for a given over-bias can be determined from the separation of the quantised charge distributions. This process can be performed using dark count events or by low intensity illumination with a light source by calculating the separation of the pedestal and first photo electron peak, Fig 4.2. A high statistics dark run of a GAMPD will provide a clear pedestal and first photoelectron peak and we call such data collection pedestal runs.

The breakdown voltage of a device increases with temperature, and hence the gain decreases when a device is run at constant bias voltage. This must be considered when running a device for a long time.

4.2 Scintillator and Wavelength Shifting (WLS) Fibre

The majority of the sub-detectors in the ND280 use plastic scintillator as their active medium. The plastic scintillator bars used in the P0D, FGD and throughout the ECal are all extruded from a common material and method at FNAL [100]. The extrusion process involves the heating of polystyrene pellets and addition of scintillation dopants.

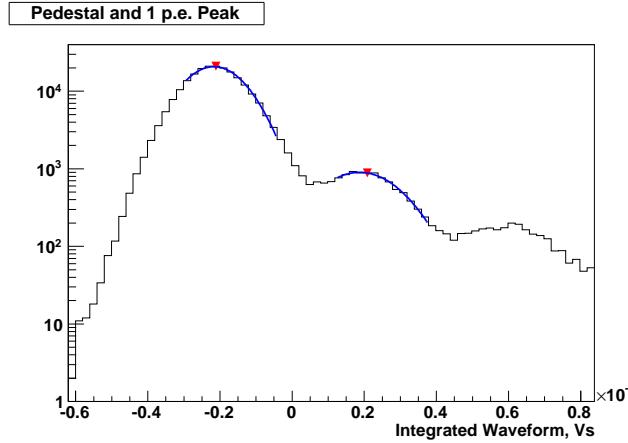


Figure 4.2: High statistics pedestal and first photo electron (p.e.) peaks are found and fit with a gaussian with the gain and associated errors calculated from their separation.

Then the resulting mixture is forced through an extrusion die, which determines the cross sectional design of individual bars. As the extrusion is a continuous process it is possible to produce many lengths of bar. A 0.25 mm thick layer of TiO_2 is co-extruded on the surface of all bars as a reflector to increase light containment. This method has been used at FNAL to produce plastic scintillator bars for many previous high energy experiments including the K2K and MINOS neutrino experiments.

There are two different scintillating dopants used. Poly(p-phenylene oxide) or PPO, included as a 1% dopant, is an organic scintillator whose purpose is to absorb the ionising energy of the passing particles which excite a π -orbital electron into a higher energy state. The de-excitation of the electron emits a photon, with maximum fluorescence at low wavelengths \sim 340-400 nm, which becomes the signal for the ND280 detector.

The second, 1,4-bis(5-phenyloxazol-2-yl) benzene or POPOP, is employed as a secondary or wavelength shifting scintillator and is included as a 0.03% dopant. This is included because the low wavelength photons produced by PPO are quickly attenuated in the scintillator. POPOP absorbs this primary scintillation light and re-emits a photon at a longer wavelength (the spectrum peaks at 410 nm). Because the ionisation energy must interact with a dopant molecule to produce light, a quenching effect has been observed in scintillators, described using a semi-empirical formula by Birks [101]. As the deposited energy grows, so will the light produced, but the lack of available scintillator molecules will limit the growth. This effect has been studied in detail for plastic scintillator [102], and is especially important for protons which can deposit a large amount of energy in a short track. The attenuation length of plastic scintillator, of the order 40 cm [103], is far shorter than the length of even the shortest bars used

in ND280. To overcome what would be a massive loss of light, large attenuation length light guides in the form of wavelength shifting (WLS) fibres are threaded through the centres of the bars to deliver light to a photosensor for readout.

The fibre selected is Kuraray Y11(200), which has an attenuation length of 350 cm, much longer than that of the extruded scintillator. The scintillation light that enters the fibre will excite a scintillating dopant. When the scintillation light is re-emitted (peak of ~ 476 nm), then some of the light will be travelling in such a direction as to be totally internally reflected within the fibre, travelling to a photosensor at one or both ends.

The die used to produce the ECal scintillator gives a bar of cross-section 40×10 mm with a central hole designed to be of diameter 1.8–2 mm into which the WLS fibre will be threaded. FNAL also produced extruded scintillator fusing the same material and method for the K2K SciBar detector. The SciBar die produced bars of cross-section 25×13 mm with a central 1.8 mm diameter hole for the insertion of WLS fibre.

Comparing the ECal test production bars with the well established SciBar scintillator will allow us to monitor the quality of the scintillator produced. The use of the MPPC photosensor and prototype housing will allow us to determine if the light yield from an incident minimally ionising particle (MIP) meets the specification laid down for the requirement of the detector physics goals.

4.3 Experimental Setup

The Muon Telescope in Sheffield consists of two light-tight scintillator paddles each coupled to a 1" EMI 9831KB photomultiplier tube (PMT). The PMTs are biased at their suggested optimal 2.2 kV and their output channels passed to a coincidencing device.

The binary coincidence output was used as an external trigger for the oscilloscope. The larger area of the muon telescope paddles (40×18 cm) with respect to the scintillator being tested meant that this triggering provided pedestal, thermal noise and MIP data in a single run. Pedestal events occurred when a muon passed through the two telescope paddles but not the test bar. Thermal events occurred as pedestal but the passage of the muon coincided with a thermal fire of one or more pixels. MIP events contained a muon which passed through all three.

For each scintillator bar a common setup was employed. Each bar was supported upon a wooden jig and held in place by foam supports. Attached to the near end of the jig

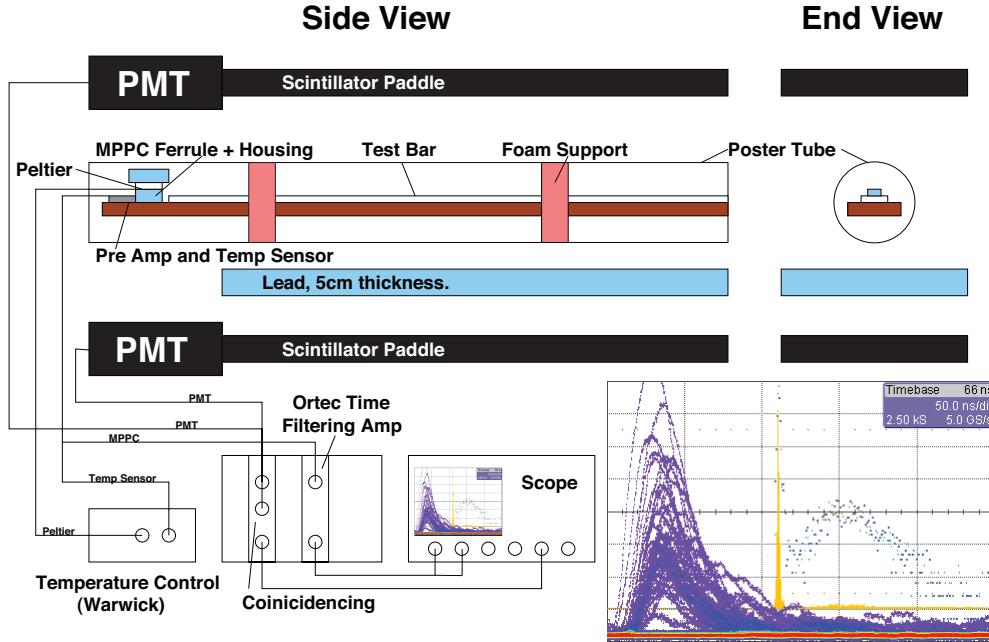


Figure 4.3: The experimental setup for the light yield studies.

was the MPPC, along with its housing, temperature control and pre-amp. The MPPC was housed inside a prototype version of the MPPC connector that is to be employed in the ECal, designed by the Colorado State University. The female half of the connector safely housed the MPPC while the WLS fibre was glued into the male half. The connector safely married the fibre end to the MPPC. This connector was then placed inside an aluminium block to which a Peltier heat pump and a temperature sensor were attached. The temperature sensor and Peltier were connected to a control box, designed and built by the University of Warwick, which monitored the current supplied to the Peltier to keep the setup at a constant temperature.

The jig and all components attached were inserted into a dark poster tube to keep the experiment light tight. At the front face of the poster tube there was an inlet, with a standard CPU fan attached, and an exhaust tube to promote air flow. This was installed as the enclosed environment was not sufficient in size for the heat sink to dissipate heat from the Peltier, which resulted in failure of the Peltier in early data runs.

The pre-amp used was a basic trans impedance amplifier. The pre-amp output was

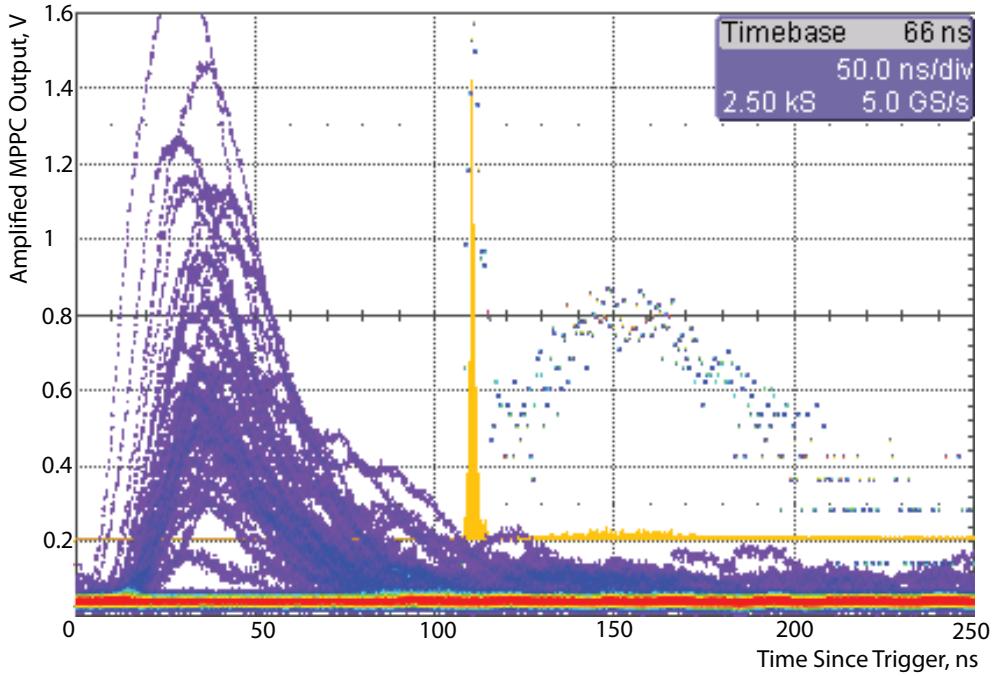


Figure 4.4: A persistency plot showing a number of waveforms received by the oscilloscope after amplification and time filtering of the MPPC signal. The yellow distribution in the background is a histogram of the integrated charge of a number of events, one can see a clear low value pedestal peak. The blue dots show this distribution in log scale, aside from the pedestal peak one can also see the MIP extended peak.

passed through an Ortec 474 timing filter amplifier, providing a coarse gain of $\times 6$ and a fine gain of $\times 2$. The amplification improved the pulse shape of the MPPC but also spread the waveform duration.

For each triggered waveform the integrated pulse area in a 150 ns gate was histogrammed on an oscilloscope. The gate was selected from observing a number of waveforms; Fig 4.4 shows that the vast majority of waveforms input into the oscilloscope, from the amplification of the MPPC signal, resided within the first three 50 ns divisions.

4.4 Analysis

The analysis of histogrammed pulse area data was done using the ROOT statistical toolkit [104]. The peaks of the pedestal and first photoelectron distributions are found using the ROOT TSpectrum::Search function. Gaussians are fitted to a small range encompassing each peak. The mean of the pedestal μ_{ped} and first photoelectron peak μ_{1pe} fits are used to calculate a translation of data $x \rightarrow x'$ from x in nV s to x' in

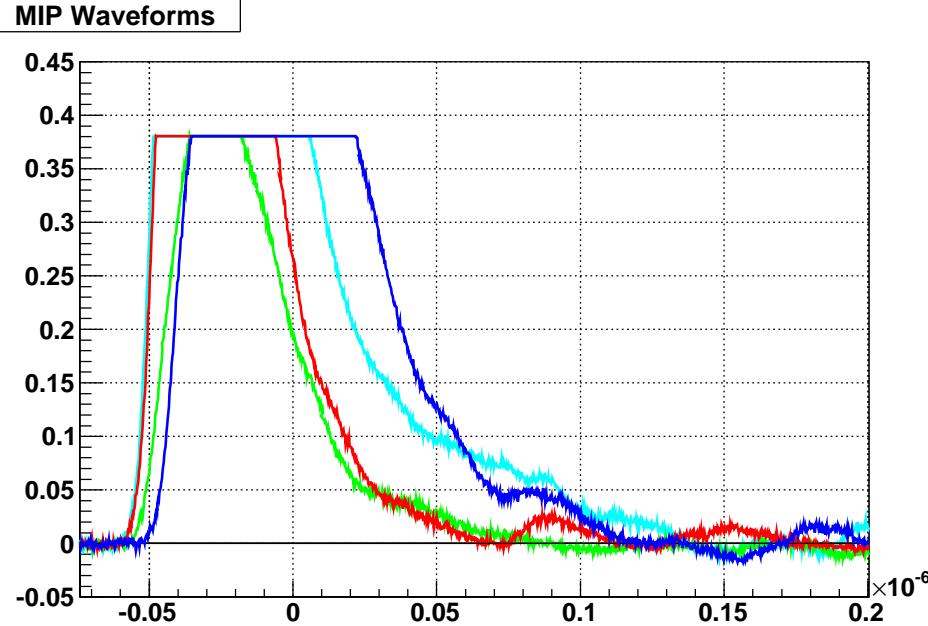


Figure 4.5: MPPC waveform truncated by the oscilloscope.

photoelectrons where

$$x' = \frac{x - \mu_{ped}}{|\mu_{1pe} - \mu_{ped}|}. \quad (4.1)$$

The rebasing of the pedestal mean to zero is performed in the numerator. The scaling of the data, in the denominator, is calculated from the distance between these two mean values, the gain. Assuming linearity, as is expected at this low end of the MPPC dynamic range, one can then determine the mean signal from a MIP by fitting a peak to the extended high p.e. distribution in the rescaled data.

Initial runs showed a nonlinearity in the p.e. peak separation, with misalignment of even low photoelectron signal peaks when transformed to photoelectrons. The MPPC devices have been shown to be linear and a charge injection investigation of the pre-amplifier showed a linear response. This nonlinearity was finally attributed to the limits of the vertical sampling performed by the oscilloscope.

The oscilloscope has a vertical accuracy of 8.5 bits and can only integrate over its display, up to a maximum 7.5 vertical divisions. For resolution of the pedestal and first p.e. peak a scale of 50mV per division is required. This allowed for a maximum signal size of just 0.375 V and it was seen that most of the MIP signals were saturating, Fig 4.5. This truncation introduced a non-linearity in the integrated values of the waveforms which in turn led to a change in peak separation at high p.e. signals.

To overcome this problem the MPPC signal was divided exactly in two using an optical splitter and entered into two separate channels of the scope. One channel had a resolution of 50 mV per division and aimed to measure the pedestal and first p.e. peak separation for calculation of scaling constants. The second channel was set at a coarser 200 mV per division resolution to ensure that the MIP signals would not saturate the vertical sampling of the scope display. Apart from the difference in vertical scale the channels' setups and histogramming were identical, with the 150 ns integration window retained from previously. The 200 mV vertical resolution for MIP signals was chosen after viewing some 2000 MIP pulse shapes.

The statistical error on the scaling factor of the data, the denominator in equation 4.1, is calculated from

$$\Delta(\mu_{1pe} - \mu_{ped}) = \sqrt{\Delta\mu_{ped}^2 + \Delta\mu_{1pe}^2}. \quad (4.2)$$

The statistical error on the entire translation is obtained from the combination in quadrature

$$\frac{\Delta scale}{scale} = \sqrt{\left(\frac{\Delta\mu_{ped}}{\mu_{ped}}\right)^2 + \left(\frac{\Delta(\mu_{1pe} - \mu_{ped})}{(\mu_{1pe} - \mu_{ped})}\right)^2}. \quad (4.3)$$

The extended, high standard deviation, MIP peak is found, again using the ROOT TSpectrum class. A Gaussian is fitted to a range about this peak with the mean taken as the average photoelectron yield per MIP. The associated errors in the MIP fit are combined with that of the scaling to give a final mean photoelectron statistical error

$$\Delta p.e. = \mu_{MIP} \sqrt{\left(\frac{\Delta\mu_{MIP}}{\mu_{MIP}}\right)^2 + \left(\frac{\Delta Scale}{Scale}\right)^2}. \quad (4.4)$$

4.5 Results

4.5.1 First Batch Prototype ECal Bar

The first batch of prototype bars from Fermilab showed troubles with the TiO_2 coating process and chosen design. The TiO_2 was only a thin, powdery and dimpled coating on the bars, one could see the WLS fibre clearly through the coating. The rectangular cross section of the bars made extrusion of a round hole impossible in the FNAL setup. The WLS would instead be threaded through a extended oval shape hole with its major axis in the bar's smallest dimension. This, coupled with the inherent weaknesses in extruded material, led stress fracturing to occur around this weakened region when the bars were machined as can be seen in Fig 4.6. This would reduce the light yield from the bars. Photons would multiply scatter and be absorbed by the fractures.

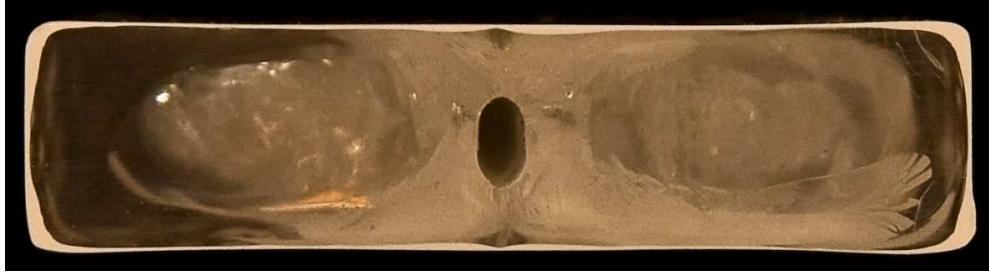


Figure 4.6: Stress fracturing around fibre hole. Note the elongated cross-section of the fibre hole.

For this initial investigation a 60 cm length of bar was cut from the end of a 2 m extruded length. The piece selected was chosen as it displayed little fracturing.

There was a large variety in p.e. yields obtained in the initial runs, which was attributed to poor light coupling between the WLS fibre and the MPPC. The initial runs employed a first prototype connector from the University of Warwick. This prototype was far from a final product and permitted a large freedom of movement between the MPPC housing and the fibre ferrule.

The second prototype connector used to obtain the data reproduced in Fig 4.7 was of a much more solid design. The MPPC was housed with a foam backing. This gentle pressure coupled the WLS fibre with the front face of the MPPC acting as a spring to gently push the MPPC into the fibre ferrule. The ferrule was held solidly in place with two prongs that clipped over the device and held it firm against the face of the MPPC.

4.5.2 Second Batch ECal Prototype

The second batch of prototype ECal bars were of production standard. The TiO_2 coating was opaque and fully glazed. The central hole was still oval with the major axis ~ 3 mm, along the bar's thickness, twice that of the minor ~ 1.5 mm, along the bar's width, but machining of these bars produced no fracturing, and a 60 cm length was again cut from a 2 m length.

The stability of the measurement was tested with separate short, overnight, and long, two day, runs. Between runs the fibre was removed from the MPPC housing and replaced. The photoelectron results agree between the two runs within statistical error when using the new connector design. This was an initial proof of concept for the new foam spring pressure coupling connector. The pressure coupling was also assumed to contribute to the increase in light yield alongside the improved TiO_2 coating. A much

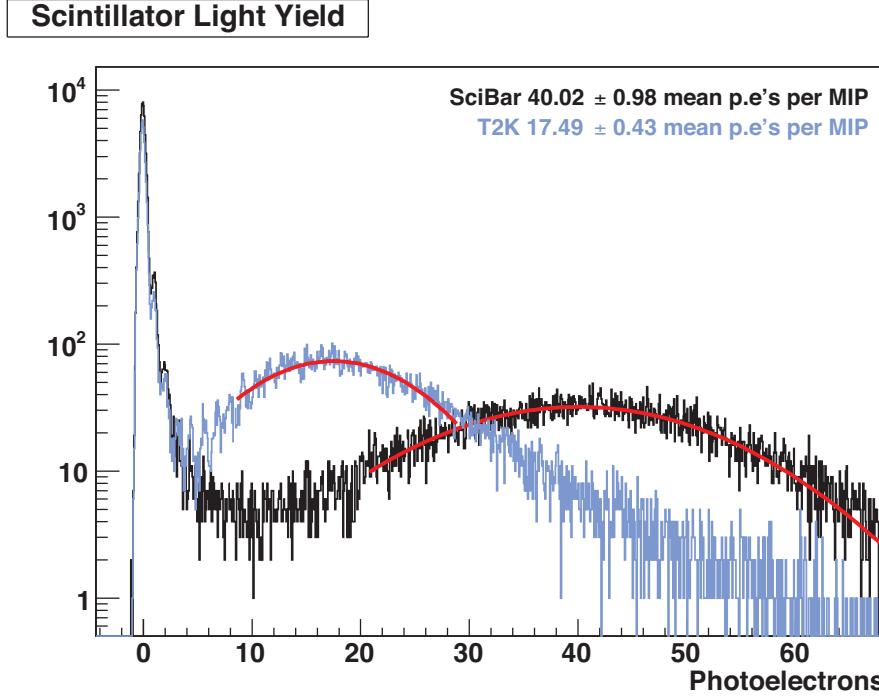


Figure 4.7: First prototype ND280 ECal scintillator light yield, and first attempt at SciBar bar light yield measurement using first prototype MPPC connector.

higher light yield of $25.25(\pm 0.51)$ p.e./MIP/cm was measured in the high statistics two day run.

4.5.3 SciBar

The SciBar [99] bar was measured with the new connector and optically split pedestal and MIP measurement setup. Before measuring it was expected that the light yield would be greater as the SciBar is 3 mm, 30%, thicker than the ECal bar, which allows muons a longer path to excite scintillation light. The smaller width of the bar required a much longer run time to achieve ECal bar comparable statistics and related errors. Again a low stats run, overnight, and a high stats, three day, run gave consistent results within error, establishing the soundness of the connector design and fibre MPPC coupling.

The 1.3 cm thick SciBar scintillator provided a much higher light yield of $32.77(\pm 0.86)$ p.e./MIP/cm.

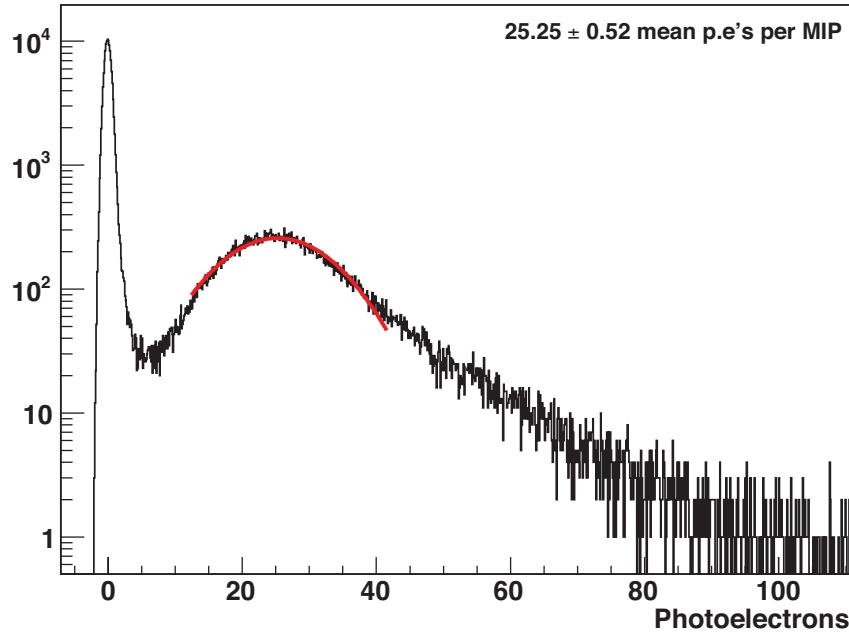
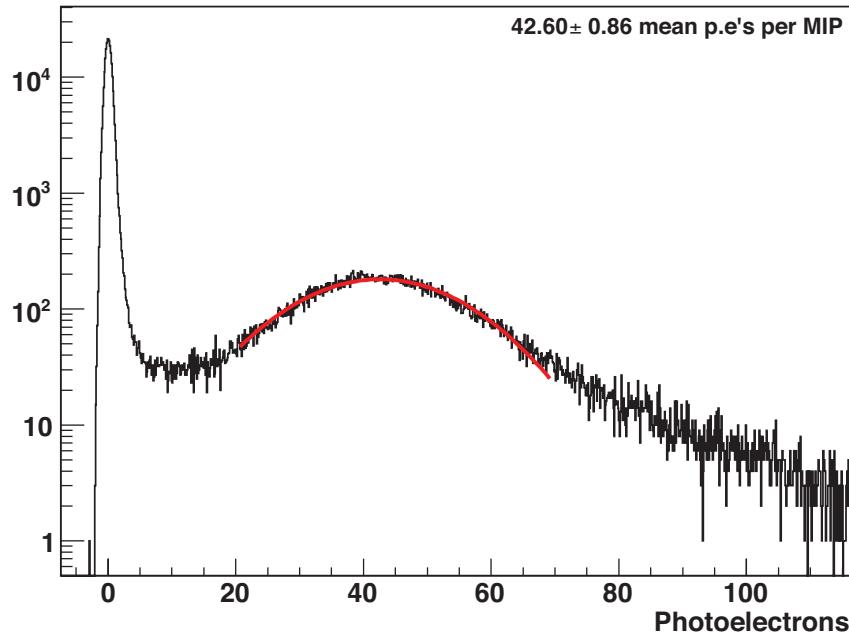
T2K ECal Scintillator Light Yield**SciBar Scintillator Light Yield**

Figure 4.8: Rescaled plot of the high statistics MIP light yield measurement of a second prototype batch ECal scintillator (top) and SciBar scintillator; measurements made with foam spring coupled MPPC connector.

4.6 Conclusions and Discussions

The light yield measured from SciBar cosmic ray data of 10.8 p.e./MIP/cm [3] was measured using the very same Kuraray Y11(200) WLS fibre but a multi anode PMT, which have a typical quantum efficiency (QE) $\sim 10\%$. MPPCs however have been measured to have a photo-detection efficiency (PDE, analogous to PMT QE) of $\sim 30\%$ [105] Taking this factor of three into account one would expected a measurement of $10.8 \times 3 \simeq 32.4$ p.e./MIP/cm from our setup.

This back of the envelope calculation agrees quite well with that measured within statistical error. This shows that despite the different technology we can obtain consistent results. It is also proof that this method of obtaining a light yield measurement from the ECal scintillator prototypes is accurate. It addition it is a powerful proof of principle for the various readout elements before front-end electronics.

The mean of 25 p.e./MIP/cm measured for the ECal scintillator is comfortably above what was deemed the minimal 14 p.e./cm (~ 6 p.e./cm after 2 m) satisfactory for physics. The bars therefore met the requirements of the ECal detector group and full scale production of DsECal bars commenced in February of 2008.

Chapter 5

ND280 Offline Software

The packages that make up the ND280 simulation and analysis suite of software are written in C++ and Python and make use of existing high energy physics software libraries. Each package is currently under development, from the Monte Carlo simulation through to analysis tools. In the following chapter I review the main packages as they existed in early 2009, and their importance to the user.

5.1 Neutrino Interaction Monte Carlo Generation

Neutrino interactions are simulated by a dedicated interaction generator. Generators take information on the flavour, energy and direction of a neutrino and the nature of the incident material as an input. The generators then use theoretical models based on data from neutrino interactions to produce an output which is essentially a list of particles created, each with an energy and a momentum. This list can then be passed to a separate simulation which propagates the particles within the detector.

There are a number of neutrino generators available, each of which has been employed by various neutrino experiments. At the moment three generators are in use in the ND280 collaboration: NEUT, NUANCE and GENIE. The three generators differ both in the models used to generate events and in their interface with detector simulations. The differences in the event generation models will not be discussed in detail here; I will instead comment upon the use of each generator.

NEUT [106, 107] was originally written to generate neutrino interactions upon oxygen nuclei within the water Čerenkov Kamiokande detector. It has since been employed in the Super-K, K2K and SciBooNE experiments. The generation model employed by NEUT has received a number of updates from these experiments and was also adapted

for an iron target. Events are generated by NEUT against a single target. The events are output as a straight list and it is then the job of the detector simulation to select the locations of interactions (vertex) within the detector. This can be implemented isotropically but is usually based on a density weighting.

The NUANCE [108] generator also produces neutrino interactions upon single nuclei. Initially written for the IMB detector, it was rewritten for use in Super-K but was not adopted as the official MC. NUANCE has been used by a number of other experiments, such as MiniBooNE [109]. The event generation is split into two tasks. The first is the most CPU intensive and produces a binary file with information on the rates of all interaction types. This is then passed to an event generator which uses MC sampling of these rates to generate each event type and final state particles. NUANCE output is again a simple list of events and final state particles with placement of the interaction vertices left to the detector simulation.

Both NEUT and NUANCE work quite separately from any detector geometry, requiring only knowledge of the target nuclei and specific fraction, with vertex placement left to the user/detector simulation. GENIE [110] on the other hand works in symbiosis with complicated detector environments and provides information on the positioning of vertices. It uses a ROOT geometry of a detector and a detailed description of the flux topology of the neutrino beam to provide an estimate of the location of neutrino interactions. As with NUANCE, the generation is split into two stages. The first calculates cross-section splines for all relevant materials, neutrino energies and interaction types. The maximum possible cross-section that neutrinos could possibly experience whilst passing through the detector is also calculated for the range of energies of the beam. Neutrinos are then propagated through the detector and the possibility of interaction determined from the cross-section experienced by each individual path re-weighted by the maximum possible. Interactions are then generated on a specific nucleus along the neutrino's path. The position of the interaction and information of the nucleus is then saved along with all other event information in ROOT tree format.

Although extremely successful in homogeneous detectors, such as the water Cherenkov Super-K and IMB, the single target interaction generation of NEUT and NUANCE are viewed to be deficient for the requirements of the ND280 detector. Composed of many different target materials, GENIE has a clear advantage when generating events in the ND280.

All three generators base calculations of quasi-elastic, exclusive resonance and DIS cross-sections around theoretical and phenomenological models from LlewellynSmith [111], Rein and Sehgal [112] and Bodek and Yang [113] respectively. There are of course

many subtle differences across the generators but the most important differences arise from their hadronic simulation of both primary generation and intra nuclear transport. The hadronisation models employed have common links in that all employ KNO scaling at low energies and move to using PYTHIA at high energies to generate primary hadronisation. The generators differ in the empirical expression they employ alongside these common methods and therefore essentially the data that they use. Hadron transport models employed are either cascade Monte Carlo simulations or in house developed data-driven models based on hadron-nucleon data.

5.2 Software Framework

5.2.1 The ND280 Data Framework: oaEvent

The `oaEvent` package is a library that defines the ND280 data structure and format, making heavy use of the ROOT [104] data analysis toolkit.

All data, simulated, run time or reconstructed, are saved within specialised objects. I will take a little time to talk about a few of these objects, shown in Fig 5.1, now as we will discuss their use in this and further chapters. The THit object is the most basic element of detector data, containing just a time, geometryID (unique to each detector element) and raw charge deposition information. There are then higher level objects which contain ever more complex data as reconstruction proceeds.

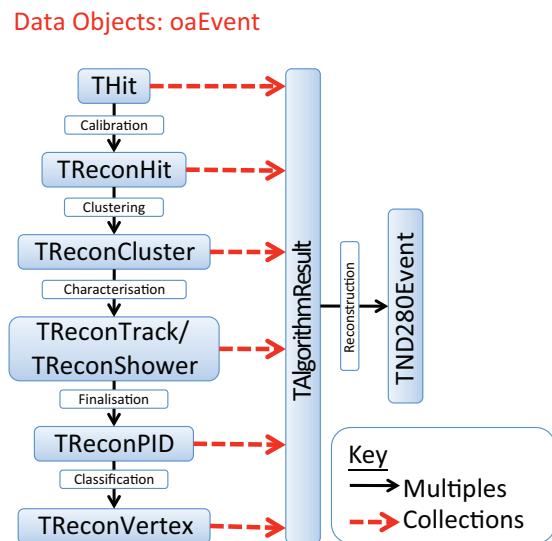


Figure 5.1: Data objects in the nd280 offline software.

TReconHit objects replace the THit objects as the basic building blocks of reconstruction; they contain extra information, such as attenuation corrected signal, reconstructed position and uncertainties of various measured quantities, and may be composed of multiple THits as is desired for dual ended readout scintillator. Multiple TReconHits are saved in TReconClusters, an object which does not assume anything of the group of hits apart from their commonality; it retains basic information such as the total charge deposit and weighted position of contributing hits.

TReconTrack and TReconShower objects assume specific classes of geometry of a group of hits and can be composed of multiple TReconClusters which may span several sub-detectors. Each has specific data relating to track and shower geometries such as curvature of tracks or opening angle of showers. Multiple showers and/or tracks may then be saved as the constituents of a TReconPID. This object allows easy access to high level information such as the most probable species of particle responsible for the contributing reconstructed showers and/or tracks.

After global reconstruction these TReconPID objects can then be saved to TReconVertex objects which represent a common neutrino interaction point. Currently global reconstruction at this level is in development and so TReconPID objects are the highest level of reconstruction data container used in this thesis.

At each stage of reconstruction one may write a collection of any number of these reconstruction objects to an TAlgorithmResult object which is identified by a name. Such objects are containers for the result of each stage of reconstruction: are used to pass data between algorithms in a single object; and may be written directly to the TND280Event.

TND280Event objects collect all information, truth or reconstructed, and are saved on an ‘event’ by ‘event’ basis to a ROOT file. The slower time response of the TPC electronics and possible interactions between beam bunches, such as (delayed) Michel electrons, have dictated that each ‘event’ that is passed into the software will initially constitute the data from an entire beam spill (section 3.2). The electronics of all the scintillator detectors have a much faster response and will reset between each of the bunches in a spill. In user defined MC data, however, an ‘event’ can range from a single particle, produced with set energy and direction, to a full beam spill, with numerous neutrino interactions throughout the ND280 basket, magnet or hall.

The oaEvent library has few dependencies, leaving just a small number of packages required to read any ND280 data file. This has been designed so that an end user with only a minimum software installation will be able to view and analyse the data. This will be covered in more detail in section 5.5.2.

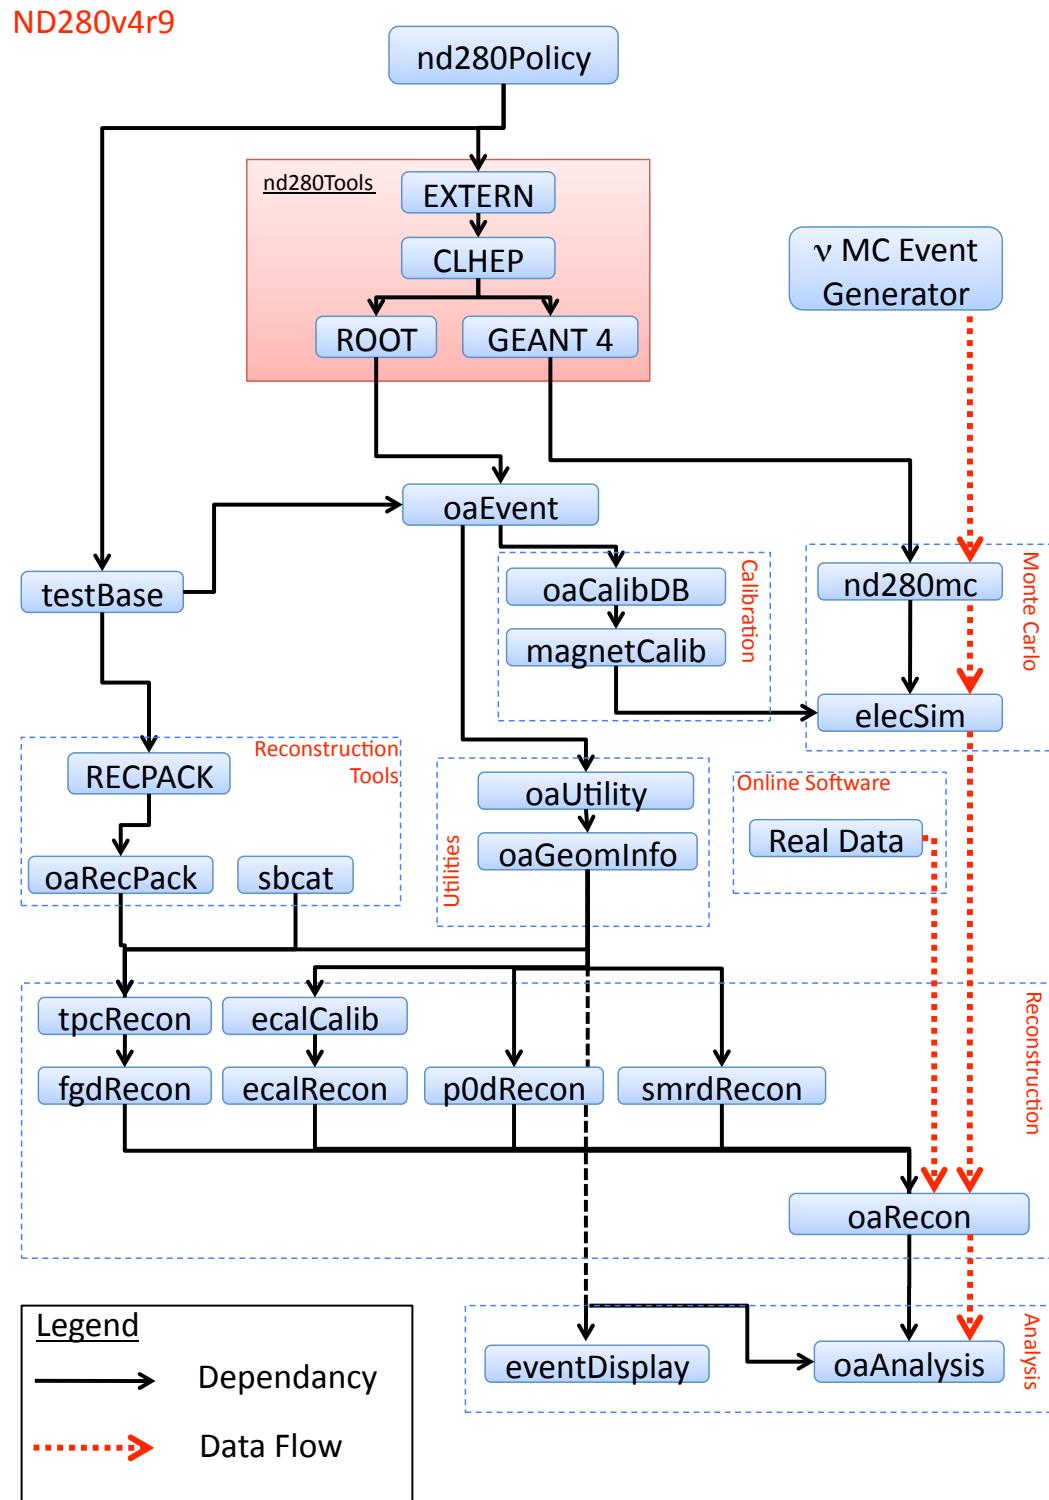


Figure 5.2: Data flow and dependancy of all nd280 offline software.

5.2.2 Offline Database

The offline database is intended to hold calibration constants and run information which will be called during event reconstruction. These calibration constants are currently hard coded into many packages but as detectors are commissioned and calibrated this database will become fully implemented.

5.2.3 User Interface: `nd280Control`

The first user release, v4r1, of the ND280 offline software contained a package aimed at steering the execution of the entire chain of ND280 software, from neutrino interaction generation right through to the production of analysis files. The `nd280Control` package also performs book-keeping of MC production, automatically generating log files to be stored in the offline database. The package is written in Python and is modular in nature with each code module responsible for one stage of the processing chain.

5.3 Monte Carlo (MC) Simulation of the ND280

The Monte Carlo (MC) method has become an essential predictive tool for particle physics experiments. The progressive increase in the complexity of particle detector geometries and associated electronics demands ever more sophisticated MC simulations.

5.3.1 Detector Monte Carlo Simulation: `nd280mc`

The GEANT4 toolkit [114, 115], maintained at CERN, is a dedicated MC package that provides tools for simulating detector geometries and propagating particles through them. The `nd280mc` package uses GEANT4 to construct the detector geometry and define particles with associated interactions.

The detector geometry is constructed from the creation of volumes which then have material properties associated to them. Volumes of the detector that actively measure a particle's passage, scintillator and TPC gas, are given sensitive properties. Whenever a particle passes through a sensitive detector region this is logged and information about the particle and its energy deposit within that material is stored.

There are two main methods of seeding particle events into the `nd280mc` simulation. The simplest insertion is from the general particle source. This is a highly customisable particle gun which introduces single particles into the detector volume with a user defined type, position, direction and energy.

The second and most important method of seeding events is to accept final state particle inputs from dedicated neutrino interaction generators (section 5.1). It has been decided that the `nd280mc` simulation should be able to accept input from as many of the available neutrino interaction generators as possible. Currently NUANCE, NEUT and GENIE are all fully supported by `nd280mc`.

Macros are given as inputs to `nd280mc` to control many aspects of the simulation. Alongside the generic GEANT4 simulation options there is an additional option list unique to `nd280mc` which allows the user to manipulate all aspects of the ND280 detector simulation. These options can be simple, such as the number of events to generate or where to seed events in the detector volume, or more complex, changes in the detector geometry or how the data is to be saved. I will discuss specific options where they have been used in the studies in this thesis.

The output of `nd280mc` is a ROOT file containing two objects. The first is a collection of `TND280Event` objects that contains all event information. The second object in the `nd280mc` output file contains the detector’s world geometry, which is translated from the construction in GEANT4 into ROOT readable format.

5.3.2 Electronics Simulation: `elecSim`

The final stage of the MC simulation is then to convert the energy deposition information saved by `nd280mc`, in sensitive detector volumes, from “true” information into a form similar to the output expected of the detector electronics. All of this is handled by the `elecSim` package.

The methods for the scintillation based detectors vary only slightly, with adjustments arising from bar size, attenuation length and analysing chip (AFTER for the FGD, TriP-T for the others). The truth hits created by GEANT4 are used to find the location of the hit in the bar and the energy which would have been deposited in the bar. The energy is converted into a number of photons produced by the WLS fibre, after accounting for Birk’s quenching effect (described in section 4.2). Bars with photosensor readout at both ends, such as in the DsECal, then have an attenuation effect applied to the photons, based on the distance the G4 hit was from the photosensor and using attenuation splines. These splines are either based on a single attenuation length or calculated from LTRI, a simple MC to estimate the light yield from MIPs at different positions along the bar.

Each hit is saved as a `THit` object, stored by sub-detector, and this information is output as it would be for real event data. All the truth information, particles created,

their paths and the hits they contributed to, is still stored, but it is hidden and only used for cross checks or analysis.

5.4 Event Reconstruction

The ND280 reconstruction follows a modular design with each sub-detector performing an initial reconstruction. The results of these independent methods are then combined across sub-detectors using methods provided by the RECPACK [116] package. RECPACK has been designed to perform fitting upon a number of different reconstruction object types, defined by a ‘state’. ND280 versions of these objects were created within the oaEvent package and are used by all the reconstruction packages.

5.4.1 Global Reconstruction: oaRecon

The oaRecon package provides global reconstruction of events in the ND280. This is essential to improve the systematics and statistics of standalone sub-detector measurements. oaRecon starts by running all independent sub-detector reconstruction methods, which are explained in the rest of this chapter, each of which returns a number of TReconObjects.

These reconstructed objects are then connected or matched by oaRecon with reconstructed objects in neighbouring sub-detectors using RECPACK tools. Currently the only successful cross sub-detector matching has been performed between the TPCs and FGDs. This is the easiest possible matching, however, as both detectors have excellent resolution and the small density of material separating them greatly reduces the likelihood of particles scattering when crossing sub-detector boundaries. In contrast, there is often dead material between an ECal and the inner basket region. The change of detector orientation from inner detectors to ECal also adds complications. The work for TPC-TECal and FGD-TECal matching is maturing fast. The poor resolution of the P0D ECal means that it cannot utilise the RECPACK tools. Matching between the P0D and P0D ECal will therefore have to be achieved via different methods.

5.4.2 Tracker Reconstruction

The two tracker sub-detectors are designed to work together during event reconstruction. They are the first working example of inter sub-detector reconstruction and use RECPACK tools in this process. Each sub-detector does an initial independent reconstruction, producing output objects which range from simple clusters to tracks or

showers with advanced PID likelihoods. These individual results are then sewn together by RECPACK.

Time Projection Chamber Reconstruction: `tpcRecon`

The first step in the TPC reconstruction chain is the application of the gain calibration constants and the removal of dead and noisy channels. The output of this process is a waveform that represents the charge acquired by each pad along the readout time.

The next step is the clustering of waveforms in the XY plane. For waveform clustering purposes the time of a charge deposit is employed as a substitute for the X position as the two are related by $x = v(t - \Delta t)$, with v the drift velocity and an unknown offset Δt . The Y axis, MICROMEGAS readout row, is selected because the positive Z axis is the preferred direction of a track; XZ methods are equivalent and will also be used.

Prior to the final track reconstruction, the drift distance along the X axis has to be reconstructed to be able to predict the size of the electron cloud due to transverse and longitudinal diffusion. This is achieved by determining the time at which the track was created (T_0 of the track) via matching with objects in the fast detectors, most probably FGD and ECAL. After this clusters are joined into tracks using a modified version of the SciBar Cellular Automaton Tracker (sbcat) [99] pattern recognition algorithm. The sbcat algorithm creates segments between any two clusters within a maximum couple of rows. The cellular automaton then works to connect track segments which share an edge and also pass a cut criterion, namely that the clusters overlap in space and time. Once a tree of segments has been constructed, the longest possible path is selected and stored as a track. Sometimes a track will branch in the TPC, normally from a hard δ ray, and if this leaves a number of independent segments another small track will be created.

The reconstruction of the track kinematics is then done using two different methods: single point reconstruction and likelihood fit. Finally the reconstructed tracks are then matched to tracks in other TPC blocks and refitted to improve the accuracy of the track parameters.

Fine Grained Detector Reconstruction: `fgdRecon`

The FGD reconstruction follows on directly from and requires `tpcRecon`. A Kalman Filter [117, 118], provided by RECPACK, is used to match TPC tracks to FGD hits. The TReconTrack from the TPC has a series of nodes, each with its own state, along the path of the track. The node closest to the FGD is selected and its state used to

define the position, direction and curvature of the track, which is then projected across the dead space into the FGD. The algorithm calculates a χ^2 for each FGD hit and this propagated track. A simple cut is then applied on this χ^2 goodness of fit to determine if the hit is to be included in the track. An additional requirement is imposed, in that each FGD module must have at least one hit in each view selected as a track candidate for the track to be extended.

Once the track hits have been found, RECPACK is used to create an FGD track, which is then connected to the TPC track. It is possible for a particle to pass through multiple sub-detectors, so a high energy muon starting in the upstream FGD may well have a track FGD1-TPC2-FGD2-TPC3.

The fgdRecon package also has a standalone reconstruction algorithm, called Iso-Recon, which is applied to any hits not clustered by the first method. This is required to find short tracks emanating from interaction vertices, which either do not leave the FGD or travel roughly perpendicular to the beam and so do not pass through, or are not reconstructed by, the TPC.

The first stage is a clustering method, based either on the TPC's cellular automaton method or the Hough transform [119] employed by the P0D. These candidate tracks are then passed through a series of cuts, testing that they are relatively straight, that there is never more than one empty layer between hits and that tracks have not been repeated. The 2D tracks passing these cuts are then matched between the XZ and YZ views and the resulting XYZ track is stored in a TRecon object.

5.4.3 P0D Reconstruction: p0dRecon

Unlike the tracker sub-detectors, the P0D currently acts as a self-contained detector, not requiring input from other sub-detectors. As such, p0dRecon is at present a fully contained process designed to reconstruct the π^0 inclusive and exclusive measurements set out in section 3.5.4. The methods for connecting P0D reconstruction objects to those in the surrounding sub-detectors are in preparation.

P0DRecon uses a two step method to reconstruct events. Firstly, tracks from MIPs are identified and removed from the hit sample, possibly also providing a vertex. Secondly, the remaining hits are clustered using a shower clustering algorithm.

The relatively straight path of charged particles within the P0D are picked out from the hit pattern of bars using a linear Hough transform, a feature extraction technique employed in digital image processing. A Hough transform converts all hit positions in XZ space into curves in $r\theta$ space. When all the curves of a set of hits have been

plotted there will be a common point at which a number of curves cross. The more lines crossing at a single point in $r\theta$ space, the more linear the feature corresponding to the unique line given by the co-ordinates of that point.

This method does not account for any anomalies that arise from the complicated geometry of the P0D detector. The majority vote result of the Hough transform is therefore just used a seed for a clustering algorithm. The clustering performed is a spatial nearest neighbour aggregation where new hits are clustered if they lie with a certain spatial distance of a cluster hit. This distance is optimised to take the geometry of the detector into account and is also weighted by the angle given by the Hough seed.

The result of this initial clustering is then used as a seed for a path-following tracking algorithm. The algorithm looks in consecutive layers in front and behind the cluster for hits to extend it. The whole process is then repeated, with the first set of tracks removed, until all significant tracks are reconstructed. MIP tracks are separated from electromagnetic showers with the intention of recognising and removing muons and protons from the sample of hits so that a later shower clustering method will only find photon showers for π^0 reconstruction. This is achieved with a basic particle identification (PID).

The track-like objects are then matched between the XZ and YZ views, to make a 3D object. A Kalman filter is used to fit the tracks and locate a vertex. A vertex from charged particles is normally more accurate than a reconstructed vertex from two photons, which is useful for the π^0 reconstruction. A further PID step is then applied to the tracks to select between muons and protons, based on the charge deposit at the end of tracks.

Once the charged track contributions have been removed, the remaining hits are passed to a shower reconstruction, to attempt a π^0 reconstruction. This is done by mapping all hits on to a sphere about a potential vertex point. The point is moved around, and when the hits form two separate groupings, each is reconstructed into a shower and an invariant mass is calculated. If this matches the mass of the π^0 , then a TReconPID is created, containing the two shower objects and recorded as a π^0 .

5.4.4 Tracker and Downstream ECal Reconstruction: `ecalRecon`

Hits are extracted from the event, placed into TReconHit objects and collated into TReconClusters according to the beam spill in which they were detected. The TReconHits are then clustered by basic spatial nearest neighbour aggregation and these collections are written into TReconClusters. At this stage a simple track or shower

may be split into a number of basic clusters. This problem is addressed by considering the direction of the primary, largest, cluster and extrapolating to consider nearby satellite clusters, merging them if they pass cuts on position and time.

After combination there are often hits remaining unclustered by the basic clustering. Principal component analysis (PCA) is used to calculate an estimator which is cut upon to determine if extraneous shower hits should be included in the cluster. The 2D TReconClusters in opposing views of each tracker ECal module are then compared to determine if they are a common object. This determination is performed by cutting on a number of variables comparing the topology of the shower and the charge weighted shower centres.

Once confident of a match, the resulting 3D clusters are passed to an energy reconstruction algorithm which uses a likelihood function to estimate the energy of clusters, assuming they are electromagnetic showers. This is performed on all clusters, as a variable calculated in this estimation is used later for particle identification purposes. An optimised hadronic energy estimator is yet to be constructed but is in the pipeline.

The result of the energy estimation, including the variables used, is saved to clusters which then get passed to a particle identification algorithm. This algorithm uses seven variables as inputs for two separate artificial neural networks (ANNs) to estimate the track-shower and EM-hadronic nature of the cluster. The clusters are then used to create TReconPid objects containing all cluster and PID info plus a TReconShower or TRecon track object, dependent on the output of the track-shower ANN. These TReconObjects are created with RECPACK fitting in mind.

The tools to match tracks and showers to inner tracker detectors are in development at the time of writing.

5.4.5 P0DECal Reconstruction: `ecalRecon`

The reconstruction chains of the ECals surrounding the tracker and P0D regions have been separated because of the significant differences between the two detectors, although initial steps of hit preparation are common. The chain of reconstruction within the P0DECal is the main topic of this thesis and is discussed in later chapters. In chapter 7 we describe the reconstruction chain in detail and evaluate its performance upon MC data. The higher level particle identification is discussed in chapter 8.

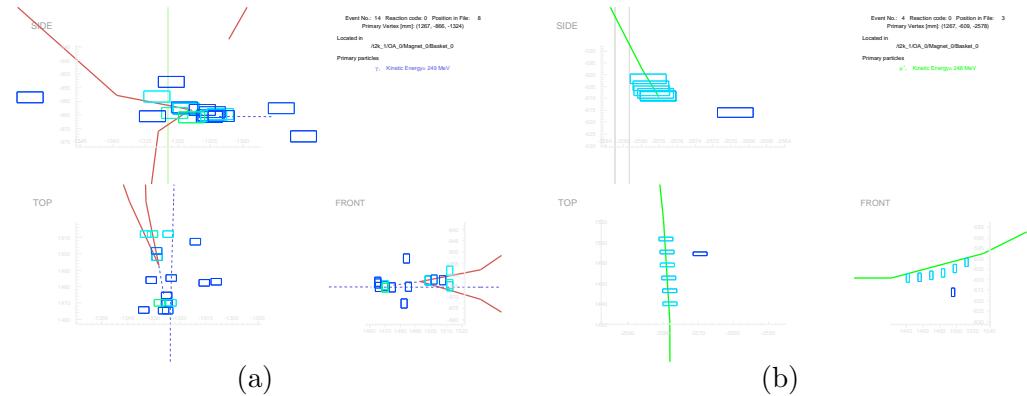


Figure 5.3: Examples of a photon shower (a) and a muon MIP (b) event, as viewed shown in the 2D events display. The three 2D views are shown along with a fourth informational panel displaying event level and primary vertex information.

5.5 Data Analysis

5.5.1 Event Display

There is a 2D, 3 view, event display and since v3r5 also a 3D event display based upon the Event Visualisation Environment (EVE) package included in the ROOT toolkit. The choice of data displayed is highly configurable in both 2D and 3D environments.

The 2D display is split into 4 sections: 3 showing each 2D view of the detector, top, side and front, with the fourth panel displaying information about the interaction vertex, incident and final state particles etc. One can choose to display truth information or just reconstructed data. True trajectories are displayed as a solid line the colour and form of which represents the particle type. Hits are shown by boxes of the physical dimensions of the bar hit and the colour corresponds to the charge deposited in the hit.

5.5.2 Analysis Framework, oaAnalysis

The Mock Data Challenge 0 (MDC0) large-scale testing of full-spill MC production ND280 offline software chain showed that after reconstruction a single event averaged 70 kB in size. The oaAnalysis package was created to distill this large amount of data into a more user friendly event size ~ 7 kB for local storage and analysis. The complex reconstructed and truth data objects are reduced to their pertinent internal variables which are then saved in ROOT trees.

oaAnalysis has a modular design with a number of separate pieces of code distilling

specific information from events and storing it in a individual ROOT tree. The modules come in three classes, each of which is saved in a separate directory of the final output ROOT file. The header directory contains event by event information such as the ID number of each event and the number of the data run. The truth directory contains the true trajectory and vertex information generated by the detector simulation. The recon directory contains modules written to distill reconstruction data. The default modules are designed to extract just the pertinent data, but users can write additional modules if they feel that the default data distilled is not what they require. I am responsible for the P0DECal module which has been used to extract data from events for all following analyses.

Because the file output by oaAnalysis is simply a collection of ROOT trees, an end user wishing to perform analyses would only need a knowledge of ROOT and no in-depth knowledge of the inner working of the ND280 offline software. With true and reconstructed data on interaction vertices a user can study efficiencies and purities allowing the optimisation of any number of analyses. An example of such an analysis will be presented in the final chapter of this thesis.

Chapter 6

Design of the P0DECal

6.1 P0DECal Overview

6.1.1 Motivation and Physics Goals

Cross sections of neutrino interactions in the energy range of the T2K beam are little known. This has been demonstrated most clearly in coherent pion production from low energy, < 2 GeV, neutrino-nucleus interactions. A lack of observed charged current coherent pion production has been remarked upon by both the K2K [120] and recently the SciBooNE collaboration. Latest results from the SciBooNE [121] experiment show a much reduced CC coherent to total CC ratio of 0.67×10^{-2} (90% C.L. upper bound) at a mean neutrino energy of 1.1 GeV compared to the prediction from the otherwise successful Rein-Sehgal model [122] of 2.04×10^{-2} . It is therefore imperative that the T2K suite of near detectors measure not only the flux of each neutrino flavour but also the cross sections of relevant neutrino interaction channels. This is essential if the systematic uncertainties of the T2K measurements are to be understood.

The P0D detector in the upstream half of the ND280 hopes to measure the cross-section of single neutral current π^0 production which is a dominant background of electron neutrino events in the Super-Kamiokande far detector (section 3.3.3). Neutral pions decay almost always into a pair of photons. Successful reconstruction of a π^0 requires conversion and containment of both of these decay photons. To improve the odds of this the P0D employs an upstream and downstream calorimeter section with a larger fraction of lead foil absorber to convert and contain photon showers. However, these sections do not form part of the fiducial volume for π^0 vertex reconstruction, reducing statistics dramatically. Despite this, many photons from π^0 decay, especially those with lower energies, still exit the P0D detector entirely without converting. This

leaves a number of π^0 events where either none or just one of the photons are converted and reconstructed by the P0D.

The P0D must also be able to tag charged current products with high efficiency, to determine if a reconstructed π^0 was indeed a product of a NC reaction or if it was emitted from a CC interaction vertex. As in the tracker the P0D can do this successfully for muons and charged pions that are travelling with an angle not too far from the beam axis. The P0D's acceptance of muon and pion tracks reduces as one increases the angle from the beam axis.

Despite a reshuffle of UK finance it was concluded that an ECal was still required to surround the P0D to attempt to address these issues. The following physics goals were identified for this π^0 detector ECal (P0D ECal) detector:

- Convert π^0 decay photons escaping detection in the P0D.
- Contain the resulting showers from photon conversion.
- Discriminate between photon showers and Muon tracks.

My work on the P0DECal design focused on optimising the details of the basic structure to ensure that these goals were met, within the constraints imposed by financial consideration.

6.1.2 Detector

Like the Tracker ECal the P0DECal is a lead sampling calorimeter consisting of a number of plastic scintillator-lead foil stacks. The scintillator is read out from wavelength shifting fibres which guide scintillation light into novel photosensors called Multi Pixel Photon Counters, MPPCs (section 4.1). There are six separate detector stacks which will surround the four outer faces of the P0D detector; two side, two top and two bottom. For simple access to inner detectors it was decided that the top and bottom modules of both the TECal and the P0DECal should be split in half. This allows the magnet yokes to be split apart with no real complications and gives easy access to the basketed inner detectors.

The tight financial envelope imposed strict constrictions on the final design, with a number of steps taken to reduce costs.

The most obvious reduction in cost would be the number of layers in each stack. A baseline design of six active plastic scintillator layers sandwiching five lead absorber

layers was envisaged to suffice in performing the basic operation of catching missing photons from the P0D [5].

Dies for scintillator bar extrusion are expensive to manufacture and refine. Scintillator used in the P0DECal was therefore restricted to the existing $4 \times 1 \text{ cm}^2$ cross-section bars used in the downstream (DsECal) and tracker (TECal) ECals.

Another cost saving can be made in minimising the number of optical readout channels required. To this end it was decided that the fibres in the P0DECal would be read out at one end, mirrored at the other end. The decision to split the top and bottom stacks leaves modules which are much longer than they are wide. It was therefore decided to align all bars with the beam axis, as opposed to two perpendicular orientations in the TECal. This reduces the required number of channels, makes clustering simpler and reduces the cost and complexity of routing power and readout to a second detector face.

With this long list of restrictions the only variable left is the combined radiation length of material through the detector; the thickness of lead absorber used between scintillator layers. In the rest of this chapter I present the results of a study which uses Monte-Carlo truth and partially reconstructed events to determine the final design of the P0D ECal.

The geometry of the baseline detector design was modified to accommodate different thicknesses of lead absorber from 1 to 5 mm in thickness. General particle source gun particles were then fired through the stack and the physics goals investigated. Neutrino beam interactions can occur at any point within the ND280 detector. This fact along with angle, energy and particle type presents an enormous phase space in which a neutrino event can lie. For the purpose of a design study it was decided to take just a small section of this enormous phase space. We considered single particle events with a flat energy spectrum over a low energy regime, 65 MeV–1 GeV. Each particle was created at the front face of a side P0DECal stack, entering at normal incidence. This was assumed to be the worst case scenario as particles at any shallower angle would traverse more of the detector. It was also the angle at which the P0D itself is most likely to miss particles, therefore making the P0DECal crucial to event reconstruction.

6.2 Photon Conversion and Containment

6.2.1 Photon Conversion

It is hoped the P0D ECal will be able to tag the second missing photon and match to a P0D vertex with just a single photon detected. It is also hoped that the P0D ECal may measure both photons from π^0 decay in the event of both photons escaping detection in the P0D. To achieve this desired improvement in efficiency we require that the P0DECal convert a large fraction of the photons escaping detection in the P0D detector.

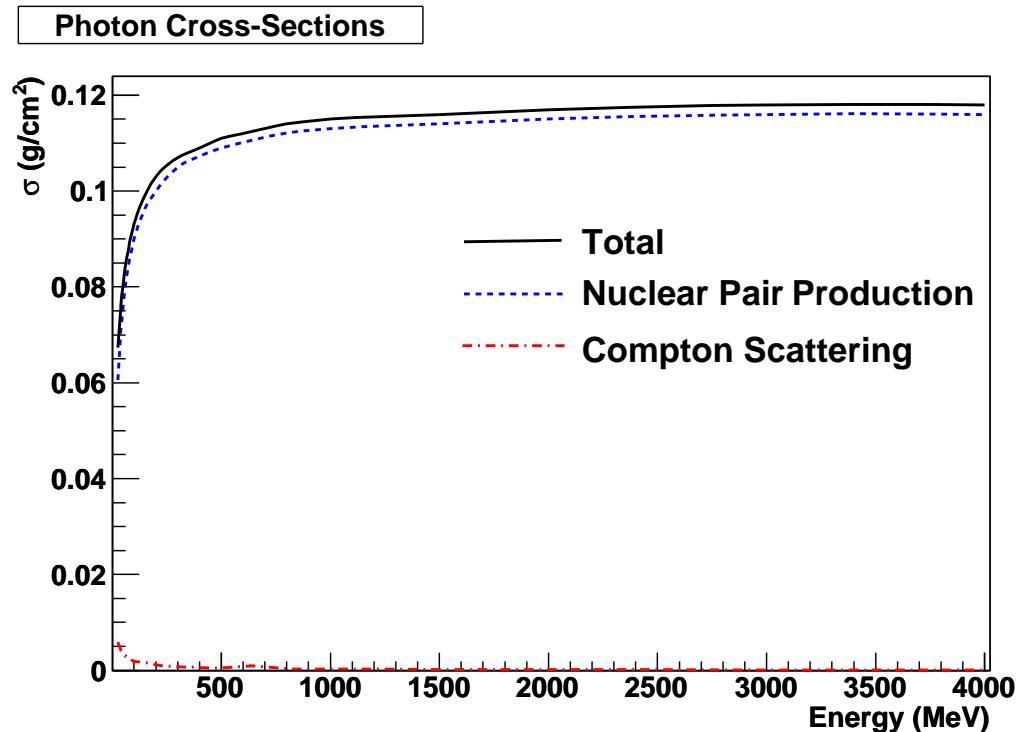


Figure 6.1: Cross-section of photon interactions in lead between 25 MeV and 4 GeV. The solid black line charts the total cross-section, the evenly dashed blue line the cross-section of the dominant nuclear field pair-production and finally the dot-dash line charts the change in Compton scattering cross-section. Data from [123].

The majority material by mass in each P0DECal stack is the lead absorber of which we are to determine the amount. The interaction cross-section of photons in lead as a function of energy, between 25 MeV and 4 GeV, is shown in Fig 6.1.

One can see that the dominant energy loss process in the energy range which bounds the π^0 decay photon energy peak and tail, 60 MeV–1 GeV, is nuclear resonant e^+e^-

pair production, with a small fraction coming from Compton scattering. If we assume the only conversion or absorption of photons in the P0DECal is via nuclear resonant pair production then from the cross section one can calculate a transmission probability

$$T = e^{-\sigma(E)\rho(\Delta x)}, \quad (6.1)$$

where $\sigma(E)$ is the interaction cross-section which is a function of photon energy, plotted in Fig 6.1 in g cm^{-2} , and ρ is the density of the interaction surface, again in g cm^{-3} . In the case of nuclear resonant pair production the density is that of the nuclei cm^{-2} of the element or, as most of the mass is concentrated in the nucleus, then this can be approximated to mass density of the element. If we take the mass density as 11.34 g cm^{-3} and choose a target of 90% conversion, 10% transmission allowed, then we can rearrange the above equation to get an estimate of the thickness of lead required to achieve this goal. For the mean expected π^0 photon energy, 200 MeV, $\sigma_{total} \simeq 0.1 \text{ g cm}^{-2}$, which yields a lead thickness of 20 mm. We therefore select a range of lead thicknesses 1–5 mm per layer as this will allow us to view all other properties up to and just beyond this roughly optimal thickness. Note that photons which enter the P0DECal obliquely will see a greater thickness than would otherwise be viewed at normal incidence.

This was an entirely Monte Carlo truth data analysis, and it used the true trajectory data saved by the GEANT 4 Monte Carlo simulation of the left side P0DECal stack in nd280mc. A particle's true Geant trajectory is terminated when a process absorbs the particle, i.e. pair production. From this logic the final points of true primary photon trajectories were used to determine the point of conversion in the ECal stack. The primary photons were tagged as exiting or converting in the P0DECal. As all photons begin life just inside the detector this tagging yields the fraction of photons converted in the P0DECal. In Fig 6.2 we plot the fraction of photons failing to convert in a P0DECal stack.

The position of conversion was also used to determine the first scintillator layer in which one could see a signal from the resulting shower. This too is an important factor as it determines the amount of information one can obtain from the resulting shower/interaction. It is most beneficial that the majority of photons convert within the first couple of absorber, or indeed scintillator, layers. This would then result in a number of active and absorber layers in which to measure signal to identify and classify the shower. The first layer that can possibly yield a signal is plotted against the initial photon energy in Fig 6.3 for each of the five lead thicknesses. The zero bin in each plot contains those photons that do not leave any signal, i.e. those shown by bin fraction

in Fig 6.3.

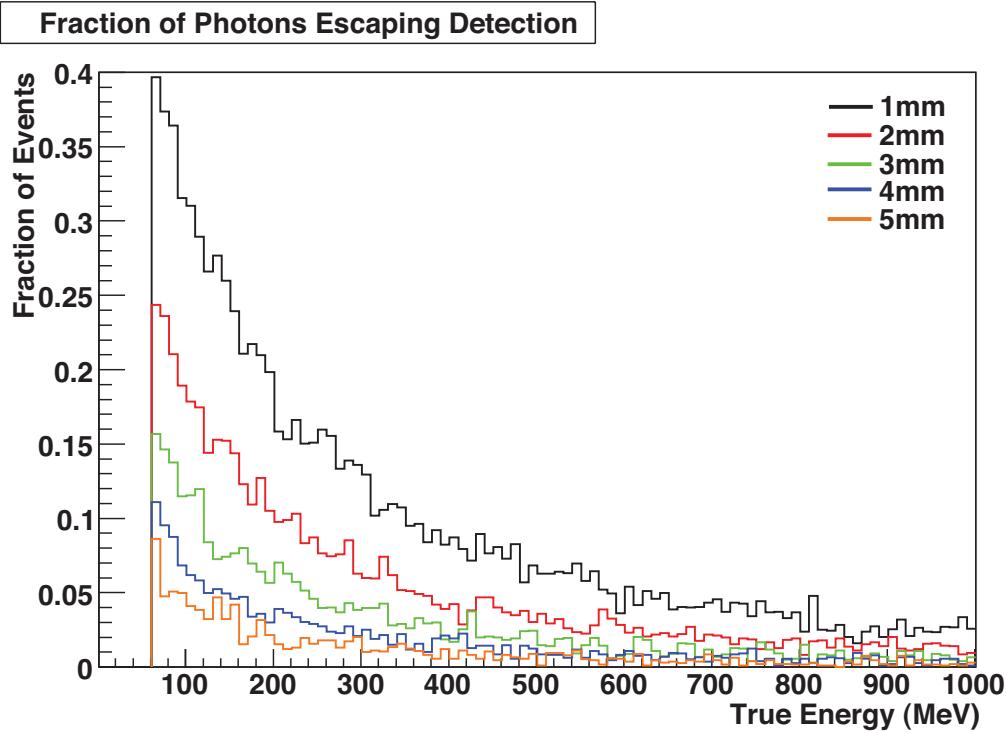


Figure 6.2: These plots show the fraction of events in each energy bin failing to interact before the final scintillator layer; these photons are essentially lost by the detector entirely. One can see an exponential decrease in the fractional loss at low energies as the lead is increased from 1 mm thickness, black, through to 5 mm, orange, as expected from the transmission equation 6.1.

6.2.2 Photon Containment

As well as converting as soon as possible, one also hope to contain a large proportion of the resulting shower's energy. This is important to gain a fair representation of the initiating photon's energy. With the true energy of the incident photons known it was just a case of quantifying the lost energy. GEANT true particle trajectories were again the only data used in this analysis. The energy of each trajectory that was found to be exiting the bounding box of the detector was summed. This true energy loss was then represented as a fraction of the incident photon's true energy. We again plot this value against the true energy of the incident photons in Fig 6.4 to gain an understanding on the shower containment power of the detector.

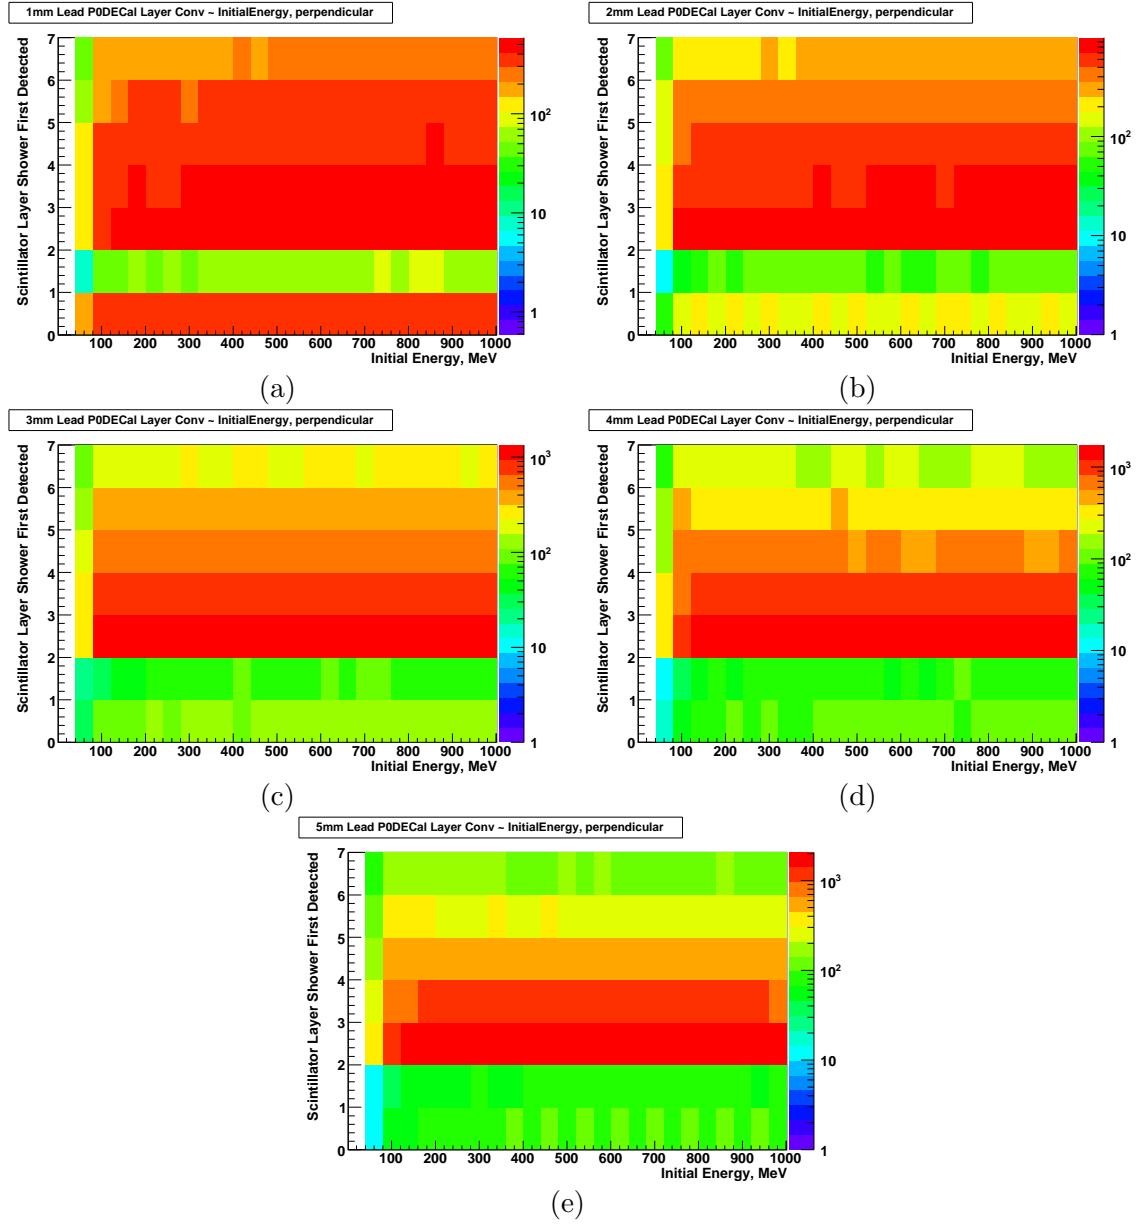


Figure 6.3: The scintillator layer in which photon showers can first be ‘seen’ by the PODECal. The photon conversion point, obtained from the final GEANT trajectory position, is plotted against incident energy. The zero bin contains the photons which escape detection entirely, i.e. those shown by fraction in Fig 6.2.

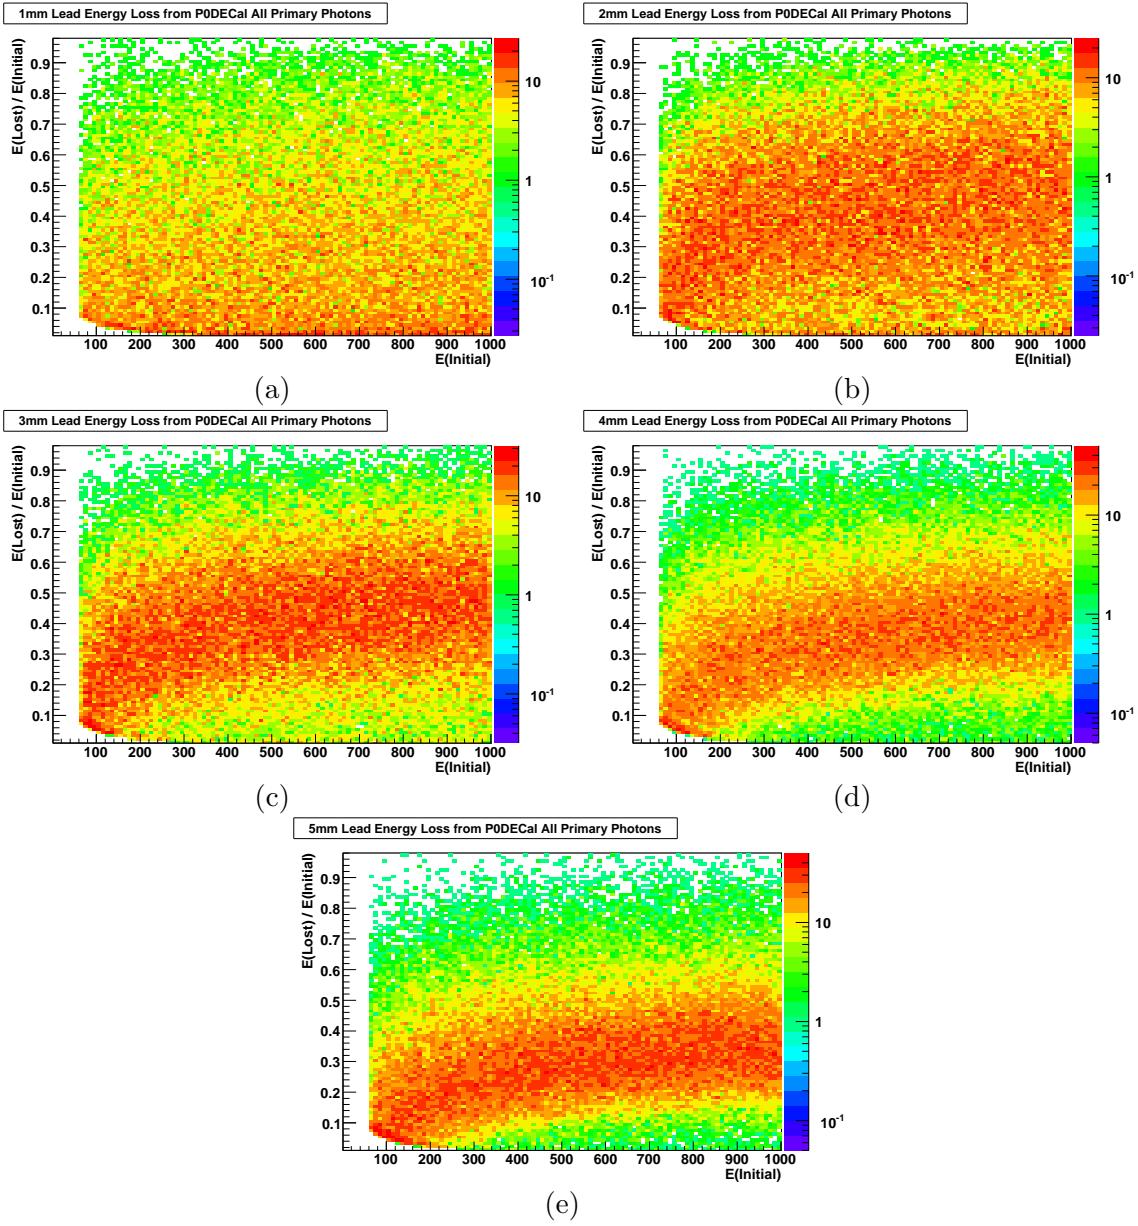


Figure 6.4: Fractional energy loss as a function of incident energy. The energies of all true GEANT trajectories exiting the P0DECal are summed, and this total is then divided by the energy of the initiating particle. The result is then plotted against the initial energy of this particle to investigate shower containment as a function of particle energy. One can see a clear dependence of containment on energy for all but the thinnest lead.

6.2.3 Discussion

Looking at the plots of fraction of photons escaping detection, Fig 6.2, one sees an exponential decrease in each individual plot with an increase in energy. This is expected from the exponential form of the transmission equation above and the increase in photon cross-section.

An increase in radiation lengths sees a suppression of the large fraction of low energy photons escaping detection, although the exponential functional form mentioned earlier does not change. Taking the lowest energy bin one observes an exponential decrease in transmission with radiation length, expected with the exponential form of equation 6.1 and the linear increase in Δx . If one tracks the fraction of escaping photons at an energy of 200 MeV we see that our previous back of the envelope calculation for this fraction was not far off that seen. We conclude from this plot alone that 4 mm+ of lead foil between each active layer in the P0DECal would suffice to stop π^0 photons.

Looking now at the layer of conversion, Fig 6.3, one observes the above decrease in the fraction of photons this time in full statistics in the zero bin of each plot. Each plot has only a small proportion of photons converting in the first active scintillator layer, which is expected because of its low density and cross-section, as shown in Fig 6.1. All plots also show a peak in photon conversion within the first layer which one would expect from the exponential nature of the photon's transmission through the stack. This exponential decrease can then be viewed as one looks at later and later active layers. As the number of radiation lengths increases then so too does the severity of the exponential drop in photons converting late in the stack. The highly spread distribution of the 1–3 mm plots would be unacceptable as a large number of photons would convert within the last two lead absorber layers, leaving only partial information about the resulting shower. The thicker 4–5 mm plots show a much more acceptable distribution with a greatly increased number of photons converting in or just after the first two lead absorber layers yielding a far greater amount of information on the resulting shower.

The scatter in layer of shower initiation can be seen clearly in the plots of fractional energy loss, Fig 6.4. The almost noise-like scatter of the 1 mm and 2 mm plots demonstrates the high spread in the layer of shower initiation. The high fractional energy loss events can be attributed to late shower initiation in the penultimate or last layer with a large amount of the shower passing into the space behind the ECal stacks. The formation of a coherent distribution in the 3–5 mm plots shows again the observed constriction in the layer of shower conversion distributions. The spread in these distributions arises from two phenomena. The first is geometrical, dependent on the layer in

which the shower is initiated, as we have discussed, as well as the shower's angle and proximity to the edge of a stack, and the second is dependent upon the physics model that is chosen by the MC to propagate the shower through the ECal stack medium. From these plots one concludes that any of the 3–5 mm geometries would suffice for containment purposes, with a preference toward the higher thicknesses because of their lower mean and standard deviation of energy loss.

The seemingly high statistics clustering towards the origin is due to the small multiplicity of low energy photon showers with many showers containing just a few trajectories. The small cutoff region towards the origin is a sampling artefact of the Monte Carlo. One cannot carry all truth data ad infinitum through a data process chain and so a photon energy cutoff must be chosen to ignore uninteresting but numerous extremely low energy photons. This artefact is not visible at higher energies, as such photons would only account for an extremely small fraction of the overall energy, but becomes visible at low energies where this assumption no longer holds.

In conclusion, the plots investigating the point of shower initiation lend support to a lead thickness of at least 4 mm.

6.3 Energy Resolution

Reconstruction of the invariant mass of a π^0 from a vertex, requires an understanding of both photons resultant from the π^0 decay.

To investigate the possible energy resolution of the P0DECal I used the charge information of all hits generated by the electronics Monte Carlo simulation elecSim 5.3.2. There were no final reconstruction algorithms applied, just a truth-based pseudo reconstruction. Hits were clustered using truth information of trajectory IDs that contributed to each hit. Trajectories were followed from the primary photon through the shower and their IDs saved and then compared to the contributors to each hit. This method is the same as that employed in the P0DECal truth clustering algorithm (section 7.4). The sum of the energy deposited by the hits in the detector was then calculated as a fraction of the truth energy of the initiating photon. This fraction was then binned as a function of true photon energy, Fig 6.5. This fraction is essentially arbitrary as an energy reconstruction has yet to be conducted on the charge measured by the electronics, but this is not important as it is only the absolute scale and deviation that are significant. The standard deviation and mean fraction were then derived for each energy bin from a Gaussian fit to the fraction distribution; an example slice with fitted Gaussian is shown in Fig 6.6(f). The resolution was then estimated as the standard

deviation divided by the mean of this fraction for each binned energy. In Fig 6.6 we plot this resolution as a function of true photon energy for each geometry.

Note that this represents the best possible resolution that could be achieved: detector effects etc. would in practice degrade this significantly.

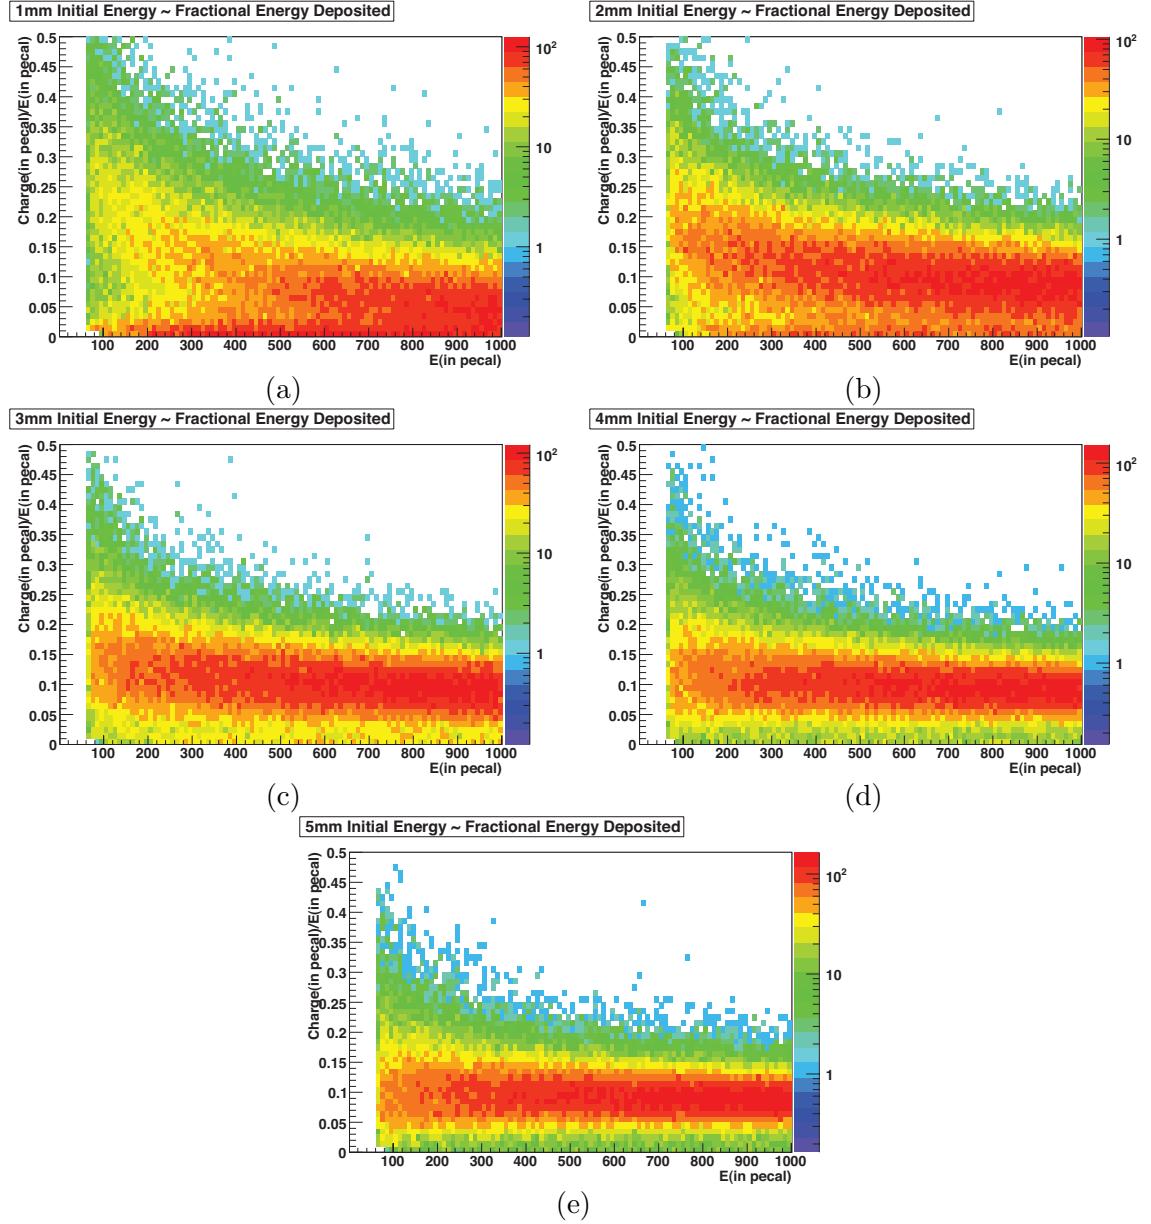


Figure 6.5: Relative energy deposition as a function of initial incident photon energy. The energies deposited in GEANT active detector hits are summed before being divided by the original energy of the initiating particle.

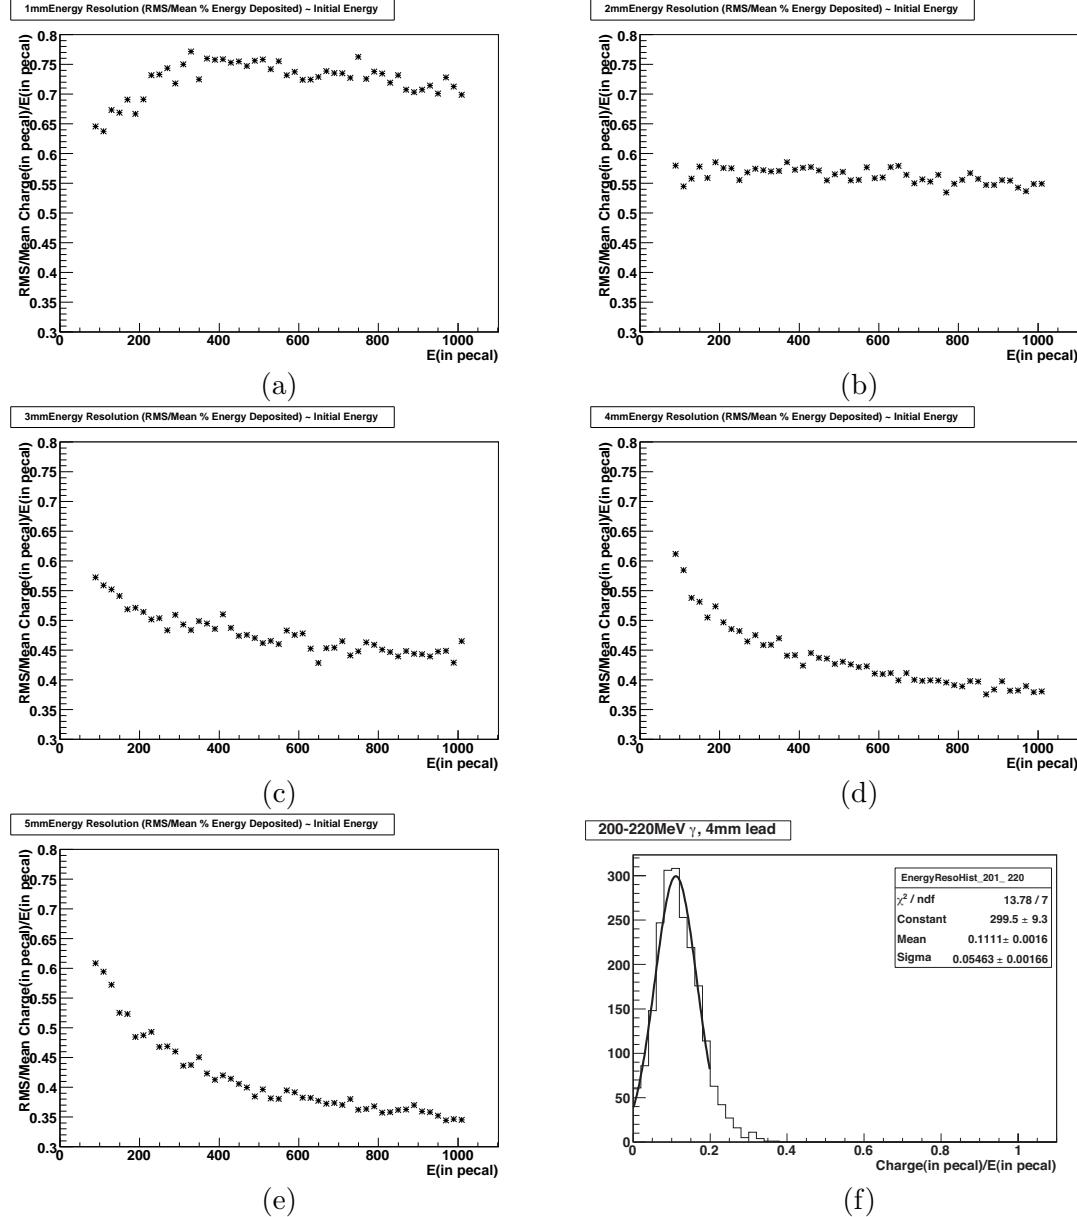


Figure 6.6: Idealised energy resolution as a function of initial photon energy. The mean, μ , divided by the standard deviation, σ of the Gaussian distributed fractional energy sampled by the P0DECal is plotted as a function of initial energy of incident photons. An example projection of the 2D histograms, in Fig 6.5, and following Gaussian fit is shown in panel (f) for the 200–220 MeV energy range with a 4 mm lead thickness. One can see from the Gaussian fit result that resolution at this energy and lead thickness is $\sigma/\mu \sim 50\%$.

6.3.1 Discussion

In the 1 mm and 2 mm lead plots one can see two distinct distributions. The upper distribution traces the distribution seen in later 4–5 mm plots for highly contained showers. The spread at each energy is down to shower containment which, as mentioned, relies on the layer in which the shower is initiated as well as the amount of material with which to absorb the shower. The small increase in fraction, and resulting decrease in resolution, at lower energies in all plots is due to the fact that some low energy photons deposit charge in just a single layer. If this is the case then the charge deposited is heavily dependent on how far through the Pb absorber the photon began to shower, i.e. the mean free path of the shower becomes comparable to the thickness of the absorber. The spread of this distribution reduces as you increase the lead thickness. This is due to a constriction in energy loss from shower initiation earlier in the ECal stacks.

The lower distribution seen in the 1–2 mm plots, with remnants still visible in the 3 mm plot, is a result of late-showering photons where only a small fraction of the shower is seen. This splitting is more obvious at low energies, as it is here that one has the largest spread in the layer of shower initiation.

The 1 and 2 mm plots do not show the same functional form as the later plots, which exhibit a decrease in resolution: they instead show a much flatter function. This functional form and larger resolution values can be attributed to the scatter in the layer of shower initiation. The higher values of resolution towards lower energies are explained, as above, by the fact that the interaction length of a photon becomes comparable to the thickness of absorber.

6.4 Discriminating Variables

With no reconstruction or particle ID algorithms in place for the P0DECal we again used truth clustered hits to investigate the third and final P0DECal physics goal of muon-photon discrimination. I selected a few variables, derived from charge and geometric distributions of detector hits, which were identified as having photon-muon discriminating power. As a first iteration we looked at the overlap in the photon and muon distribution of each variable to assess their discriminating power. The key thing that we were identifying in this study was the change in discrimination with radiation lengths in the detector.

Muons are minimally ionising particles, MIPs, which lose very little energy when traversing matter. From event displays one can clearly see this behaviour as the muon

passes straight through the P0DECal stack depositing just a small amount of energy in each detector layer. Conversely photons lose energy primarily through initiation of an electromagnetic particle shower from $e^+ e^-$ pair production. This is an exponential energy loss which results in a large deposit of charge in most layers and a shorter distance travelled through the detector.

These characteristics bring to the fore a number of possible discriminating variables, from which I chose the two variables track length (the number of layers travelled through the stack), and mean charge deposited per layer. Here I show just 1 mm and 4 mm extremes of the variables to illustrate the change in these variables with increasing lead thickness.

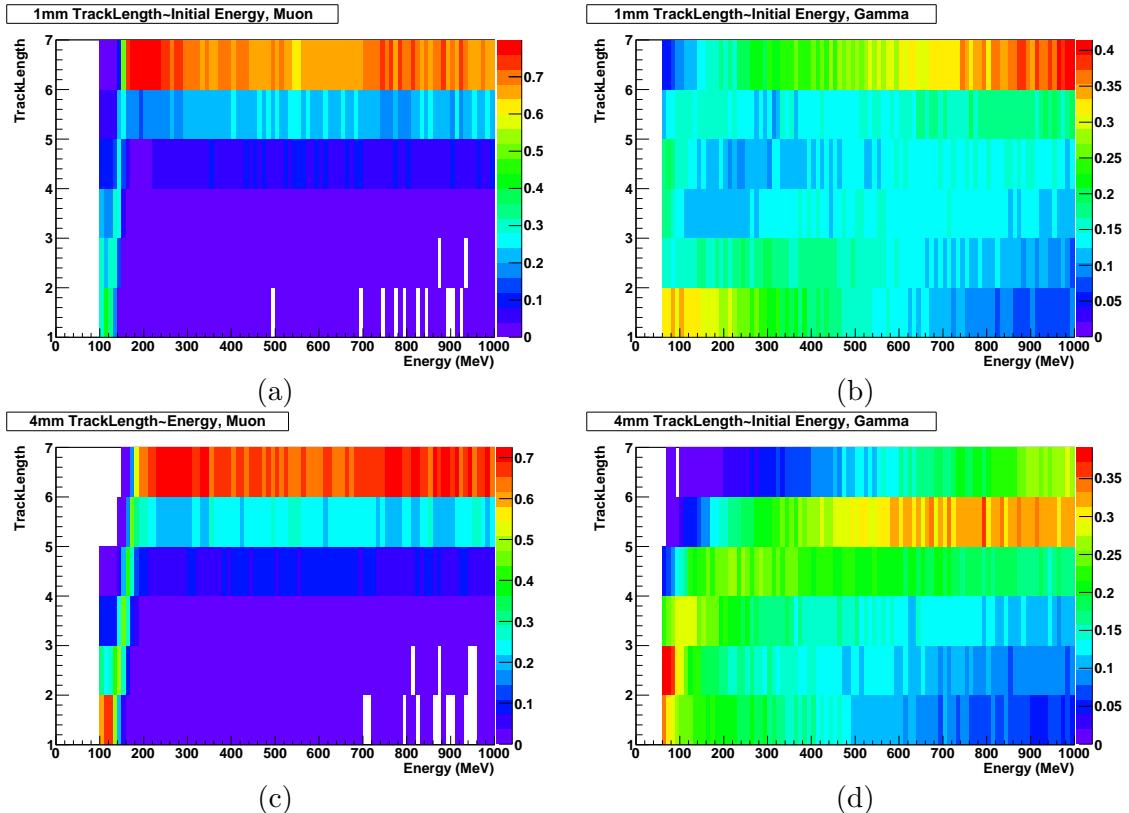


Figure 6.7: Plots of track length, expressed as the number of layers traversed between first and last layer with a signal, as a function of energy, for 1 mm (top) and 4 mm (bottom) lead thicknesses. The left-hand panels show muons and the right-hand panels photons. The z, colour, scale shows the fraction of events in each track length bin, normalised to a total of 1 for each energy bin.

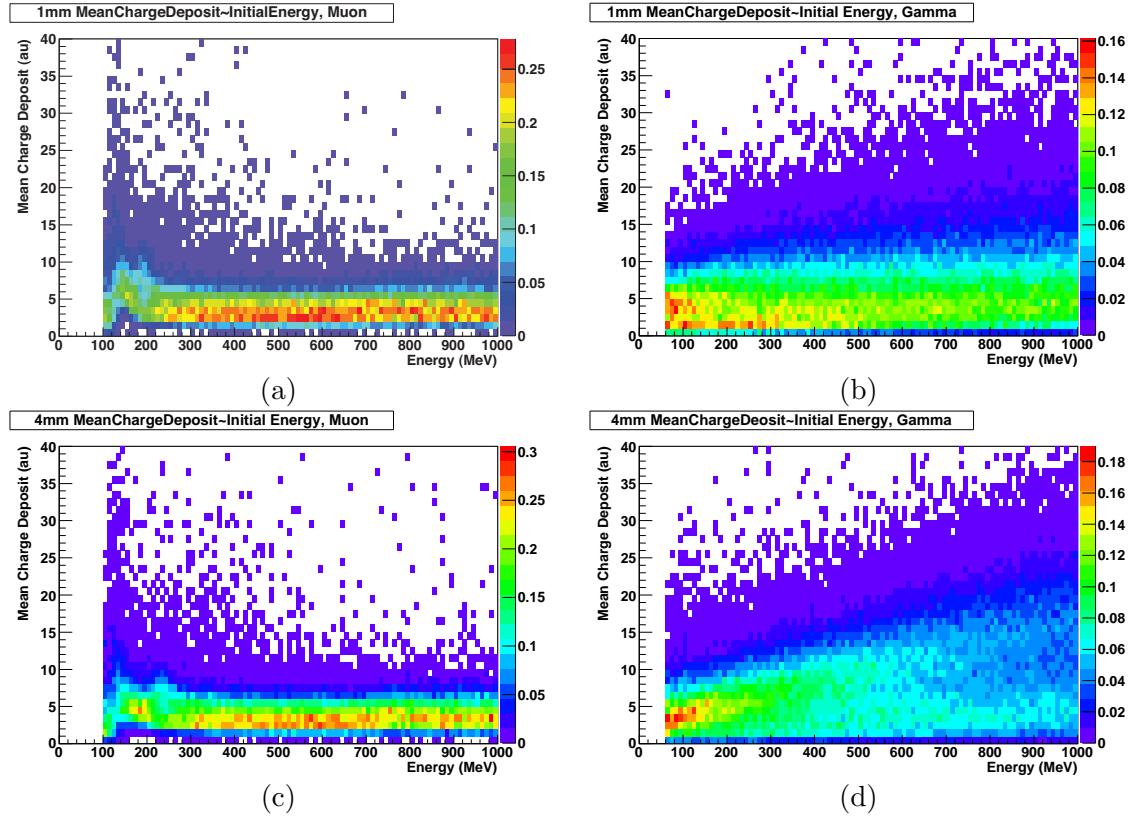


Figure 6.8: Plots of mean charge deposited per layer. Shown here 1 mm (top) and 4 mm (bottom) lead thicknesses. The left-hand panels show muons and the right-hand panels photons. The z, colour, scale shows the fraction of events in each track length bin, normalised to a total of 1 for each energy bin.

6.4.1 Discussion

Looking at the 1 mm plots of track length variable, Fig 6.7(b), we see that many of the photon showers are traversing the maximum 6 detector layers. This peaked distribution is smeared because of the spread in the location of conversion. This trend only alters when the energy of photons is such that transmission becomes the most probable outcome for the photon (section 6.2). The muons perform as expected with a peak at six layers for the through-going muons. The small but not insignificant distributions at four and five layers hit are a consequence of the low ionisation of the muon and its extremely localised nature. Here the muon in question may have only glanced one or two bars and not passed through enough scintillator to leave a significant ionisation signal.

Looking now at the 4 mm plot of track length, Fig 6.7(d), one sees a dramatic change

in the photon distributions but little change between the muon distributions, Fig 6.7(a) & (c). The peak track length for photon showers has reduced across almost the entire energy range. Photons below 400 MeV predominantly leave signal in just four layers or fewer. One can also see the reduction in the spread of this variable from the constriction of the layer in which the photons converted. The flat nature of the plot is a consequence of only partial sampling of photon showers, the latter portion of the shower. The muon distribution has changed only mildly. The steep linear increase to the maximal 6 layers traversed, located at very low energies, is a consequence of stopping muons that decay before passing through all 6 layers. This has increased slightly due to the higher number of radiation lengths being able to stop higher energy muons.

From these plots it is clear that the increase in radiation lengths of absorber leads to a separation of these two distributions.

The second set of plots, Fig 6.8, shows mean charge deposited per layer and demonstrates a similar trend as track length. The photon distributions, (b) & (d), are again highly spread. This time not only due to the layer of conversion, but primarily because of the multitude of possible ways in which a shower may progress through matter. The 1 mm plots, (a) & (b), show a large overlap at all energies, with the peaks of the 1D projections just a small number of charge units apart. The distribution of mean charge deposited by photons is a very broad spectrum as expected from a highly ionising EM shower. The muons, (a) & (c), display an expected uniform distribution, with a tight Gaussian spread, at all energies from their minimally ionising nature. Comparing back to the track length plots one can see that the small increase at low energies is a consequence of stopping muons depositing a larger amount of energy in their stopped layer.

The 4 mm plots, Fig 6.8(c) & (d), show the photon distribution gradient increase and rise from out of the shadow of the muon distributions, this is clearest at higher energies. This is coupled with an increase in the peak of the mean energy distribution. This shift is a consequence of a greater sampling of the photon shower, i.e. less shower energy escaping the P0DECal. The muon plot changes little, as one simply samples a MIP track at different intervals, which should not change the measurements one receives, with an exception for those muons which can now be stopped by the increased radiation lengths.

From these plots one can see that the particle identification improves primarily because one gains more information about the photon showers with increasing radiation lengths. This arises from an earlier conversion and a greater containment. This benefit will however cease when reducing the sampling again with a continuing increase

in containment. If one continues to increase the radiation lengths the sampling begins once more to decrease as photon showers are contained over just a few bars.

6.5 Conclusions

I have studied the key physics goals of the P0DECal detector which have been investigated using a mixture of true trajectory and partially reconstructed hit data. Calculations based on the known radiation length of lead suggest that at least 20 mm total is required for good photon containment, and the Monte Carlo bears this out. Higher lead thicknesses also improve the separation between photon and muon distributions in possible particle-identification variables such as track length and mean deposited energy. On the other hand it was felt unwise to consider thicker lead layers as this would increase the number of photons which deposit energy in just one layer leaving a single unreconstructable hit.

As a direct result of these studies, a lead thickness of 4 mm was selected for the P0DECal design.

Chapter 7

Monte Carlo Analysis of P0DECal Reconstruction

From the schedule of ND280 construction and commissioning (section 3.6) it can be seen that construction of the first P0DECal module will begin in late April 2009 at earliest. This has dictated that the efficiency studies of all P0DECal reconstruction methods presented in this chapter are all Monte Carlo (MC) based. After the design study, chapter 6, the geometry of the P0DECal was finalised and was used in all of the following MC analyses.

Fig 7.1 shows a flow diagram of the P0DECal reconstruction chain; on the left we see the flow of event data through the various reconstruction algorithms and on the right the `TAlgorithmResult` objects which contain the final result of each algorithm and are written back to the event. The `TECalGetHits` and `TECalPrepareHits` algorithms are common between tracker and P0DECal. The `TP0DECalBasicClustering` algorithm performs the collation of detector hits into clusters. The `TP0DECalCreatePid` algorithm employs two stages: `TP0DECalPidVarCalc` calculates discriminating variables and track-shower discrimination values, and `TP0DECalCreateOutputObjects` creates final `TRecon Tracks` and `Showers` which are then saved into the `P0DECalPID` result. The particle identification is discussed in the next chapter.

7.1 Hit Preparation: `TECALPrepareHits`

Hits are first extracted from the `TND280Event` by subdetector, in the `TECalGetHits` algorithm. The result is a single collection of `THits` for all P0DECal modules covering

P0DECal Reconstruction

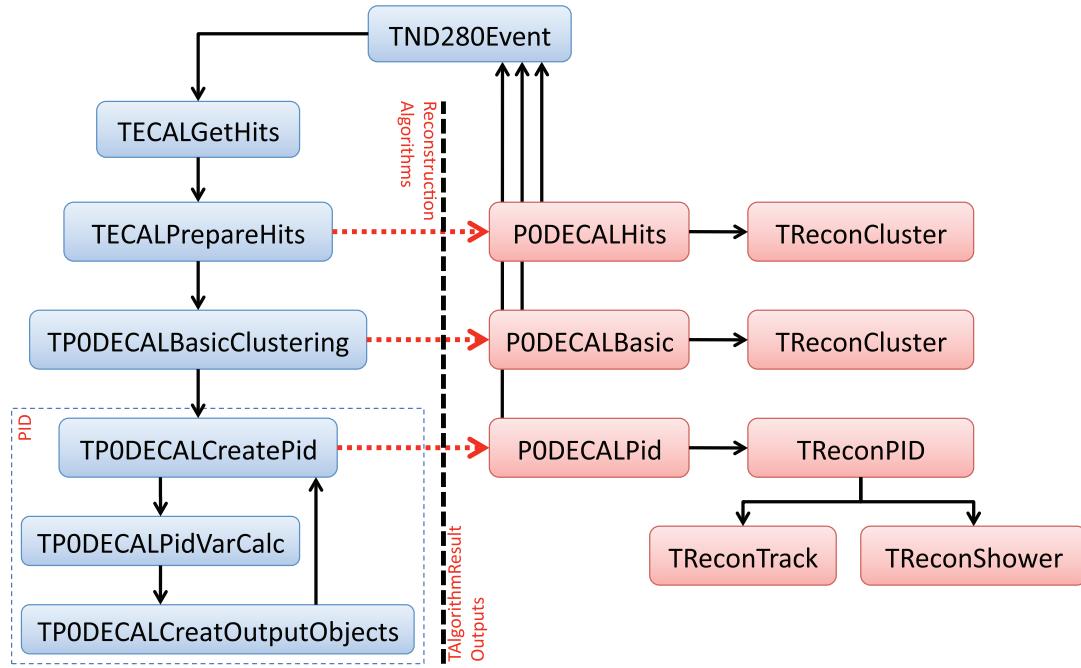


Figure 7.1: Flow of event data in P0DECal Reconstruction.

the time of an entire beam spill. Before reconstruction of events can begin in earnest, one must prepare these hits for the downstream reconstruction algorithms.

The first stage of this preparation involves the separation of hits into separate P0DECal stacks; top, bottom, side on both the left and right hand side of the detector. The hits are also segmented into beam bunches, as the fast response of the scintillator detector electronics allows data from each bunch to be stored before resetting. Information on spill and bunch structuring is taken from the offline database and as such is dynamic.

After separation comes a first pass of signal calibration. The elecSim package (section 5.3.2) takes the energy deposited in a scintillator bar and simulates each stage until the photoelectron signal saved by the DAQ, as will be seen in real data. Each stage of signal calibration, starting with the first pass described here, works to reverse this and give an energy deposition measurement which is to be given in MIP units. A MIP unit is defined as the mean number of scintillation photons created by a minimum ionising particle, such as a muon, travelling a distance of 1 cm through ECal scintillator.

The initial stage of this calibration procedure aims to convert the p.e. signal to the number of scintillation photons that reach the face of the photosensor. This is calculated using the photo detection efficiency (PDE) of the MPPC and the poissonian probability

of two photons triggering just a single pixel. This number of photons is then converted from NPhotons to MIPs, at this stage assuming that the particle passed through the scintillator at the centre of the bar¹.

Once the ECal is installed and commissioned, calibration data from cosmic muons in inter-spill periods will be plentiful. These data will be written to the offline calibration database on a regular basis and the attenuation parameters and photosensor response for each bar will be used by the reconstruction to calibrate hits. Calibration such as this is repeated wherever information on reconstructed objects is updated.

The results of this calibration and all other hit information are then saved into TReconHit objects before being placed into time bunched TReconClusters for reconstruction to begin in earnest. For each event the TAlgorithmResult titled *P0DECalHits* is written to the TND280Event object and passed to the next stage of reconstruction, TP0DECalBasicClustering.

7.2 TP0DECALBasicClustering

The entire clustering of P0D ECal hits is done in just a single stage using spatial nearest neighbour aggregation, employed in the TP0DECALBasicClustering algorithm.

The algorithm is applied to each spill collection of hits and aims to collate hits originating from a common incident particle. The bar with the largest charge deposition, in each P0DECal stack and bunch collection of hits, is taken as a seed from which a cluster can grow. From the centre of this seed bar, each other bar receiving signal within a fixed horizontal ‘width’ and longitudinal ‘depth’ distance are included as a member of the cluster. This process is then repeated, centring upon each newly collated hit, until no new hits are found within the constraints. If there are more than a defined minimum number of hits collated then the hits are saved as a new TReconCluster.

The highest signal hit in any remaining uncollated hits is then selected as a new seed and the process repeated. This continues until no further clusters pass the minimum hit constraint. Any hits remaining unclustered are saved as a separate collection and passed through the reconstruction chain for possible inclusion in later reconstruction methods.

In the tracker ECal, it is found that the analogous basic clustering algorithm often splits hits from single particles into several clusters, and there is therefore a “combine” method designed to reunite such fragments. This is found to have little effect on

¹In the tracker ECal dual readout bars a position of the event along the bar is also estimated from the relative signals and attenuation correction applied.

P0DECal reconstruction: with only six layers and an increased radiation length of absorber between there is little scope for tracks or showers to fragment. The small percentage of events that do fragment tend to be at oblique angles, where a horizontal neighbour is missing.

The uncollated hits and all TReconClusters are written to the *P0DECalBasic* TAlgorithmResult each event and passed to the next level of reconstruction, TP0DECalCreatePid.

7.3 TP0DECALCreatePid

The TP0DECALCreatePid algorithm has the following principal tasks:

- Identify clusters as track-like or shower-like.
- Determine cluster position.
- Determine cluster direction.

7.3.1 Identify Clusters as Track-like or Shower-like

Initial trials of linear discrimination techniques proved unsatisfactory, so particle identification in the P0DECal is achieved using non-linear multivariate analysis (MVA) techniques. The basic clusters are passed to TP0DECalPidVarCalc which first calculates a number of discriminating variables. These are then passed to a Multilayer Perceptron (MLP), a type of Artificial Neural Network (ANN), which determines the track-like or shower-like nature of the cluster. This is the topic of the next chapter and will be discussed further there.

This stage of the algorithm then prompts TP0DECalCreateOutput to create TReconTrack or TReconShower objects, and relevant data, according to the results of the MLP.

7.3.2 Determine Cluster Position

The front positions of both tracks and showers are calculated from a charge weighting of the hits in the innermost two layers with signal. For each layer the position and charge deposited in each hit bar are multiplied and summed before being divided by the total charge deposit in the layer. This is repeated for the next innermost layer, if available, and the average of these charge weighted positions taken as the front position of the track or shower.

The performance of this algorithm is discussed in section 7.7.

7.3.3 Determine Direction

The directions of tracks and showers are reconstructed using principal component analysis. The major axis of the clustered hits is found and its direction taken as that of the first principal component. If this fails then the direction is set perpendicular to the face of the stack. The final direction performance of tracks and showers is also discussed in section 7.7.

7.3.4 Output Results

Once the TReconTrack or TReconShower is returned by TP0DECalCreateOutputObjects the final TReconPID object is created; this not only contains the solution selected by the MVA output, track or shower, but also the alternative, shower or track. These higher level reconstruction objects also carry RECPACK node information which could in principle be used for inter-detector fitting but which are crudely generated due to the limited information from the P0DECal. It is likely that a simple energy flow algorithm may turn out to be more appropriate for matching P0D and P0DECal objects. In the NC1 π^0 analysis presented in chapter 9 we are interested in particles which remain unnoticed in the P0D, leaving matching a moot point. These tools are still in development by the global reconstruction group.

7.4 True Clustering

The reconstructed clusters will be compared to clusters of hits collated using true trajectory information. Points along each true trajectory are checked to test if it at any time enters the P0DECal. The first trajectory found entering, or primary trajectory starting therein, is taken as a seed for clustering. All subsequent trajectories that are deemed to be a product of the seed, or a previously collated trajectory, are attached to this seed's collection. True trajectories contributing to each hit, produced by elecSim and passed to the reconstruction, are compared to trajectory collections. If a hit is not already accounted for then it is associated with that particular trajectory group. These grouped hits result in ‘idealised’ clusters containing only hits associated with particles entering the P0DECal, or created from interaction vertices therein.

7.5 Monte Carlo Test Samples

Before presenting studies of the performance of P0DECALRecon it is essential to describe the method(s) by which events were inserted into the Monte Carlo simulation of the detector and electronics. The Monte Carlo simulation chain of the ND280 off-line software, described in section 5.3.1, is capable of generating and simulating a wide range of event types within the entire ND280; hall, magnet and subdetectors.

The closest approach to real data is provided by “full bunch” events in which there are multiple neutrino interaction vertices within the ND280 each beam bunch. These were considered too complicated to use in training and optimisation of the P0DECal reconstruction. Most of what follows is therefore based on particle gun samples which consist of just one particle emanating from a single vertex with each event containing just one bunch.

7.5.1 Particle Gun Monte Carlo

Single particle events are inserted into the detector Monte Carlo using the particle gun method. Various particle types were fired from a starting point at the face of the left side P0D ECal over a range of energies and incident angles.

Two classes of particle were investigated, minimally ionising particles (MIPs) and electromagnetically (EM) showering particles. Muons are selected to represent the MIP class of particles as they best exhibit MIP properties and they will also be the dominant MIP particle incident upon the P0D ECal.

Photons were selected to represent EM showering particles in the following study. Photons and electrons (positrons) both lose energy through photon and electron-positron pair-production cascades, called EM showers. The only discernible difference between these two particles is the point at which this cascading of EM particles begins. For the charged electron it is almost an immediate effect as its electric field interacts with that of ambient electrons. Photons, however, can travel some distance before pair producing, generally when in the presence of a nuclear field. Low energy electrons in this method can therefore lose a large proportion of their energy in dead material in front of the first layer of scintillator, dramatically affecting energy reconstruction of the EM shower. Photons however are seen to initiate EM showering in the first couple of lead absorber layers (see section 6.2) leaving the detector privy to the early stage of the EM shower.

Neutrino interactions may occur at any point within the ND280 detector, and each

particle type can exit neutrino interaction vertices at a wide range of energies and angles. This leaves an enormous region of phase space to investigate for each particle type.

When investigating optimisation of an algorithm one would hope to use the particle type, angle and energy spectrum expected in the detector. These variables can be extracted from a large sample of neutrino interaction Monte Carlo data. The energy spectrum of each particle type is histogrammed and used as an input for particle gun energies. This allows a user defined position and direction of single ‘real’ energy spectrum particles, enabling easy production of large scale Monte Carlo statistics from which to optimise reconstruction algorithms.

For the purpose of efficiency studies, however, one would prefer a flat spectrum of energy; for the ND280 this would cover the low kinematic energy region 0-1 GeV where the energies of neutrinos and their products peak and detectors are predicted to find reconstruction difficulties. This allows an unbiased measure of efficiency to be constructed.

Because of the isotropic nature of events in the plane perpendicular to the beam axis a flat distribution of particle incidence angle, θ , on the front face of the module was selected for both optimisation and efficiency studies. Throughout this chapter θ is defined as the angle from the minor perpendicular axis of the ECal stack and ϕ the rotation about this axis, Fig 7.2. A single octant of $\theta = 0 - 90^\circ$, $\phi = 0 - 90^\circ$ was chosen; the other three octants of this upper hemisphere will be symmetries of this.

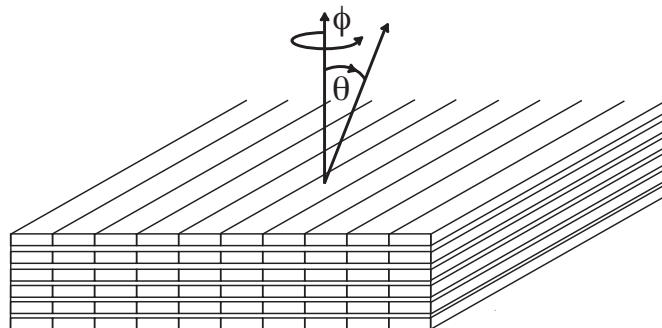


Figure 7.2: Angular convention in the following tests: θ is the angle from the perpendicular through the stack and ϕ the rotation about this axis.

7.5.2 Background

Testing the purity of clustering is a more difficult affair as one requires a great deal of knowledge of all forms of background: electronics noise or additional particles.

Full simulations of the photosensor and electronics readout are used in generating all Monte Carlo samples. As well as determining the response of the detector readout, from energy deposition to DAQ readout, elecSim also generates photosensor and electronics noise. This simulation will be tuned as calibration data becomes available during cosmic testing, particle test beam running and in-situ commissioning. As such the reconstruction described here may need re-optimisation and perhaps the inclusion of a pre-clustering noise filter algorithm.

The timing cuts set by the beam's bunch structure give a powerful separation technique to exclude cosmic muon events. There will still be a possibility, however, of contamination from hits associated with particles from other interaction vertices within the detector. It has been estimated that at full intensity the full mass of the ND280 detector will receive a mean of eight interactions per bunch [3]. If particles pass through a P0DECal stack at close quarters then there may be some cross contamination or possibility of defining just one cluster for the two particles.

Background hits must not be collated as part of a cluster as they will affect all higher reconstruction processes, e.g. particle identification. It is therefore important that the purity of clustering is also investigated. The purity of reconstructed clusters is obtained from comparison of true clustered hits, using particle trajectories to determine the contributors to each detector hit.

Although eight interactions per spill seems large, the probability of such contamination is expected to be low. Still, to investigate the effect the backgrounds may have on the purity of clustering, we produce Monte Carlo with a test particle fired into a P0DECal stack, using the procedure described above, and full bunch background added as follows. Interaction vertices are generated using the NUANCE neutrino generator with the T2K neutrino beam flux as input. The vertices are then randomly distributed throughout the detector weighted by density of material.

7.6 Clustering Efficiency and Purity

7.6.1 Cluster Finding

This basic clustering algorithm has three free parameters: horizontal distance cut, transverse/depth distance cut and the minimum number of hits that constitute a cluster. The 40 mm width of the ECal bar was selected after considering the Molière radius of electromagnetic showers in plastic scintillator, ~ 4 cm, so that there will be hits within multiple adjacent bars within any ECal detector. For this reason the width cut was set as the bar width. A grid search was then performed using the other two free parameters to maximise efficiency and purity of clustering. For this it was deemed necessary to use energy spectra akin to that expected in running, so as to gain a realistic performance efficiency, as well as full spill/electronics background, to test the purity of such clustering.

The purity of clustering was affected little by an increase in transverse distance cut, but a marked increase in efficiency was seen. The efficiency improvement arose from the collection of numerous low energy photons which can travel some distance through the stack before depositing energy. The most dramatic effect on the efficiency of clustering arose from a change in the minimum number of hits required to constitute a cluster. The expected energy spectra of both muons and photons of interest for event reconstruction purposes peak at low energies. Such particles tend to leave very few hits in the P0DECal, which makes the efficiency loss caused by minimum hit number more marked. The result of this grid search was a transverse cut of 50 mm and a minimum hit limit of just 2 hits. Once calibration data of real world electronics noise are collected and analysed this optimisation may have to be reconsidered.

The Venn diagram in Fig 7.3 shows the region of detector hits corresponding to truth clusters, TC, reconstructed clusters, RC, and correctly reconstructed true hits, TR. The efficiency and purity are the ratio of this overlapping TR region with TC and RC respectively. This can be applied to the charge accumulated or to the count of hits in each cluster. The hits within the TR region are selected by comparing the unique detector volume ID of individual hits within the detector.

The performance of the chosen clustering was then investigated using a flat energy spectrum of particles isotropically fired in one octant. This gave an unbiased evaluation of the detector's ability to cluster at a range of angles and energies. A total of 400,000 events were generated over the energy range 0–1000 MeV and angular range $\theta, \phi = 0 - 75^\circ$. A cut was applied to select only events which have a single true cluster of

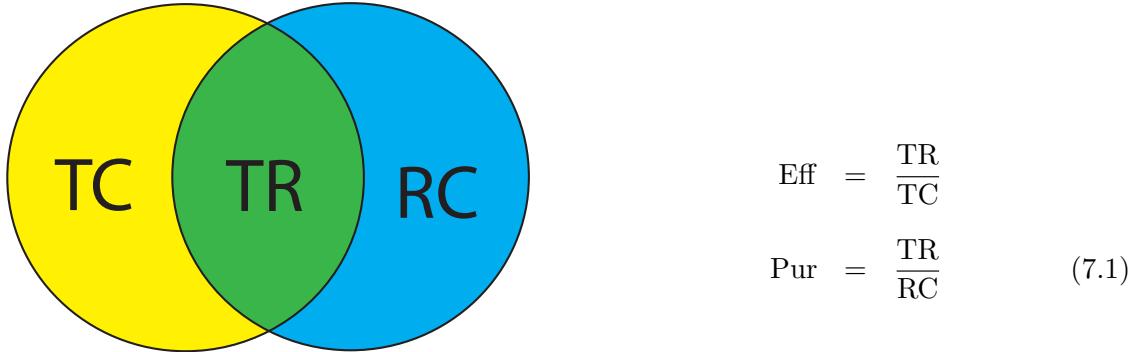


Figure 7.3: Venn diagram of true clustered hits, TC, reconstructed clusters, RC, and reconstructed true hits, TR.

hits as desired². These events were then used to quantify the efficiency of clustering by investigating the number of reconstructed clusters and their hit completeness.

Both the kinetic energy and angle of incidence have an important effect on cluster finding for both classes of particle; the effect of each is shown on the top and bottom of figures 7.4 and 7.5 for photons and muons respectively. Although θ and ϕ both vary we display data for θ only; ϕ is unreconstructable owing to the geometry of the detector. The histograms' x axes are split into bins of 50 MeV in energy and 15° for θ . Each bin is normalised to 100%. The blue hatched area denotes those truth clusters with just a single hit, therefore where reconstruction would always be unsuccessful. The red hatched histogram shows the percentage of events that reconstruct a single cluster, the optimal for single particle events. The green hatched histogram shows the percentage of events that reconstruct two clusters. Practically invisible, at less than the 1% level, is a final distribution which contains all events with greater than two clusters reconstructed. The white area between the blue and red hatched histograms is where the reconstruction algorithm can be said to have failed, as there were enough hits in a truth cluster but they were not successfully reconstructed.

Looking at the reconstruction of photons as a function of energy (Top Fig 7.4) one first notices there is a clear downturn in reconstruction efficiency at low energies. Below 50 MeV around 34% deposit just a single hit in the detector and could therefore not be reconstructed, although these hits may have an important part to play in more developed analyses. Of the remaining just $\sim 31\%$ of these are successfully reconstructed. This is attributed to the sparse distribution of hits from low energy photons in the

²Some single particle events contained 2 true clusters from back-scattered or externally produced particles re-entering an ECal stack

shower, separating hits beyond the limits of the reconstruction algorithm. The percentage rises dramatically in the next, 50–100 MeV, energy bin where over 79% of the eligible 88% of events are reconstructed.

Beyond 150 MeV there is a clear band of less than 5% of events that are not eligible for reconstruction; this consists of late converting photons where just a single hit is registered in a later layer. Of the remaining events $\geq 90\%$ have at least one cluster, rising sharply and plateauing around 400 MeV with a cluster finding success of 96%.

The angular dependence of photon cluster finding is much less marked. The band of un-reconstructable events, late converting events and failed sparse clusters remains. The small increase in percentage of reconstructable clusters arises from the increase in radiation lengths traversed by the photons; those photons with perpendicular incidence to the ECal stack (low θ) traverse less material and have a higher probability of converting late and leaving just a single hit. The lowest reconstruction efficiency is also seen at low angles of incidence for photons, i.e. directly through the stack, as this maximises the possibility of disjointed side photon initiated showers which may not be clustered because of the tight nearest neighbour only cut in the horizontal direction. This improves from $\sim 90\%$ at $0\text{--}15^\circ$ to greater than 96% in the most oblique $60\text{--}75^\circ$ bin.

One can see that fragmentation of clusters within the P0DECal is not an important issue; fragmentation occurs on the few percent level with a small increase in energy as showers grow bigger in size.

Muons of energy 0–50 MeV lose their energy after traversing just a small amount of material, such as the dead space in front of an ECal stack. Once stopped the muon decays $\sim 100\%$ of the time via $\mu^- \rightarrow \nu_\mu \bar{\nu}_e e^-$; events in this bin essentially represent the ability for the P0DECal to reconstruct the resultant Michel electrons incident upon the stack. This leaves the reconstruction in this bin very much akin to that for low energy photons, to which it is comparable with just 38% of events reconstructed out of a possible 76%. The percentage of reconstructable muons rises sharply to almost 98% of the possible 97% of reconstructable events in the energy range 50–100 MeV as the majority of these muons, even though they may stop in the ECal stack, leave a reconstructable track signal over two or more layers. The small failure in this bin is attributed again to Michel electron clustering rather than a genuine muon track. This is logical if one then looks at all higher energies where the reconstructable percentage and those events reconstructed are indistinguishable from 100%; 97% of events are reconstructable in the 150–200 MeV rising to 100% to 2 decimal places after that, and the percentage of these reconstructed agrees with 100% to 4 decimal places for all bins

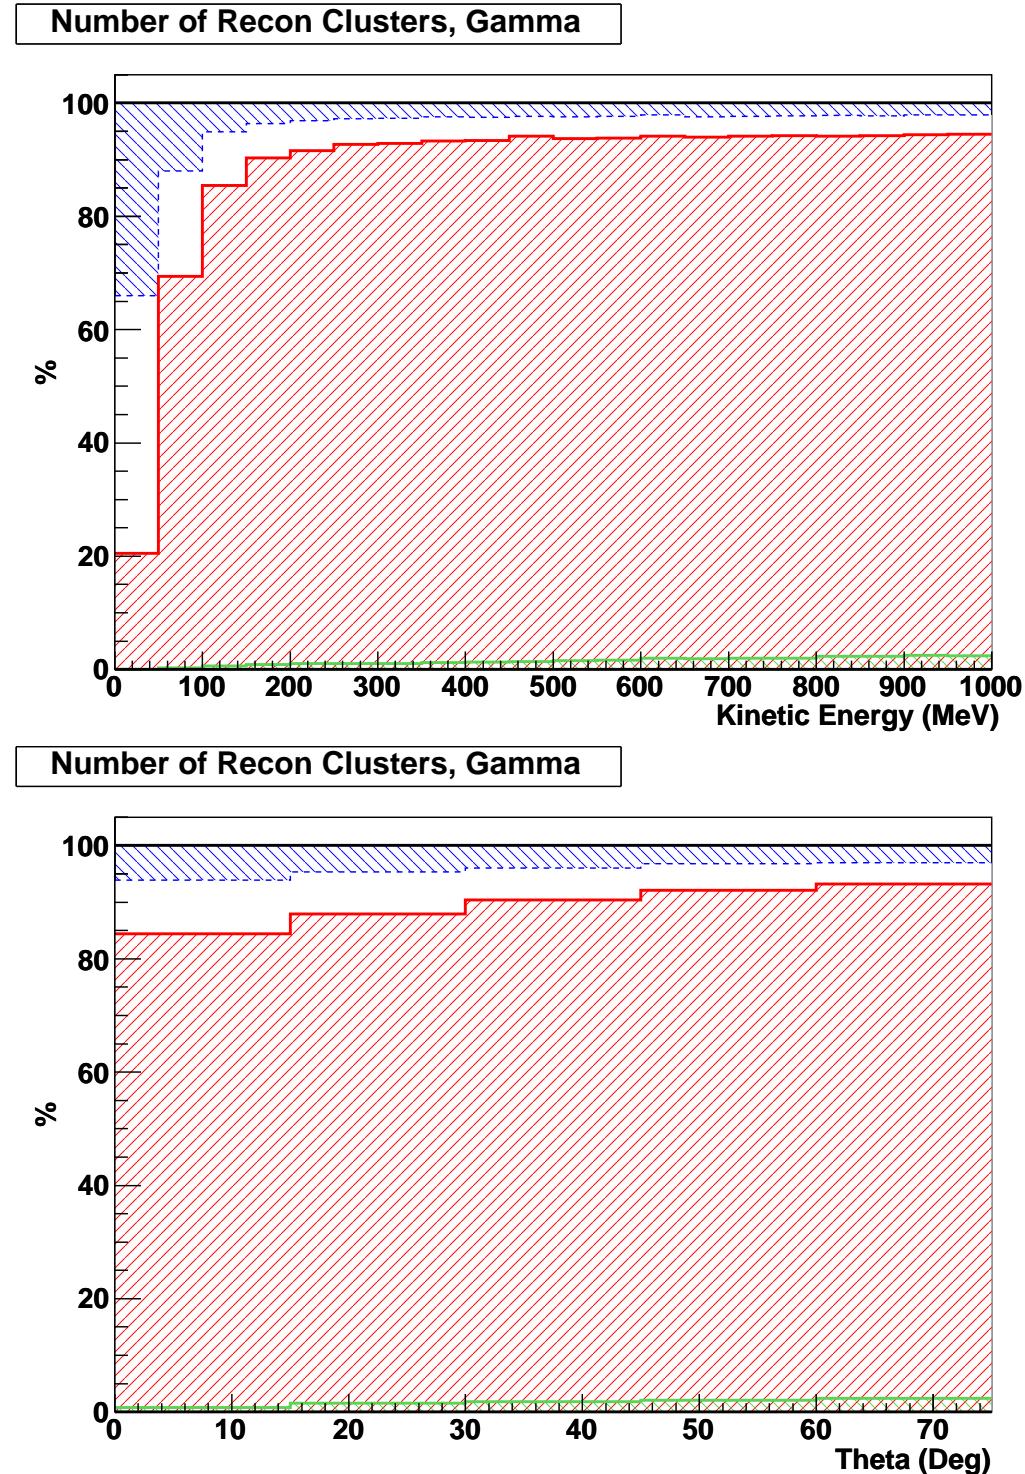


Figure 7.4: The effect of energy and angle of incidence upon photon reconstruction. In each plot the number of reconstructed clusters is shown: Red hatch = 1, green hatch = 2. Also shown are the events with less than the minimum 2 hits required to form a cluster, blue hatch. The white space between the red and blue hatch therefore represents the clustering inefficiency – i.e. potentially reconstructable events in which no cluster is found.

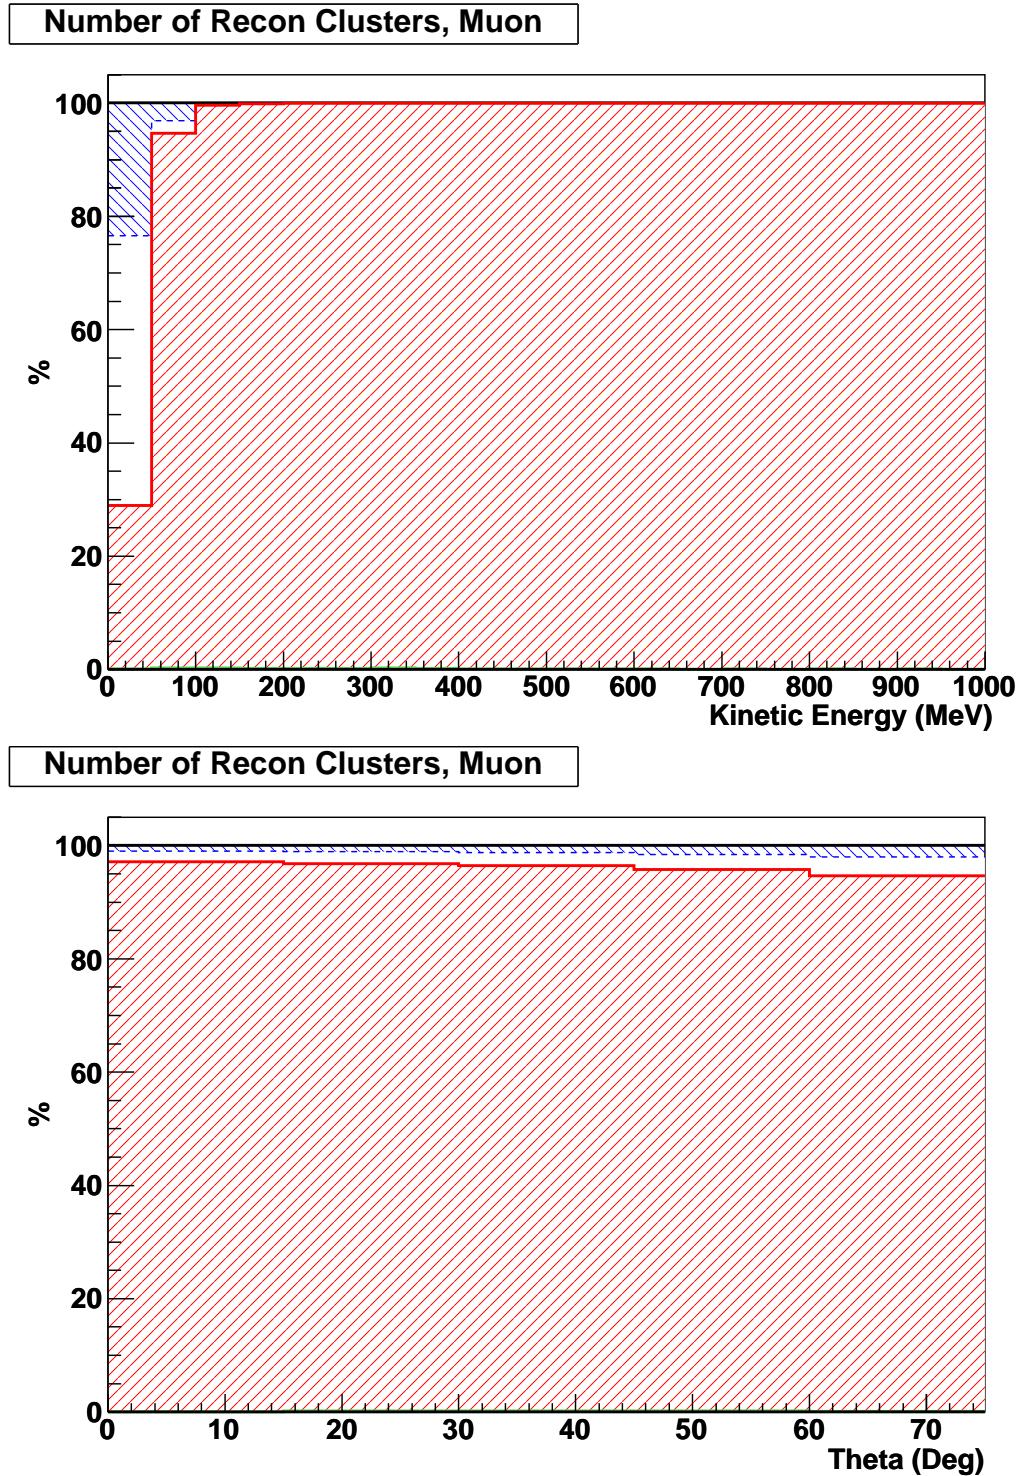


Figure 7.5: The effect of energy and angle of incidence upon muon reconstruction. In each plot the number of reconstructed clusters is shown: Red hatch = 1, green hatch = 2. Also shown are the events with less than the minimum 2 hits required to form a cluster, blue hatch.

≥ 150 MeV.

Again the angle of particle incidence is seen to affect the cluster finding efficiency little. Un-reconstructable low energy events along with associated low energy failed reconstruction form an almost uniform band. There is a slight increase in both un-reconstructable events and failure due to increased radiation lengths and horizontal reconstruction cuts respectively.

7.6.2 Clustering Efficiency

Once a cluster is found one needs to consider how much of the true cluster is reconstructed, and if there is any external contamination from noise or other particles. The former we describe as the efficiency of the reconstruction and the latter the purity. Both quantities can be considered as a bulk number of hits or as charge deposition. Both of these efficiency and purity measures are required for different aspects of higher reconstruction.

For energy reconstruction of EM showers it is true that one requires an optimal clustering of charge, but for hadronic energy estimation one must also include geometric measures. For particle identification purposes not only charge but spatial dependent variables are used to discriminate between types of particle. The position and direction of a cluster are determined by fitting to charge weighted distributions of hits and therefore also convolve these two quantities.

It is therefore necessary to investigate the efficiency and purity of both hit and charge deposition measures when considering the effectiveness of clustering. Again the efficiency and purity are obtained from comparison with true clusters collated using truth clustering. The efficiency is obtained from single particle gun Monte Carlo and the purity from particle gun MC with additional full bunch background.

The plots in figures 7.6 and 7.7 display the clustering efficiency of photon (top) and muon (bottom) events in the P0DECal with change in energy (left) and incidence angle (right). Each energy or angle bin is normalised to one to demonstrate an unbiased performance of charge clustering, Fig 7.6, and bulk hit clustering, Fig 7.7.

It can be seen in Fig 7.6 that the efficiency of clustering charge deposited by both muons and photons is high with the vast majority of events collating $> 95\%$ of the charge deposited. At low photon energies there is a slight decrease in efficiency for some events as charge spreads farther due to low energy photons in the EM cascade and these low energy photons becoming a larger part of the whole. Overall the vast majority of events cluster $> 90\%$ of the charge deposited by photons; this agrees nicely

with the definition of the Molière radius which states that 90% of the EM shower will lie within the Molière radius.

The band of poor efficiency events at low muon energies is again primarily a sign of Michel electron events. The electrons may initiate an EM shower before reaching the P0DECal and as such spread wide before depositing charge in the detector. This can happen in varying degrees, hence the low efficiency band; still note that for the majority of these events that do indeed deposit signal in the detector $> 95\%$ of the charge is clustered.

The top right plot of Fig 7.6 shows us that the clustering of charge deposited by photon interactions is independent of the particle's initial angle of incidence. Muons (bottom, right), on the other hand, show a clear increase in charge clustering efficiency at higher angles; this can be explained best by now turning to look at the bulk efficiency of hit clustering.

The bulk hit clustering efficiency is displayed with a coarse binning due to the discrete nature of the number of hits; 12 bins represent an average of two hits per layer of the detector. The bottom right plot of Fig 7.7 shows the same increase in clustering efficiency, this time for the bulk number of hits. This is attributed to the clustering of low energy delta rays, electrons knocked free of detector material. These electrons leave a small but non-negligible signal in bars to the side of a muon track. A number of these events will see such a delta ray deposit energy in the next nearest neighbour. Perpendicular tracks are therefore unlikely to cluster such hits as part of the track but as one increases the angle there is a greater opportunity for track hits to lie in front of or behind such a delta ray hit. It can be noticed, however, that the majority of muon events do indeed cluster all associated hits.

Apart from the low energy Michel electron events it can be seen that the bulk hit clustering of muons is, as with the charge clustering case, independent of energy. This is expected as MIPs should leave a consistent signal within the detector.

The hit efficiency for photons is far more spread out. This can be explained by concentrating on the small decrease in efficiency as the energy increases. The spread in efficiency arises from low energy photons, produced in the EM shower, travelling beyond nearest transverse neighbour and being excluded by the clustering. These hits are clearly low energy, as the charge clustering efficiency of the detector is high, and small in number. The decrease with energy arises from an increase in multiplicity of such low energy photons and hence a larger fraction of the total hit count. The decrease in the corresponding angle plot is due to a similar effect: as the cluster becomes more dispersed horizontally the likelihood of missed next nearest neighbour hits increases.

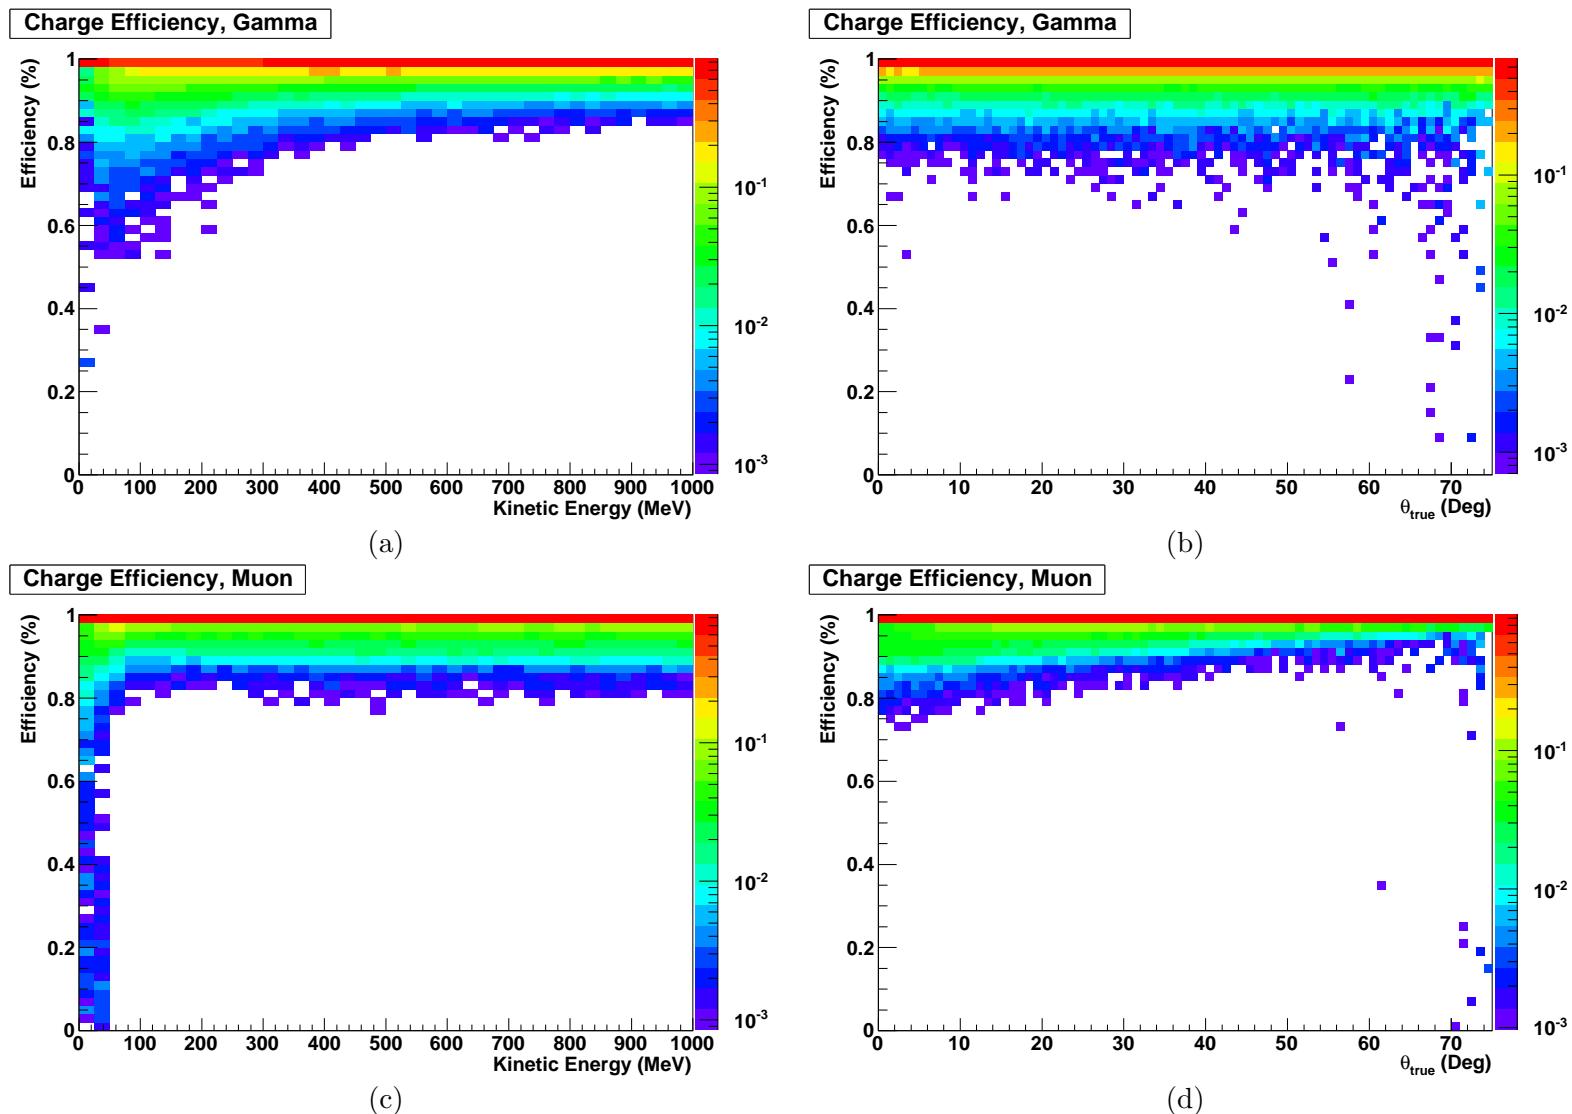


Figure 7.6: The charge clustering efficiency of P0DECal Recon as a function of energy (a & c) and angle (b & d) for both photons (a & b) and muons (c & d). The colour z-axis on these plots denotes the fraction of events normalised to each energy or angular bin.

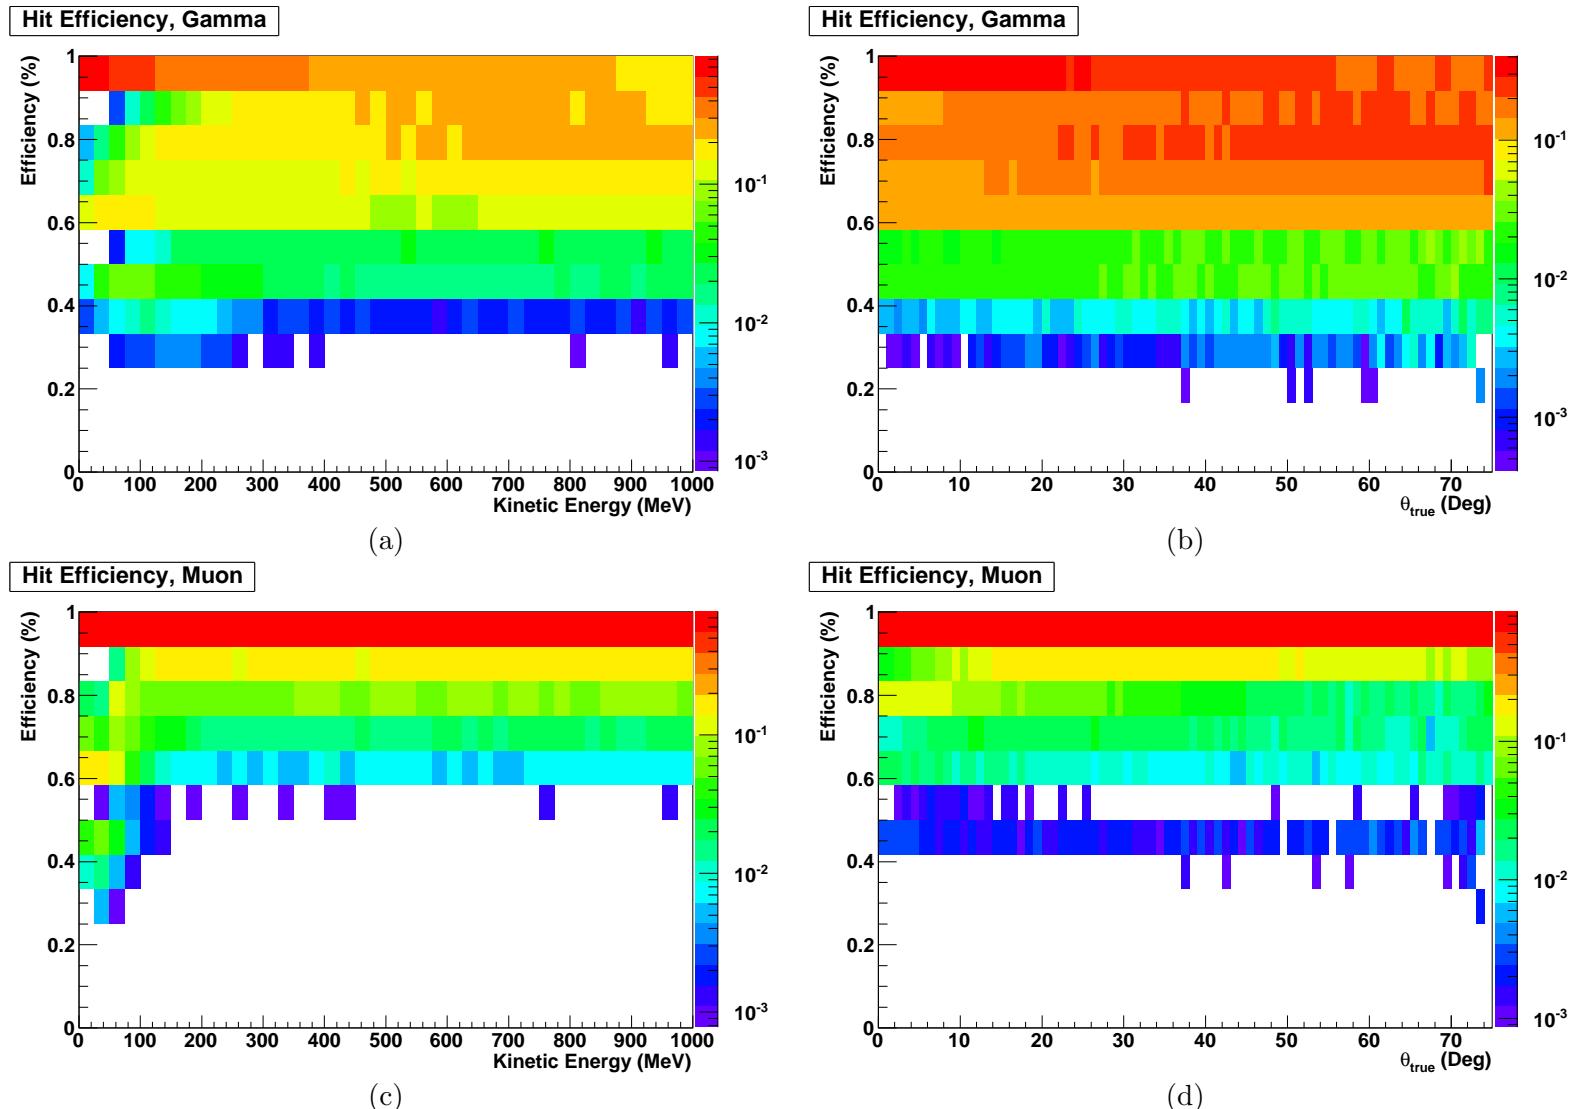


Figure 7.7: The bulk hit clustering efficiency of P0DECal Recon as a function of energy (a & c) and angle (b & d) for both photons (a & b) and muons (c & d). The colour z-axis on these plots denotes the fraction of events normalised to each energy or angular bin.

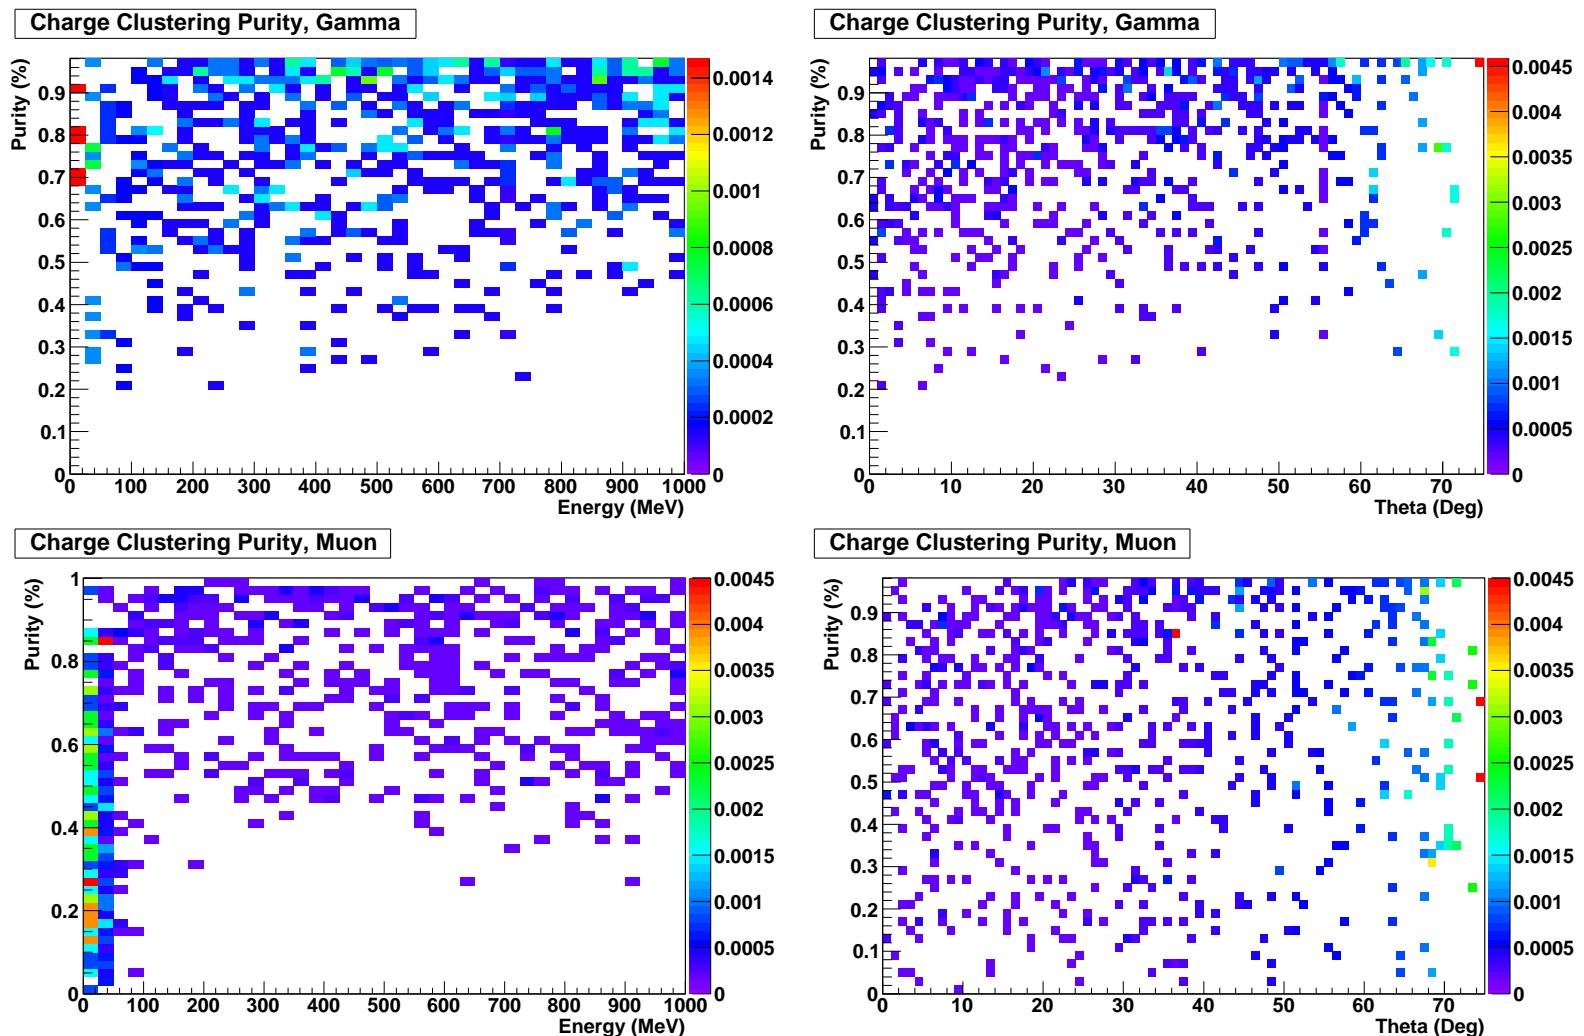


Figure 7.8: The charge clustering purity of P0DECal Recon as a function of energy (a & c) and angle (b & d) for both photons (a & b) and muons (c & d). The colour z-axis on these plots denotes the fraction of events normalised to each energy or angular bin.

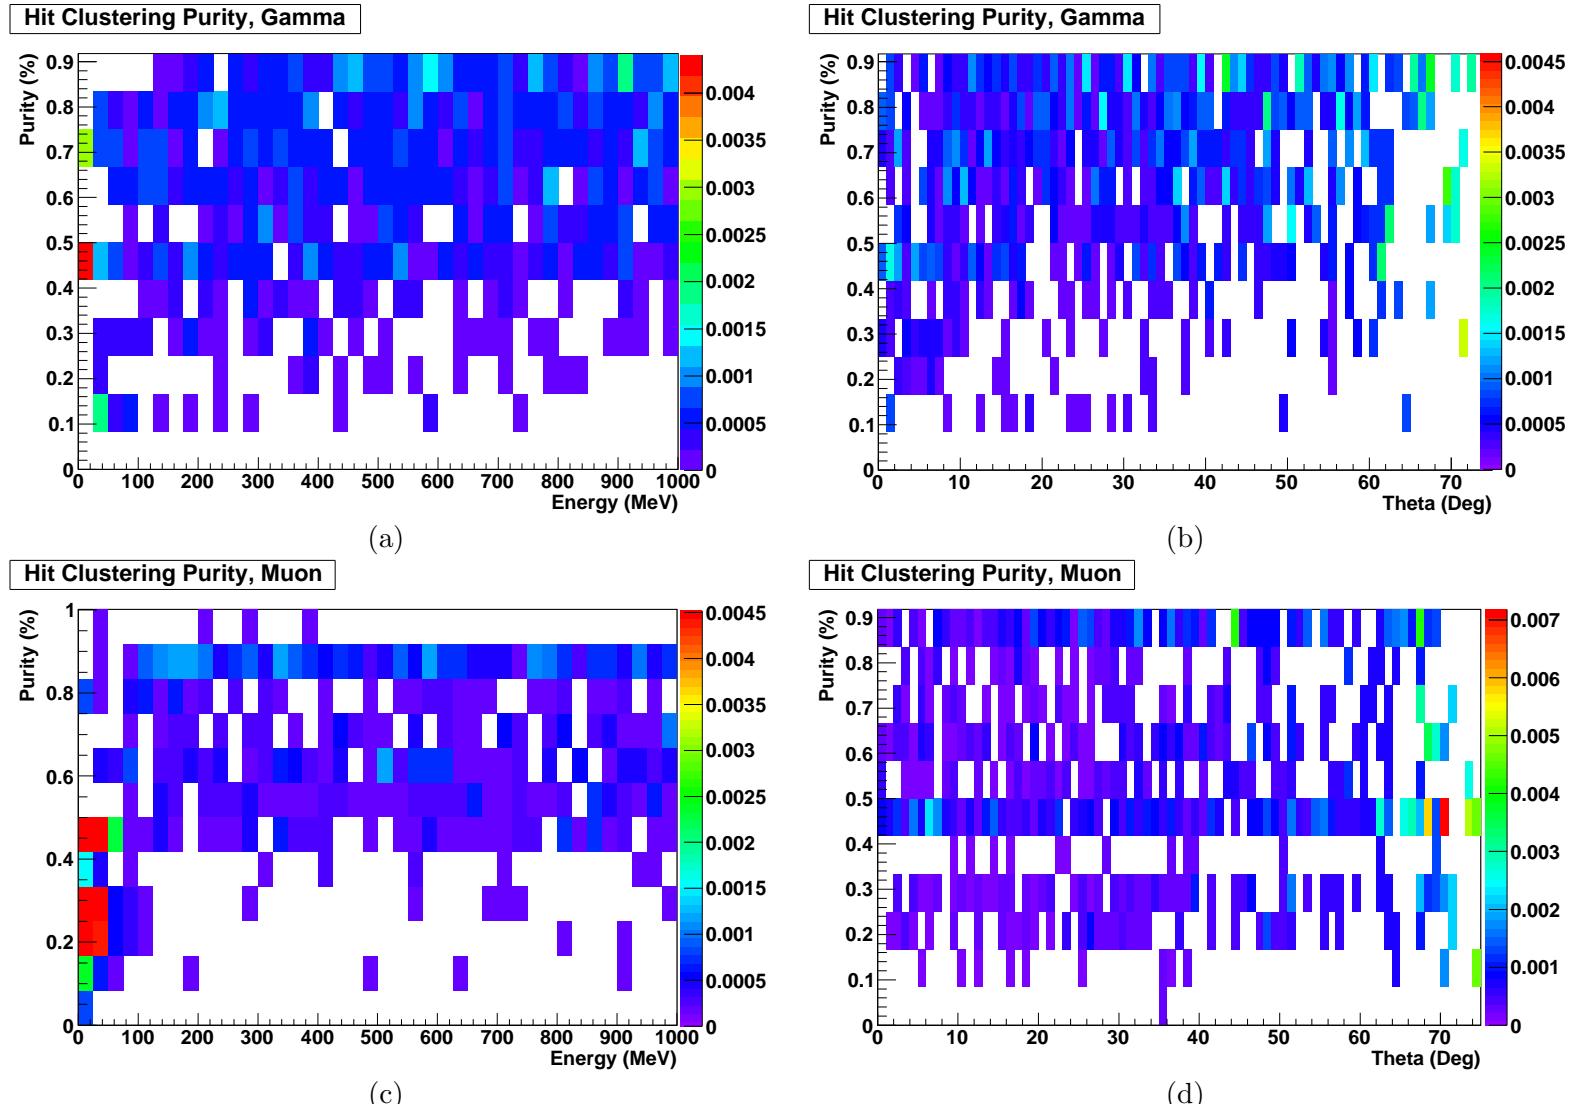


Figure 7.9: The bulk hit clustering purity of P0DECal Recon as a function of energy (a & c) and angle (b & d) for both photons (a & b) and muons (c & d). The colour z-axis on these plots denotes the fraction of events normalised to each energy or angular bin.

7.6.3 Clustering Purity

The purity of clustering is important as any incorrectly associated hits will modify any higher level reconstruction. Looking at the z-axes of figures 7.8 and 7.9, which display the fraction of events in each energy bin, one can see that even in a full bunch environment the contamination of clusters from additional particle events and electronics noise is at an extremely small, $> 0.5\%$, level. Almost 100% of events experience zero contamination with just a few events affected at low energies and high angles.

The events affected most are the extremely low energy events in which the few hits tend to be clustered with either noise or diffuse hits from other particles in the P0DECal. One should, of course, note that this study is based on Monte Carlo, and noise levels in the actual P0DECal stacks have not been measured as the detector has yet to be built. It will be necessary to check the conclusions of this study with real data when it becomes available.

7.6.4 Conclusion

Although the clustering efficiency of bulk number of hits seems poor, this is attributed primarily to the few hits received by the P0DECal in each event. The high efficiency of the charge clustering demonstrates that these hits form just a small fraction, with the vast majority clustering over 90% of the deposited charge. This is not expected to affect particle identification techniques as, although all hits are not collated, the different classes of event still display noticeably different geometric topologies. The low energy of these hits also means that charge weighted position and direction algorithms are not expected to be affected dramatically. The large charge clustering efficiency should therefore not detrimentally affect any charge dependent energy estimation.

7.7 Position and Direction Reconstruction

To match P0DECal events to vertices in the P0D it is important to understand the position at which tracks/showers enter the P0DECal and their direction. This information is obtained from fits to the hit distributions within the reconstructed cluster.

7.7.1 Position Resolution

The position of entry is determined by calculating the charge weighted position of hits in the innermost layer that receives signal. Each hit position is multiplied by the

charge deposited in that particular bar and summed before being finally divided by the total charge in that layer. The construction of the P0DECAL leaves it only sensitive to position in the Y-direction (X-direction) on the front face of the side (top/bottom) modules.

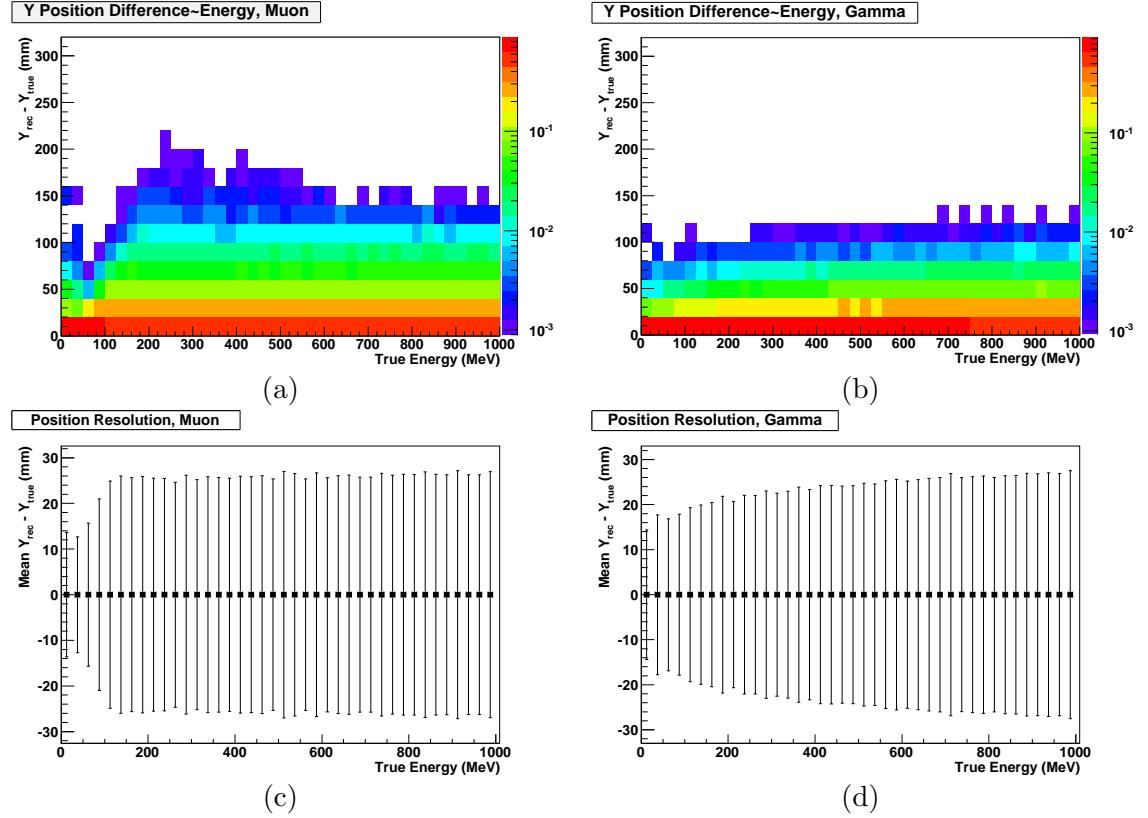


Figure 7.10: The position resolution of the P0DECAL reconstruction. Plots (a) and (b) show the difference in reconstructed and true front position for muons and photons respectively. The colour z-axis denotes the fraction of events in each energy bin. The error bars on plots (c) and (d) show the standard deviation from the true position, obtained from a gaussian fit to 1D projections of (a) and (b) on an energy bin by bin basis.

In Fig 7.10 we display the difference of reconstructed Y position to the true position of the trajectory as it enters a side P0DECAL stack as a function of energy (top) and a resulting Gaussian fit to each energy bin (bottom).

The resolution is seen to be slightly worse than the bar width for both photons and muons with a standard deviation around $\sim \pm 25$ mm. This is expected when using just a single layer estimate as events hitting just one bar will have a resolution of bar width at best and those events which share charge between bars will add further smear to this measure.

7.7.2 Direction Resolution

Two methods of determining the direction of a cluster have been considered: thrust (a variable originally devised for hadronic jet reconstruction, which chooses an axis to maximise longitudinal momentum) and principal component analysis (PCA). Both methods have limited effectiveness in the P0DECal because of the small number of layers, but PCA is found to perform better than thrust and is the method adopted in this study.

The primary principal component, the major axis, is taken as the direction of the cluster. Note that if only one layer is hit then the method returns a direction perpendicular to the face of the stack; there is therefore a constant undertow of events reconstructed at $\theta = 0$.

It can be seen in Fig 7.12 that this method is most successful for muon tracks. This is because muon events typically traverse the entire detector and give a clear primary axis, features which, combined, yield more directional information. The converse are the reasons for the difficulty in reconstructing photon direction; just a few layers traversed and a rounder hit pattern.

The left of Fig 7.13 shows the fraction of these events classed as ‘failed’ by the reconstruction. The blue distribution shows the fraction of events incorrectly reconstructed to be traversing the stack vertically. The stacked black distribution shows the fraction of events within the failed bins, where failed is defined here as reconstructing a mean angular difference greater than one standard deviation from zero.

It can be seen from the photon plot that the primary reason for failure is due to photon showers hitting just a single layer (or not leaving much information in a second). This effect increases with incidence angle due to two factors: increased fragmentation/missing hits and an increase in radiation length, of dead material in front of the stack and absorber between layers. There is a small increase in reconstruction failures at high angles again associated with missing late or early hits.

The top right plot shows the angular resolution for photon showers which remains around $\pm 10^\circ$. The apparent increase in resolution as one reaches higher angles is false as it is the angle that is aligning with the short-comings of the reconstruction; i.e. if just one layer is predominantly hit then the reconstruction naturally prefers large angles.

The ‘successful’ region in the muon plot is seen to be wider and shows a marked improvement upon the photon case. This is expected for the long straight track left by a muon event. The right plot shows that the resolution of the muon track in this region is much improved on the photon events and shows a one standard deviation difference

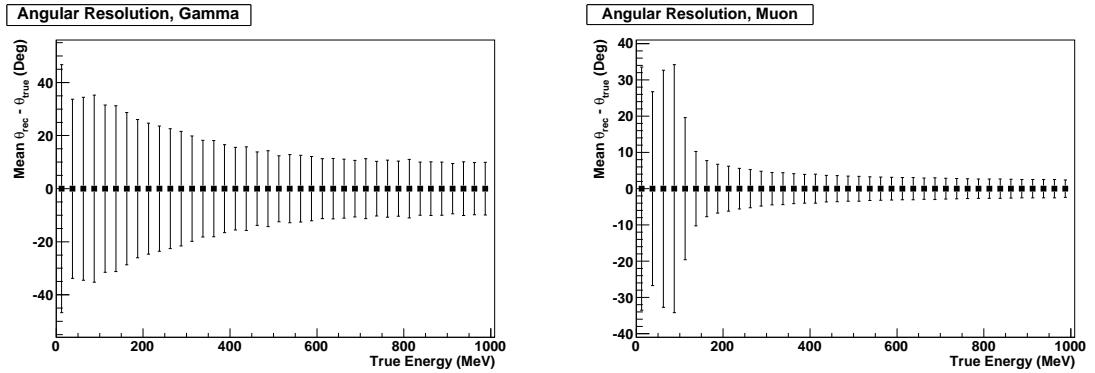


Figure 7.11: Angular resolution as a function of true energy; the error bars show the RMS of the difference between reconstructed and true angle of particle incidence.

of $\pm 2 - 4^\circ$. Those events assigned to the 0 bin or failing reconstruction are below the 10% level for the majority of the successful region rising only at the extremes. At low angles the increase is due to the finite width of the ECal bars, all low angles being indistinguishable from zero. At higher angles the fraction of events reconstructed as perpendicular arise from the increased probability of creating, and collecting, delta ray hits. These are the same reasons for the decrease in angular resolution with increase of true angle in the right hand plot. The failure at high angles is again associated with the geometry of the detector; at 75° a horizontal distance of 4 cm gives a vertical distance of ~ 1 cm and so at high angles there is a greater chance of muons depositing signal in just a couple of layers before exiting the detector.

We now plot the angular resolution with true energy in Fig 7.11, here we have included all results, both ‘failed’ and ‘successful’. It can be seen that the worst resolution is at low energies for both type of particle. The sharp decrease in the muon, right-hand, plot is associated with the lower energy events being dominated by Michel electrons rather than a muon track. These events and the low energy photon events are of low resolution due to perpendicular reconstruction, as just one layer has been hit, or failed reconstruction. The resolution for events which are predominantly tracks, above 150 MeV, is $< \pm 10^\circ$.

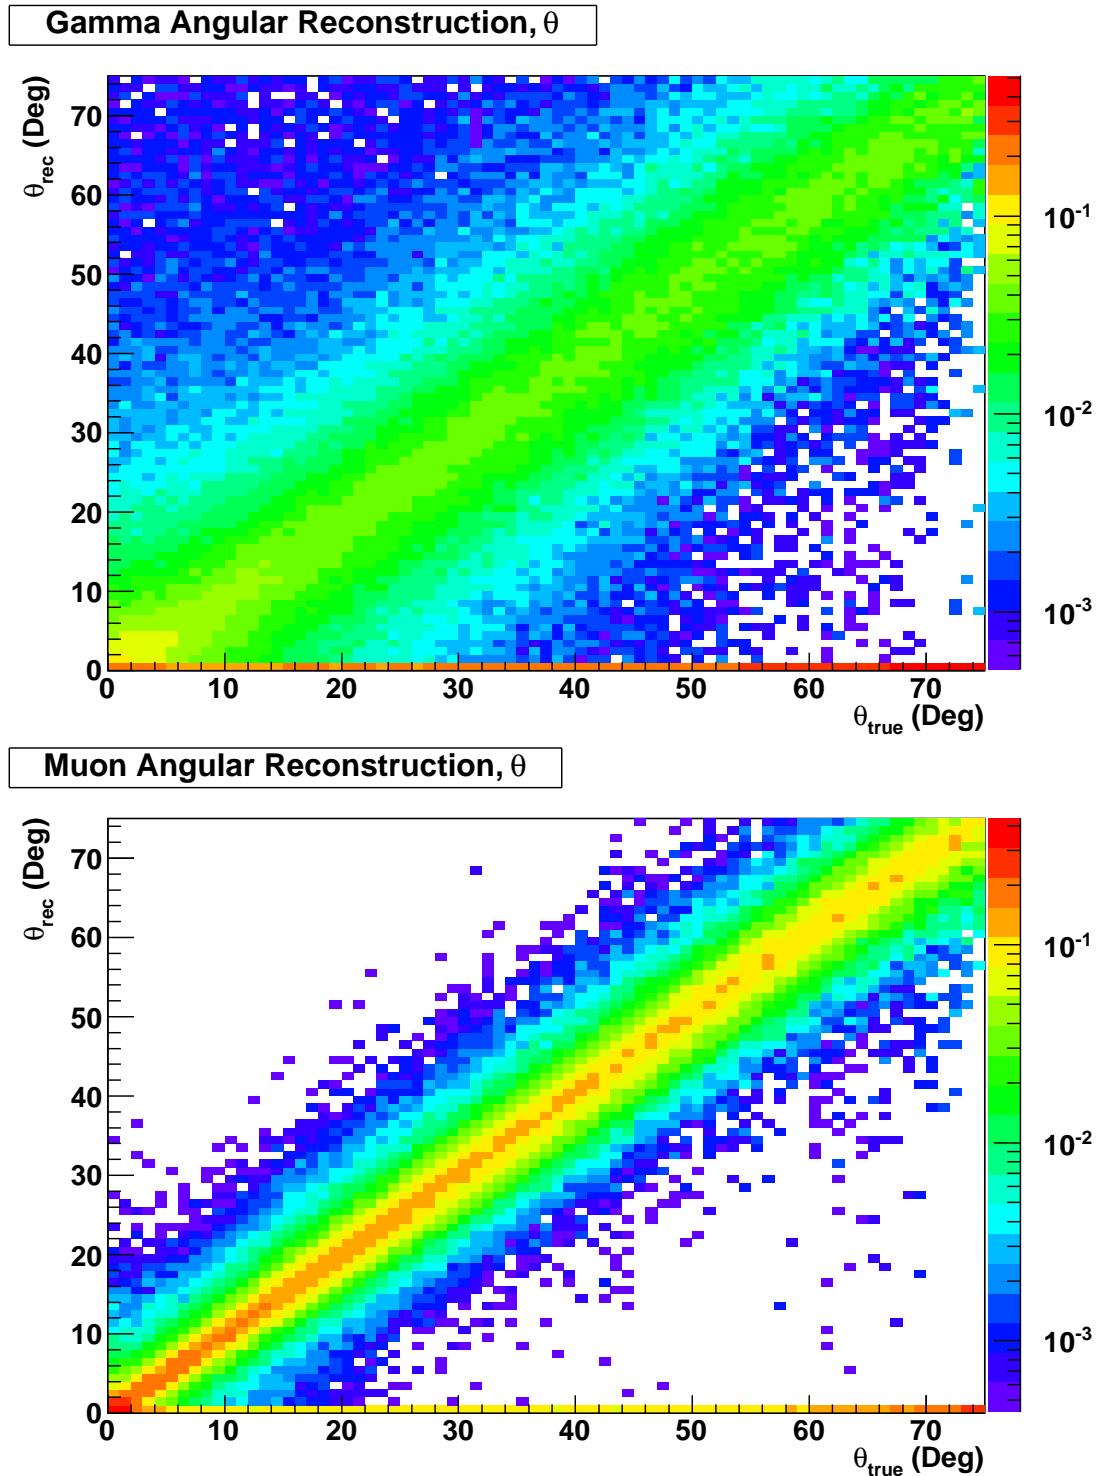


Figure 7.12: The PCA direction reconstruction of P0DECalRecon as a function of true incident direction, for photons (top) and muons (bottom). The z-axes denote the fraction of events within each angular bin.

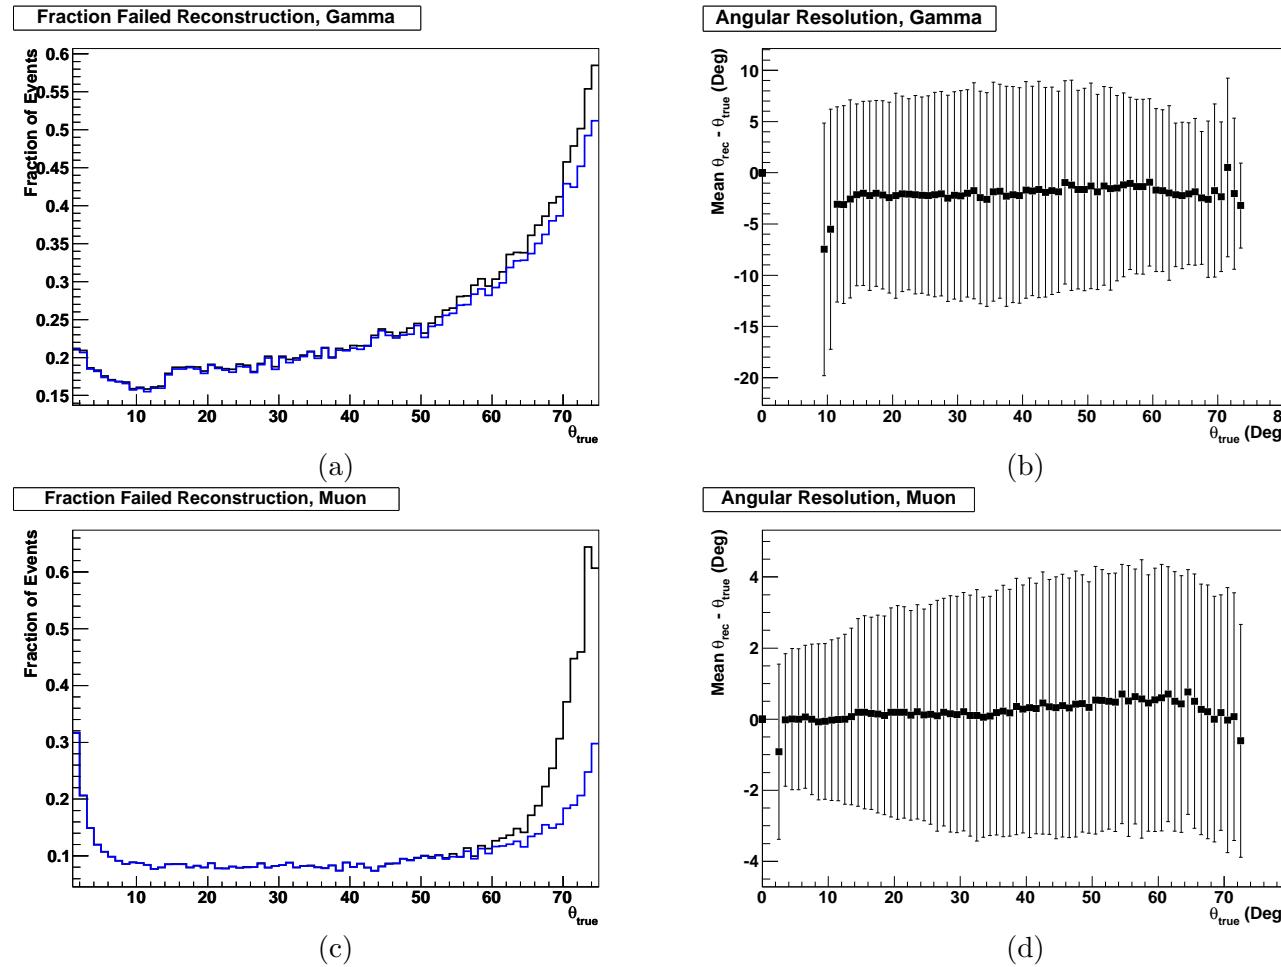


Figure 7.13: The fraction of failed direction reconstructions for photons (a) and muons (c). The blue line charts the fraction of events with no handle upon direction and the black charts those events which lie within a bin with mean greater than 1 standard deviation from the true value. The direction resolution of P0DECalRecon for successfully reconstructed directions is shown in (b) for photons and (d) for muons.

7.7.3 Conclusion

Angular reconstruction with just a six layer detector was always going to be difficult. When combined with the successful position resolution the accuracy of pointing into the P0D, combined with the expected low number of interactions per spill therein, should be sufficient to confidently match tracks and even photon showers to P0D vertices. This is an important factor when selecting π^0 candidate photons and excluding charged current events when tagging muon tracks.

7.8 Energy Reconstruction

In a calorimeter the simplest energy reconstruction can be performed for electromagnetic showers. Assuming containment, the charge deposited by the hits of an EM shower directly relates to the energy of the EM particle that initiated it. Here I present the results of an attempt to use a likelihood function method to reconstruct the energy of EM showers in the P0DECal.

7.8.1 Reconstruction Algorithm

A maximum likelihood fit was performed to different charge deposition variables: the total charge (QSum), the maximum charge deposited in a bar (QMax), and the mean (QMean) and population standard deviation (QRMS) of the charge per bar. Geometrical variables were considered, but had been found to detrimentally affect the energy estimation in the Tracker ECal [124] and were therefore disregarded for this investigation.

Each of these parameters is well represented by a skewed Gaussian function,

$$\frac{A}{\sigma_s \sqrt{2\pi}} e^{-\frac{(E-\mu)^2}{2\sigma_s^2}} \quad \text{where} \quad \sigma_s = \sigma + s(E - \mu) \quad (7.2)$$

which is fitted to each of these variables at a range of energies, E . The mean (μ), skew (s) and standard deviation (σ) of these fits are then plotted as a function of energy and fitted with polynomials of different order. The fits to these variables x_i are then used to create a negative log likelihood function evaluated at each energy over a desired working range.

Probability is calculated from the functional form of the various variables x_i . At each energy E the inverse covariance matrix C_{ij}^{-1} , corresponding determinant and vector of difference $D_i = x_i - \mu_i$ of each variable and corresponding mean are calculated and

constructed. A single real variable measure of matching success is then calculated, $M = D_i^T C_{ij}^{-1} D_j$, which is essentially the mean of the distance of each value from its corresponding variable mean, normalised to the standard deviation. When divided by the determinant of the original covariance matrix we get a normalised measure (probability) that such a vector of data belonged to each energy.

Any energy values outside of the desired working range have their probability penalised by an exponentially decreasing factor, dependent upon their distance outside of this range. This limits the bounds of the fit in a continuous way. The working range of 25 MeV–7 GeV used by the TECal is employed in the following investigation.

The tails of skewed Gaussians do not vanish far from the mean value. This is fixed by again penalising those values which lie in the tails, defined as the regions $\pm 5\sigma$ of the mean. A second Gaussian is superimposed upon the first to push the values in the tails towards negligible.

The final probability is minimised using the TMinuit gradient minimisation technique and the minimum chosen as the respective shower energy.

7.8.2 Calibration

From the Monte Carlo truth design study it was understood that the best attainable energy resolution of the P0DECal was quite poor, see section 6.3, due to the coarse sampling of particle showers. Any accuracy of energy reconstruction will reduce still further when this resolution effect is coupled with the response of the electronics.

There is also a competing detrimental effect arising from attenuation correction of the measured signal. Each P0DECal scintillator is read out from a single end of a bar and all bars are parallel, with no crossing geometry. These two facts dictate that one cannot determine with any accuracy, without external information, the position of an event along a bar’s length. Without this information one cannot calibrate any signal attenuation loss along the bar; for this reason this step is currently not employed in P0DECALRecon.

To investigate how attenuation calibration affects possible energy reconstruction of EM showers in the P0DECal the reconstructed Z position, along the bar, was replaced in each hit with: the bar centre position (worst case reference), true position and gaussian smeared true position of standard deviation 25, 50, 75 and 100 cm. The likelihood function was trained and tested in each scenario to obtain a measure of energy reconstruction.

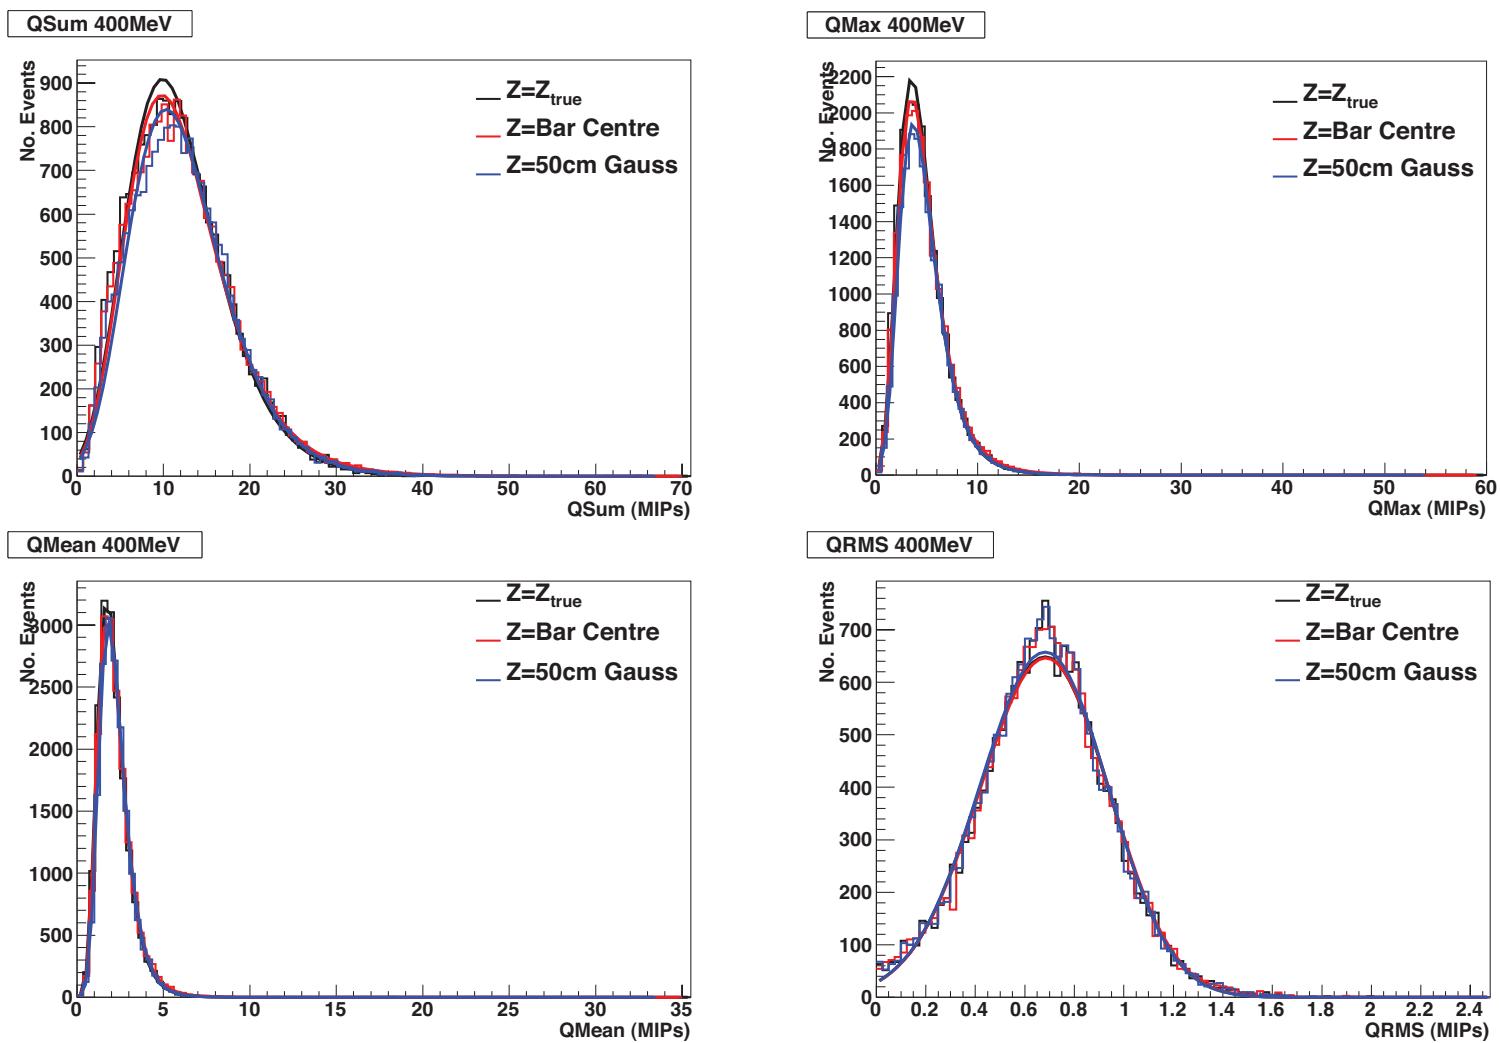


Figure 7.14: Variables used in the energy reconstruction likelihood function. These distributions are for 400 MeV gammas. All variables are well fitted by a skewed Gaussian function and there is little change in these variables with varying accuracy of attenuation correction: there is little difference between the attenuation corrected signal using the real, centre and $\sigma = 50\text{ cm}$ gaussian smeared real position along the detector z-axis, length of the scintillator bars. This suggests this is not the limiting factor of the energy estimation.

To train the estimator particle gun MC data were used. Mono-energetic photons were fired perpendicularly from a small distance in front of the left side PODECal stack, with an isotropic position, over the front face of the detector. Photons were chosen, as commonly they do not lose energy within the dead material in front of an ECal stack, which would affect the energy estimation. For the same reason though, photons also give the most varied distribution of each variable as one samples different percentages of showers, although from the design study one can see that the majority of photons convert in the first two layers of absorber with the majority of shower energy contained. Fifty energies were selected, at different intervals, covering the range 75 MeV–10 GeV, with a greater density chosen at low energies.

If we look at the distribution of the variables used in this likelihood estimator, Fig 7.14, we see little effect from altering the Z position, and therefore accuracy of attenuation correction, of the hits. The estimation power, as explained in describing the method above, comes from the normalised deviation of a variable from its respective mean. One can see that the standard deviation divided by the mean of any of these similar distributions will be extremely small.

Energy Recon Result 400MeV Gamma

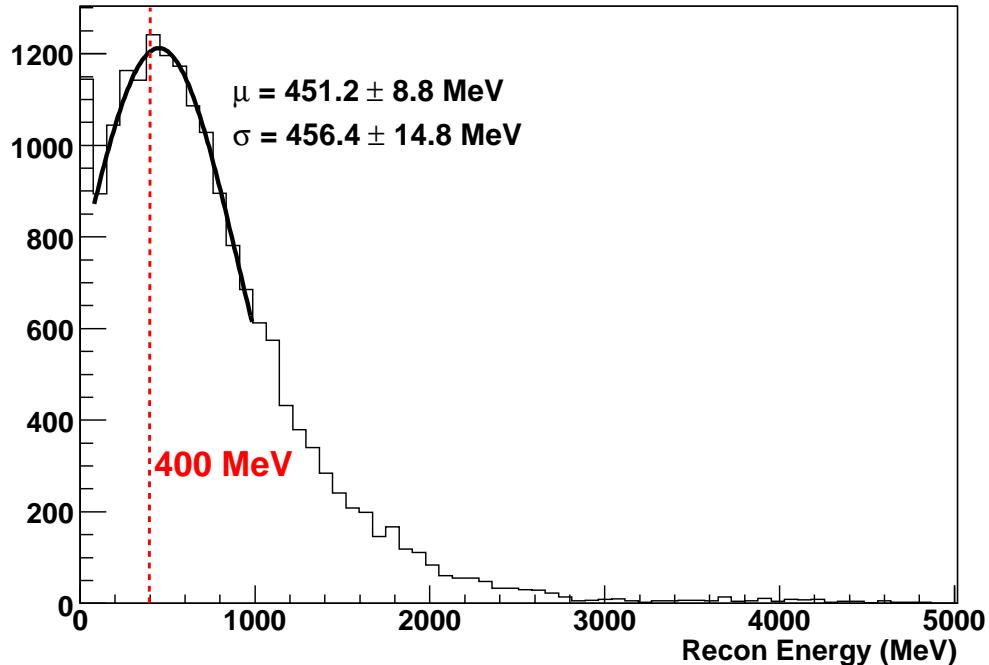


Figure 7.15: Reconstructed energies of true 400 MeV energy photons.

What we indeed find is that this method fails to estimate the energy of EM showers.

Fig 7.15 shows the reconstructed energies of 400 MeV photons fired perpendicularly into a P0DECal stack. As can be seen the distribution is, to a good approximation, random.

7.8.3 Conclusion

In conclusion, the use of a likelihood function estimator to reconstruct the energy of photon showers is unsuccessful in the coarse P0DECal detector. This failure is attributed to the inadequate sampling of the shower which is affected by the few active elements and the relatively large amount of inactive radiation lengths between each layer.

The clustering has been seen to perform well in the vast majority of events and so we conclude that the choice of estimator used or the overall sampling resolution of the detector is the weakest link in reconstructing the energy of EM showers. It can be seen in section 8.2 that some of these variables do possess some degree of correlation with true energy, suggesting that it may be still be possible to gain a handle on showers in the P0DECal. Further work is required and will be continued by myself or collaborators in the months before detector commissioning.

7.9 Conclusion

In conclusion the P0DECal demonstrates sufficient performance, in almost all aspects of reconstruction, to satisfy the requirements of higher level reconstruction techniques. The high cluster finding and hit clustering efficiencies produce representative and pure sample of each class of particle; this is essential for particle identification performed on such clusters. The position and direction resolution are sufficient to point to the regions of the inner P0D detector where vertices may exist, allowing for charged current discrimination and π^0 photon tagging. The weakest link in the entire reconstruction chain is the energy reconstruction which requires more work to produce usable results.

Chapter 8

P0DECal Particle Identification

The coarse granularity of the P0DECal detector limits the number of different geometric and charge deposition event topologies. To meet the physics requirements of the P0DECal these topologies are split into two classes; track and shower. Tracks are defined by minimally ionising particles (MIPs). Showers are traditionally split into two further classes of electromagnetic (EM) and hadronic, defined by the particle initiating the shower. Distinguishing the different topologies arising from these two shower classes would be too great a challenge using the discretised data given by the six single ended readout scintillator active layers of the P0D ECal. For this reason they are grouped into a single shower event class.

Particle gun MC is used to simulate numerous examples of each class of event in a P0DECal stack. Events were generated for each particle type with their respective expected ND280 energy range, gleaned from large scale neutrino interaction MC. These MC data are then used to train and test linear (Fisher discriminant and projective likelihood) and non-linear (artificial neural networks and boosted decision trees) multivariate analysis (MVA) techniques for the discrimination of track and shower events.

8.1 Training Samples

To train MVA techniques it is necessary to select reference (training) samples of each class of event – in this case “track” and “shower” events. This cannot be done purely on the identity of the initiating particle (π^\pm can be either MIPs or showers; decaying muons may look shower-like). Therefore “good” examples of each class were first selected using true trajectory information.

8.1.1 Tracks

Minimally ionising particles leave narrow geometric and charge deposition distributions. Tracks typically hit no more than one bar per active layer in the P0DECal and the charge deposit per bar is Landau distributed with a centre at the MIP charge. These minimally ionising tracks are formed by heavy charged particles, almost entirely muons and non-interacting charged pions in the ND280. MIP tracks typically traverse many radiation lengths of material because of their minimal energy loss: above 150 MeV $> 99\%$ of muons traverse the entire P0DECal stack. Therefore the narrow drawn out hit pattern characterises the geometric topology of a track. The relatively stable and low energy energy deposition, close to the MIP charge, characterises the charge deposition distribution of track events.

As the kinetic energy of muons decreases they may lose enough energy to ionisation to stop in the detector. These events with stunted tracks, over just a few P0DECal layers, lose some amount of the characteristic long narrow geometric hit distribution. Pions may interact at any point in their journey through a detector, more likely at lower energies as they spend a greater amount of time in the field of detector matter. These interactions can be simple capture by a nucleus, in which case the track stops short much like low energy muons, or they may initiate a hadronic shower. A pion track that initiates a hadronic shower has been dubbed a lollipop event, with a track stick and a shower sweet.

All muons were assumed track-like if their trajectories traverse the entire P0DECal stack. All pions doing the same are accepted as tracks, as also are those pions which shower in just the final absorber or active layer. This was decided upon as for the vast majority of such a pion's journey it was indeed a MIP and the anomalous signal in the final layer will do little to alter this.

8.1.2 Showers

EM showers are initiated by photons or electrons whilst hadronic showers are initiated primarily by charged pions. The geometric topologies of showers are widely spread distributions. The Molière radius of EM showers in plastic scintillator and lead (~ 12 cm) dictates that showers should hit multiple bars in each active layer. The charge deposition of shower-like particles is also variable, with high energy deposits greatly above that of MIP (for the majority of ND280 energies) from primary particles and lower charge depositions from secondary, tertiary, etc., particles.

The radius of an EM shower is determined primarily by multiple scattering of non-

radiating electrons, which is characterised by the Molière radius

$$R_M = 21X_0/\epsilon_c \simeq 7A/Z \text{ g cm}^{-2}. \quad (8.1)$$

The radius of a hadronic shower is determined approximately by the scale of the hadronic interaction length because of hadronic secondary, tertiary, etc., particles produced transverse to the initial particle momentum,

$$\lambda_I = \frac{A}{N_A \sigma_I} \simeq 35A^{1/3} \text{ g cm}^{-2}, \quad (8.2)$$

where $\sigma_I \propto A^{2/3}$ is the interaction cross-section which we approximate as proportional to $A^{2/3}$ from basic geometrical assumptions. This interaction length is much greater than the Molière radius both in the plastic scintillator and lead absorber. This increases the radius of hadronic showers when compared to EM showers. Hadronic showers are not in fact purely hadronic, as one-third of the time neutral pions can be produced, decaying into photons which can then initiate EM showers. Hadronic showers are often therefore a mix of hadronic and EM components which cannot be separated. The large amount of inactive lead absorber in the P0DECal will stunt the progression of both EM and hadronic shower over just a few layers; this fact alone makes distinguishing between the two class of shower a difficult task in the P0DECal.

Showering pions were selected as trajectories stopping in or before the penultimate absorber layer of the P0DECal stack which produce hadronic daughters. Photons with trajectories converting in the P0DECal and all electrons are selected as EM showers.

8.2 Discriminating Variables

Once fully reconstructed, the collection of hits is passed to a particle identification algorithm which proceeds to characterise the nature of the cluster. The first stage in this process is the calculation of a number of discriminating variables which will be used for classification.

Within the P0DECal we have two distinct groups of variable; geometric and charge deposition. Geometric variables depend upon the pattern of scintillator bars that receive a signal from a traversing particle. Charge deposition variables rely on the size of the signal received by scintillator bars. These groups of variable are deconvolved by the granulated nature of an ECal. Geometric variables are highly discrete because of the geometry of the ECal bars while charge variables display greater continuity. In the limit of infinite sampling the two groups of variables would be perfectly correlated.

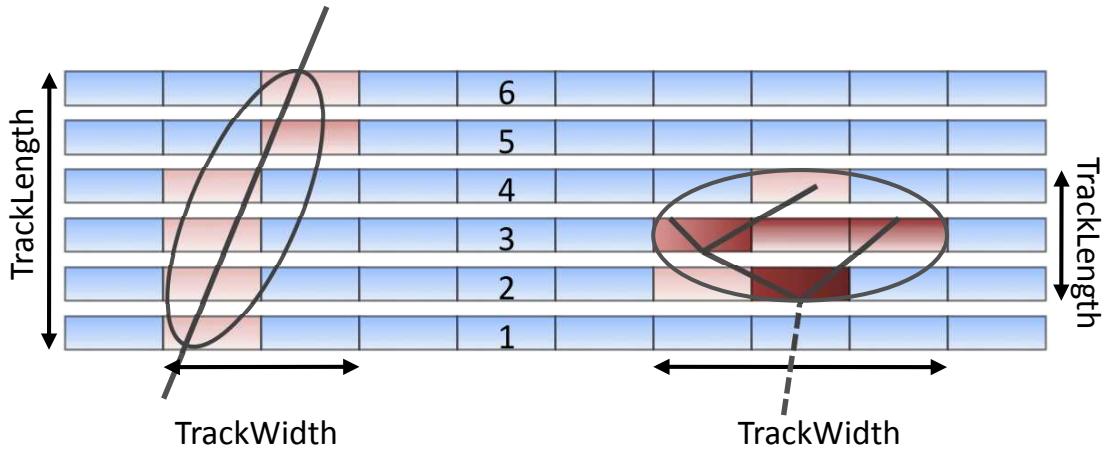


Figure 8.1: .
Schematic diagram of a typical track and shower event.

When selecting variables from these groups we must not only take into account their discriminating power but also their angular and energy correlation. This is important as the range of energies and angles of particles incident upon the P0DECal presents a massive phase space for which one would need training data and which would most likely result in a number of versions of the MVA technique being employed in discrete regions of this massive phase space. It will be possible to cover a wide range of angular phase space if required during test beam studies at CERN, but the same cannot be said for energy. The testbeam will be operating at energies above the peak range expected in the ND280. Although tuning of MC could allow extrapolation to lower energies, this is not the optimal method, so one would ideally select variables which present little correlation with particle energy.

We will also discuss the correlation of variables with other discriminators. This is important to consider as highly correlated variables can reduce the effectiveness of a chosen MVA technique. Some MVA techniques, such as projective likelihood, are ignorant of variable correlation, which drastically reduces the combined discriminating power of correlated variables. Most other techniques, e.g. boosted decision trees, also exhibit marked underperformance in the presence of highly correlated data.

8.2.1 Track Length

The first discriminating variable that sprang to mind, and was used as a measure in the P0DECal design study, was the number of layers traversed by the cluster or TrackLength. It is expected to peak at the full six layers for tracks and at just a few

for less penetrating showers, see Fig 8.1. This is indeed what we see in Fig 8.3i, an impressive separation. When looking at the correlations of this variable we see that it is highly correlated with the energy of a shower, as one would expect as energy dictates a showers penetration depth, but remains uncorrelated with incidence angle.

8.2.2 Axis Max Ratio, AMR

The AMR variable is a geometrical variable that gives a measure of the sphericity of a cluster. It is calculated using principal component analysis (PCA) which, given a set of variables, calculates a set of principal components, eigenvalues and vectors, which track the variance of the set of variables. The principal components are listed in order of ascending variance. The eigenvector of the first principal component therefore gives the major axis of a cluster of hits and the second eigenvector the minor axis. This is shown schematically in Fig 8.1 by encompassing the hit bar centres inside ellipses: the major axis of the ellipse corresponds to the the primary principal component, and its minor axis the secondary principal component.

The axis max ratio is the ratio of these first two principal components. The value of the AMR is expected to be largest for tracks, with a long major axis and short minor axis as the narrow track traverses the entire distance through the detector stack. The transverse spread seen in shower class events should produce smaller values as the topologies of the hit pattern become more spherical.

This is indeed what we see in Fig 8.3a. To avoid large ratios the AMR variable is capped at 300, any events with a value greater than this are simply set at 300. In this bin one can clearly see the track events far outweighing the shower events. If we look at the main body of the distribution one can see that the track distribution has a higher mean and long high AMR tail whilst the shower event distribution is heavily peaked toward one.

AMR is almost entirely uncorrelated with energy and incidence angle of track class particles, but demonstrates a relatively strong correlation with energy for shower events. This is a negative correlation and is attributed to the penetration of showers through the P0DECal increasing with energy and becoming more spherical, as the width of showers remains largely constant. The correlation is however less than the similar TrackLength variable.

8.2.3 Path Charge Ratio

The path charge ratio takes the ratio of the charge deposited in the first few layers and back few layers of a cluster. If a cluster traverses 6 layers and deposits charge enough to register a hit in each then this variable is analogous to the charge ratio used in the TECal [125]. When a cluster traverses four or five layers and again registers hits in consecutive front and back two layers then the ratio is taken of the charge deposited in these. All other times the ratio is just between the charge deposited in the first and last layer in which hits are registered.

As an example the track on the left in Fig 8.1 would have a path ratio of the sum of charge in the first three to the last three layers as it traverses the entire stack and leaves a hit in each layer. The three layers hit by the photon shower, right Fig 8.1, would mean that this event's path charge ratio is the ratio of charge deposited in layers 2 and 4 respectively.

Fig 8.3(h) shows that track events have a tight distribution which peaks around 1, which one would expect as a track deposits charge uniformly along its trajectory. The small high ratio tail is most likely from delta ray production by both muons and pions and the very start of a shower in a small fraction of pion events. The shower distribution has a much larger standard deviation and long high value tail. The peak at low values arises from the highly asymmetric events where low energy photons penetrate to leave charge in later layers.

Again this variable shows no correlation with energy or incidence angle of track events and remains largely uncorrelated with energy for shower events.

8.2.4 Max Charge Ratio

The Max Charge Ratio is the ratio of the charge deposited in the layer which receives the largest signal and the charge deposited in the layer with the smallest signal. In our schematic in Fig 8.1 this would be between the slightly higher deposition in layer 5 and the lowest of the other layers for the track and clearly between layers 3 and 4 of the shower.

From data it is seen that the MaxRatio is larger for showers, whose charge deposition is highly asymmetric, and peaked at one for uniform tracks; this can be seen in Fig 8.3(d).

Fig 8.2 shows that although this is a charge variable it is relatively uncorrelated to the energy of a track or shower incident particle. This is because of the relative scaling of the charge deposited in all layers of a shower with energy. It is also fairly uncorrelated

to the incident angle of a particle with the tracks.

8.2.5 Charge Standard Deviation

The charge standard deviation variable is the population standard deviation of the charge deposited in each bar. This is expected to be larger for showers which leave a disparate deposition of energy and low for the uniform charge deposition of tracks, which can be seen from data in Fig 8.3(b).

This is quite highly correlated with the energy of shower events, growing larger with energy, as one would expect as larger showers result in more particles and therefore more varied hit patterns. The correlation is reduced by dividing by the MeanCharge variable, to which it is highly correlated and which is also correlated strongly to energy, this is the normalised charge standard deviation described below.

The mild correlation of this variable with angle, seen in the track events, arises from tracks increasingly passing between bars in a layer at higher angles. As expected the variance of this variable distribution is small for track events and large in shower events with the peak of the distributions following this same pattern.

8.2.6 Charge Mean

The charge mean is the mean charge deposited within an ECal bar. This is expected to show a similar trend to the standard deviation, with large means for high energy deposition from showers and low values from minimally ionising track. It indeed shows much the same correlation and distribution shape, Fig 8.3(f), as the above charge standard deviation. This is a major drawback, as these variables show potential to be highly successful discriminators.

8.2.7 Normalised Charged Standard Deviation

The normalised charge standard deviation combines the previous two variables: it is the charge standard deviation divided by the mean charge. Its distribution can be seen in Fig 8.3(g). This variable dramatically reduces the correlation with energy in shower events, by around 20% compared to the two variables displayed individually. The correlation with angle remains at the same level.

For track class events there is a small increase in energy correlation but on the 5% level whilst the angular correlation decreases from 59% and 33%, for mean charge and charge standard deviation respectively, to just 6%. This decrease in angular correlation

arises again from multiple bars traversed by oblique angle tracks which smears charge distributions.

8.2.8 Max Width

The MaxWidth variable is the number of bars that constitute the maximum horizontal spread in a certain layer. This variable is hoped to be small for thin tracks and large for diffuse showers, as shown in Fig 8.1. Fig 8.3(e) shows that this is not as successful as hoped. Both distributions peak at a width of two bars, but there still remain regions of separation which non-linear discriminating techniques may make use of.

This variable is relatively uncorrelated to both energy and angle for both classes of event. One expects this behaviour as shower width is determined primarily by the detector material and tracks are uniform.

8.2.9 Charge Weighted Track Width

This variable was devised to overcome the shortcomings of the MaxWidth variable. The width of a cluster in each layer is calculated as a number of bars and multiplied by the charge deposited in that layer. Each charge weighted width is then summed before being divided by the total charge in the cluster. For shower events this gives the layer in which a particle deposited the most charge the dominant vote in how spread a cluster is. For track events it is primarily a majority vote as charge deposition between layers is expected to be uniform.

The variable is therefore expected to have low values for tracks and larger values for more diffuse showers; this can be seen in Fig 8.3(c). This variable also shows low correlation with both energy and angle for tracks and showers.

8.2.10 Variable Correlations

The varied correlations between discriminating variables do not pose as great a problem as correlation with particle incident energy or angle but can have a marked effect on the performance of MVA algorithms.

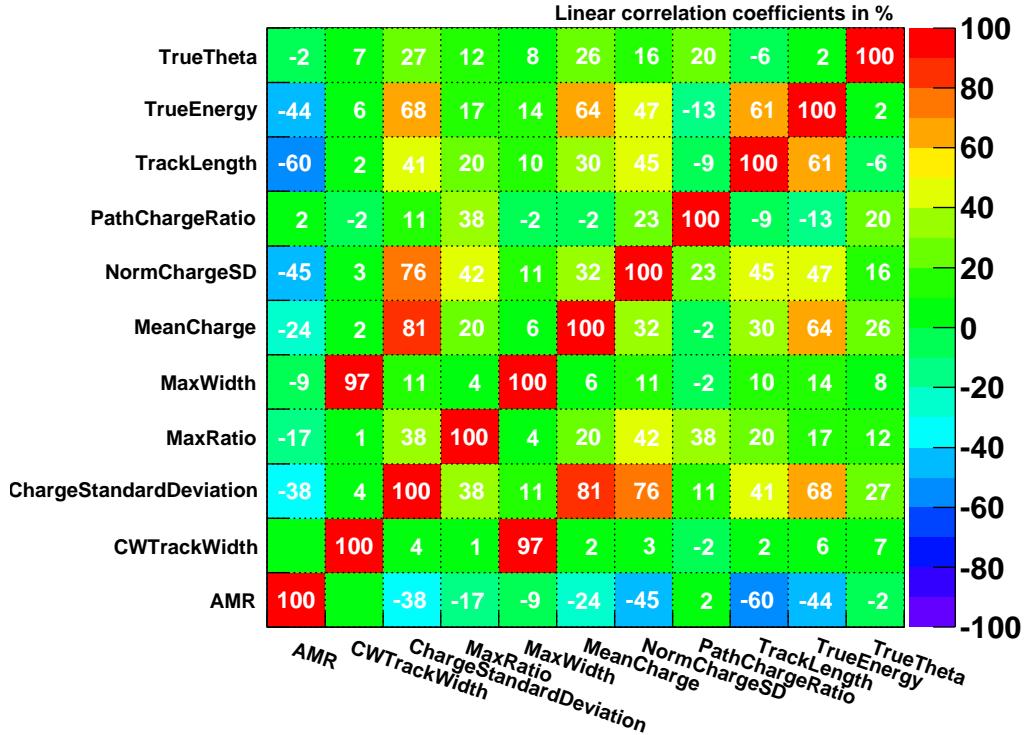
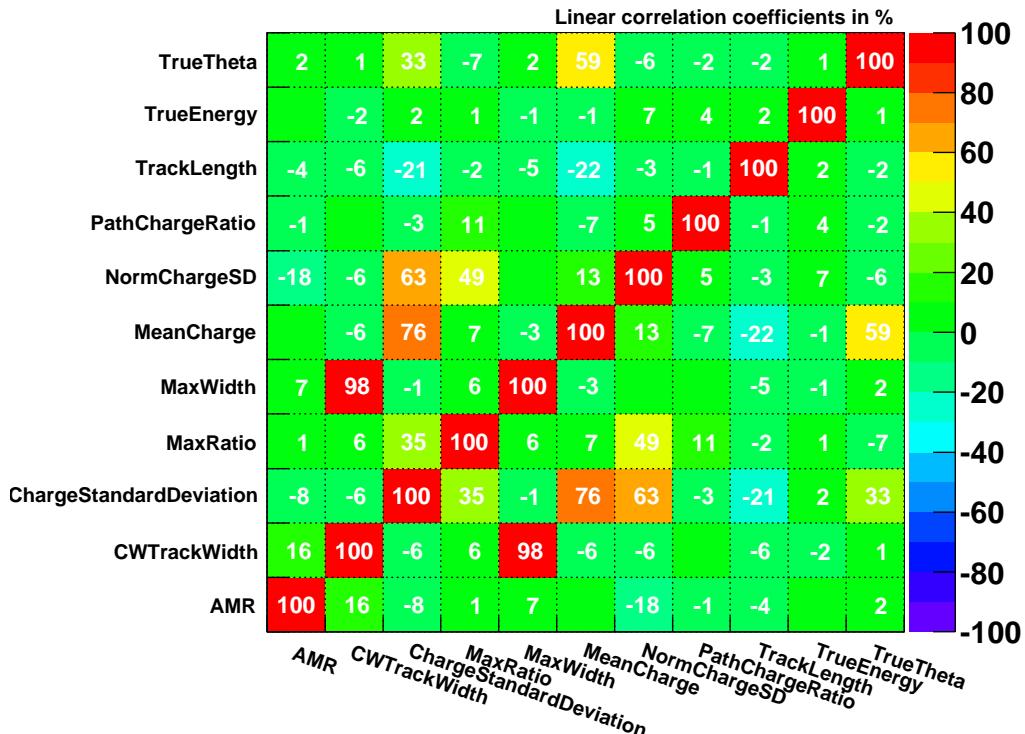
Correlation Matrix Shower**Correlation Matrix Track**

Figure 8.2: Correlation of PID variables with other PID variables and true energy and angle of incidence.

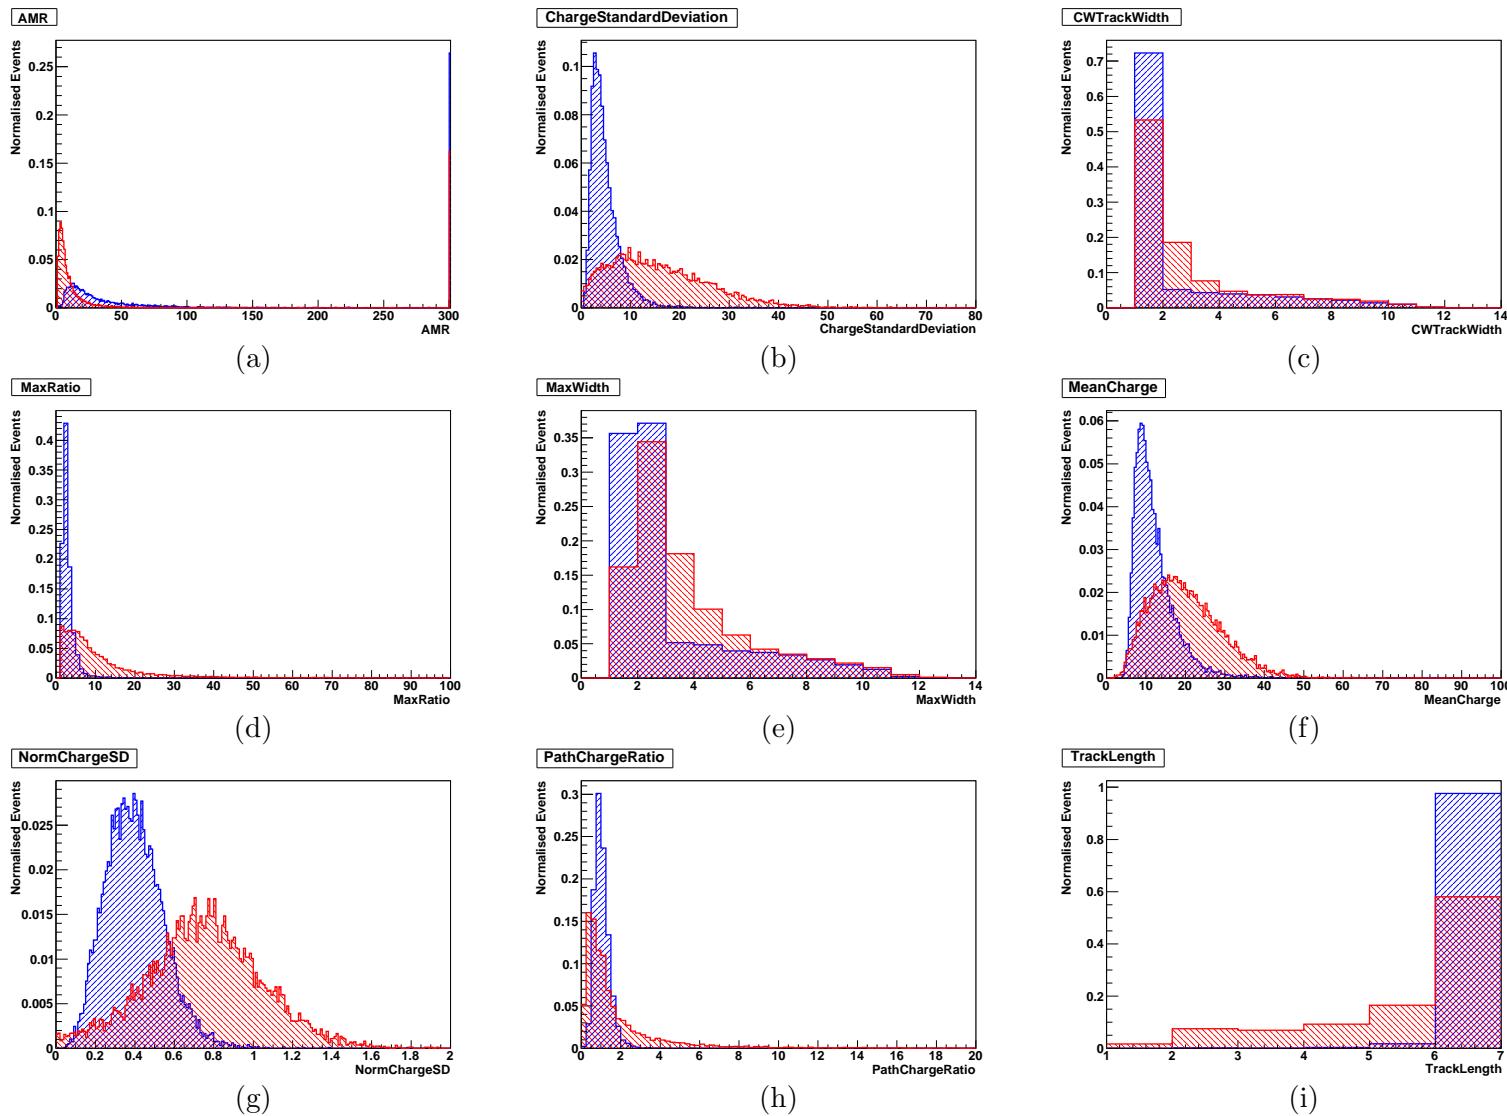


Figure 8.3: Normalised distributions of each discriminating variable considered; presented in blue for track and red for shower class particles.

Variable Pre-Processing

One can pre-process data before training an MVA classifier to try to optimise the discriminating power using decorrelation techniques. Preprocessing relies on information of the variable distribution and as such will be different for each class of event. When implementing a trained MVA method one does not know the class to which the event belongs (if we did there would be no point in all this!).

For some MVA techniques this is not a problem as each class is treated separately. An example is projective likelihood (section 8.5) which can perform transformation \hat{T} of each event class independently. Most MVA techniques, however, must employ a common transformation, constructed from either background pre-processing, signal pre-processing or some mixture of the two.

Linear correlations such as those shown in Fig 8.2 can be most simply accounted for by multiplying the data vector with the square root of the covariance matrix.

Another method of decorrelation is Principal Component Decomposition. This method rotates a data set to attain the maximal variation. The axis over which the data show maximum variance is obtained from principal component analysis. The data are then rotated so that this axis lies along the x-axis of the transformed correlation plot. This then in effect leaves the data with just the spread of points about the linear correlation.

Both of these methods are available as part of the TMVA package which is used to construct train and test the MVA techniques presented in this chapter. It must be noted that these methods only completely remove correlations when they are perfectly linear and the variables gaussian distributed. In real world scenarios this is usually not the case.

8.2.11 Variable Choice

When choosing variables to use in a discrimination algorithm we consider not only the discriminating power but also correlation with energy¹. The discriminating power of a variable is a difficult thing to quantify as different MVA methods utilise areas of the multi-variate phase space in different ways. A quick check, however, can be made on the bulk of each variable by calculating the separation of the distributions of each class. This separation is given by

$$\langle S^2 \rangle = \frac{1}{2} \int \frac{(\hat{y}_S(y) - \hat{y}_B(y))^2}{\hat{y}_S(y) + \hat{y}_B(y)} dy \quad (8.3)$$

¹I only mention energy as angular dependence is smaller and can be tackled with various setups at test beam.

Variable	Separation			Rank		
	Raw	Decor	PCA	Separation	Energy	Overall
MaxRatio	41.4%	29.5%	43.7%	2	4	1
NormChargeSD	46.4%	17.4%	48.7%	1	6	2
CWTrackWidth	16.3%	12.0%	37.1%	7	1	3
PathChargeRatio	23.2%	25.0%	23.2%	6	2	4
AMR	27.6%	26.1%	29.0%	4	5	5
MaxWidth	8.3%	10.6%	26.8%	9	3	6
TrackLength	17.8%	37.2%	1.1%	5	7	7
ChargeStandardDeviation	46.3%	36.0%	27.0%	3	9	8
MeanCharge	27.8%	14.3%	21.7%	8	8	9

Table 8.1: Separation of discriminating variables with and without decorrelation.

where \hat{y}_S and \hat{y}_B are the PDFs of the signal and background respectively. This measure is 1 for PDFs showing no overlap and 0 for PDFs with perfect overlap. The overlaps of the 9 variables investigated, for a representative training sample, are shown in table 8.1 for the raw data as well as the inverse covariance matrix and PCA linear decorrelation methods.

The separation between track and shower class distributions for each variable, unprocessed, covariance matrix decorrelated and PCA decorrelated, is used to rank the variables in order of ‘raw’ discriminating power. The addition of these rankings determined a global separation ranking, with any ties determined by the variable’s separation with no pre-processing. This rank was then added to each variable’s ranking of energy correlation to give an overall rank for each variable; any ties here were decided by the energy correlation ranking.

The five highest ranked variables, which span both geometric – AMR and CWTrackWidth – and charge – MaxRatio, NormChargeSD and PathChargeRatio – were then used to train different multivariate discriminators.

Although a weaker bulk discriminator than some of its other competitors the CWTrackWidth comes out above the MaxWidth and the low correlation with energy sees it better the TrackLength geometric variables. The other variables shine predominantly for their discriminating power but in total they comprise 5 out of the top six ranked variables for low energy correlation.

It was essentially the energy correlation that excluded the TrackLength, ChargeStandardDeviation and Mean Charge whilst it was the poor discriminating power of MaxWidth which led to its demise. This is a good outcome as one would not wish to use the MeanCharge or ChargeStandardDeviation with the NormChargeSD, which is after all

just a combination of the two so one would not expect a marked gain in discrimination power. The MaxWidth variable seems to be almost randomly selected with a large overlap due to tracks hitting 2 bars when passing through a layer and/or knocking electrons free whilst low energy showers fail to expand much. The TrackLength is not the best of discriminators and has one of the highest correlations with energy of the variables investigated, a combination which makes it unsatisfactory.

8.3 Analysing MVA Output

Because many MVA methods are non-linear and rely on stochastic “training” procedures rather than deterministic cuts, they are inherently difficult to compare and evaluate. A consistent, well defined and robust evaluation method is essential both to determine an MVA’s overall performance and to assess its generality and robustness when presented with new data.

8.3.1 Performance and The Receiver Operation Characteristic (ROC) Curve

The most popular independent method for evaluating the performance of a binary discriminator is through construction of its Receiver Operation Characteristic (ROC) curve, used for some years in signal detection theory. It has been demonstrated [126] that ROC curves can be used to evaluate and compare discrimination algorithms.

I will now discuss ROC curves using the classification example presented in this chapter. Selected data sets of track and shower events are passed through the classifier and produce independent regions t and s in the MVA output space. Classification is then obtained by placing a single linear cut c dividing the output space into the two regions track, T , and shower, S . There are four possible outcomes; true or false, shower or track, Fig 8.4.

In discrimination problems one is generally interested in just one class, usually termed signal. In classification, as is the target here, one is interested in both classes; we therefore generate ROC curves for both classes independently.

In the form of the ROC curve used in this thesis, I plot $1 - \text{falseT}_i(S_i)$, the purity, against $\text{trueT}_i(S_i)$, the efficiency, as the position of the cut c_i is scanned across the MVA output space y ; an example of an MVA output is shown in Fig 8.5(a). The area under such a curve gives an indication of the performance of classification, with the most successful classifiers having a large area: this indicates that high efficiencies are

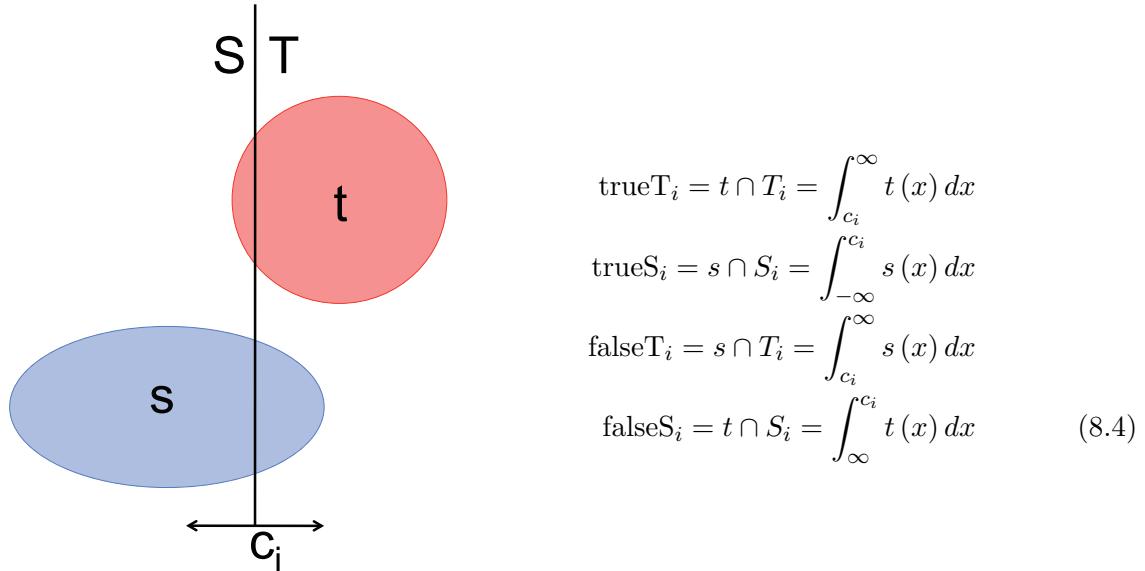


Figure 8.4: Venn diagram of ROC curve construction.

retained even at high purity.

Examples of such curves are drawn in Fig 8.5(b). The black curve represents an ideal classifier in which $t \cap S$ and $s \cap T$ are zero: the purity of the selected sample is always 100%. The blue curve represents a good classifier, which can produce relatively high purity samples with good efficiency. The red curve is a poor classifier: efficiency can only be increased by accepting significant amounts of background. The dashed line shows the worst possible performance where the two distributions are entirely overlapping i.e. have zero separation.

These curves provide a method of independently comparing performance of various MVA techniques and also different training schemes within a particular MVA choice; more specifically we use the integral of the curve, above this middle worst performance line.

As a metric of performance we also consider two other measures determined from a fixed central cut; the first is the successful classification of track and shower events:

$$\text{Success} = \frac{t \cap T + s \cap S}{t + s} \quad (8.5)$$

where $t \cap T$ are the track events which lie within the track cut region T and $s \cap S$ are those shower events in the shower cut region S . The cut applied here is at zero which is expected to be close to optimal.

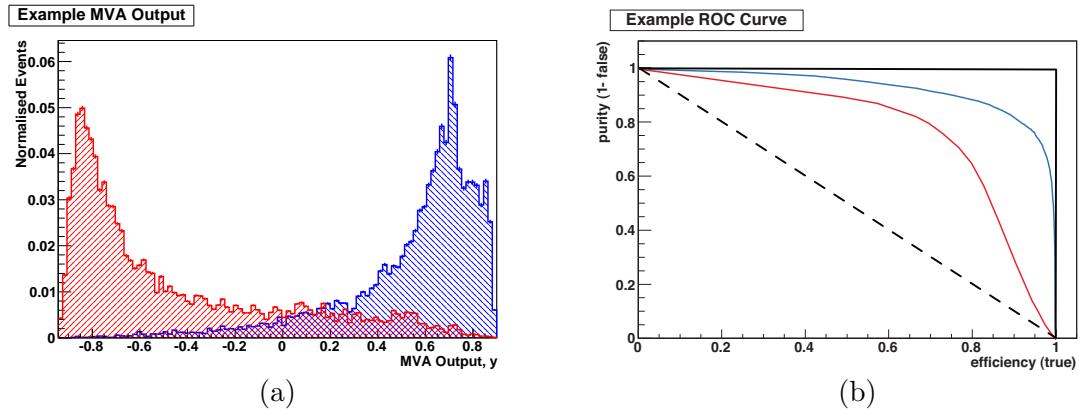


Figure 8.5: An example output of a MVA discriminator (a) the red and blue represent the training data of the two different classes. An example of a Receiver Operating Characteristic curve is shown in (b), the black line shows the perfect maximum, blue a good discriminator, red a poor discriminator and the dashed lines shows the minimum.

The second measure was the statistical significance

$$\text{Significance} = \frac{t \cap T}{\sqrt{s \cap T + t \cap T}} + \frac{s \cap S}{\sqrt{s \cap S + t \cap S}} \quad (8.6)$$

which is a measure of the purity as well as the performance of a discriminator.

8.3.2 Robustness: Overtraining Checks

Overtraining occurs when an MVA method loses its generality of evaluation. This occurs when too many model parameters are adjusted to too few data points. The result is an MVA method that can give the best possible discrimination for the sample used to train it but performs poorly on any new data presented to it.

One can therefore search for overtraining by comparing the response of a fully trained MVA to the training sample and an independent test sample of data. This is performed separately for track and shower events for the MVAs discussed here, using large statistics test samples of $> 100,000$ events.

Since ROC curves are a standard method for evaluating MVA performance, it is natural to consider the difference between training and test sample ROC curves as a potential indicator of overtraining. After considerable experimentation, the integrated difference between test and training ROC curves, described below, was selected as the most satisfactory metric.

The integral is numerically performed on a bin by bin basis. For each bin the difference

between training and test curves is multiplied by the bin width. The values are then summed to obtain an integral. As it is a measure of difference, it is expected that this will be smallest for optimally trained and larger for under/overtrained MVA techniques.

8.4 Fisher Linear Discriminant

Fisher linear discriminants are the simplest of the MVA techniques covered here and posses the fewest degrees of freedom. Despite this fact Fisher discriminants can in certain cases out-perform linear projective likelihood methods and, in a case where we have variables with linear correlations, equivalent RMS but different means, then they can out-perform even non-linear MVA methods.

8.4.1 Description

The input data are transformed into a single statistic, the Fisher linear discriminant, which maximises variable mean separation between classes, described by the between-class covariance matrix B , whilst minimising the scatter of the variables relative to their respective means within each class, described by the within-class covariance matrix W . This is achieved by maximising

$$J(w) = \frac{f^T B f}{f^T W f} \quad (8.7)$$

where the within and between-class covariance matrices W and B are transformed by the vector of Fisher coefficients f .

$$W_{kl} = \sum_{U=S,B} \langle x_{U,k} - \bar{x}_{U,k} \rangle \langle x_{U,l} - \bar{x}_{U,l} \rangle = C_{S,kl} + C_{B,kl} \quad (8.8)$$

$$B_{kl} = \sum_{U=S,B} (x_{U,k} - \bar{x}_{U,k}) (x_{U,l} - \bar{x}_{U,l}) \quad (8.9)$$

This is found to be maximal when the Fisher coefficients are equivalent to

$$f_k = \frac{\sqrt{N_T N_S}}{N_T + N_S} \sum_{l=1}^{n_{var}} W_{kl}^{-1} (\bar{x}_{T,l} - \bar{x}_{S,l}); \quad (8.10)$$

the first term is a normalisation factor over the N_T track events and the N_S shower events, and the $\bar{x}_{C,l}$ are the sample means of the track T and S shower class C

statistics x_l respectively.

The Fisher discriminant is then given each i^{th} event by a linear combination of these coefficients and their respective statistic

$$y_F(i) = f_0 + \sum_{k=1}^{n_{\text{var}}} f_k x_k(i); \quad (8.11)$$

here the offset f_0 centres the sample mean of all $N_T + N_S$ events.

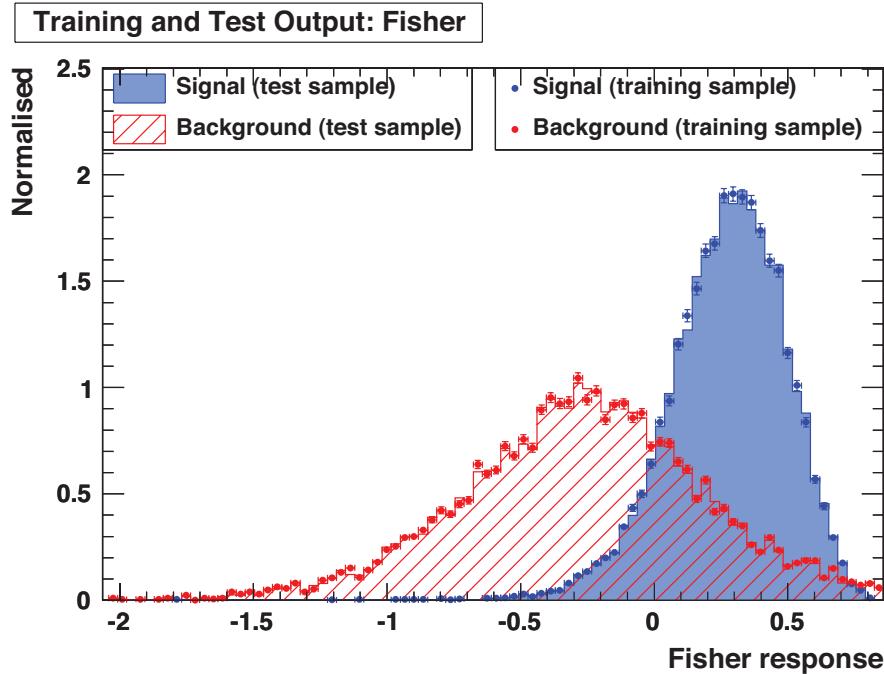


Figure 8.6: Example output of a Fisher discriminant using the five variables chosen above.

8.4.2 Training

From the above one can see that the success of Fisher discriminant training relies upon the degree to which the input data represents the mean and covariance of data of each class. Therefore with a representative sample it is impossible to wrongly train a Fisher discriminant.

8.5 Maximum Likelihood

The method of maximum likelihood has become an integral part of almost all high energy physics analyses. The discriminator is built out of probability density functions (PDF) which are designed to accurately reproduce the input variables for each class of event. These PDFs can be represented by original histograms or may have their functional form estimated.

8.5.1 Description

For any given event there will be a vector of discriminating variables x_i ; these are compared to the probability densities of each variable. These probabilities are then multiplied before being normalised through division by the sum of the signal and background likelihoods.

Mathematically the discriminator output $y_{\mathcal{L}}$ is written as

$$y_{\mathcal{L}}(i) = \frac{\mathcal{L}_S(i)}{\mathcal{L}_S(i) + \mathcal{L}_B(i)}, \quad (8.12)$$

where

$$\mathcal{L}_{S(B)}(i) = \prod_{k=1}^{n_{var}} p_{S(B),k}(x_k(i)) \quad (8.13)$$

is the likelihood of the signal and background for a given event i , the product of the PDF values of each variable x_k .

As each variable is treated individually, any correlations among the variables are ignored. As mentioned above, linear correlation may be accounted for with preprocessing of event data. As the likelihood estimator treats each class of event separately then linear transformations \hat{T} can be performed on signal and background independently:

$$y_{\mathcal{L}}(i_{event}) = \frac{\prod_{k=1}^{n_{var}} p_k^S(\hat{T}^S x_k(i_{event}))}{\prod_{k=1}^{n_{var}} p_k^S(\hat{T}^S x_k(i_{event})) + \prod_{k=1}^{n_{var}} p_k^B(\hat{T}^B x_k(i_{event}))}. \quad (8.14)$$

In the majority of cases likelihood discriminators benefit from an improvement of classification after decorrelation.

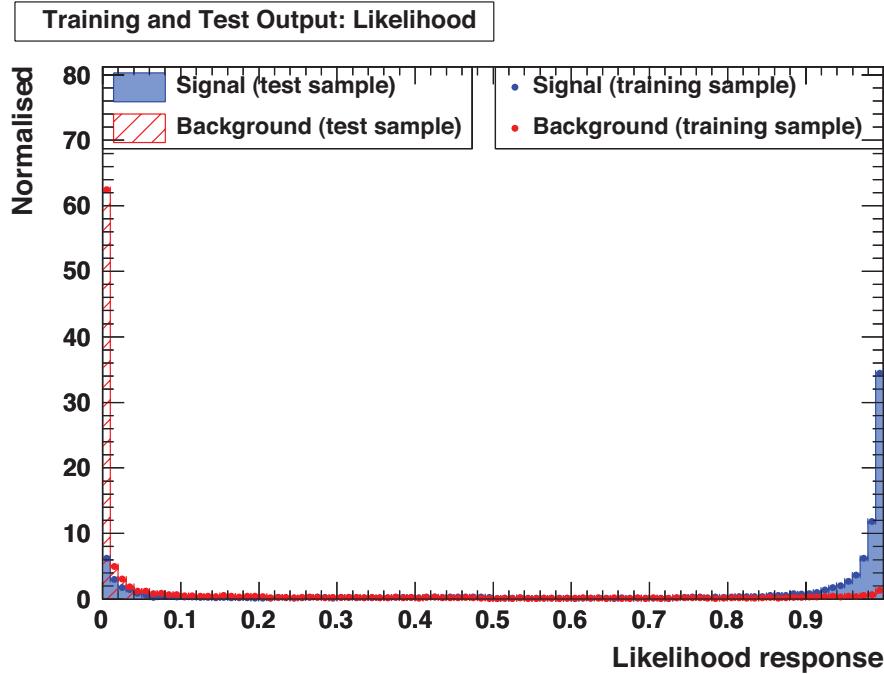


Figure 8.7: Example output of a likelihood discriminant using the five variables chosen above.

Probability Density Function (PDF) Parametrisation

Probability density estimation is a field in its own right and there have been many methods developed. Here we will discuss three which are implemented inside TMVA/in the most common use. Most simply, the form of a probability density function (PDF) can be estimated by a normalisation of the histogrammed data. More sophisticated non-parametric techniques parameterise input data, examples including spline and kernel density estimation.

TMVA employs spline estimation by first binning data in bins containing equal average numbers of events. The histogram is then smoothed a user defined number of times. Care must be taken not to smooth too much; highly peaked distributions can suffer less smoothing than broad distributions before being distorted. Each histogram is then parameterised by fitting a chosen degree of polynomial between each histogram bin centre. The functions of these splines are then saved in finely binned histograms for easy access on implementation.

The second type of PDF parameterisation is in using a Gaussian Kernel Density Estimator (KDE). The Gaussian KDE estimates the shape of a PDF $p(x)$ of variable x

from a sum over smeared raw data training events x_i

$$p(x) = \frac{1}{Nh} \sum_{i=1}^N K\left(\frac{\bar{x} - x_i}{h}\right), \quad (8.15)$$

where the N is the number of events, the function K is gaussian and h is the bandwidth of the kernel. In the following results the bandwidth is applied as a fixed constant for each variable over the entire training set. The optimal bandwidth for independent gaussian variables has been shown to be [127]

$$h = \left(\frac{4}{3}\right)^{1/5} \sigma_x N^{1/5}, \quad (8.16)$$

where σ_x is the RMS of variable x . This bandwidth and KDE method were employed in the following likelihood discrimination.

KDE methods can suffer from edge effects, in that data is smeared into regions where no data existed before, i.e. beyond the boundary of the variable limits. This results in probability leakage from the PDF, leaving the total probability below unity. This is particularly problematic for those variables that peak close to the limits of the variable range, which is true for the majority of variables used here. Two methods of overcoming this are renormalisation of the PDF or mirroring of a PDF at its boundaries. Renormalisation can leave the extremes of the distribution with smaller probabilities than expected, which will alter the applicability of that variable to the task in hand. Mirroring on the other hand compensates for this effect almost perfectly. The PDF is mirrored at its boundaries and the kernel applied to the original and mirrored PDFs. The PDF is then reselected, by region, after the process and any loss outside of this boundary due to smearing would effectively be reflected back into the original variable domain by KDE smoothing of the mirrored PDF.

One sees this when training with no correction and with mirrored correction. The significance of variables within the classifier changes and so too does the performance. In the following investigation I use KDE parameterisation with mirroring as this produces the best and most stable result and is important for the variables used, a number of which have peaks at or near such boundaries, Fig 8.3.

8.6 Boosted Decision Trees

8.6.1 Theory

A decision tree is a binary tree structured classifier like the one sketched in Fig 8.8. Repeated decisions (cuts) are performed, on a single event variable at any one time, and passed to further nodes until a termination criterion is reached. This dissects the multivariate phase space into regions of two classes, dependent on the composition of training events in the terminated nodes. The multiple paths down such a tree provide a very individual sequence of cuts.

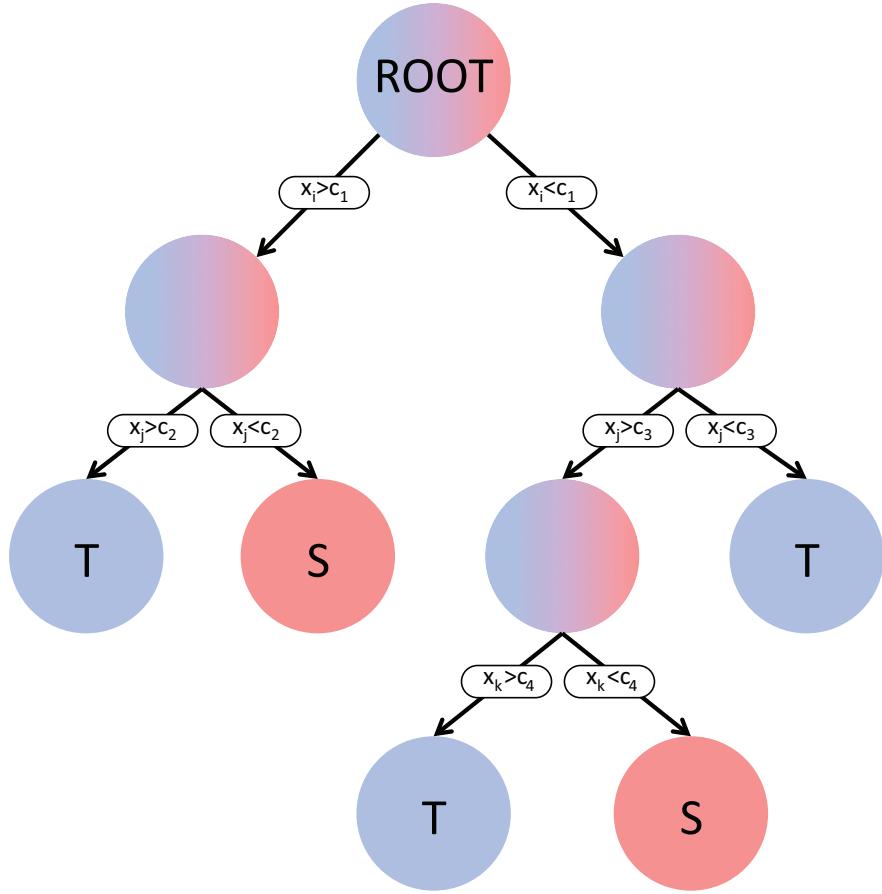


Figure 8.8: Schematic diagram of a decision tree. At each node the optimal cut c is employed on the event variables x .

The decision employed at each tree node in the analyses presented here is the Gini index of diversity $p(1 - p)$, where p is the probability of either class. This yields a maximum where the classes are fully mixed and a minimum when just one class of

event is present. At each node all variables and possible cut values are tested and the variable and cut which minimise this criterion is the one selected. Cut values are equi-spaced values over the range of each variable and the number of such cuts is a user defined factor that must optimise performance and computation. There was no discernable change in performance over the range of 10–100 cuts employed although there is a marked increase in processing time.

Decision trees are susceptible to training upon statistical fluctuations in data. If a fluctuation makes a certain variable appear a stronger discriminator than it really is, a node may choose to cut on this variable in preference to the “real” best discriminator, and this will in general modify the whole structure of a tree beyond this node, resulting in an inefficient classifier.

This problem is addressed by instead creating a forest of decision trees, all trained on the same data but with events re-weighted by what is called boosting. Classification is then obtained from a majority vote of all trees within the forest.

Boosting is a general re-weighting scheme that can be employed in other multivariate analyses, not just decision trees. The most popular boosting method, and the one employed in this analysis, is adaptive boosting or AdaBoost. In this scheme, events that are misclassified in one tree have their weight boosted in the next tree that is trained in the forest. These events are multiplied by a constant boost weight α derived from the misclassification rate err of the previous tree,

$$\alpha = \frac{1 - \text{err}}{\text{err}}. \quad (8.17)$$

This boost weight is also used to modify the weight of each trees vote, $h(x)$, when quantifying the final classifier output, y_{BDT} , for each event, x :

$$y_{BDT}(x) = \sum_i^N \ln(\alpha_i) \dot{h}_i(x), \quad (8.18)$$

where the sum is over the N trees in the forest.

In principle each tree may continue until each termination node contains just a single class of event. This gives huge discriminating power, but only for the test sample which is never a perfect representation of all data. This condition is called overtraining and is the one major flaw of boosted decision trees. This is addressed by pruning a tree once trained to remove any training on unique events.

This essentially amounts to removing statistically insignificant nodes, and there are again a number of methods of achieving this. Cost complexity pruning [128] is used

in this investigation. The cost of each node $R(\text{node})$ is chosen as the misclassification rate $1 - \max(p, 1 - p)$ and calculated for each node in the tree. The tree is broken down into a number of sub-trees of which each node, apart from those terminating, is a root. The cost-complexity of each sub-tree is then defined by

$$\rho = \frac{R(\text{node}) - R(\text{subtree})}{N_{\text{nodes}} - 1}, \quad (8.19)$$

where N_{nodes} is the number of nodes in the subtree and $R(\text{subtree})$ is the cost of the subtree. Each node with the smallest cost-complexity, ρ , is recursively pruned until a user-defined pruning limit, known as the prune strength, is reached.

We use the integrated difference method described above to determine overtraining as a function of prune strength, Fig 8.9. At each prune strength value 10–100 we train 30 BDTs each with 600 trees and 16000 training events. The data points plot the mean and the error bars chart the RMS of the integrated difference of all 30 trees at each prune strength.

As one would expect the difference between test and training ROC curves increases at low prune strengths where overtraining is expected to be most probable. This rapidly decreases until a minimum around a prune strength of ~ 3 ; the difference then rises as the trees become undertrained. The error bars plot the RMS of the values.

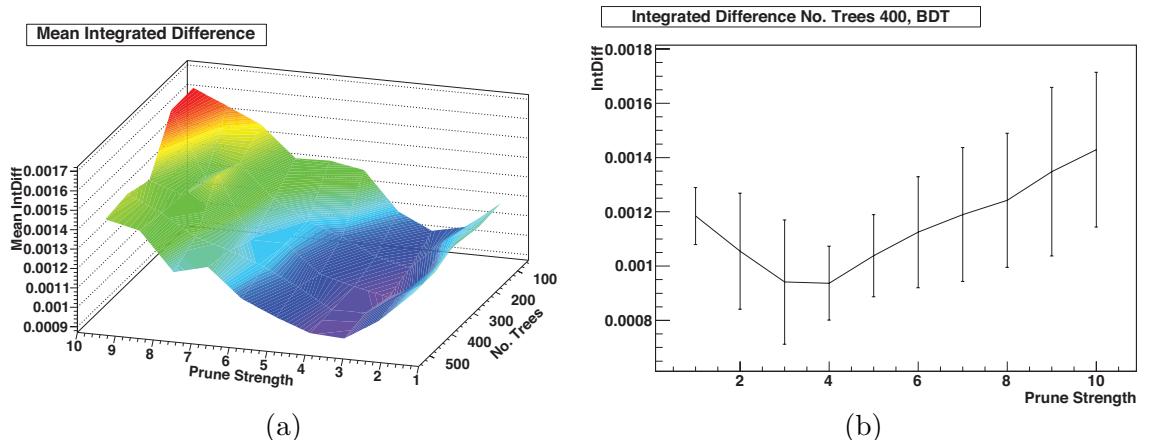


Figure 8.9: Overtraining check of BDT at various numbers of trees and prune strengths.

I also note here that it was seen that if one increases the number of training events then one must reduce the prune strength accordingly. We have not performed this study in great detail, only examining the success of training at various prune strengths with 16,000 and 50,000 training events because of processing time (Fig 8.9(a) alone requiring over 2,000 hours of CPU time). Between these two we notice that one requires a

significantly smaller prune strength of ~ 3 when training on 50,000 events compared to ~ 50 for just 16,000 events. This is attributed to the reduction of statistical fluctuations in the data as one increases the number of events: the PDFs which are cut upon become more representative of the ‘real’ distributions.

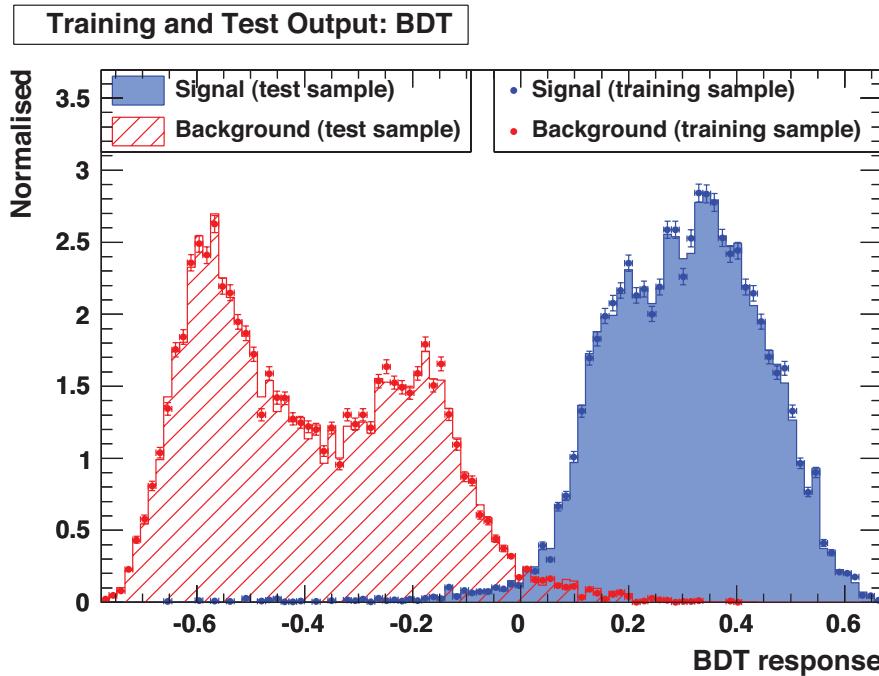


Figure 8.10: Example output of a BDT discriminant using the five variables chosen above.

8.7 Multilayer Perceptron

Artificial neural networks (ANNs) were originally developed to model cognitive processes by simulating the behaviour of neurons in the human brain using a network of perceptrons. Perceptrons are idealised neurons with multiple input and output synapses. Each input synapse of a perceptron reacts with a different weighting, dependent upon the input stimulus strength and type. The perceptron then activates, through one of a number of possible activation functions, which determine the magnitude of output.

A generalised ANN with n perceptrons can have $n(n - 1)$ directional connections. This complexity can be greatly simplified by restricting connections to a single direction between layers of perceptrons. Such a network is called a multilayer perceptron (MLP).

8.7.1 Theory

Perceptron

The building block of an MLP is the perceptron which can take multiple inputs and produce various responses for a number of output synapses. The mapping of perceptron input to output is defined by the response function ρ which can be viewed as a convolution of an $\mathcal{R}^n \rightarrow \mathcal{R}$ synapse function and an $\mathcal{R} \rightarrow \mathcal{R}$ activation function α . The synapse function determines how a perceptron combines the inputs received whilst the activation function decides how the perceptron is to respond to the input stimuli. Examples of synapse functions are summations of each weighted input, the weighted inputs squared or of the absolute value of the weighted input. The synapse function, κ , employed in the following investigation is a simple summation of weighted, w_{ij} , inputs, y_i :

$$\kappa(y_1^{(l)}, \dots, y_n^{(l)} | w_{0j}^{(l)}, \dots, w_{nj}^{(l)}) = w_{0j}^{(l)} + \sum_{i=1}^n y_i^{(l)} w_{ij}^{(l)}. \quad (8.20)$$

The activation function decides how the perceptron is to respond to the input stimuli and can take almost any functional form. For the purpose of training it is best to select a functional form which is continuous and differentiable in the weights. With classification problems one is ideally hoping for a definite one way or another decision; yes or no, or in this case track or shower. One can select a form of the synapse function that matches this desire. The tanh function,

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}, \quad (8.21)$$

is an example of such a function and is the one used in the following investigation; x is replaced by the result of the synapse function, κ .

Choosing the Architecture of a Network

The first layer of a MLP is called the input layer. This layer contains a perceptron for each of the n input variables; x_1, x_2, \dots, x_n . For classification problems, such as those discussed herein, there is just a single perceptron in the final layer which outputs as single number, the neural network estimator y_{MLP} . Any layers between the input and final layer are called hidden layers.

The number of hidden layers and the number of perceptrons they contain are free parameters and should be selected carefully for the job in hand. It has been noted [129] that an arbitrarily large accuracy of approximation of any continuous correlation func-

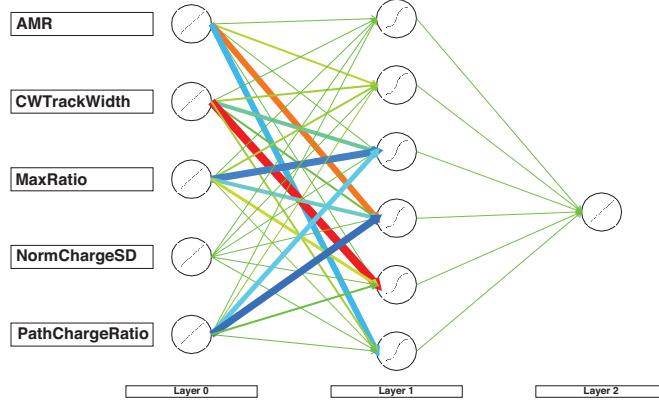


Figure 8.11: Example of an MLP network with the five selected variables as input nodes, a single hidden layer, layer 1, and a single discriminating output node, layer 2. The thickness of the lines represent the weight of each synapse connection.

tion can be achieved with a single hidden layer, if the layer contains an arbitrarily large number of perceptrons. One can achieve comparable performance from a single hidden layer MLP and a second MLP with more hidden layers but relatively fewer perceptrons. As the non-linearity increases so to does the complexity of error estimation as more hidden layers are added to the network.

It can therefore be deduced that for a physics application, where good error estimation is a key requirement, one would usually choose to use just a single layer MLP. A single layer perceptron requires more training events and consequent computing time than its multiple hidden layer counterpart, but for particle identification purposes, as it is used here, this is not an issue. One uses large statistics from Monte Carlo, cosmic muon or known particle test beam data from which to select a training set.

The output y_{MLP} of a single hidden layer MLP with n input variables and m perceptrons in the hidden layer can be written as

$$y_{ANN} = \sum_{j=0}^n y_j^{(2)} w_{j1}^{(2)} = \sum_{j=0}^n \tanh \left(\sum_{i=0}^m y_i^{(1)} w_{ij}^{(1)} \right) w_{j1}^{(2)}, \quad (8.22)$$

where each node has a \tanh activation function and straight addition synapse function.

Increasing the number of perceptrons in the single hidden layer one would expect to see an improvement in performance as the number of free parameters increases. However, no marked improvement in performance is seen, Fig 8.12(a), beyond 6 neurons in the hidden layer, in our five variable neural network. As the possibility of overtraining is understood to scale with the number of degrees of freedom for a given data set, this

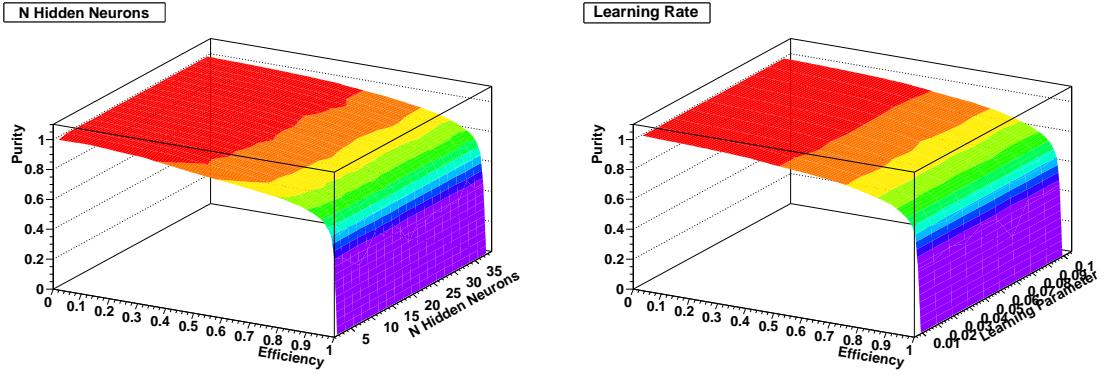


Figure 8.12: ROC curve performance of MLP as function of the number of neurons in the single hidden layer (left) and the learning parameter η (right).

value of 6 was chosen for the networks employed hereafter.

Training

The training of a multilayer perceptron amounts to tweaking the weight of each synapse connection between perceptrons. The most popular method of training, and the one employed here, is supervised back propagation. The neural network estimator y_{MLP} output by the network is compared to the desired output \hat{y} for event a in some error function. The sum of squares error

$$E_a(x_a|w) = \frac{1}{2} (y_{ANN,a} - \hat{y}_a)^2 \quad (8.23)$$

is simple and works with feed forward networks, and is the one employed herein.

The method of steepest descent can be employed to find the synapse weights corresponding to a minimum in this error function. For this to be possible the selected activation function must be differentiable with respect to the weights, a criterion mentioned above.

Starting with an initial guess w^0 the weights are updated, at each step of the minimisation s , by stepping a small distance in weight space in the direction where the error function decreases most rapidly,

$$w_{ij}^{s+1} = w_{ij}^s - \eta \Delta_{w_{ij}} E, \quad (8.24)$$

where η , the learning parameter, determines the step size of this minimisation and must be selected carefully; too large and the minimisation may oscillate, too small and con-

vergence will be unnecessarily slow. Compared to boosted decision trees, MLPs require about a quarter of the CPU time (around 20 mins on a machine with a SPECint [130] of ~ 1500) with a learning parameter of 0.02. Nonetheless a quick check of performance over a range of learning parameters was carried out. One can see in Fig 8.12(b) that over an order of magnitude in the learning parameter, from 0.01 to 0.1, there is no effect seen on the average MLP performance. There was little time gain observed and so a relatively low 0.02 was chosen to err on the side of caution.

The weights of synapses connected to the output layer are updated first as the values of these are required to then update the preceding layer and so on until the synapses relating to the input layer are updated. This updating of weights from output to input gives the method its name of back propagation.

In the method described here the weights are updated on an event by event basis, otherwise known as online learning. The contrasting method, known as bulk learning, updates network weights after each training cycle using the sum of errors from all events. The benefit of online learning is the ability to train a network using a limited number of events, over fewer epochs, which results in a faster training. The events are entirely randomised.

The optimal training of such a network is dependent not only upon the number of epochs one chooses to train over but also the number of events available for training. To investigate overtraining we therefore chart the mean integrated difference as a function of epochs (training cycles) and number of training events. We train networks with 6 neurons in a single hidden layer, as explained above, with learning rate of 0.02. A grid of epochs, 50–600, and number of training events, 10000–100000, was then selected and 30 networks trained at each grid point giving a total of around 4,000 networks. The mean of the integrated difference was then calculated at each grid point and plotted as a surface, Fig 8.13(a). There is clearly a strong dependence on the number of training events, which is expected from online learning. There is also a more subtle dependence on the number of training epochs. In Fig 8.13(b) we show a slice of the surface plot for a fixed 50000 training events. Here one can see that at the low number of epochs there is a larger difference due to slight undertraining. The square difference reaches a minimum around 300 epochs before gently rising once again due to overtraining.

This gentle rise is a far cry from the step descent seen in the overtraining in BDT. This indicates that MLPs are more robust to training, less susceptible to overtraining. In the present case, repeat trainings of discriminators will undoubtedly be necessary as the current simulated data are replaced by testbeam data, T2K data, new calibrations, etc. This robustness is therefore an important advantage of MLPs.

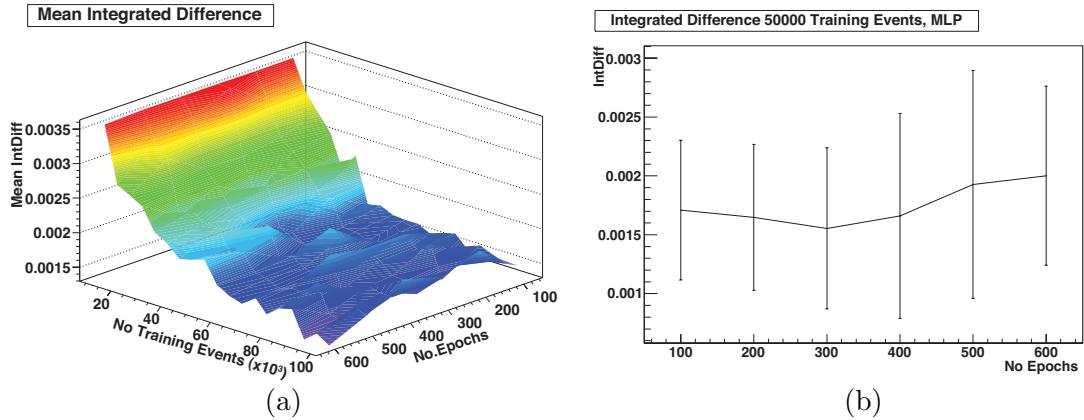


Figure 8.13: Overtraining check of MLP at various numbers of training epochs and events.

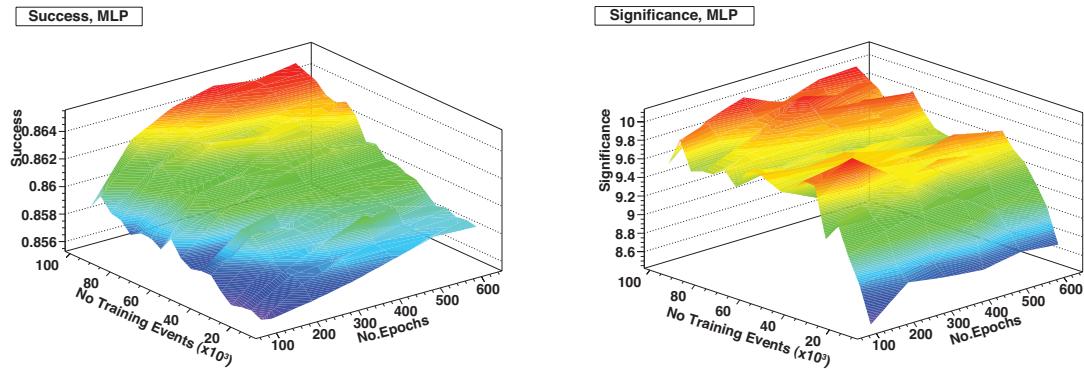


Figure 8.14: Performance of MLP as function of the number of epochs and training events.

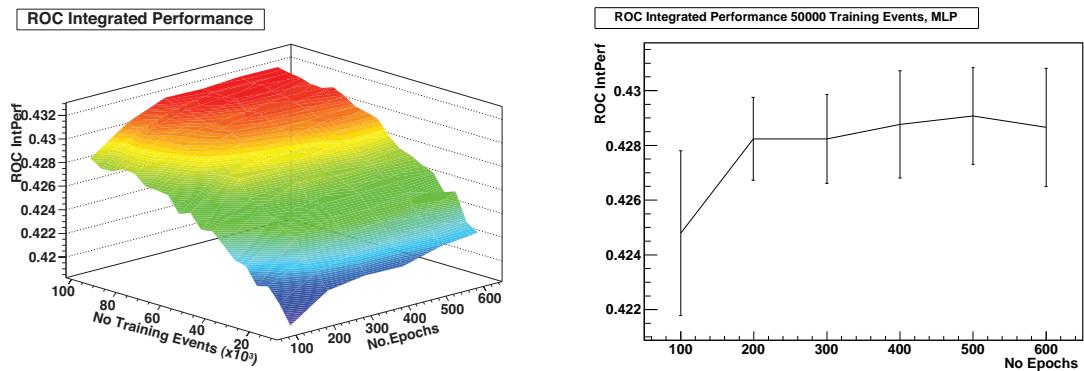


Figure 8.15: Performance of MLP as function of the number of epochs and training events.

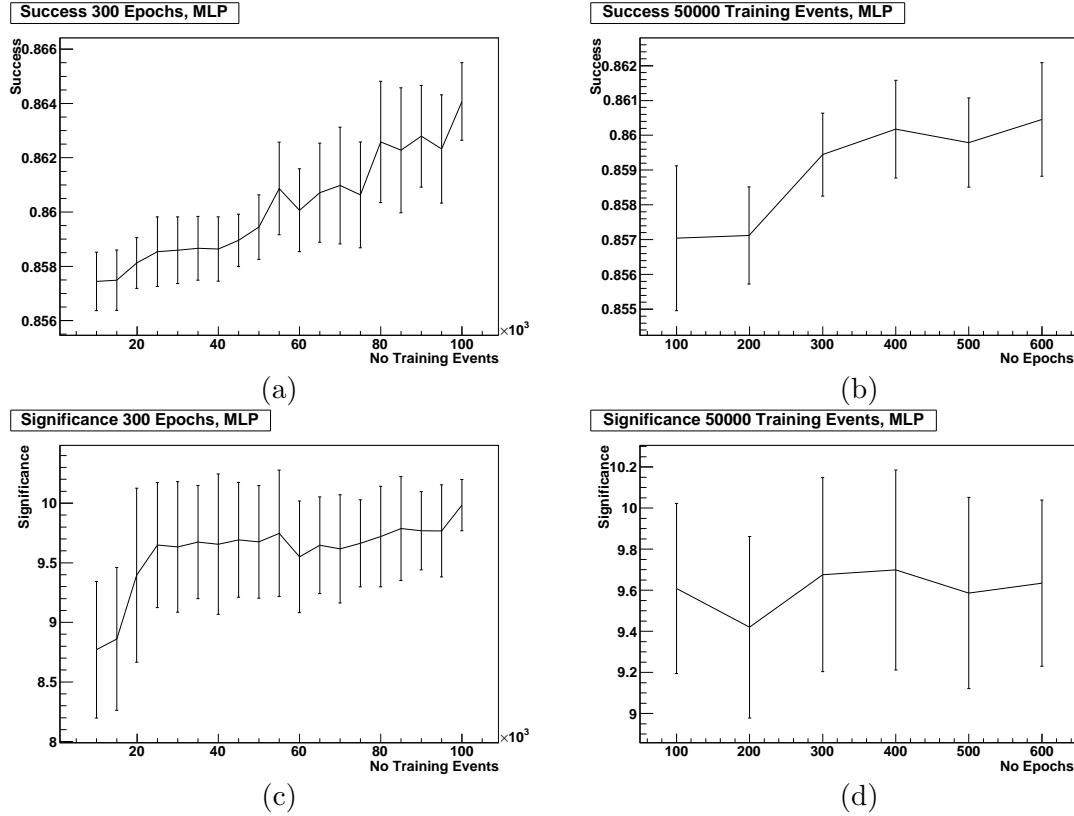


Figure 8.16: Performance of MLP as function of the number of epochs and training events.

In Fig 8.14 the performance of the neural networks are charted as a function of event number and training epochs. As one would expect the success, Fig 8.14(a), of track and shower identification rises with increase in both training variables; more information leads to a greater classification power. When considering the significance of the discrimination, Fig 8.14(b), we see that there is a sharp increase at low numbers of training events before a plateau after 30,000 events. The significance remains relatively insensitive to the number of epochs used in training the network. This demonstrates that although the fraction of events successfully classified continues to rise, so too does the inclusion of background.

The performance increase can also be seen in the integrated value of the ROC curve and in Fig 8.15 we show that this follows much the same pattern as the success.

Choosing the values of 50,000 events and 300 epochs as above leaves us in a region firmly in the plateau of significance and toward the upper end of success. From Fig 8.16(b) one can see that a choice of 300 epochs at 50,000 training events is a shrewd one as there is little gain in success by increasing either value further. Although there is an

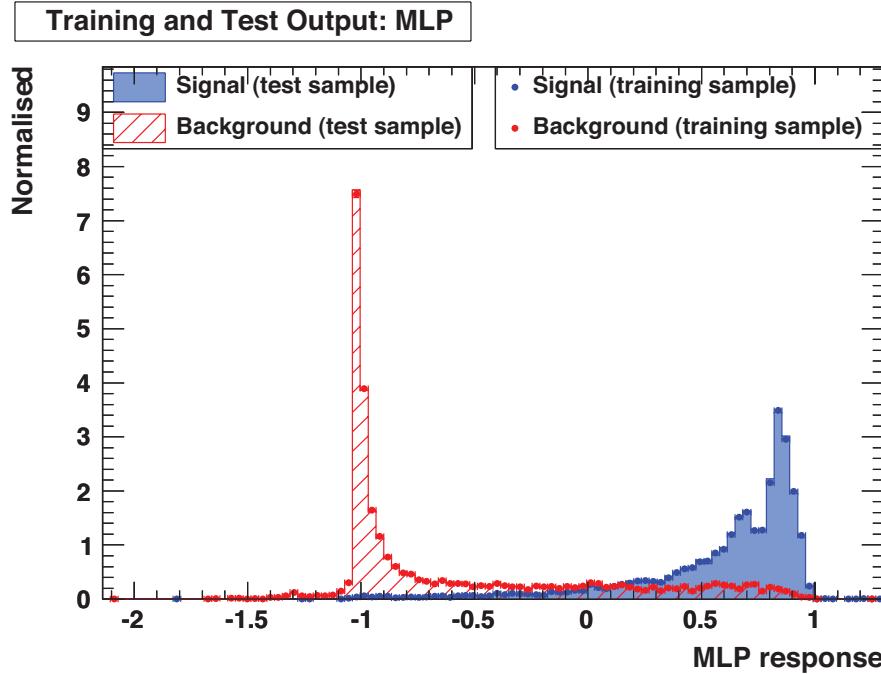


Figure 8.17: Example output of a MLP discriminant using the five variables chosen above.

increase in success with ever more training events, Fig 8.16(a), the increase is relatively slight with a change of just $\sim 0.7\%$ over the entire range 10,000–100,000.

The indifference of the significance to the number of training epochs can be seen clearly in Fig 8.16(d) with a little changing mean value. The plateauing seen in Fig 8.14 can also be seen clearly in the cut through at 300 epochs in Fig 8.16(c).

8.7.2 Conclusion

MLPs have shown themselves to be complex to understand and a black art to select optimal training for. Despite this they are far less sensitive to training anomalies than BDTs, the other non-linear method considered here. For the numerous trainings that one may wish to make for various analyses during the lifetime of the ND280 detector it is felt that the MLP is a more sensible choice.

8.8 Comparison and Conclusion

From Fig 8.18 it can be seen that the linear discrimination functions were not sufficient for the task in hand. BDTs show great promise but further understanding is required before one can confidently employ them in analysis techniques.

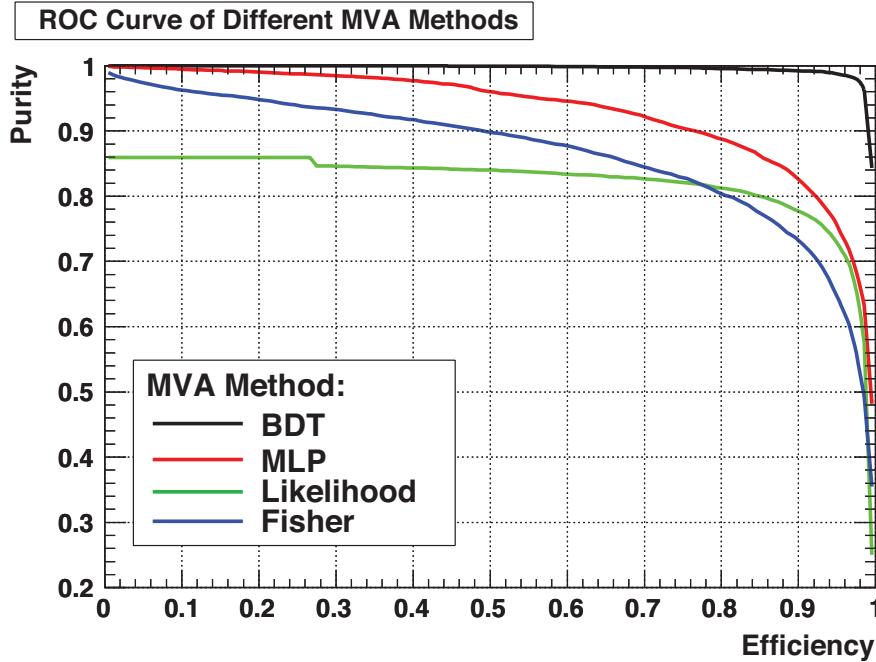


Figure 8.18: Comparison ROC curves for different MVA techniques.

For these reasons we conclude that an MLP, with network architecture of just a single hidden layer, with six nodes, is the correct choice of discriminator at this stage. This network is trained over 300 epochs with 50000 training events per track and shower sample. This network is chosen in order to perform competently and consistently.

Chapter 9

P0D + P0DECal NC1 π^0 Event Selection

The P0D subdetector of the T2K ND280 is a dedicated π^0 detector which aims to measure inclusive and exclusive neutral pion production channels. Its large mass and water target in and out running modes allow the detector to collect a large set of events from which to achieve these goals.

The P0D π^0 reconstruction however suffers statistically and systematically from the entire and partial loss of showers initiated by π^0 decay photons. The P0D incorporates upstream and downstream calorimeter regions to attempt to address this problem for photons travelling at small angles with respect to the beam axis. However, the isotropic nature and relatively low energy of π^0 decays as seen in the ND280 leaves a number of photons with much higher angles. The P0DECal is primarily a veto detector intended to tag particles missed by the P0D proper.

The P0DECal can in principle improve the P0D π^0 identification in a number of ways:

- reducing the misidentification of CC1 π^0 events as NC1 π^0 , by detecting muons missed by the P0D;
- improving efficiency by detecting photons that did not shower in the P0D;
- “tail catching” particles which leave only short tracks/showers within the P0D.

The final item requires P0D-P0DECal cluster matching which is not yet implemented and therefore is not considered here. This chapter presents a feasibility study which investigates the effect the P0DECal will have on the efficiency and purity of NC1 π^0

events selection, by employing a series of single cuts on data from separately reconstructed P0D and P0DECal objects. The creation of the final reconstructed P0DECal objects used in this investigation have been the topic of this thesis until this point.

9.1 Monte Carlo Generation

The generation of primary neutrino interactions was performed using GENIE and NUANCE. Single vertices were selected to be located within the general volume of the P0D with all interaction types produced according to their respective cross-sections as defined by the default settings of both generators. The events were then passed through the MC simulation chain of nd280mc and elecSim; all MC steps were automated using nd280Control.

Events were then reconstructed by p0dRecon (v3r9) and the latest P0DECalRecon algorithms. The final data were distilled into ROOT trees using oaAnalysis, with updated P0DECal module, and the analysis presented here uses this final output.

The MC was generated half of the time in P0D water in mode and the other half with water out; this altered the active portion of the detector and total mass, which one would expect to affect both event selection efficiency and purity. It is important to consider both modes as the cross-section of NC1 π^0 upon water will be determined through comparison of water in and out data.

Vertices are classified firstly using truth information: oaAnalysis saves true vertices with associated final state particles alongside reconstructed data. The first selection criterion applied to the truth data is that there be just a single vertex in the P0D. Next the vertex is tested to determine if it is indeed NC1 π^0 by cutting on the PDG codes of primary particles exiting the vertex.

This poses a problem when using GENIE in standard mode as the majority of π^0 produced decay or interact within the nucleus falling under the jurisdiction of GENIE's intranuclear interaction. To overcome this problem the final state photons are recombined [131] to once again form a π^0 . NUANCE on the other hand does not display this feature and the decay of pions is left to nd280mc. Finally events are rejected if there are any charged primary particles and if there is more or less than one π^0 exiting the vertex.

Dalitz decay mode π^0 are accounted for in GENIE data by also considering events with electron-positron pairs in the final state of a GENIE interaction vertex. This is expected to have little effect upon the NC1 π^0 event selection however as they represent

< 5% of all π^0 decays and are expected to produce three EM showers rather than the the signature of just two showers for which we are searching.

9.2 Analysis

This analysis is an extension of the NC1 π^0 selection written for the P0D by Le Phoc Trung. This method applies a series of single cuts to a number of high level reconstructed variables, the structure of which is set out in Fig 9.1. The blue boxes outline cuts that are affected only by P0D data and are left fixed for this investigation. The red boxes represent cuts that arise from, or are affected by, the addition of P0DECal data; these cuts will be altered and the optimal values chosen.

The first cut restricts reconstructed vertices to lie within the fiducial region set by TP0DGeom in the oaGeometry package; 30 cm from the P0D edge in x and y plus 6 cm in z . The P0D NC1 π^0 analysis ignores all vertices reconstructed outside of this region, but it is hoped that the P0DECal will be able to extend this volume to include more events and improve the NC1 π^0 statistics. The increase in global acceptance is especially important in the early phases of data taking during which every measurement is statistically limited.

The fiducial statistic, d_{fid} , associated with each vertex gives the shortest distance to the edge of the fiducial volume and is positive inside and negative outside. We will alter the cut on this variable from the P0D chosen zero to the very edge of the P0D; this change will apply only in the x and y directions, as any increase in fiducial volume along the z direction could not be affected by the P0DECal.

The second stage of the selection considers clusters of hits within the P0D which are unmatched to a vertex. All such events are rejected by the P0D NC1 π^0 selection to remove ambiguity. Here we hope to match outlying unmatched clusters with objects in the P0DECal using position information only. The P0D hits provide a charge weighted position to which the informational P0DECal position is compared and the difference cut upon.

The third stage of the selection rejects charged current (CC) events by tagging MIP tracks. This is split into two separate stages: the first uses just P0D tracks and rejects events with at least one muon-like track, where muon-like is determined by likelihood; the second looks for tracks within the P0DECal.

There is currently no full matching algorithm between P0D and P0DECal, so any comparison of vertices or unmatched P0D clusters and P0DECal reconstructed objects

POD + P0DECal NC1 π^0 Event Selection

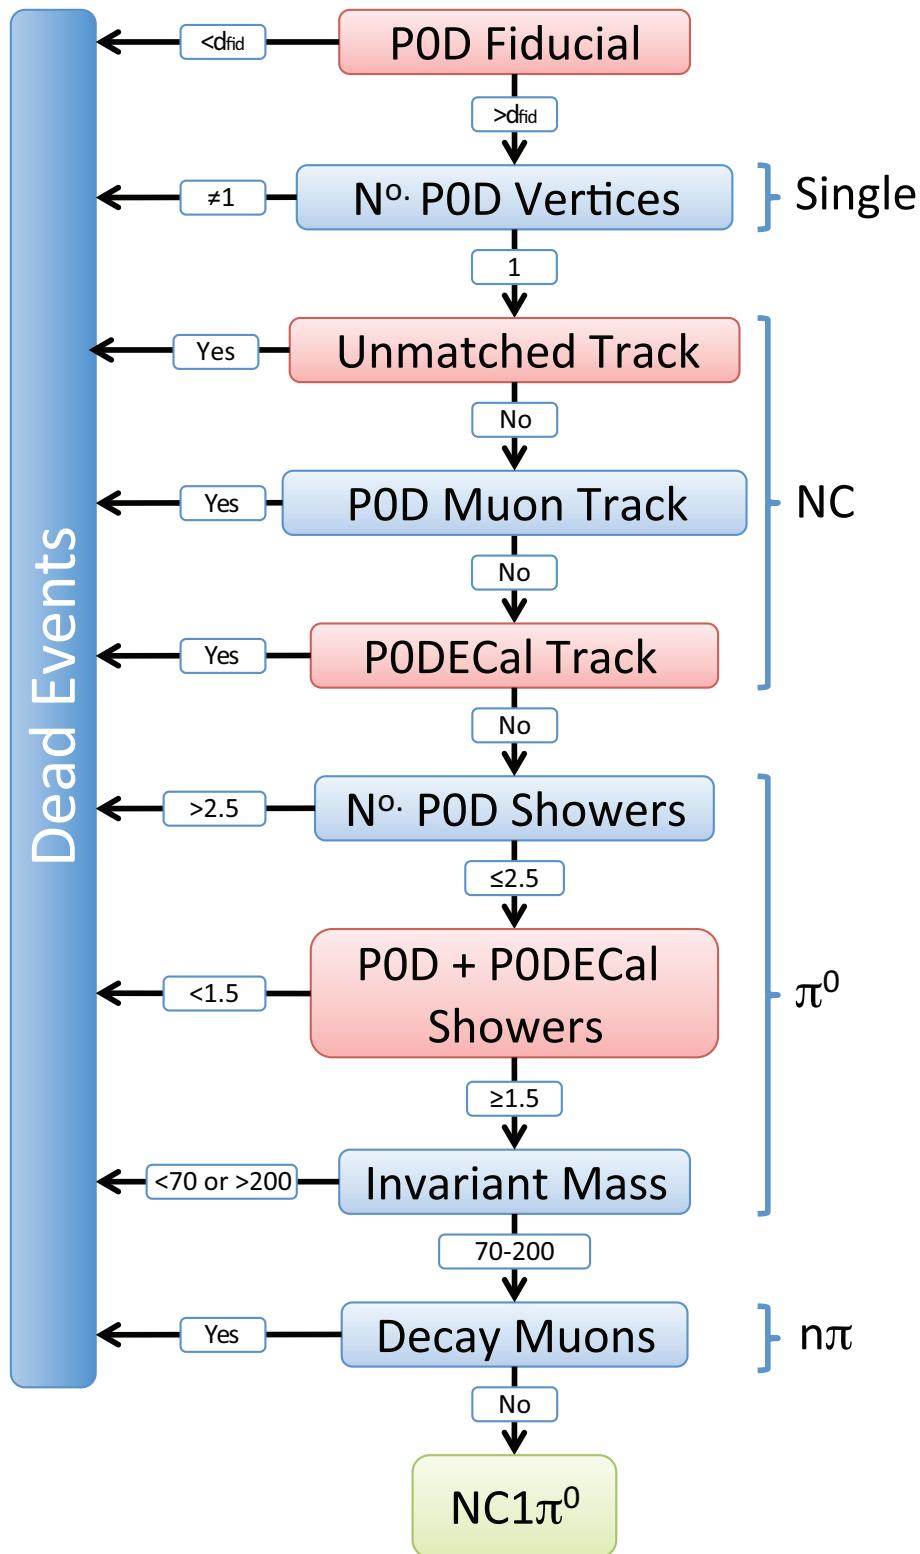
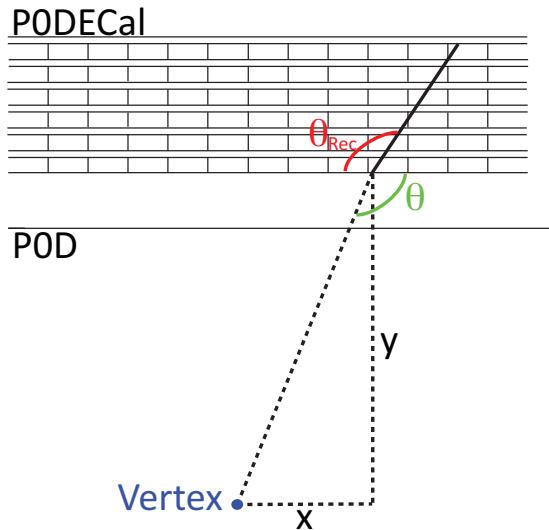


Figure 9.1: Flow of P0D + P0DECal NC1 π^0 event selection. Blue boxes denote P0D only cuts, red boxes those cuts which use P0D and P0DECal reconstructed events, and the green box the selection of events.

is achieved using cuts upon a number of variables. Timing is the first and most obvious cut that is made, with common objects depositing signal in a very tight time window within the similar scintillator MPPC readout and electronics of the two detectors.

P0DECal tracks are also compared with vertices, and inner unmatched clusters, within the P0D through comparison of angles. Given the reconstructed position of a vertex in the P0D and the reconstructed front position of the track one can calculate an expected angle of intercept θ . This is then compared to the track angle reconstructed by the P0DECal, θ_{Rec} . The errors on such measurements must be accounted for and we divide through by the total error σ_T to give an unbiased measure of angular difference. Our statistic which we cut upon is therefore



$$\begin{aligned}
 \Delta\theta &= \frac{\theta_{Rec} - \theta}{\sigma_T} \\
 \sigma_T &= \sqrt{\sigma_{\theta_{Rec}}^2 + \sigma_\theta^2} \\
 \sigma_\theta &= \sigma_x \left| \frac{d\theta}{dx} \right| \\
 &= \frac{\sigma_x}{y \left(1 + \frac{x^2}{y^2} \right)}. \quad (9.1)
 \end{aligned}$$

Figure 9.2: Theta comparison cut, P0DECal track and P0D vertex position.

The cut on this variable is another we alter to optimise event selection.

The fourth stage, and the final in which P0DECal information is used, determines the number of showers resulting from this event. The P0D has two views and so counts shower number in halves; an event is accepted if there are 1.5–2.5 showers, which allows for successful reconstruction in just one view and also fragmentation of a shower in one view. All events with > 2.5 showers are also rejected by the P0DECal analysis.

Events with P0D shower numbers below 1.5, however, are passed to the P0DECal to search for showers. The reconstruction analysis (chapter 7) shows that fragmentation does not occur in the P0DECal and so we can say with confidence that we have an accurate shower count. If there are zero P0D showers we require two P0DECal showers

and if there are 0.5 or 1 P0D showers then we require just one P0DECal shower. Due to the poor angular reconstruction of showers in the P0DECal we are generally unable to repeat the matching method employed with the tracks above but do attempt to do so for events with angular information. The main impact on the shower cuts is timing, and this will be the third cut parameter that we will adjust when searching for optimal cuts.

Stage five makes a cut on the reconstructed invariant mass of the P0D photon showers. It is hoped that the P0DECal will soon be able to have an impact here but as the energy reconstruction stands we can offer no additional information.

The sixth and final stage searches for Michel electrons from pion–muon decay to tag multi-pion events where the charged pion is not seen. The P0DECal may again be able to supply additional measures of such events but we will not discuss this here. This cut is, in any case, found to have a very marginal effect, so any gain or loss from including P0DECal data is expected to be minimal.

9.3 Results

In searching for the optimal cuts we first alter the fiducial cut to see the extent to which the P0DECal could extend the P0D fiducial volume. In Fig 9.3 we show the effect of altering this cut in selecting NC1 π^0 vertices from the NUANCE and GENIE water in and out MC. The efficiency plots are similar in shape, each showing a 5–6% increase in global efficiency, which means a larger increase with respect to P0D-only efficiency for the water out than the water in running mode.

The purity plots, however, display marked differences between the water in and out running modes. The water in modes show a slight, but constant, decreased purity until 150 mm where the decrease begins to grow at a linear rate. The water out data display a noticeable difference from water in. The P0DECal is seen to improve the purity of events as the fiducial volume is increased. This arises because of the change in radiation lengths of material in the P0D; particles traverse a greater distance and as such a greater fraction interact with the P0DECal and more photons will fail to convert within the P0D.

Concentrating once more on the water in running mode we see that for both GENIE and NUANCE the purity loss remains below 5% at a fiducial cut out to ~ 200 mm. We therefore choose this fiducial cut as a base for the next stage of optimisation; grid search of the track theta matching cut and the shower and track timing cut.

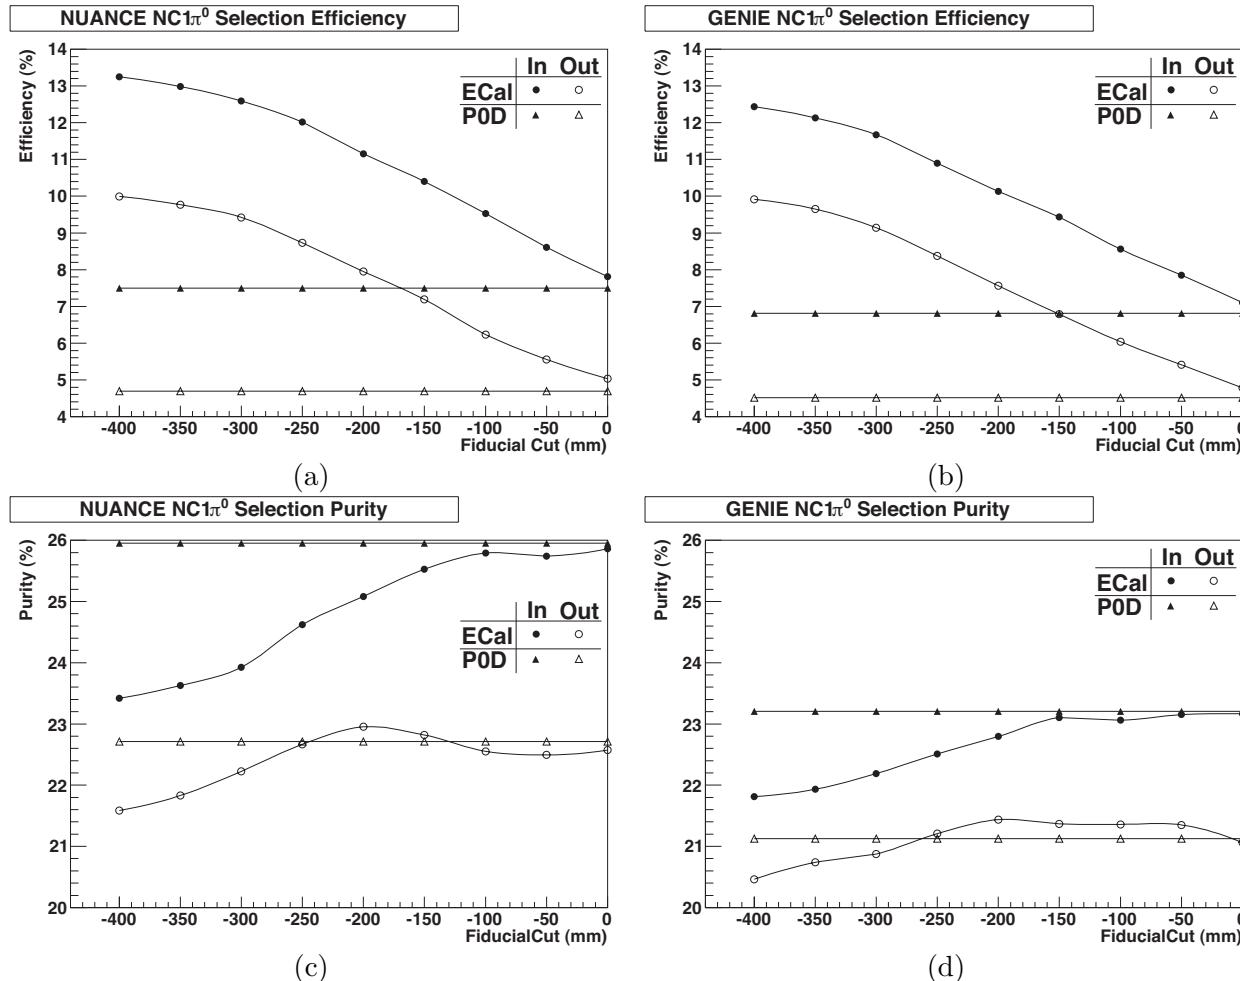


Figure 9.3: Efficiency (a & b) and purity (c & d) of NC1 π^0 event selection for large statistics MC of all processes generated by the NUANCE (a & c) and GENIE (b & d) neutrino generator for both P0D water in and out running modes.

We choose a range of theta cut values of $0\text{--}5\sigma$ and step of 0.5σ . For the timing cut we investigate the range 100–1000 ns taking steps of 100 ns. The multiplication of efficiency and purity is then plotted as a function of the cuts, shown in Fig 9.4.

We see a clear peak in the track cuts around 3σ for water in running mode using both GENIE and NUANCE MC. The peak in timing is less prominent but peaks at 500 ns for GENIE data and a more restrictive 100 ns for NUANCE water in data.

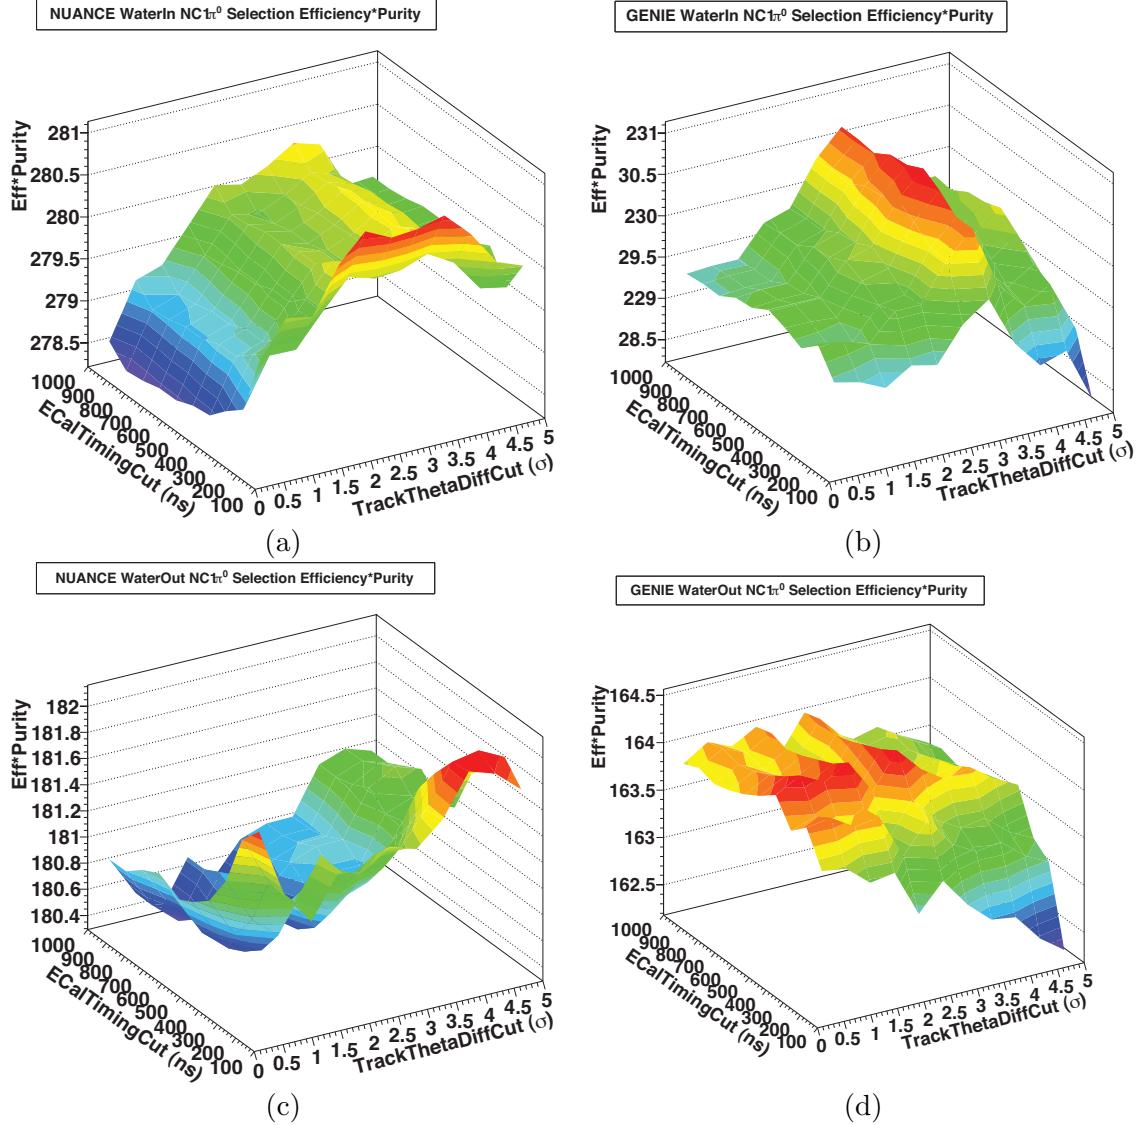


Figure 9.4: The efficiency multiplied by the purity of P0D+P0DECal NC1 π^0 event selection for water in (a & b) and water out (c & d) P0D running modes with data from the NUANCE (a & c) and GENIE (b & d) neutrino MC generators.

The water out running mode data are very different between the two generator datasets.

This difference was not expected and is not understood; further study is required. For this analysis we use the GENIE results as they are felt to represent a more accurate representation of the numerous interaction modes and distribution throughout the P0D. In Fig 9.5 we chart the efficiency and purity of the cuts employed.

It is clear from these plots that the weakest link in both efficiency and purity is cut number 6, which is the invariant mass cut. With a lack of successful energy reconstruction the P0DECal cannot supply information to the total invariant mass cut which therefore relies solely on the shower viewed by the P0D. This efficiency is hit with the total loss of two P0DECal shower, zero P0D shower events, and some one P0D one P0DECal shower events selected by the previous cut 5. The purity reduces due to some events with large shower energies which do not arise from a π^0 photon but lie within the correct invariant mass window.

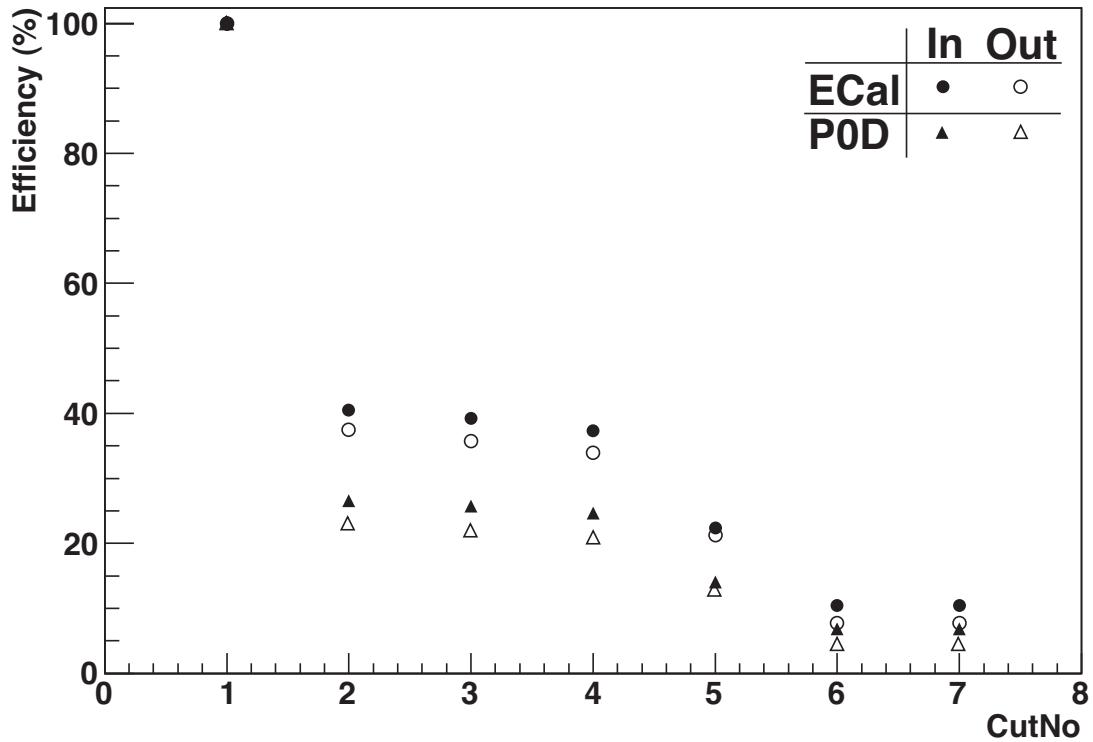
In tables 9.1 and 9.2 we breakdown the effect of each cut on various interaction channels. One can see that the most persistent background to the NC1 π^0 process in this selection is that of CC1 π^0 arising from overlapping tracks and/or misreconstruction of the muon produced.

In water out running mode there is a notable increase in acceptance of CC and NC single charged pion production (NC1Pi and CC1Pi) events. This can only be attributed to the P0DECals current inability to distinguish between hadronic and EM showers and the most likely route for improvement. Electron neutrino (nuE) events are rejected with higher efficiency during water out running mode suggesting sucessful tagging of such events by the P0DECal. All other interaction channels remain similarly supressed during both running modes at $\pm 1\%$ level.

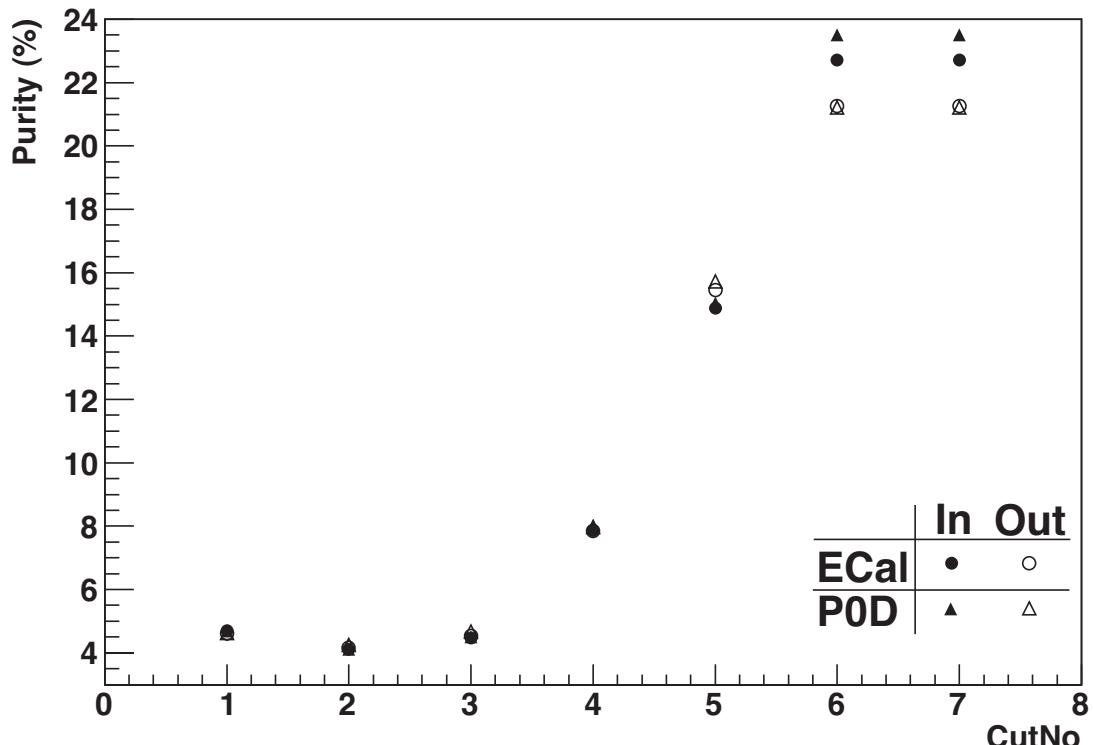
9.4 Conclusion and Discussion

This initial study has shown the P0DECal to be a powerful addition to the ND280. During the early stages of the experiment statistics will be limited and focus will be on efficiency of NC1 π^0 event selection. The P0DECal is shown to increase the efficiency, for the cuts selected, by 52% and 73% with a small change in purity of -3.4% and +0.1% for GENIE data of water in and water out running P0D modes when compared to P0D alone.

From this initial study it can be seen that the P0DECal will increase statistics in the early life of the experiment and, one hopes, with further refinement of the reconstruction, including possible energy information, it will be able to improve the purity of NC1 π^0 samples in the latter stages of the experiment. It is shown to come into its own

GENIE NC1 π^0 Event Selection Efficiency~CutNo

(a)

GENIE NC1 π^0 Event Selection Purity~CutNo

(b)

Figure 9.5: The efficiency (a) and purity (b) of NC1 π^0 event selection using the P0D and P0D+P0DECal methods at the chosen cuts (see text).

when the number of radiation lengths are reduced during water out running mode; the selection efficiency and purity are dramatically affected making it an essential tool.

GENIE Water In: Fiducial Cut=200 mm, Theta Cut=3 σ and Timing Cut=500 ns

P0D												
Cut	Eff	Pur	NC1Pi0	All	CC1Pi0	nuE	nPiRes	CCQE	NCQE	CC1Pi	NC1Pi	Other
0. None	100.00	4.69	20454	435727	19135	5923	56101	169725	66168	74016	15180	9025
1. Fiducial	26.58	4.11	5436	132163	5898	1775	17434	50007	23659	21243	3950	2761
2. Unmatched P0D	25.71	4.50	5259	116770	4886	1567	12907	44582	23482	18068	3567	2452
3. CC MIP Track	24.63	8.02	5038	62836	2339	1452	7451	12481	23263	7390	2843	579
4. ~2 Showers	14.09	15.04	2882	19160	1613	729	4350	3253	635	4511	915	272
5. Inv. Mass	6.82	23.51	1395	5934	493	112	435	1078	166	1876	338	41
6. Decay e ⁻	6.82	23.51	1395	5934	493	112	435	1078	166	1876	338	41
P0D+P0DECal												
Cut	Eff	Pur	NC1Pi0	All	CC1Pi0	nuE	nPiRes	CCQE	NCQE	CC1Pi	NC1Pi	Other
0. None	100.00	4.69	20454	435727	19135	5923	56101	169725	66168	74016	15180	9025
1. Fiducial	40.48	4.12	8280	200955	8910	2691	26717	76875	34335	32537	6169	4441
2. Unmatched P0D	39.21	4.48	8020	179083	7482	2420	20035	69546	34099	27889	5640	3952
3. CC MIP Track	37.32	7.87	7633	96982	3626	2220	11686	21136	33629	11576	4492	984
4. ~2 Showers	22.30	14.87	4561	30675	2578	1178	7147	5336	1053	6900	1453	469
5. Inv. Mass	10.37	22.71	2121	9341	842	203	729	1735	266	2854	518	73
6. Decay e ⁻	10.37	22.71	2121	9341	842	203	729	1735	266	2854	518	73

Table 9.1: Breakdown of NC1 π^0 selection from GENIE water in P0D data.

GENIE Water Out:Fiducial Cut=200 mm, Theta Cut=3 σ and Timing Cut=500 ns

P0D												
Cut	Eff	Pur	NC1Pi0	All	CC1Pi0	nuE	nPiRes	CCQE	NCQE	CC1Pi	NC1Pi	Other
0. None	100.00	4.63	17217	372233	16434	5156	47788	145744	56229	63020	12860	7785
1. Fiducial	23.11	4.25	3978	93619	4037	1257	12051	34556	17792	14995	2918	2035
2. Unmatched P0D	21.90	4.67	3770	80695	3278	1084	8590	29805	17642	12240	2544	1742
3. CC MIP Track	20.93	7.97	3603	45204	1684	1009	4936	8976	17444	5165	1921	466
4. \sim 2 Showers	12.74	15.73	2193	13940	1129	510	2785	2689	554	3192	675	213
5. Inv. Mass	4.48	21.21	771	3635	271	54	251	743	140	1136	220	49
6. Decay e^-	4.48	21.21	771	3635	271	54	251	743	140	1136	220	49

P0D+P0DECal												
Cut	Eff	Pur	NC1Pi0	All	CC1Pi0	nuE	nPiRes	CCQE	NCQE	CC1Pi	NC1Pi	Other
0. None	100.00	4.63	17217	372233	16434	5156	47788	145744	56229	63020	12860	7785
1. Fiducial	37.47	4.17	6452	154670	6813	2054	20231	58846	26927	25129	4809	3409
2. Unmatched P0D	35.75	4.53	6155	135787	5643	1818	14830	52348	26712	21039	4291	2951
3. CC MIP Track	33.89	7.87	5834	74139	2832	1667	8458	16418	26281	8636	3260	753
4. \sim 2 Showers	21.23	15.45	3655	23652	1959	889	5120	4486	927	5130	1122	364
5. Inv. Mass	7.73	21.24	1331	6266	533	114	489	1300	227	1822	365	85
6. Decay e^-	7.73	21.24	1331	6266	533	114	489	1300	227	1822	365	85

Table 9.2: Breakdown of NC1 π^0 selection from GENIE water out P0D data.

Chapter 10

Summary and Conclusions

This thesis has described the implementation of the T2K ND280 P0DECal software, from initial design studies to a preliminary investigation of its likely physics impact. Despite the relative simplicity and crude nature of the P0DECal, it has been shown that it is capable of achieving its primary physics goals: particles passing through the detector can be reconstructed and identified with good efficiency, and the efficiency of π^0 reconstruction in the P0D can be significantly increased.

In Chapter 4, the results of light yield measurements from cosmic ray muons incident on ECal scintillator bars was presented. An average light yield of 25 photoelectrons/MIP was observed using the second generation of MPPC connector. This is in excess of the minimum required for the successful running of the ECal, therefore validating the basic design parameters assumed in the initial ECal design proposal. It is also consistent with the expected and measured light yield in the scintillator bars of the existing SciBar detector, which were studied in the same experimental set-up.

Budget considerations dictated that the P0DECal, whose primary function is shower tagging rather than precision measurement, should be a considerably cruder device than the rest of the ND280 electromagnetic calorimeter. In Chapter 6 I presented the results of my design study optimising the thickness of lead foil in this downscoped detector. The choice of 4 mm taken as a result of this study is fully validated by the much more detailed performance studies described in susequent chapters.

Chapters 7 and 8 present the development and performance of the ECal reconstruction software. In Chapter 7, the basic clustering algorithm is described, and its performance evaluated. It is shown that muons above 50 MeV in kinetic energy are reconstructed with efficiency close to 100%, while photons are reconstructed with efficiency of 70% in the range 50–100 MeV rising to >90% above 400 MeV. The small number of active

layers in the P0DECal does, however, adversely affect the position, direction and, especially, energy resolution of the reconstructed clusters. The position resolution is found to be ± 25 mm in the appropriate coordinate (only the X(Y) coordinate can be reconstructed by the top/bottom(side) P0DECal modules owing to the orientation of the bars). The angular resolution is $\sim 10^\circ$ for photons and $\sim 2\text{--}4^\circ$ for muons, but with a substantial population, especially for photons, where no angular information can be extracted; and attempts to reconstruct the energy using the likelihood fitter developed for the Tracker ECal proved unsuccessful. It is possible that some degree of energy resolution could be achieved using a multivariate analysis analogous to that used for particle identification, and this approach is worthy of further study.

Chapter 8 considers the classification of P0DEcal clusters into tracks (muons and non-showering charged pions) and showers (photons, electrons and showering pions). Several charge-based and geometric discriminating variables are discussed. Owing to the lack of energy resolution in the P0DECal, it is essential to minimise any correlations between particle identification variables and energy, and this requirement is used to select the most appropriate set of variables. Standard linear techniques (Fisher discriminant and maximum likelihood) were found to give inadequate separation, but good results were obtained using non-linear multivariate techniques: boosted decision trees (BDTs) and multilayer perceptrons (MLPs). While BDTs seem to be capable of the best discrimination, they require much more CPU time to develop, are susceptible to overtraining, and have a much less extensive knowledge base among particle physicists than MLPs. It was therefore concluded that at present an MLP architecture with 6 neurons in a single hidden layer represents the best choice.

Finally in Chapter 9 I present the results of a study which takes the standard P0D π^0 event selection analysis and modifies it to make use of the additional information provided by the P0DECal. Although the lack of a useful P0DECal energy reconstruction remains the limiting factor, it is shown that the inclusion of P0DECal data allows us to relax the P0D fiducial volume cuts, increasing statistics by 52% (73%) whilst the purity changed little, $-3.4\% (+0.1\%)$ for water in(out) P0D running modes. Progress on energy reconstruction and/or on cluster matching between P0D and P0DECal might allow still further improvements to be made.

In conclusion, this thesis has demonstrated that, despite its simplicity and limited performance, Monte Carlo studies show that the P0DECal can detect particles with kinetic energies above 50 MeV with high efficiency, can discriminate effectively between shower-like and track-like topologies, and can substantially enhance the efficiency of π^0 reconstruction in the P0D. It is hoped that these expectations will be fully realised when the P0DEcal is installed in the T2K near detector later this year.

Bibliography

- [1] Schwetz, T., Tortola, M., and Valle, J. W. F. (2008) *New J. Phys.* **10**, 113011.
- [2] Cottingham, W. and Greenwood, D. (1998) An Introduction to the Standard Model of Particle Physics, Cambridge University Press, .
- [3] T2K ND280 Collaboration Nov 2005 <http://www.nd280.org>.
- [4] Fukuda, Y. et al. (2003) *Nucl. Instrum. Meth.* **A501**, 418–462.
- [5] T2K Collaboration November 2006 <http://www.nd280.org>.
- [6] J. Chadwick, F. (1932) *Proc. Roy. Soc., A* **136**, 692–708.
- [7] Fermi, E. (1934) *Z. Phys.* **88**, 161–177.
- [8] Bethe, H. and Peierls, R. (1934) *Nature* **133**, 532.
- [9] Reines, F. and Cowan, C. L. Nov 1953 *Phys. Rev.* **92(3)**, 830–831.
- [10] Davis, R. Feb 1955 *Phys. Rev.* **97(3)**, 766–769.
- [11] Cowan, Jr., C. L., Reines, F., Harrison, F. B., Kruse, H. W., and McGuire, A. D. July 1956 *Science* **124**, 103–104.
- [12] Sakata, S. and Inoue, T. November 1946 *Progress of Theoretical Physics* **1**, 143–150.
- [13] Marshak, R. E. and Bethe, H. A. Sep 1947 *Phys. Rev.* **72(6)**, 506–509.
- [14] Yukawa, H. (1935) *Proc.Phys.Math.Soc.Jap.* **17**, 48.
- [15] Lattes, C. M. G., Muirhead, H., Occhialini, G. P. S., and Powell, C. F. (1947) *Nature* **159**, 694–697.
- [16] Lee, T. D. and Yang, C. N. March 1960 *Physical Review Letters* **4**, 307–311.

- [17] Danby, G., Gaillard, J.-M., Goulianos, K., Lederman, L. M., Mistry, N., Schwartz, M., and Steinberger, J. July 1962 *Physical Review Letters* **9**, 36–44.
- [18] DONUT Collaboration (2001) *Phys. Lett.* **B504**, 218–224.
- [19] Lee, T. D. and Yang, C. N. Oct 1956 *Phys. Rev.* **104**(1), 254–258.
- [20] Wu, C. S., Ambler, E., Hayward, R. W., Hoppes, D. D., and Hudson, R. P. Feb 1957 *Phys. Rev.* **105**(4), 1413–1415.
- [21] Garwin, R. L., Lederman, L. M., and Weinrich, M. Feb 1957 *Phys. Rev.* **105**(4), 1415–1417.
- [22] Lee, T. D., Oehme, R., and Yang, C. N. Apr 1957 *Phys. Rev.* **106**(2), 340–345.
- [23] Goldhaber, M., Grodzins, L., and Sunyar, A. W. Feb 1958 *Phys. Rev.* **109**(3), 1015–1017.
- [24] Sudarshan, E. C. G. and Marshak, R. E. Mar 1958 *Phys. Rev.* **109**(5), 1860–1862.
- [25] Feynman, R. P. and Gell-Mann, M. Jan 1958 *Phys. Rev.* **109**(1), 193–198.
- [26] Glashow, S. L. (1961) *Nuclear Physics* **22**(4), 579 – 588.
- [27] Salam, A. and Ward, J. (1964) *Physics Letters* **13**(2), 168 – 171.
- [28] Weinberg, S. Nov 1967 *Phys. Rev. Lett.* **19**(21), 1264–1266.
- [29] tHooft, G. (1971) *Nuclear Physics B* **35**(1), 167 – 188.
- [30] Gargamelle Collaboration (1973) *Physics Letters B* **46**(1), 121 – 124.
- [31] Gargamelle Collaboration (1973) *Physics Letters B* **46**(1), 138 – 140.
- [32] Majorana, E. (1937) *Nuovo Cimento* **14**, 171.
- [33] COBRA Collaboration (2007) *Phys. Rev.* **C76**, 025501.
- [34] CUORICINO Collaboration (2008) *Phys. Rev.* **C78**, 035502.
- [35] Bakalyarov, A. M., Balysh, A. Y., Belyaev, S. T., Lebedev, V. I., and Zhukov, S. V. (2005) *Phys. Part. Nucl. Lett.* **2**, 77–81.
- [36] IGEX Collaboration (2002) *Phys. Rev.* **D65**, 092007.
- [37] NEMO Collaboration (2008) *arXiv* p. 0810.0248.
- [38] Particle Data Group (2008) *Phys. Lett.* **B667**, 1.

[39] Assamagan, K. et al. Jun 1996 *Phys. Rev. D* **53**, 6065–6077.

[40] ALEPH Collaboration (1998) *Eur. Phys. J. C* **2**, 395–406.

[41] Loredo, T. J. and Lamb, D. Q. (2002) *Phys. Rev. D* **65**, 063002.

[42] MINOS Collaboration (2007) *Phys. Rev. D* **76**, 072005.

[43] Spergel, D. N. et al. (2007) *The Astrophysical Journal Supplement Series* **170**(2), 377–408.

[44] Davis, R., Harmer, D. S., and Hoffman, K. C. May 1968 *Phys. Rev. Lett.* **20**(21), 1205–1209.

[45] Kamiokande Collaboration Jul 1989 *Phys. Rev. Lett.* **63**(1), 16–19.

[46] Abazov, A. I. et al. Dec 1991 *Phys. Rev. Lett.* **67**(24), 3332–3335.

[47] Gallex Colaboration (1992) *Physics Letters B* **285**(4), 376 – 389.

[48] Bahcall, J. N., Serenelli, A. M., and Basu, S. (2005) *Astrophys. J.* **621**, L85–L88.

[49] Hirata, K. S., Kajita, T., Koshiba, M., et al. (1988) *Physics Letters B* **205**(2-3), 416 – 420.

[50] IMB Collaboration Oct 1986 *Phys. Rev. Lett.* **57**(16), 1986–1989.

[51] IMB Collaboration May 1991 *Phys. Rev. Lett.* **66**(20), 2561–2564.

[52] Cabibbo, N. Jun 1963 *Phys. Rev. Lett.* **10**(12), 531–533.

[53] Mki, Z., Nakagawa, M., and Sakata, S. (1962) *Prog. Theor. Phys.* **28**, 870–880.

[54] Glashow, S. L., Iliopoulos, J., and Maiani, L. Oct 1970 *Phys. Rev. D* **2**(7), 1285–1292.

[55] Kobayashi, M. and Maskawa, T. (1973) *Progress of Theoretical Physics* **49**(2), 652–657.

[56] Christenson, J. H., Cronin, J. W., Fitch, V. L., and Turlay, R. Jul 1964 *Phys. Rev. Lett.* **13**(4), 138–140.

[57] Pontecorvo, B. May 1968 *Soviet Journal of Experimental and Theoretical Physics* **26**, 984–.

[58] Chau, L.-L. and Keung, W.-Y. Nov 1984 *Phys. Rev. Lett.* **53**(19), 1802–1805.

- [59] Gell-Mann, M. and Pais, A. Mar 1955 *Phys. Rev.* **97**(5), 1387–1389.
- [60] Pais, A. and Piccioni, O. Dec 1955 *Phys. Rev.* **100**(5), 1487–1489.
- [61] Pontecorvo, B. (1958) *Soviet Journal of Experimental and Theoretical Physics* **6**, 429–+.
- [62] Pontecorvo, B. (1958) *Soviet Journal of Experimental and Theoretical Physics* **7**, 172–+.
- [63] Lipkin, H. J. (2006) *Physics Letters B* **642**(4), 366 – 371.
- [64] SNO Collaboration (2002) *Phys. Rev. Lett.* **89**, 011301.
- [65] CHOOZ Collaboration (2003) *Eur. Phys. J.* **C27**, 331–374.
- [66] Boehm, F. et al. (2001) *Phys. Rev.* **D64**, 112001.
- [67] KamLAND Collaboration (2005) *Phys. Rev. Lett.* **94**, 081801.
- [68] KamLAND Collaboration (2008) *Phys. Rev. Lett.* **100**, 221803.
- [69] Allison, W. W. M. et al. (1997) *Phys. Lett.* **B391**, 491–500.
- [70] Super-Kamiokande Collaboration (1998) *Phys. Rev. Lett.* **81**, 1562–1567.
- [71] Super-Kamiokande Collaboration (2006) *Phys. Rev.* **D74**, 032002.
- [72] Super-Kamiokande Collaboration (2004) *Phys. Rev. Lett.* **93**, 101801.
- [73] MINOS Collaboration (2008) *Phys. Rev. Lett.* **101**, 131802.
- [74] Wolfenstein, L. May 1978 *Phys. Rev. D* **17**(9), 2369–2374.
- [75] Mikheev, S. P. and Smirnov, A. Y. (1987) *Sov. Phys. Usp.* **30**, 759–790.
- [76] SNO Collaboration (2008) *Phys. Rev. Lett.* **101**, 111301.
- [77] MACRO Collaboration (2004) *European Physical Journal C* **36**, 323–339.
- [78] K2K Collaboration (2006) *Phys. Rev.* **D74**, 072003.
- [79] KARMEN Collaboration (2002) *Phys. Rev.* **D65**, 112001.
- [80] NOMAD Collaboration (2003) *Phys. Lett.* **B570**, 19–31.
- [81] MiniBOONE Collaboration (2007) *Physical Review Letters* **98**(23), 231801.
- [82] MiniBooNE Collaboration (2008) *Phys. Rev.* **D78**, 012007.

[83] Fogli, G. L., Lisi, E., Marrone, A., Palazzo, A., and Rotunno, A. M. (2008) *Phys. Rev. Lett.* **101**, 141801.

[84] NOvA Collaboration (2004) *arXiv* pp. hep-ex/0503053.

[85] Fogli, G. L., Lisi, E., Marrone, A., and Palazzo, A. (2006) *Prog. Part. Nucl. Phys.* **57**, 742–795.

[86] Albert, J. April 2009 *T2K Internal Note*.

[87] Cervera, A., Donini, A., Gavela, M. B., Cdenas, J. J. G., Hernndez, P., Mena, O., and Rigolin, S. (2000) *Nuclear Physics B* **579(1-2)**, 17 – 55.

[88] Double Chooz Collaboration (2006) *arXiv* pp. hep-ex/0606025.

[89] Daya-Bay Collaboration (2007) *arXiv* pp. hep-ex/0701029.

[90] IMB Collaboration (1993) *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **324(1-2)**, 363 – 382.

[91] Shiozawa, M. August 1999 *Nuclear Instruments and Methods in Physics Research A* **433**, 240–246.

[92] Super-Kamiokande Collaboration (2000) *Phys. Rev. Lett.* **85**, 3999–4003.

[93] (for theT2K TPC groups), T. L. (2007) *Journal of Physics: Conference Series* **65**, 012018 (6pp).

[94] (for theT2K TPC groups), M. D. M. (2007) *Journal of Physics: Conference Series* **65**, 012019 (6pp).

[95] K2K Collaboration (2005) *Physics Letters B* **619(3-4)**, 255 – 262.

[96] Vacheret, A., Noy, M., Raymond, M., and Weber, A. (2006) *PoS PD07*, 027.

[97] Retière, F. o. b. o. t. T. N. F. g. (2006) *PoS PD07*, 017.

[98] Baron, P., Delagnes, E., Calvet, D., de laBroise, X., Delbart, A., Druillole, F., Fallou, J.-L., Mazzucato, E., Monmarthe, E., Pierre, F., Sarrat, A., Zonca, E., and Zito, M. 26 2007-Nov. 3 2007 *Nuclear Science Symposium Conference Record, 2007. NSS '07. IEEE* **3**, 1865–1872.

[99] Yamamoto, S. Oct. 2004 *Nuclear Science Symposium Conference Record, 2004 IEEE* **2**, 1034–1038 Vol. 2.

- [100] Pla-Dalmau, A., Bross, A. D., and Rykalin, V. V. FERMILAB-CONF-03-318-E.
- [101] Birks, J. B. (1951) *Proceedings of the Physical Society. Section A* **64**(10), 874–877.
- [102] Hirschberg, M., Beckmann, R., Brandenburg, U., Bruckmann, H., and Wick, K. Aug 1992 *Nuclear Science, IEEE Transactions on* **39**(4), 511–514.
- [103] Pla-Dalmau, A., Bross, A., Rykalin, V., and Wood, B. Oct. 2005 *Nuclear Science Symposium Conference Record, 2005 IEEE* **3**, 1298–1300.
- [104] Brun, R. and Rademakers, F. (1997) *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **389**(1-2), 81 – 86 New Computing Techniques in Physics Research V.
- [105] Ward, M. Mar 2009 *Private Communication*.
- [106] Hayato, Y. (2002) *Nuclear Physics B - Proceedings Supplements* **112**(1-3), 171 – 176.
- [107] Mitsuka, G. (2008) volume **981**, : AIP pp. 262–264.
- [108] Casper, D. (2002) *Nucl. Phys. Proc. Suppl.* **112**, 161–170.
- [109] Hawker, E. (2005) *Nuclear Physics B - Proceedings Supplements* **139**, 260 – 265 Proceedings of the Third International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region.
- [110] Andreopoulos, C. (2006) *Nuclear Physics B - Proceedings Supplements* **159**, 217 – 222 Proceedings of the 4th International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region, Proceedings of the 4th International Workshop on Neutrino-Nucleus Interactions in the Few-GeV Region.
- [111] Smith, C. H. L. (1972) *Physics Reports* **3**(5), 261 – 379.
- [112] Rein, D. and Sehgal, L. M. (1981) *Annals of Physics* **133**(1), 79 – 153.
- [113] Bodek, A., Park, I., and Yang, U.-k. (2005) *Nucl. Phys. Proc. Suppl.* **139**, 113–118.
- [114] GEANT 4 Collaboration (2003) *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**(3), 250 – 303.

[115] GEANT 4 Collaboration Feb. 2006 *Nuclear Science, IEEE Transactions on* **53(1)**, 270–278.

[116] Cervera-Villanueva, A., Gmez-Cadenas, J., and Hernando, J. (2004) *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **534(1-2)**, 180 – 183 Proceedings of the IXth International Workshop on Advanced Computing and Analysis Techniques in Physics Research.

[117] Kalman, R. E. (1960) *AMES Journal of Basic Engineering* **82(1)**, 35–45.

[118] Kalman, R. E. and Bucy, R. S. (1961) *AMES Journal of Basic Engineering* **83(1)**, 95–108.

[119] Duda, R. O. and Hart, P. E. (1972) *Commun. ACM* **15(1)**, 11–15.

[120] K2K Collaboration Dec 2005 *Phys. Rev. Lett.* **95(25)**, 252301.

[121] SciBooNE Collaboration (2008) *Phys. Rev.* **D78**, 112004.

[122] Rein, D. and Sehgal, L. M. (1983) *Nuclear Physics B* **223(1)**, 29 – 44.

[123] XCOM Photon Cross-Section Database.

[124] McCauley, N. Aug 2008 *Private Communication*.

[125] Carver, A. Feb 2009 *T2K Internal Note*.

[126] Spackman, K. A. (1989) In Proceedings of the sixth international workshop on Machine learning San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.. pp. 160–163.

[127] Scott, D. (1992) *Multivariate Density Estimation, Theory, Practice, and Visualization*, Wiley Interscience, New York, .

[128] L. Breiman, J. Friedman, R. O. and Stone, C. (1984) *Classification and Regression Trees*, Wadsworth, .

[129] Hornik, K., Stinchcombe, M., and White, H. (1989) *Neural Networks* **2(5)**, 359 – 366.

[130] SPEC <http://www.spec.org/>.

[131] Guzowski, P. Feb 2009 *Private Communication*.