Evaluating the Low-energy Response of the ProtoDUNE–SP Detector using Michel Electrons

Aidan Reynolds

University College
University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Trinity 2020

Abstract

This thesis presents the results of a study of electromagnetic interactions in the ProtoDUNE–SP liquid argon time projection chamber (LArTPC) detector. The LArTPC detector technology provides high spatial resolution on the final states of charged particle interactions, which allows for different interaction modes to be distinguished based on the geometry of the ionisation energy deposition in the event. As a result, the technology has been used in a number of recent neutrino experiments, and is the technology of choice for the Deep Underground Neutrino Experiment. In order to perform high precision measurements of neutrinos in LArTPC detectors, final state particles need to be effectively identified, and their energy accurately reconstructed. This thesis focusses on these challenges with two studies on data from the ProtoDUNE–SP LArTPC: a study of track–shower classification, and a study of Michel electron energy reconstruction. A track–shower classification algorithm was developed based on the use of convolutional neural networks, and its performance was compared to the current track–shower classification algorithm in ProtoDUNE–SP. The results of this network were used to select a sample of Michel electron events, which were then used to develop a Michel electron reconstruction algorithm based on semantic segmentation with a convolutional neural network. This sample of Michel electron events were then used to estimate the energy resolution and bias for low–energy electrons in ProtoDUNE–SP based on this algorithm.

Evaluating the Low-energy Response of the ProtoDUNE–SP Detector using Michel Electrons



Aidan Reynolds

University College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2020

To my parents

Abstract

This thesis presents the results of a study of electromagnetic interactions in the ProtoDUNE–SP liquid argon time projection chamber (LArTPC) detector. The LArTPC detector technology provides high spatial resolution on the final states of charged particle interactions, which allows for different interaction modes to be distinguished based on the geometry of the ionisation energy deposition in the event. As a result, the technology has been used in a number of recent neutrino experiments, and is the technology of choice for the Deep Underground Neutrino Experiment. In order to perform high precision measurements of neutrinos in LArTPC detectors, final state particles need to be effectively identified, and their energy accurately reconstructed. This thesis focusses on these challenges with two studies on data from the ProtoDUNE–SP LArTPC: a study of track–shower classification, and a study of Michel electron energy reconstruction. A track–shower classification algorithm was developed based on the use of convolutional neural networks, and its performance was compared to the current track–shower classification algorithm in ProtoDUNE–SP. The results of this network were used to select a sample of Michel electron events, which were then used to develop a Michel electron reconstruction algorithm based on semantic segmentation with a convolutional neural network. This sample of Michel electron events were then used to estimate the energy resolution and bias for low–energy electrons in ProtoDUNE–SP based on this algorithm.

Statement of Originality

This thesis and the work presented in it are my own and were produced by me as a result of my own original research. Results and figures from published works by others have been clearly attributed, and this work has not been submitted for another qualification at this or any other university.

After a brief introduction in Chapter 1, Chapter 2 gives an overview of the theory motivating the work presented in this thesis, which represents a summary of the work of others, for which all relevant sources have been referenced. Similarly, Chapters 4 and 5 represent a summary of the theory and tools that are the foundations of the work in this thesis, and are the work of others.

Chapter 3 describes the ProtoDUNE–SP experiment, whose data was analysed in this thesis. This chapter is predominantly a review of the work of others; however, I made contributions to the development of the online monitoring system, for which I have included a statement outlining my contributions in Section 3.8.

The development of the hit tagging algorithm based on a convolutional neural network in Chapter 6 is my own work, as well as the analysis of the results for the ProtoDUNE–SP beam particles and cosmic–rays. The hit reconstruction and clustering algorithms that led to the clustering of hits into tracks and showers were developed by others, and the relevant documentation has been referenced. In the same way, Section 6.4 details the use of my hit tagging algorithm by others, whose work is referenced.

The work on Michel electron reconstruction in Chapter 7 and Appendix A is my own. As with Chapter 6, this work relies on the clustering algorithms developed by others. In addition, the energy calibration factors used in the Michel electron energy reconstruction were calculated in the work of others.

Finally, the conclusions presented in Chapter 8 are my own.

Acknowledgements

This thesis would not have been possible without the advice and support of many people. I would like to start by thanking my supervisor, Alfons Weber. He has been a brilliant advisor, and I would like to thank him for all the opportunities he has given me. Thanks also to the rest of the Oxford Neutrino Physics group, particularly Justo Martin–Albo, whose support during my first year helped me to settle into the DPhil, and my fellow students – Fabio, Alex, and Ciaran – who have made the office, both physical and virtual, an enjoyable place to work.

I have worked with many great collaborators from the DUNE experiment as part of my DPhil. I'd particularly like to thank the members of the ProtoDUNE–SP reconstruction and analysis group, who have always given excellent feedback on my work. Special thanks go to Dorota Stefan and Robert Sulej, who were incredibly supportive during the early years of my DPhil, and to Leigh Whitehead and Tingjun Yang, whose insightful conversations have been invaluable.

I am very fortunate to have had the chance to work at CERN for part of my DPhil, and I would like to thank all of the members of the on–site ProtoDUNE–SP team, especially Alex, Chris, Geoff, James, Milo, and Seb, who I thoroughly enjoyed working with during my time at CERN, and from whom I have learned so much.

My greatest thanks go to my friends and family for their continued support over the years. First, I'd like to thank Babs, whose proof reading has considerably improved the clarity of this thesis. Next to my friends in Oxford: to Amy, Rory, and Helena, who have always been there to help me relax at the end of the day, and to the members of Oxford Ultimate, who have been so welcoming over the past year, you have kept me going during the final stages. Finally, to my family, who fostered a love of learning which will always be with me, and to Ellie, who could never understand how much her support has meant to me over the years.

Contents

List of Figures viii

List of Tables xii

List of Abbreviations xiii

1 Introduction 1

2 Neutrino Physics 5

2.1 A Brief History of Neutrino Physics 6

2.2 Neutrinos in the Standard Model 12

2.3 Neutrino Oscillations 14

2.3.1 Neutrino Oscillations in Vacuum 14

2.3.2 Neutrino Oscillations in Matter 17

2.3.3 Current Knowledge and Open Questions 18

2.4 Neutrino Interactions 22

2.5 Supernova Neutrinos 26

2.5.1 Core–collapse Supernova Dynamics 27

2.5.2 SN1987A 28

2.5.2.1 Kamiokande–II 29

2.5.2.2 IMB 29

2.5.2.3 Baksan 29

2.5.3 Supernova Neutrino Prospects in DUNE 31

2.5.3.1 Supernova Neutrino Interactions in Liquid Argon . 32

2.5.3.2 Supernova Neutrino Events in DUNE 33

3 The ProtoDUNE–SP Experiment 35

3.1 ProtoDUNE–SP in the Context of DUNE 36

3.1.1 Liquid Argon Time Projection Chambers 37

3.1.2 ProtoDUNE–SP Goals 39

3.2 The ProtoDUNE–SP Detector 40

3.2.1 The Liquid Argon TPC 41

3.2.2 The Photon Detection System 43

3.3 The H4 beamline . . . 45
3.4 The Cosmic–Ray Tagger . . . 46
3.5 The Data Acquisition System . . . 48
3.6 Simulation . . . 50
3.7 Reconstruction . . . 51
3.8 ProtoDUNE–SP Online Monitoring System . . . 57
3.8.1 Data Processing . . . 58
3.8.2 Monitoring Web Interface . . . 64
3.8.3 System Performance and Possible Improvements . . . 69

4 Electromagnetic Energy Loss in Liquid Argon 73
4.1 Electromagnetic Energy Loss in Matter . . . 73
4.1.1 Energy Loss for Heavy Charged Particles . . . 74
4.1.2 Energy Loss for Electrons . . . 75
4.1.3 Energy Loss for Photons . . . 78
4.2 Interaction Signatures in ProtoDUNE–SP . . . 79
4.3 Scintillation Light in Liquid Argon . . . 82

5 Neural Networks 83
5.1 Artificial Neural Networks . . . 84
5.1.1 Convolutional Neural Networks . . . 86
5.1.2 Residual Neural Networks . . . 88
5.2 Supervised Learning . . . 89
5.2.1 The Backpropagation Algorithm . . . 90
5.2.1.1 The Choice of Activation Function . . . 92
5.2.2 Regularisation . . . 93

6 Hit Classification with Convolutional Neural Networks 95
6.1 Hit Classification with Convolutional Neural Networks . . . 96
6.1.1 Data Preparation . . . 97
6.1.2 Network Architecture . . . 99
6.1.3 Training and Validation . . . 102
6.2 Performance on ProtoDUNE–SP Simulation . . . 103
6.2.1 Comparison with Pandora . . . 107
6.3 Validation on ProtoDUNE–SP Data . . . 110
6.3.1 Systematic Uncertainties . . . 119
6.4 Application in ProtoDUNE–SP Analyses . . . 121

7 Study of Michel Electron Reconstruction in ProtoDUNE–SP 123
7.1 Michel Electrons in Liquid Argon 124
7.2 Michel Electron Event Selection 131
7.3 Michel Electron Energy Reconstruction 134
7.3.1 Michel Electron Hit Tagging with a U–ResNet 134
7.3.2 Michel Electron Reconstruction 139
7.3.2.1 Hit Selection 140
7.3.2.2 Ionisation Energy Reconstruction 142
7.4 Conclusion 153

8 Conclusion 155

Appendices

A Additional Details for Michel Electron Reconstruction 159

References 161

List of Figures

2.1 Ratio of data to Monte Carlo for electron and muon neutrino fluxes measured by the Super Kamiokande experiment as a function of L/E_ν. 10
2.2 Solar neutrino flux composition as measured by the SNO experiment. 11
2.3 Ratio of observed neutrino flux to the predicted flux in the absence of neutrino oscillations in the KamLAND experiment. 12
2.4 Feynman diagrams for neutrino scattering in matter. 17
2.5 The two possible neutrino mass orderings. 18
2.6 90% confidence intervals for $\sin^2(\theta_{23})$–Δm^2_{32} from a number of leading neutrino oscillation experiments. 21
2.7 Confidence intervals for δ_{CP} from the T2K experiment. 23
2.8 Confidence intervals for δ_{CP} from the NOνA experiment. 24
2.9 Muon neutrino cross section as a function of energy. 26
2.10 Measured supernova neutrino events from SN1987A in Kamiokande–II. 30
2.11 Supernova neutrino flux at earth as a function of neutrino energy. . 31
2.12 Neutrino cross sections in liquid argon as a function of neutrino energy, in the energy regime relevant for supernova neutrinos. . . . 32
2.13 Predicted supernova neutrino event rate and energy spectra for a 10 kpc core collapse supernova measured by DUNE. 34

3.1 The Deep Underground Neutrino Experiment. 36
3.2 A schematic of the LArTPC detection princile. 38
3.3 A 3D rendering of the main components of the ProtoDUNE–SP TPC. 41
3.4 The operating principle of the photon detector modules in ProtoDUNE–SP. 43
3.5 Location of the H4 beamline and cosmic–ray taggers in relation to the ProtoDUNE–SP detector. 45
3.6 A schematic of the H4 beamline magnets and instrumentation. . . . 46
3.7 Time of flight vs momentum distributions from the H4 beamline instrumentation. 47
3.8 Outline of the data acquisition system in ProtoDUNE–SP. 49
3.9 Example of noise filtering and 2D deconvolution in ProtoDUNE–SP data. 52

3.10 An example of reconstructed hits from ProtoDUNE–SP data. 53
3.11 Diagram demonstrating the track stitching algorithm in ProtoDUNE–SP. 54
3.12 An example of a reconstructed particle hierarchy from Pandora. . . 55
3.13 dE/dx vs residual range for a stopping muon sample in ProtoDUNE–SP data. 57
3.14 A diagram representing the data flow in the ProtoDUNE–SP online monitoring system. 59
3.15 Examples of online monitoring plots for the TPC data. 61
3.16 An example of a beam window event display from the ProtoDUNE–SP online monitoring system. 62
3.17 Examples of online monitoring plots for the photon detection system data. 63
3.18 Examples of online monitoring plots for the timing and trigger data. 65
3.19 Examples of online monitoring plots for the beam instrumentation. 66
3.20 Examples of online monitoring plots for the cosmic–ray tagger. . . . 67
3.21 Data flow diagram for the web interface in the ProtoDUNE–SP online monitoring system. 68
3.22 An example of an online monitoring page. 68
3.23 Example of the dynamic binning algorithm in the ProtoDUNE–SP online monitoring system. 70

4.1 Stopping power as a function of energy for muons in liquid argon. . 75
4.2 Stopping power as a function of energy for electrons in liquid argon. 77
4.3 Photon interaction cross sections in liquid argon. 79
4.4 Photon mean free path in liquid argon. 80
4.5 Typical particle signatures in ProtoDUNE–SP for different particle species. 81

5.1 Graphical representation of an individual neuron in an artificial neural network. 84
5.2 A graphical representation of a multi–layer perceptron. 85
5.3 Application of a Sobel edge detector to an image. 88
5.4 A graphical representation of a convolutional neural network. 89
5.5 An example of a shortcut connection in a convolutional neural network. 90

6.1 Example CNN input images for the track, shower, and Michel classes. 98
6.2 Network architecture used for hit classification. 101
6.3 Evolution of the validation and training set losses during the training process for the hit tagging CNN. 103

6.4 Shower score output distributions for the TSE classifier. 105
6.5 ROC curves for the shower score from the TSE classifier. 106
6.6 Michel electron score distributions for the Michel electron classifier. 107
6.7 ROC curves for the Michel electron classifier. 108
6.8 Comparison of track and shower classifications by Pandora and the CNN. 109
6.9 Track classifier scores for a reconstructed event in ProtoDUNE–SP data. 111
6.10 Quality cuts on reconstructed hits for CNN comparison. 112
6.11 Overall shower score distribution from the CNN for all hits in data and simulation. 113
6.12 Data quality cuts for CNN comparison. 116
6.13 Shower classifier scores for different particle species from the ProtoDUNE–SP beam. 117
6.14 Shower classifier score for hits from cosmic–rays. 118
6.15 CNN shower score distribution for secondary particles in neutral pion event selection sample. 122

7.1 Michel electron energy spectra for free muons and captured muons in liquid argon from ProtoDUNE–SP simulation. 125
7.2 Michel electron candidate event from ProtoDUNE–SP data. 126
7.3 Properties of radiated energy deposits from Michel electrons. 127
7.4 Spatial distribution of radiated ionisation deposits from Michel electrons. 128
7.5 Available ionisation energy vs true Michel electron energy for two different collection radii. 129
7.6 Fraction of Michel electron energy collected as ionisation vs collection radius. 130
7.7 Total true ionisation energy recovered for a 40 cm collection radius, and an 80 cm collection radius. 131
7.8 Fraction of Michel–like hits in the best Michel electron candidate. . 133
7.9 Efficiency of Michel electron event selection as a function of energy. 134
7.10 Spatial and angular distributions for primary muons associated with selected Michel electrons. 135
7.11 U–ResNet CNN architecture used to select Michel electron ionisation energy deposits. 138
7.12 U–ResNet training and validation loss as a function of epoch. 139
7.13 Example input, truth, and predicted images for U–ResNet. 140
7.14 U–ResNet predicted score distribution for hits in Michel electron ionisation images. 141

7.15 Number of selected hits in reconstructed Michel electron events. . . 143
7.16 Reconstructed ionisation energy for hits in U–ResNet input images. 144
7.17 Reconstructed ionisation energy for Michel electron candidates in ProtoDUNE–SP data and simulation. 145
7.18 Reconstructed ionisation energy vs total deposited ionisation energy for Michel electron events in ProtoDUNE–SP simulation. 146
7.19 Fractional energy difference between reconstructed and true Michel electron energy. 147
7.20 Energy resolution and bias for Michel electrons as a function of true ionisation energy deposition. 148
7.21 Energy uncertainty contribution from hit energy reconstruction for Michel electron ionisation energy reconstruction. 150

A.1 Gaussian fits to the fractional energy difference bewteen the reconstructed ionisation energy and the true ionisation energy for Michel electron events. 160

List of Tables

3.1 Summary of expected and recorded beam triggers in ProtoDUNE–SP. 47

6.1 Truth labels for the CNN output for different particle species. . . . 99
6.2 Number of input images with each truth label. 99
6.3 False–positive rates for track shower classification with Pandora and the CNN. 109
6.4 Fraction of hits in correct class for different samples in ProtoDUNE–SP data and simulation. 119

7.1 Comparison of energy resolution estimates for Michel electron ionisation energy. 149

List of Abbreviations

Adam . . . . Adaptive Momentum Estimation

ANN . . . . . Artificial Neural Network

APA . . . . . Anode–plane Assembly

BI . . . . . Beam Instrumentation

CC . . . . . Charged Current

CE . . . . . Cold Electronics

CNN . . . . . Convolutional Neural Network

CP . . . . . Charge–Parity

CPT . . . . . Charge–Parity–Time

CPA . . . . . Cathode–plane Assembly

CRT . . . . . Cosmic–ray Tagger

CTB . . . . . Central Trigger Board

DAQ . . . . . Data Acquisition System

DIS . . . . . Deep Inelastic Scattering

DUNE . . . . Deep Underground Neutrino Experiment

ES . . . . . Elastic Scattering

FELIX . . . . Front–end Link Exchange

FEMB . . . . Front–end Mother Board

FFT . . . . . Fast Fourier Transform

FNAL . . . . Fermilab National Accelerator Laboratory

IO . . . . . Inverted Ordering

KamLAND . Kamioka Liquid Scintillator Anti–neutrino Detector

LArTPC . . Liquid Argon Time Projection Chamber

LEP . . . . . Large Electron–Positron Collider

LHC . . . . . Large Hadron Collider

LIGO . . . . Laser Interferometer Gravitational–Wave Observatory
LNGS . . . . Laboratori Nazionali del Gran Sasso
MC . . . . . Monte–carlo Simulation
MIP . . . . . Minimum Ionising Particle
ML . . . . . Machine Learning
MLP . . . . . Multi–layer Perceptron
MSW . . . . Mikheyev–Smirnov–Wolfenstein
NC . . . . . Neutral Current
NN . . . . . Neural Network
NO . . . . . Normal Ordering
OM . . . . . Online Monitoring
PDS . . . . . Photon Detection System
PFParticle . Paricle Flow Particle
PID . . . . . Particle Identification
PMNS . . . . Pontecorvo–Maki–Nakagawa–Sakata
QE . . . . . Quasi–elastic
RCE . . . . . Reconfigurable Computing Element
ReLU . . . . Rectified Linear Unit
RES . . . . . Resonance
ResNet . . . Residual Neural Network
RMS . . . . . Root Mean Square Deviation
ROC . . . . . Reciever Operator Characteristic
SCE . . . . . Space Charge Effect
SGD . . . . . Stochastic Gradient Descent
SiPM . . . . Silicon Photo–multiplier
SM . . . . . Standard Model
SNO . . . . . Sudbury Neutrino Observatory
SP . . . . . Single Phase
SSM . . . . . Standard Solar Model
SSP . . . . . SiPM Signal Processor
SURF . . . . Sanford Underground Research Facility

TOF . . . . . Time of Flight

TPB . . . . . Tetraphenyl-butadiene

TPC . . . . . Time Projection Chamber

TSE . . . . . Track–Shower–Empty

WIB . . . . . Warm Interface Board

1 Introduction

Since the discovery of neutrino flavour oscillations, which implies that neutrinos have mass, neutrino physics has enjoyed a period of rapid development. The field has begun to transition into an era of precision, with many of the parameters governing these oscillations having been well constrained. The fact that neutrinos have mass, and the success of the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) theory in describing neutrino oscillations, leads to a number of fundamental questions that have important implications for both particle physics and cosmology:

- What is the mechanism giving rise to neutrino mass?
- Are neutrinos Dirac or Majorana particles?
- What is the absolute scale and ordering of the neutrino masses?
- Do neutrinos and anti–neutrinos oscillate differently, and would this help to explain the matter anti–matter asymmetry in the universe?
- Are there any sterile neutrinos?

The high resolution and large masses of modern neutrino detectors also make them useful tools for both astronomy and astrophysics. 2017 has widely been considered as the dawn of multi–messenger astronomy, with a measurement of gravitational waves at the Laser Interferometer Gravitational–Wave Observatory (LIGO) being correlated with measurements of a neutron star merger from electromagnetic

telescopes[1]. This measurement was shortly followed by a similar correlation, but in the neutrino sector, between a high energy neutrino event in the IceCube Neutrino Observatory and a number of traditional telescopes[2]. Within our galaxy, neutrino detectors provide a unique opportunity to understand the underlying mechanisms in supernovae; in the case of such a supernova, the structure of the neutrino flux at earth provides a mechanism to measure effects in the early stages of the supernova burst, which are inaccessible with electromagnetic measurements[3].

Each of these questions places constraints on the design of an appropriate neutrino detector. The discovery of a matter anti–matter asymmetry in neutrino oscillations could be answered by making precise measurements of neutrino oscillations. This requires reliably identifying the flavour and energy of neutrinos, in order to measure the appearance and disappearance spectra associated with neutrinos produced in long baseline neutrino experiments. To identify the low energy electrons produced in supernova neutrino interactions, a detector with low thresholds and low backgrounds is required. The Deep Underground Neutrino Experiment (DUNE) aims to tackle these challenges by utilising the liquid argon time projection chamber (LArTPC) technology, whose high spatial and calorimetric resolution allows for accurate geometric classification of neutrino interactions. To achieve these goals, a significant programme of LArTPC research is ongoing with construction, reconstruction, and analysis methods all under development in a number of LArTPC based experiments[4–7].

This thesis presents analyses of charged particle interactions in the ProtoDUNE–SP LArTPC detector, using data collected with a test beam between August and November 2018. Hit classification and Michel electron reconstruction are investigated, and a sample of Michel electrons is used to provide a measurement of the energy resolution and bias for low energy electrons in ProtoDUNE–SP.

Particle classification plays an important role in event reconstruction in LArTPC detectors. In particular, the clustering of hits into tracks and showers is an important step in reconstructing events in a LArTPC. After tracks and showers have been reconstructed, they can be combined to build up a picture of the full particle

interaction. This thesis presents the results of the development of a hit classification algorithm based on convolutional neural networks. The primary goal of this algorithm is to classify hits as track–like or shower–like on a hit–by–hit basis. The output from the hit classification has been applied to analyses of beam particle interactions, such as pion cross–section analyses, and cosmic–ray interactions, which are used to calibrate the ProtoDUNE–SP detector. These outputs are also used to select a sample of Michel electron candidates that are also analysed in this thesis.

Michel electrons have an energy spectrum spanning 0–60 MeV. Understanding electrons in this energy range is important, as they have a similar energy to the electrons produced when neutrinos from supernova bursts interact. In a LArTPC at these energies, the energy deposition of electrons transitions between ionisation dominated and radiation dominated regimes, making for a characteristic event signature.

The work presented in this thesis also details the selection and reconstruction of Michel electron events in ProtoDUNE–SP, based on machine learning algorithms. Analysis of the reconstructed Michel electron events is used to quantify the energy resolution and bias for low energy electrons in ProtoDUNE–SP based on this approach. The results of this analysis provide valuable inputs to studies of supernova burst neutrinos in LArTPC detectors.

In this thesis, Chapter 2 provides a theoretical overview of neutrinos within the standard model. Interactions, oscillations, and production will be discussed summarising the current knowledge in this field, as well as open questions that will be studied in ongoing and upcoming experiments. The role of neutrinos in supernova bursts, and the detection of such neutrinos in a LArTPC detector, will also be discussed in more detail.

The ProtoDUNE–SP experiment is described in Chapter 3, including details of the beam line, detector, cosmic–ray flux, and simulations. An overview of the LArTPC detection principle will be given with specific details of the ProtoDUNE–SP design. Some details of detector operations will be discussed, paying particular

attention to the monitoring of the detector via the online data quality monitoring system, which has been the author's major contribution to the detector operations.

Chapter 4 will cover details of electromagnetic energy loss in liquid argon. Energy loss for muons, electrons, and photons will be discussed as well as the production of scintillation light in liquid argon. The impacts of these effects on the event signatures in liquid argon will be highlighted.

The main analyses of this thesis, which are discussed in Chapters 6 and 7, make use of neural network algorithms for reconstruction, therefore, Chapter 5 will briefly outline the relevant details of these algorithms.

Chapter 6 will detail the development of a hit classification algorithm based on convolutional neural networks. This algorithm is used to classify hits as track–like or shower–like in ProtoDUNE–SP, as well as to identify Michel electron hits. This algorithm and the results it produces are currently being used in a number of analyses of ProtoDUNE–SP data, including the selection of Michel electron events, which is discussed in Chapter 7.

In Chapter 7, Michel electron reconstruction in ProtoDUNE–SP will be discussed. First, a brief overview of Michel electron production and energy loss in liquid argon will be given. This will be followed by a discussion of event selection and energy reconstruction algorithms, which were developed to study Michel electron events in ProtoDUNE–SP. The reconstructed Michel electron spectrum will be compared between data and simulation, and the energy resolution and bias for low energy electrons in the ProtoDUNE–SP detector will be estimated. Finally, the results will be compared to those from similar experiments, and their implications for DUNE will be highlighted.

A summary of the results will be given in Chapter 8, along with a discussion of the implications of these results for physics in LArTPC detectors.

2
Neutrino Physics

Contents

2.1 A Brief History of Neutrino Physics 6
2.2 Neutrinos in the Standard Model 12
2.3 Neutrino Oscillations 14
2.3.1 Neutrino Oscillations in Vacuum 14
2.3.2 Neutrino Oscillations in Matter 17
2.3.3 Current Knowledge and Open Questions 18
2.4 Neutrino Interactions 22
2.5 Supernova Neutrinos 26
2.5.1 Core–collapse Supernova Dynamics 27
2.5.2 SN1987A 28
2.5.3 Supernova Neutrino Prospects in DUNE 31

Despite being one of the most abundant particles in the universe, neutrinos are some of the most elusive; due to the fact that they can only interact via the weak interaction. The history of neutrino physics is, therefore, strongly connected to the discovery and study of weak interactions.

Measurements by Chadwick in 1914 showed that the energy spectrum of electrons released in β–decays was continuous[8], this is in contrast to discrete spectra observed in α and γ decays, and seemingly violates conservation of energy under the assumption of a two–body final state, which was expected at the time. In order to solve this problem, Pauli postulated that the continuous energy spectrum could

be explained if the energy released in a β–decay is shared with an additional neutral weakly interacting fermion[9], which Pauli named the neutron. Fermi later renamed Pauli's fermion to the neutrino, after Chadwick's discovery of the neutron in 1932. Despite claims that neutrinos might never be detected, not only have they now been discovered but they have also been found to have a number of interesting properties, which were not anticipated when they were first postulated.

This chapter will detail some of the history and theory of neutrinos and their interactions. Section 2.1 will give a brief historical overview of neutrino physics. Section 2.2 will introduce the theory of neutrinos in the Standard Model, followed by a discussion of neutrino oscillations in Section 2.3. Neutrino interactions will be discussed briefly in Section 2.4, and finally Section 2.5 will discuss the production and measurement of neutrinos from supernovae.

2.1 A Brief History of Neutrino Physics

The first attempt to incorporate the neutrino into a theoretical model came in 1934, in Fermi's theory of β–decay. In this theory, the neutrino takes part in a four–point interaction with the other components of the β–decay interaction[10]. The incredible success of this theory in explaining the observed properties of β–decays provided strong evidence for the neutrinos existence, however, in 1934 after using Fermi's theory to predict the strength of neutrino interactions, H. Bethe and R. Peierls found that the interactions were so weak that they might never be observed, a hypothesis that wouldn't be disproven for over 20 years[11].

The first breakthrough in experimental neutrino physics would come in 1956, when F. Reines and C. Cowan were attempting to measure positrons produced in inverse β–decay interactions,

$$\overline{\nu_e} + p \rightarrow n + e^+.$$

A detector containing 1400 litres of liquid scintillator was used to measure the large flux of electron anti–neutrinos in the vicinity of the Savannah River nuclear reactor. They observed a large increase in the rate of positron events when the

reactor was on compared to when the reactor was switched off. This was the first experimental evidence for the existence of neutrinos[12].

This discovery gave rise to questions of neutrino flavour. As neutrinos are produced alongside a charged lepton, it is natural to compare the properties of neutrinos with their partners in the weak interaction. At the time of the discovery of the neutrino, there were two known charged leptons, the electron and the muon, and physicists questioned whether the neutrinos produced alongside muons are different from those produced alongside electrons. In 1962, Lederman et al discovered the muon neutrino at Brookhaven National Laboratory; by creating a beam of muon–associated neutrinos from decaying pions, and observing the leptons produced in neutrino interactions after all other particles had been absorbed. They found that only muons were produced in the resulting neutrino interactions and, therefore, the neutrinos were only ever associated with a muon. This showed that neutrinos are produced with a distinct flavour in weak interactions[13].

In 1973, the Gargamelle experiment at the European Organisation for Nuclear Research (CERN) released results on the measurement of neutrino interactions[14]. They observed a new type of interaction, neutral–current (NC) interactions,

$$\nu_l + N \rightarrow \nu_l + X,$$

which are characterised by the lack of an observable charged lepton in the final state. Unlike charged–current (CC) interactions, which are mediated by the charged W boson, these NC interactions are mediated by the neutral Z^0 boson.

With the discovery of the tau–lepton in 1977, it was expected that there should be an associated tau neutrino, however, it wouldn't be detected until 2001 by the DONUT experiment[15]. In this experiment, tau neutrinos were produced from the decay of charmed mesons produced in collisions between protons and a stationary target. The neutrino interactions were detected in emulsion detectors, where the unique geometry of the interaction, in which a short tau track is produced at the vertex followed by a long muon track, allowed them to be distinguished from other decays.

While additional neutrino species are possible, data from measurements of the Z boson line–shape at the Large Electron–Positron Collider (LEP) in 1992 restricts the number of active light neutrino species to three[16]. An active light neutrino is any neutrino with $m_\nu < \frac{m_Z}{2}$ that can interact with the Z boson, such that the decay $Z \rightarrow \nu\bar{\nu}$ is possible.

Alongside the discovery of three different types of neutrino, there were interesting results when observing neutrinos produced in the Sun. The flux of neutrinos from the Sun at the Earth's surface had been predicted with Bachall's Standard Solar Model (SSM). However, in 1968 when Davis et al measured the flux in the Homestake experiment, they found a deficit with respect to the prediction of the SSM[17, 18], the so called solar neutrino problem. In the Homestake experiment, electron neutrinos were being measured via their inverse beta decay interactions with the chlorine in the target,

$$v_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e^-.$$

The neutrino interaction rate was measured by counting the number of argon atoms in the chlorine tank, which was achieved by periodically bubbling helium gas through the chamber to capture the argon atoms.

In addition to the solar neutrino problem, a similar deficit was observed in 1988 for muon neutrinos produced during cosmic–ray showers. The Kamiokande experiment was able to measure both electron and muon neutrino interactions, via the Cerenkov radiation produced by the charged leptons in water. The data from Kamiokande was consistent with the expected rate of electron neutrinos from the atmosphere, however, a deficit of muon neutrinos was observed[19].

Super–Kamiokande, the next generation of the Kamiokande experiment, aimed to understand the observed deficit of atmospheric muon neutrinos, by using a larger water Cerenkov detector capable of resolving the angular distribution of atmospheric neutrino interactions. Super–Kamiokande consists of a cylindrical vessel containing 50 kt of ultra pure water, surrounded by an array of around 13,000 photomultiplier tubes to detect the Cerenkov light. Electron and muon neutrinos

can be distinguished based on the pattern of Cerenkov light that is left in the detector. Muons leave clear Cerenkov rings in the detector due to their higher mass, while electrons, which can scatter and shower, tend to leave diffuse "fuzzy" rings on the wall of the detector. In 1998, Super–Kamiokande published measurements of the flux of atmospheric muon neutrinos as a function of azimuthal angle[20]. Since these neutrinos are created a short distance from the Earth's surface, their incoming angle can be used to estimate the distance travelled before they arrive at the detector. The down–going neutrinos have only travelled a short distance in the atmosphere ($\sim$ 10km), while the up–going neutrinos have travelled through the entire Earth to reach the detector ($\sim 13,000$km). Figure 2.1 shows the ratio of the measured neutrino flux in Super–Kamiokande to the prediction in the absence of oscillations, as a function of L/E_ν. The muon neutrino flux is consistent with the no–oscillation prediction at small L/E_ν, however, for large L/E_ν a clear deficit is observed for ν_μ.

While it couldn't isolate the exact cause of the solar neutrino problem, the Sudbury Neutrino Observatory (SNO) was able to provide unique insight into the observed solar neutrino fluxes in 2002. Unlike other water Cerenkov detectors, SNO was filled with heavy water, D_2O, instead of its lighter isotope. The use of heavy water gives rise to additional neutrino interactions, which allowed the SNO experiment to distinguish between three different interaction modes: charged–current (CC), neutral–current (NC), and elastic scattering (ES). Each mode is sensitive to different parts of the solar neutrino flux, including some sensitivity to the muon neutrino and tau neutrino fluxes via the NC and ES interactions. Analysis of the data for each of the three unique interaction modes led to a measurement of the flavour composition of the solar neutrino flux at Earth, whilst also finding the overall neutrino flux at Earth to be consistent with the SSM. Figure 2.2 shows the composition of the solar neutrino flux as measured in the SNO experiment[21], the flux prediction based on the measured rate of NC events is consistent with the predictions of the SSM. The composition of solar neutrinos measured in the SNO experiment is not a result of simple neutrino oscillations, it also depends on the effect of matter on the neutrino propagation in the Sun, via

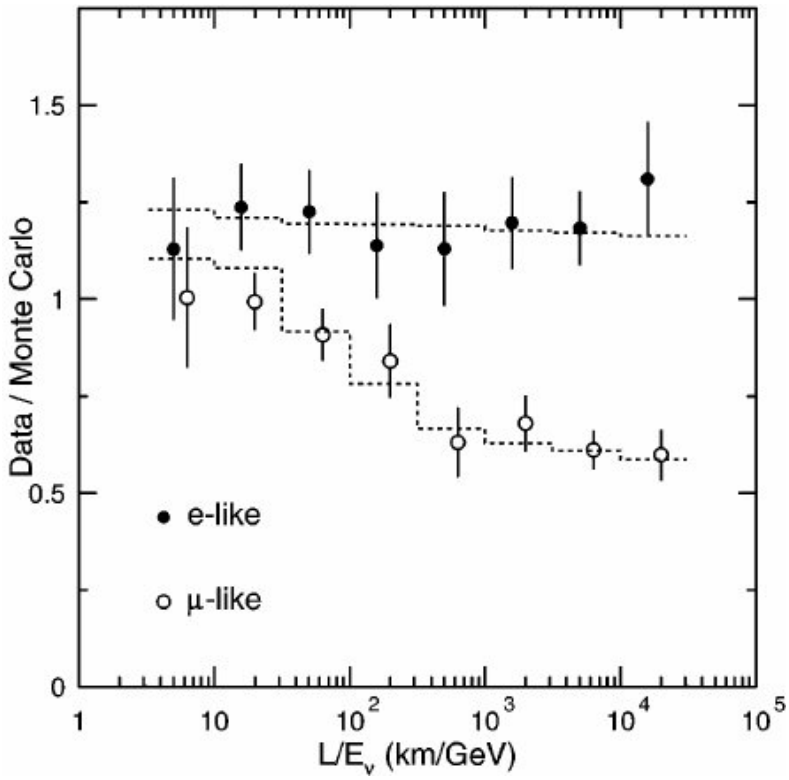


Figure 2.1: Ratio of data to Monte Carlo for electron and muon neutrino fluxes measured by the Super Kamiokande experiment as a function of L/E_ν. The Monte Carlo prediction is based on the assumption of no oscillations. The muon neutrino flux is consistent with the no oscillation prediction at small L/E_ν, however, for large L/E_ν a clear deficit is observed. The best fit under the assumption of atmospheric ($\nu_\mu \rightarrow \nu_\tau$) oscillations is shown, the best fit parameters are $\Delta m^2 = 2.2 \times 10^{-3}\text{eV}^2$, and $\sin^2 2\theta = 1$. Figure from [20].

the Mikheyev–Smirnov–Wolfenstein (MSW) effect. However, at the time a number of solutions to the solar neutrino problem were still considered possible: MSW conversion, decoherence, neutrino decay, and others[22].

An L/E_ν dependence in the neutrino flux would have to be measured in order to prove that neutrino oscillations were a unique solution to the problem. This measurement would require a much shorter baseline, and a detector with a good energy resolution, which would be provided by the Kamioka Liquid Scintillator Anti–neutrino Detector (KamLAND). In 2002, the experiment measured $\overline{\nu_e}$ oscillations from a number of nuclear reactors, which produce neutrinos at the MeV scale[23, 24]. Along with an overall deficit of neutrino events, they were able to use the high energy resolution of the KamLAND detector to measure an L/E_ν dependence of

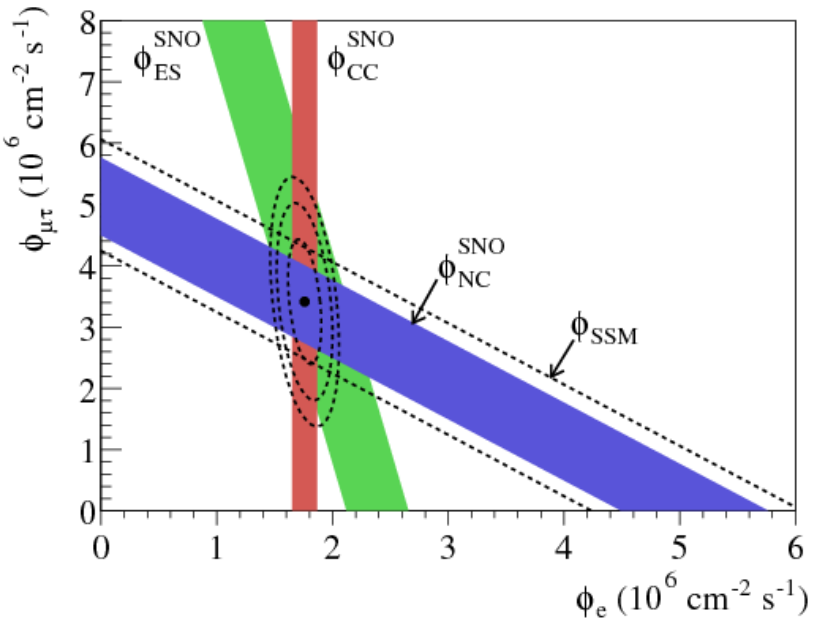


Figure 2.2: Solar neutrino flux composition as measured by the SNO experiment. The coloured bands represent the measured flux of charged–current (CC), neutral–current (NC), and elastic scattering (ES) events, including a $\pm 1\sigma$ spread. The central contours represent 68%, 95%, and 99% probability contours for the joint ϕ_e and $\phi_{\mu\tau}$ fit. The dashed lines represent the predicted flux of ^{8}B neutrinos based on the standard solar model. Figure from [21].

the $\overline{\nu_e}$ survival probability. Figure 2.3 shows the ratio of the observed neutrino flux to the no–oscillation predicted flux as a function of L/E_ν. A clear dependence can be seen, and these data meant that neutrino oscillations were the only solution to the solar neutrino problem with experimental support.

Based on the results of the above experiments, it was assumed that electron and muon type neutrinos oscillate into tau type neutrinos, which go undetected. The first evidence of tau neutrino production in oscillations wouldn't come until 2010, when the OPERA experiment measured a ν_τ candidate in a ν_μ beam. They used similar emulsion detectors to those used to discover the ν_τ in DONUT, and a muon neutrino beam on a 730 km baseline, from CERN to Laboratori Nazionali del Gran Sasso (LNGS). By the end of the experiment a total of 10 candidate events had been observed, 6.1 σ above the expected background[25, 26].

Since the discovery of neutrino oscillations, many more experiments have made measurements of oscillations, and the majority of the parameters of the neutrino

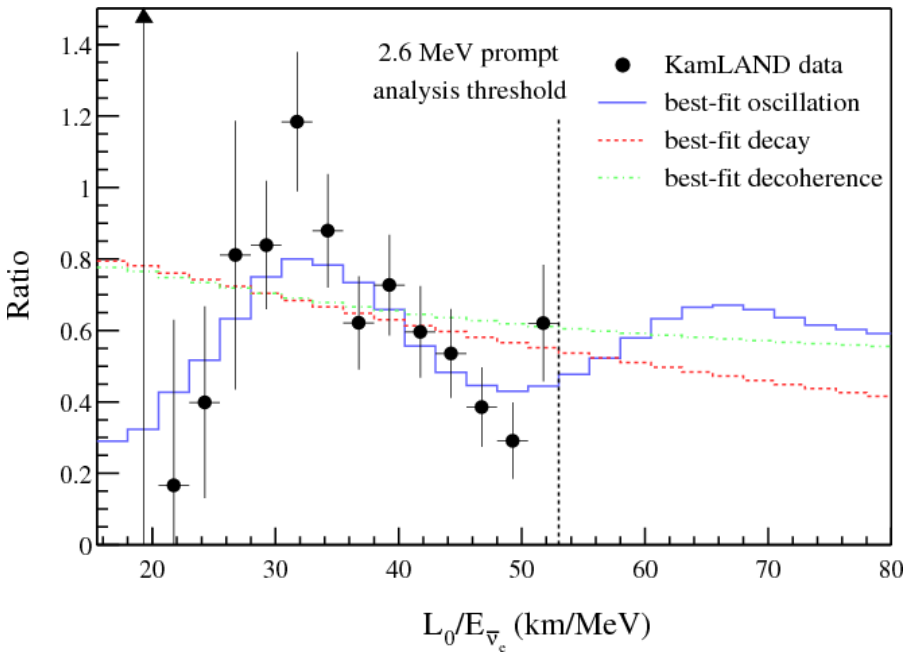


Figure 2.3: Ratio of observed neutrino flux to the predicted flux in the absence of neutrino oscillations in the KamLAND experiment as a function of L/E_ν. The data is fit with three different models: the red dashed line represents the best fit to a neutrino decay model, the green dashed line is for a decoherence model, and the solid blue line represents the best fit of the data to a neutrino oscillation model. The neutrino oscillation model, which has a different shape to the other two models, is found to give the best fit to the data. Figure from [24].

oscillation models have been constrained. Important results of these experiments for constraining the parameters of the neutrino oscillation models will be highlighted in Section 2.3, along with a theoretical overview of neutrino oscillations.

2.2 Neutrinos in the Standard Model

In the standard model, neutrinos form part of the left handed fermion doublets,

$$\psi_i = \begin{pmatrix} \nu_i \\ l_i^- \end{pmatrix},$$

where they are paired with a charged lepton of the same flavour in CC interactions, and i represents any of the three known generations of leptons. Their interaction with other particles in the standard model is determined by the electroweak theory, which is derived from the $SU(2) \times U(1)$ gauge group. The neutrino fields enter

into the SM Lagrangian in the CC and NC interactions:

$$\begin{aligned}\mathcal{L}^{CC} &= -\frac{g_W}{\sqrt{2}}\, j_\alpha^{CC}(x)\, W^\alpha(x) \;+\; h.c. \\ \mathcal{L}^{NC} &= -\frac{g_W}{cos\theta_W}\, j_\alpha^{NC}(x)\, Z^\alpha(x) \;+\; h.c.\end{aligned} \tag{2.1}$$

Here

$$j_\alpha^{CC}(x) = \sum_{\beta=e,\mu,\tau} \overline{\nu_\beta}(x)\, \gamma_\alpha\, \frac{1}{2}\left(1-\gamma^5\right) l_\beta(x)$$

is the leptonic charged–current, and

$$j_\alpha^{NC}(x) = \frac{1}{2}\sum_{\beta=e,\mu,\tau} \overline{\nu_\beta}(x)\, \gamma_\alpha\, \frac{1}{2}\left(1-\gamma^5\right) \nu_\beta(x)$$

is the neutrino neutral–current, $W^\alpha(x)$ and $Z^\alpha(x)$ are the vector boson fields for the $W^\pm$ and Z^0 bosons respectively, g_W is the electroweak coupling constant, and θ_W is the Weinberg angle.

Mass is included in the standard model through the Dirac mass term in the Lagrangian,

$$\begin{aligned}\mathcal{L}^D &= m_D\overline{\psi}\psi \\ &= m_D\overline{(\psi_L+\psi_R)}(\psi_L+\psi_R) \\ &= m_D(\overline{\psi}_L\psi_R + \overline{\psi}_R\psi_L),\end{aligned}$$

where L and R represent the left and right handed components of the field. The lack of right handed neutrino states means that neutrinos are assumed to be massless in the standard model. For massive neutrinos to exists in the standard model, right handed neutrino fields need to be introduced. In addition, it is still not known whether neutrinos are Dirac or Majorana particles, meaning that additional Majorana mass terms are possible. A more general neutrino mass term, including both Dirac and Majorana components is

$$\mathcal{L}^{D+M} = \begin{pmatrix}\bar{\nu}_L\ \bar{\nu}_R\end{pmatrix}\begin{pmatrix}m_L\ m_D \\ m_D\ m_R\end{pmatrix}\begin{pmatrix}\nu_L \\ \nu_R\end{pmatrix}. \tag{2.2}$$

2.3 Neutrino Oscillations

Neutrino oscillations are a result of quantum mechanical interference between different massive neutrino eigenstates. Neutrinos are produced in a state of definite flavour, $\alpha = e, \mu, \tau$, in charged–current (CC) and neutral–current (NC) weak interactions,

$$W^+ \to l^+_\alpha \nu_\alpha, \quad W^- \to l^-_\alpha \overline{\nu_\alpha}, \quad Z \to \nu_\alpha \overline{\nu_\alpha}.$$

The CC processes are used in neutrino oscillation experiments because they give information about the initial flavour state of the neutrinos. These processes are governed by the Lagrangian of the CC leptonic interactions, as in Equation 2.1.

Neutrino flavour states, ν_α, can be represented as a superposition of neutrino mass eigenstates in the form,

$$\nu_\alpha = \sum_k U^*_{\alpha k} \nu_k, \tag{2.3}$$

where ν_k are the neutrino mass eigenstates, and U is a unitary mixing matrix.

The representation of neutrino flavour states as a superposition of mass eigenstates gives rise to the phenomenon of neutrino oscillations. Consider a neutrino produced in a CC weak interaction with flavour α. This neutrino flavour state is described by Equation 2.3, where U is a unitary mixing matrix called the PMNS (Pontecorvo, Maki, Nakagawa, and Sakata) matrix. For three flavour mixing, the PMNS matrix takes the form:

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}.$$

2.3.1 Neutrino Oscillations in Vacuum

In a vacuum, the neutrino mass states are eigenstates of the free particle Hamiltonian

$$\mathcal{H} \left| \nu_k \right\rangle = E_k \left| \nu_k \right\rangle,$$

with energy

$$E_k = \sqrt{\mathbf{p}^2 + m_k^2}.$$

So the solutions to the time dependent Schrodinger equation are plane waves,

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k\rangle .$$

The time evolution of the initial flavour state is:

$$|\nu_\alpha(t)\rangle = \sum_k U^*_{\alpha k} e^{-iE_k t} |\nu_k\rangle . \tag{2.4}$$

The mass states can be written in terms of the flavour states by inverting Equation 2.3,

$$|\nu_k\rangle = \sum_\alpha U_{\alpha k} |\nu_\alpha\rangle , \tag{2.5}$$

where we have used the fact that the states form an orthonormal basis, $\langle\nu_\alpha|\nu_\beta\rangle = \delta_{\alpha\beta}$, and that the transformation matrix is unitary, $UU^\dagger = \mathbf{1}$.

Substituting Equation 2.5 into the time evolution of the flavour state, Equation 2.4, gives:

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_k U^*_{\alpha k} e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle$$

So as the initial flavour state evolves with time, it becomes a superposition of different flavour states; this process is known as neutrino oscillation. The probability of finding the initial neutrino in flavour state ν_β as a function of time is:

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta}(t) &= |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 \\ &= \sum_{kj} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} e^{-i(E_k - E_j)t}. \end{aligned} \tag{2.6}$$

All neutrino oscillation experiments to date operate in a regime where $E \gg m$, in this regime the relativistic energy relation for neutrinos can be expanded as $E_k \simeq E + \frac{m_k^2}{2E}$, where $E = |p|$. Hence,

$$E_k - E_j \simeq \frac{m_k^2 - m_j^2}{2E} = \frac{\Delta m^2_{kj}}{2E}.$$

Another modification to the formula arises because neutrino oscillation experiments do not measure the neutrino propagation time, t. Instead, they measure the propagation distance, L, also known as the baseline. In the ultra–relativistic limit, $t \simeq L$ in natural units, and as a result, Equation 2.6 can be approximated as:

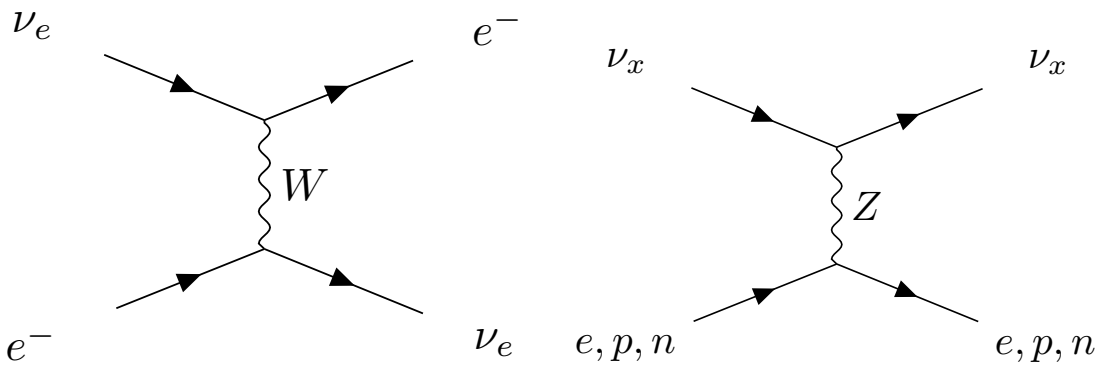
$$P_{\nu_\alpha \to \nu_\beta}(t) = \sum_{kj} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} e^{-i\frac{\Delta m^2_{kj} L}{2E}}.$$

Splitting the sum into its real and imaginary parts emphasises the possibility of charge–parity (CP) violation in neutrino oscillations. CP–violation occurs if the imaginary component of the matrix product is non–zero.

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta}(t) = \delta_{\alpha\beta} &- 4 \sum_{j>k} \mathrm{Re}(U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}) \sin^2(\frac{\Delta m^2_{kj} L}{2E}) \\ &\pm 2 \sum_{j>k} \mathrm{Im}(U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}) \sin^2(\frac{\Delta m^2_{kj} L}{2E}). \end{aligned}$$

The third term here is responsible for CP violation in neutrino oscillations, the positive case corresponds to neutrinos and the negative case to anti–neutrinos.

The probability of oscillation is, therefore, dependent on properties determined by nature, in the form of PMNS matrix elements and mass squared differences, and those which can be chosen by experiments, the distance travelled, L, and the neutrino energy, E. In a vacuum, the oscillation probability is only dependent on the magnitude of the squared mass difference between the neutrino mass eigenstates. There is no information about the absolute masses of the neutrino eigenstates and, therefore, the question of the absolute neutrino masses cannot be answered with neutrino oscillation experiments. To get access to the absolute masses other experiments are required, e.g. direct mass measurements with tritium beta decay experiments such as KATRIN[27]. Neutrino oscillation experiments can give insight on the ordering of the neutrino masses, but the oscillations of neutrinos in matter must be considered in order to achieve this.



(a) Charged–current scattering. **(b)** Neutral–current scattering.

Figure 2.4: Feynman diagrams for neutrino scattering in matter via a charged–current and a neutral–current process.

2.3.2 Neutrino Oscillations in Matter

In real neutrino oscillation experiments neutrinos travel through matter, whether it be the matter in a star or the Earth's crust. Neutrinos propagating in matter are subject to an additional potential due to the interaction of the neutrinos with the electrons and nucleons in the medium. This interaction can significantly modify the flavour of the beam relative to oscillations in vacuum.

In matter, neutrinos can scatter off particles by interacting with either the charged–current or neutral–current interactions, as shown in Figure 2.4. These interactions give rise to additional effective potentials that the neutrinos experience while they are travelling,

$$V_{CC} = \sqrt{2} G_F N_e,$$
$$V_{NC}^f = \sqrt{2} G_F N_f g_V^f,$$

In these formulae, G_F is the Fermi weak coupling constant, N_e and N_f are the number densities of electrons and fermions respectively, and g_V^f is the vector coupling constant for a given fermion f.

Neutral–current scattering is independent of neutrino flavour, meaning that the additional potential does not affect oscillations. However, the charged–current potential is only present for ν_e and, therefore, it will impact the mass eigenstates by different amounts depending on their relative ν_e component.

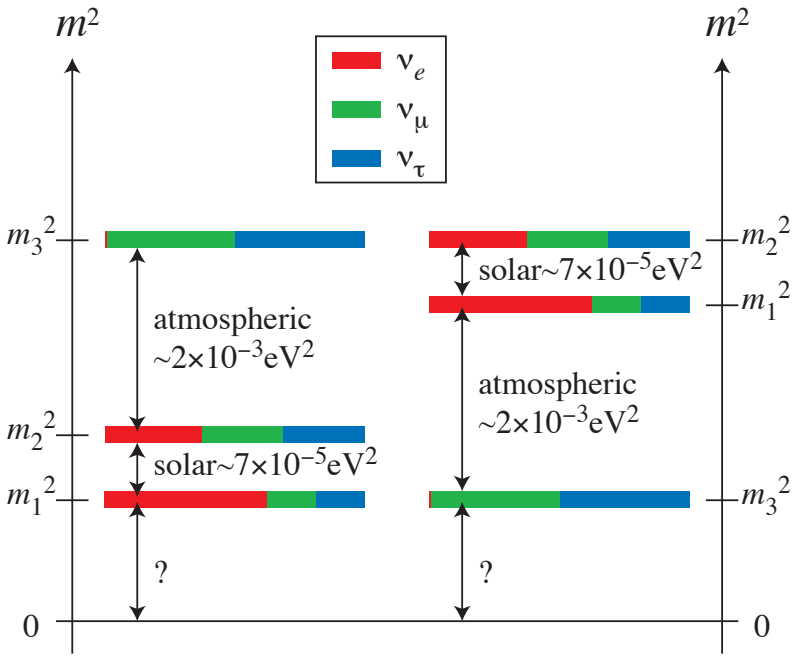


Figure 2.5: The two possible neutrino mass orderings. Left: normal ordering. Right: inverted ordering. Figure from [30].

A full description of the effects of matter on neutrino oscillations is outside the scope of this thesis, although it is discussed at length in other sources such as [28]. For the purposes of this thesis, it is sufficient to note that the propagation of neutrinos in matter changes the oscillation probabilities, and that this must be taken into account by oscillation experiments.

One implication of the effects of matter on oscillations is that it introduces sensitivity to the sign of the mass splittings[28], allowing experiments to determine the relative ordering of the neutrino masses. By considering the oscillation of neutrinos in the sun, it is possible to determine that $\Delta m^2_{21} > 0$[29]. The sign of the remaining mass splitting, Δm^2_{32}, remains unknown, which leaves two possibilities for the ordering of neutrino masses. Normal ordering (NO), in which $m_3 > m_2 > m_1$, and inverted ordering (IO), where $m_2 > m_1 > m_3$. These two possibilities are depicted in Figure 2.5.

2.3.3 Current Knowledge and Open Questions

At the time of writing, the most widely accepted model of neutrino oscillations involves three neutrino mass eigenstates. In this model, the PMNS matrix is

often parametrised in terms of three mixing angles θ_{12}, θ_{13}, and θ_{23} and three CP–violating phases δ_{CP}, α_1, and α_2:

$$U = \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Solar}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix}}_{\text{Cross–mixing}} \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{Atmospheric}} \underbrace{\begin{pmatrix} e^{i\frac{\alpha_1}{2}} & 0 & 0 \\ 0 & e^{i\frac{\alpha_2}{2}} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Majorana}},$$

where $c_{ij} = \cos(\theta_{ij})$ and $s_{ij} = \sin(\theta_{ij})$. Expressing the mixing matrix like this, factorises the matrix into its components, which are responsible for oscillations in different regimes. The solar mixing matrix contains only θ_{12}, which is dominant in solar neutrino oscillations, and the atmospheric mixing matrix, which dominates in the mixing of atmospheric neutrinos, is a function of θ_{23}. The final component that has an impact on neutrino oscillations, is the cross–mixing matrix or reactor matrix, which depends on the final mixing angle, θ_{13}, and one of the CP–violating phases, δ_{CP}. If δ_{CP} is non–zero, then U will have complex components in off diagonal elements, leading to different probabilities for CP flipped oscillations, $P(\nu_\alpha \to \nu_\beta) \neq P(\overline{\nu_\alpha} \to \overline{\nu_\beta})$. Discovery of this effect, which is known as CP–violation, is one of the major goals of the next generation of neutrino oscillation experiments.

The final matrix in the factorised version of the PMNS matrix is called the Majorana component. The CP–violating phases in this matrix cancel out in the oscillation probability, and so they can't be measured in neutrino oscillation experiments. In fact, they only lead to physical effects if neutrinos are Majorana particles (i.e. if they are their own antiparticle). Other experiments are required to determine if neutrinos are Majorana particles, for example, neutrinoless double beta decay experiments such as CUORE[31], NEXT[32], and SNO+[33]. The question of the nature of neutrinos has implications on neutrino mass generation, as mentioned in Equation 2.2.

A large number of neutrino oscillation measurements have now been made by solar, reactor, atmospheric, and accelerator neutrino experiments. When combined, the results of these experiments give us our best estimates of neutrino oscillation

parameters. The latest combined results, from the 2019 Nu–Fit global neutrino oscillation analysis[34] are summarised below. Each measurement is given for both the normal hierarchy and the inverted hierarchy cases, and the major contributing experiments are highlighted in each case.

θ_{12}

The constraints on the solar mixing angle, θ_{12}, are dominated by a combination of data from solar neutrino experiments, such as SNO[21] and Super Kamiokande[35], with data from the KamLAND experiment[24]. The current constraint is,

$$\begin{aligned} \sin^2(\theta_{12}) &= 0.310\ ^{+0.013}_{-0.012} \quad \text{(NO)} \\ &= 0.310\ ^{+0.013}_{-0.012} \quad \text{(IO)} \end{aligned}$$

Δm^2_{21}

The best measurement of Δm^2_{21} is also dominated by a combination of solar neutrino data with the results from KamLAND. The measured value is,

$$\begin{aligned} \Delta m^2_{21} &= (7.39\ ^{+0.21}_{-0.20}) \times 10^{-5}\ \text{eV}^2 \quad \text{(NO)} \\ &= (7.39\ ^{+0.21}_{-0.20}) \times 10^{-5}\ \text{eV}^2 \quad \text{(IO)} \end{aligned}$$

θ_{23} and Δm^2_{32}

There is a strong correlation between θ_{23} and Δm^2_{32} and, therefore, their measurements are usually presented as a two–dimensional contour. Figure 2.6 shows a comparison of the world–leading contours for $\sin^2(\theta_{23})$–Δm^2_{32}, with the tightest error bands coming from long baseline accelerator experiments such as T2K, MINOS, and NOνA[36–38]. The results are dependent on the neutrino mass ordering; based on the Nu–Fit global three neutrino oscillation analysis[34],

$$\begin{aligned} |\Delta m^2_{32}| &= (+2.525\ ^{+0.033}_{-0.031}) \times 10^{-3}\ \text{eV}^2 \quad \text{(NO)} \\ &= (-2.512\ ^{+0.034}_{-0.031}) \times 10^{-3}\ \text{eV}^2 \quad \text{(IO)}, \end{aligned}$$

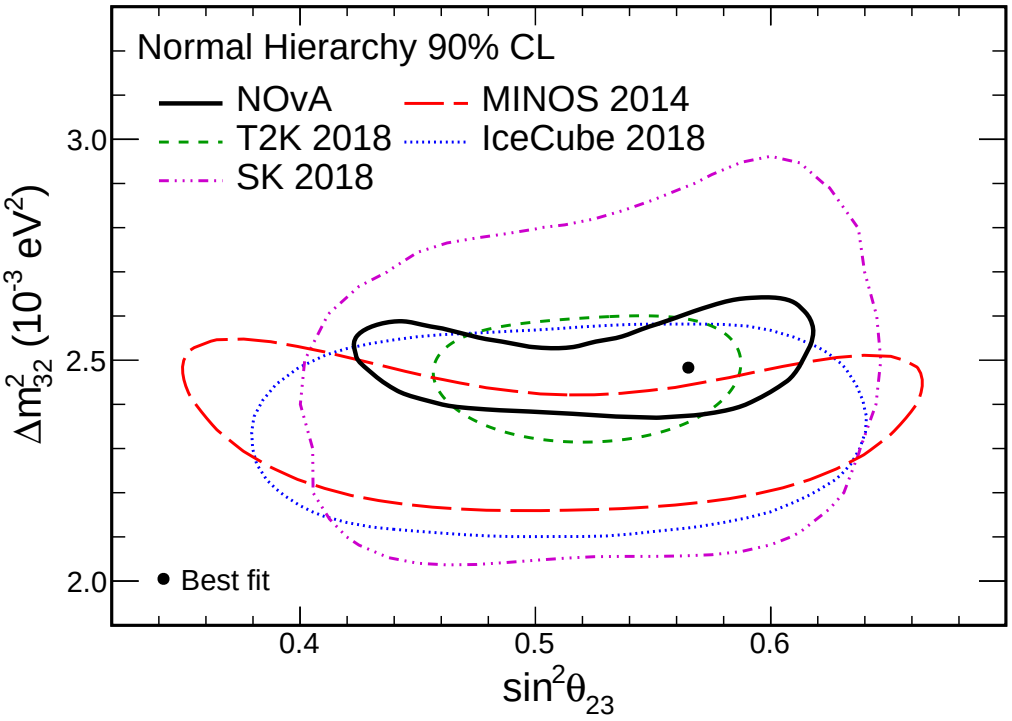


Figure 2.6: The 90% confidence region contours for $\sin^2(\theta_{23})$–Δm^2_{32} from a number of the leading neutrino oscillation experiments[36–38]. The best fit point for the NOνA experiment is shown as a black dot, and is based on data up to 2019. Figure from [38].

and

$$\sin^2(\theta_{23}) = 0.582\ ^{+0.015}_{-0.019} \quad \text{(NO)}$$
$$= 0.582\ ^{+0.015}_{-0.018} \quad \text{(IO)}.$$

θ_{13}

Reactor neutrino experiments such as Daya Bay[39], Double Chooz[40], and RENO[41] have made precise measurements of θ_{13}, as well as some long baseline oscillation experiments, such as T2K and NOνA. The Nu–Fit three neutrino global fit to the reactor data gives,

$$\sin^2\theta_{13} = 0.02240^{+0.00065}_{-0.00066} \quad \text{(NO)}$$
$$= 0.02263^{+0.00065}_{-0.00066} \quad \text{(IO)}$$

Mass Ordering

The long baseline oscillation experiments T2K and NOνA have both shown a preference for normal mass ordering[36, 38]. When these results are combined in the Nu–Fit three–neutrino oscillation analysis, inverted ordering is currently disfavoured with a $\Delta\chi^2$ of 9.3.

δ_{CP}

While there are no accurate measurements of δ_{CP}, there are some hints from long baseline accelerator neutrino experiments T2K and NOνA, that it may be non–zero. These limits are based on joint fits to four data samples: electron neutrino appearance, electron anti–neutrino appearance, muon neutrino disappearance, and muon anti–neutrino disappearance.

The T2K experiment's joint fit shows an excess of electron neutrino events, and a deficit of anti–electron neutrino events, when compared to the predictions for $\delta_{CP} = 0$. This results in a preference for negative values of δ_{CP} with a 3σ confidence interval of $[-3.41, -0.03]$ in the case of normal neutrino mass ordering. The results of the T2K fit are shown in Figure 2.7[42].

The NOνA experiment has also performed joint fits to neutrino and anti–neutrino oscillation data, the results of the fit are shown in Fig. 2.8. As with T2K, the normal mass ordering is preferred, however, for NOνA the full range of δ_{CP} is covered at 3σ, highlighting the need for further study of CP–violation in neutrino oscillations[38].

2.4 Neutrino Interactions

To perform a neutrino oscillation experiment the composition of the beam needs to be measured, and the change in beam composition as a function of energy and distance travelled is analysed. In practice, this relies on measuring neutrino events and categorising them by flavour in order to compare the measurement to prediction. However, because the detectors can only see the neutrinos that interact, the result is actually a convolution of the neutrino beam composition with the

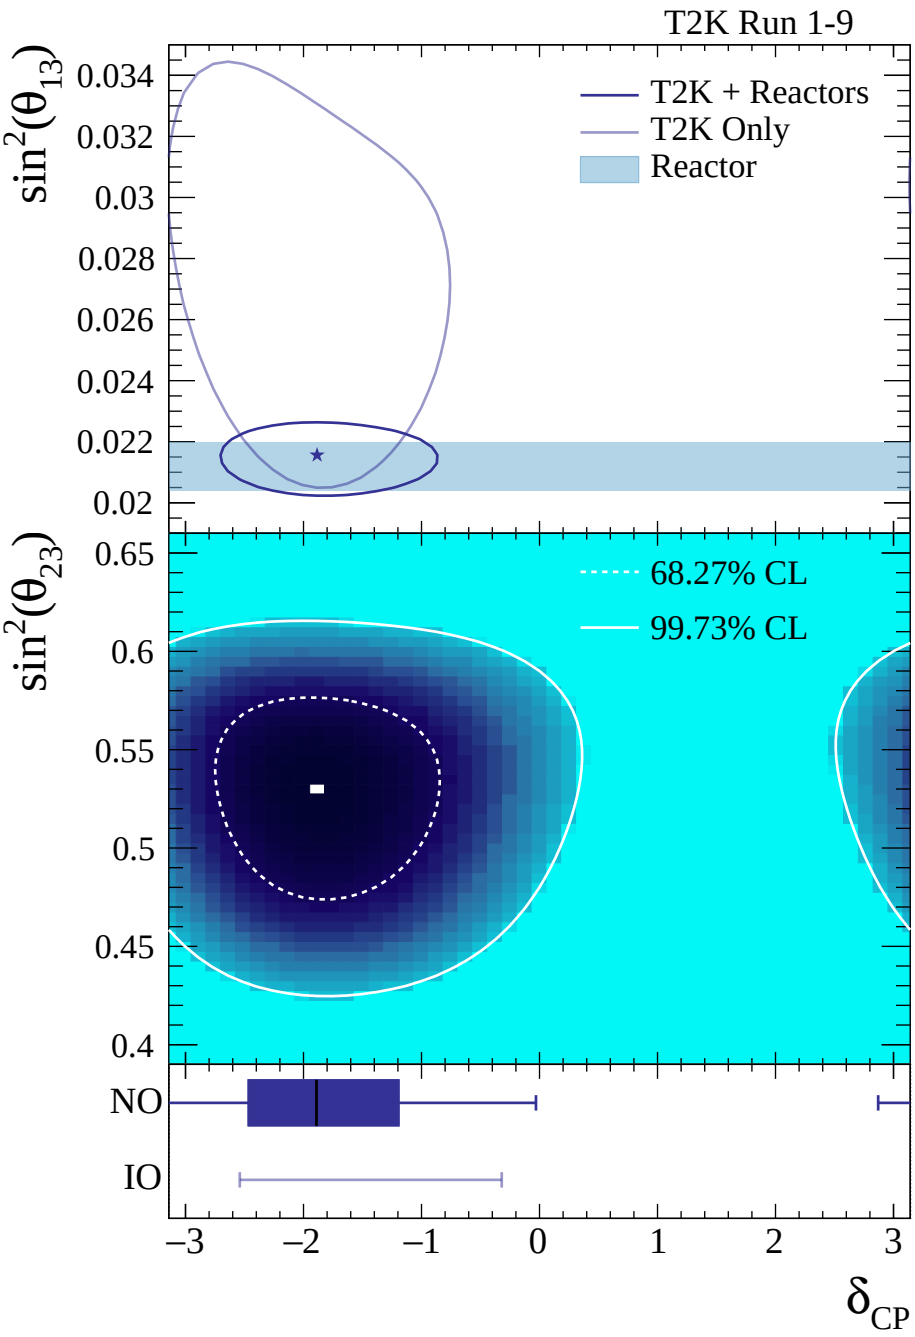


Figure 2.7: Confidence intervals for δ_{CP} from the T2K experiment. Figure from [42].

Top: 68.27% confidence level contours for δ_{CP} versus $\sin^2\theta_{13}$ under the assumption of normal ordering.

Middle: Confidence intervals at the 68.27% and 99.73% confidence level for δ_{CP} versus $sin^2\theta_{23}$ from a fit to T2K and reactor data under the assumption of normal ordering.

Bottom: Confidence intervals for δ_{CP} from a fit to T2K and reactor data for both the normal and inverted orderings. The vertical line in the shaded box shows the best-fit value of δ_{CP}, the shaded box shows the 68.27% confidence interval, and the error bar shows the 99.73% confidence interval.

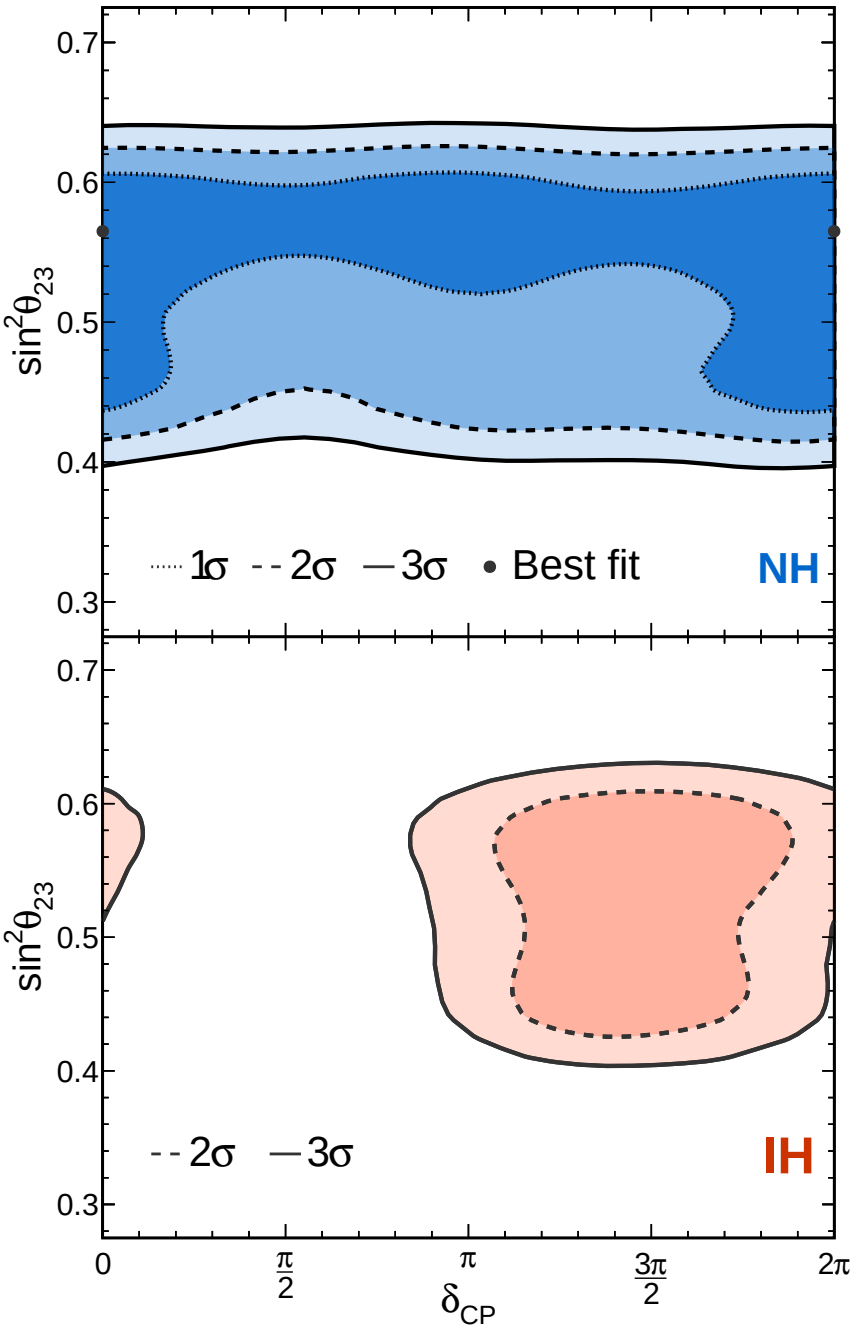


Figure 2.8: Confidence intervals for δ_{CP} from the NOνA experiment. Figure from [38].

Top: Confidence interval for δ_{CP} versus $\sin^2\theta_{23}$ under the assumption of normal ordering.

Bottom: Confidence interval for δ_{CP} versus $\sin^2\theta_{23}$ under the assumption of inverted ordering.

neutrino interaction cross section. Therefore, it is important to understand neutrino interaction cross sections in order to make accurate oscillation predictions.

Broadly speaking, there are two major types of neutrino interactions: charged–current (CC) and neutral–current (NC). CC interactions are usually used in oscillation experiments because they are the only type of interaction that allows the flavour of the neutrino to be determined. Some NC interactions, e.g. scattering from an electron, produce high energy leptons in the detector and, therefore, form an irreducible background. These backgrounds must be modelled as part of the experiment's simulation.

The CC interactions are further split into three main categories.

Quasi–elastic (QE)

A neutrino scatters off an individual nucleon. Depending on the energy transfer, the nucleon may be liberated from the nucleus.

Resonance (RES)

A neutrino excites the target nucleon into a resonance that then decays resulting in the possibility of mesons in the final state.

Deep Inelastic Scattering (DIS)

The neutrino has enough energy to resolve the individual quarks within the target nucleon, liberating a quark and resulting in a hadronic final state.

The cross sections for these processes vary as a function of neutrino energy and flavour, and they are each dominant in different energy regimes. Predictions and measurements of the charged–current cross section for ν_μ are shown in Figure 2.9. The three main components of the cross section are shown, and the energy range relevant for DUNE is highlighted[43]. Within the energy range of the DUNE beam, all three of the major neutrino cross section components have a region in which they dominate. This means that understanding each of these cross sections is crucial for any neutrino oscillation measurement in DUNE.

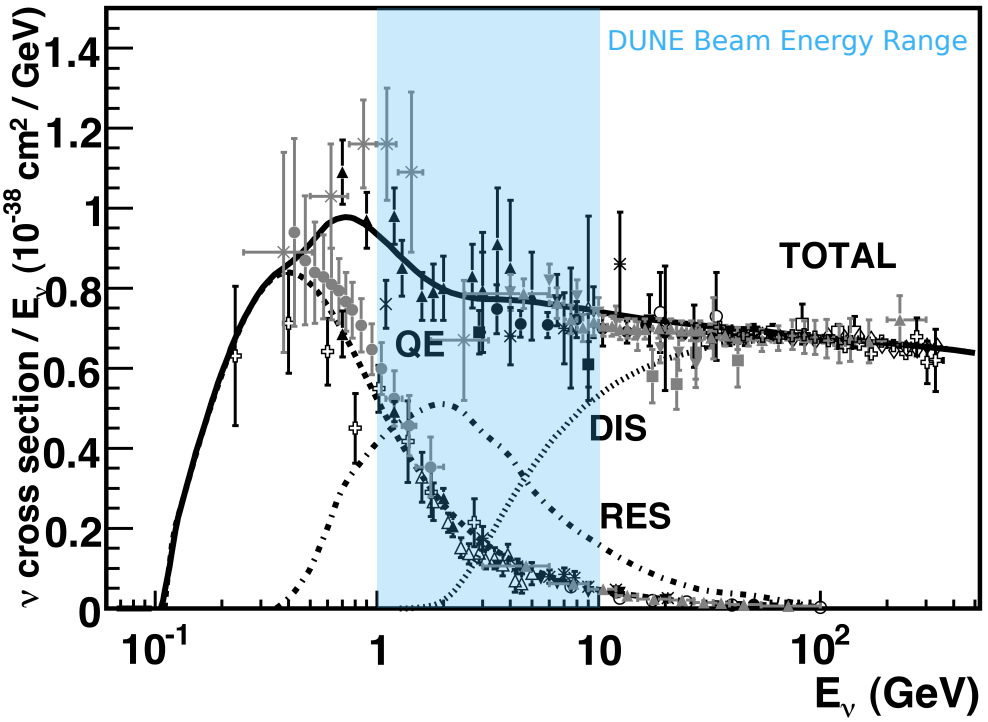


Figure 2.9: Muon neutrino cross section per nucleon, per neutrino energy for an isoscalar target as a function of neutrino energy. The light blue region represents the region of interest for the DUNE neutrino beam. Original figure from [43].

The array of interaction modes available in DUNE means that there are many particles that can be produced in the final state; understanding the composition of the final state is essential to the neutrino oscillation analysis. Part of this thesis, Chapter 6, looks at the identification of charge deposits in a LArTPC. The results provide input in analyses of ProtoDUNE–SP data, and the methods used could be adapted and developed for application in DUNE.

2.5 Supernova Neutrinos

Supernovae are extremely violent explosions that are undergone by certain types of stars at the end of their life. These explosions can emit on the order of 10^{46} Joules (10^{65} eV) of energy, and in certain cases, known as core–collapse supernovae, about 99% of this energy is carried away by neutrinos. Measurements of these neutrinos can provide insight into the mechanism involved in supernova bursts, as well as the study of neutrino masses[28].

2.5.1 Core–collapse Supernova Dynamics

Core–collapse supernovae occur in stars with masses of around 10–60 solar masses. These stars will have undergone all stages of nuclear fusion during their life, but since iron is the most tightly bound nucleus there is no fuel left to burn after it has been produced. At this stage, the iron core of the star, which has a mass on the order of 1 solar mass, begins to collapse as the pressure produced by nuclear fusion is no longer enough to counter the force of gravity.

During the collapse of the core, electron neutrinos are produced through electron capture on both nuclei and free protons in the plasma at the core of the star,

$$e^- + N(Z, A) \rightarrow N(Z-1, A) + \nu_e$$
$$e^- + p \rightarrow n + \nu_e.$$

At first, the mean free path of these neutrinos is much longer than the size of the core, and the neutrinos leave the core carrying away their energy. This phase of the collapse is known as the infall phase, it lasts on the order of 10 ms and releases neutrinos with energies of around 12–16 MeV.

Once the density of the core increases beyond around 3×10^{11} g / cm^3 neutrinos become trapped in the core, as the cross section for coherent scattering becomes large enough to prevent neutrinos from passing through the core. Neutrinos are still being produced by capture processes at this point, but they are unable to leave the core.

The core–collapse comes to a sudden halt about 1 second after it begins, when the density of the inner core reaches that of nuclear matter. At this stage, the core settles into equilibrium as a proto–neutron star while a shock–wave caused by this sudden halt propagates out through the layers of the star. Behind the shock, neutrino production is accelerated as nuclei are dissociated and the free protons capture electrons. Neutrinos begin to build–up behind the opaque shock. A few milliseconds after the bounce, the shock reaches a region of low enough density and becomes transparent, releasing the build–up of neutrinos in just a few milliseconds, this is known as the neutronisation burst.

One aspect of supernova dynamics that is still unclear is the so–called revival of the shock. As the shock dissipates through nucleon dissociation and neutrino emission it becomes weakened, and eventually stalls around 100 ms after the initial bounce. During this time, known as the accretion phase, matter falls onto the collapsed core. If the shock cannot be revived a supernova will not occur. It is thought that neutrino flux produced in the proto–neutron star is able to revive the shock, but the precise mechanics of this revival are still unknown. Studies have shown that the impact of spherical asymmetries might play an important role in allowing the burst to take place[44].

After the neutronisation burst, the main source of neutrino production comes from the proto–neutron star. Neutrinos of all flavours are produced in the core at a temperature of around 40 MeV. The release of neutrinos at this stage, known as the neutrino cooling phase, is significantly slower than that of the neutronisation burst and can last tens of seconds. During this phase, much like photons within a star, the neutrinos are mostly contained within the opaque environment of the proto–neutron star. They can only escape if they travel far enough from the core to a region where the opacity is sufficiently low. This region is called the neutrinosphere, it depends on flavour, and it is different for neutrinos and anti–neutrinos. These differences result in different energy distributions for each neutrino type, owing to the relative temperature of the star at the radius of each neutrinosphere,

$$\langle E_{\nu_e} \rangle \approx 10\text{MeV}$$
$$\langle E_{\overline{\nu_e}} \rangle \approx 15\text{MeV}$$
$$\langle E_{\nu_x} \rangle \approx 20\text{MeV}$$

where $\nu_x \in [\nu_\mu, \overline{\nu_\mu}, \nu_\tau, \overline{\nu_\tau}]$.

2.5.2 SN1987A

The first, and currently only, experimental observation of neutrinos from a supernova burst occurred in 1987. A small number of low–energy neutrino events were detected in coincidence with a supernova burst from the Large Magellanic Cloud, referred to

as SN1987A. Three neutrino detectors reported an excess of low energy events in coincidence with the supernova: Kamiokande–II, IMB, and Baksan[45–47]. These events were primarily produced by two types of interaction, inverse beta decay (IBD),

$$n + \nu_e \rightarrow p + e^+,$$

and elastic scattering,

$$\nu_e + e^- \rightarrow \nu_e + e^-.$$

For supernova neutrinos, the inverse beta decay cross section is significantly higher than that of elastic scattering, therefore, most events are likely to be due to inverse beta decay.

2.5.2.1 Kamiokande–II

Kamiokande–II was a water Cerenkov detector containing 2.1 kt of water. A significant increase in the rate of electron events with respect to background was observed, for a 10 second window coincident with SN1987A[45]. Around 12 events were observed, their time sequence and amplitudes are shown in Figure 2.10. Unfortunately, due to an inaccurate detector clock, the absolute time of these events is only accurate to about one minute and, therefore, a time coincidence check with the other neutrino observations cannot be made.

2.5.2.2 IMB

IMB was another water Cerenkov detector, with a fiducial mass of 3.3 kt. They observed eight neutrino candidate events, with energies in the range of 20–40 MeV within a six second window. The background rate was estimated to be around 2 events per day[46].

2.5.2.3 Baksan

Unlike Kamiokande–II and IMB, the Baksan detector was a segmented liquid scintillator detector with a total mass of around 330 tons. However, only a limited fiducial mass of around 200 tons was used for the supernova neutrino measurements.

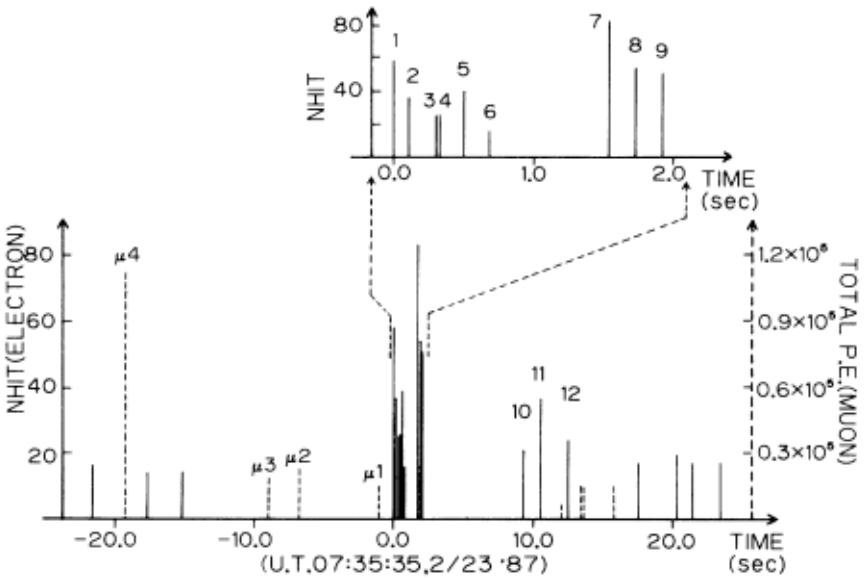


Figure 2.10: Measured time of supernova neutrino events from SN1987A in Kamiokande–II. Electron events are plotted as solid vertical lines, where the height of the line represents the amplitude of the event in terms of number of PMT hits. Muons are represented as dashed vertical lines, and the height of the line represents the total number of measured photo–electrons. A clear excess of electron events can be seen around the centre, this region is expanded above the main axis of the figure. The supernova neutrino candidates are labelled with a number above the solid vertical line. Figure from [45].

No excess above background could be observed by Baksan in isolation; however, when input from Kamiokande–II and IMB is used, a cluster of 5 events within a 10 s window was observed coincident with the measurements from IMB[47].

Since SN1987A, significant progress has been made for both experimental and theoretical aspects of neutrino physics. Progress has been made in modelling supernova explosions, and the impacts of neutrino oscillations on the supernova environment have been considered[48]. The current and next generation of neutrino experiments expect to measure thousands of neutrino events for a galactic supernova, instead of the tens measured by the previous generation. The flux of neutrinos at earth has three components, ν_e, $\overline{\nu_e}$, and ν_x, the predicted energy spectra of each component is shown in Figure 2.11. Measuring the spectrum and time structure of a supernova neutrino burst could provide insights into the dynamics of supernovae, as well as neutrino physics.

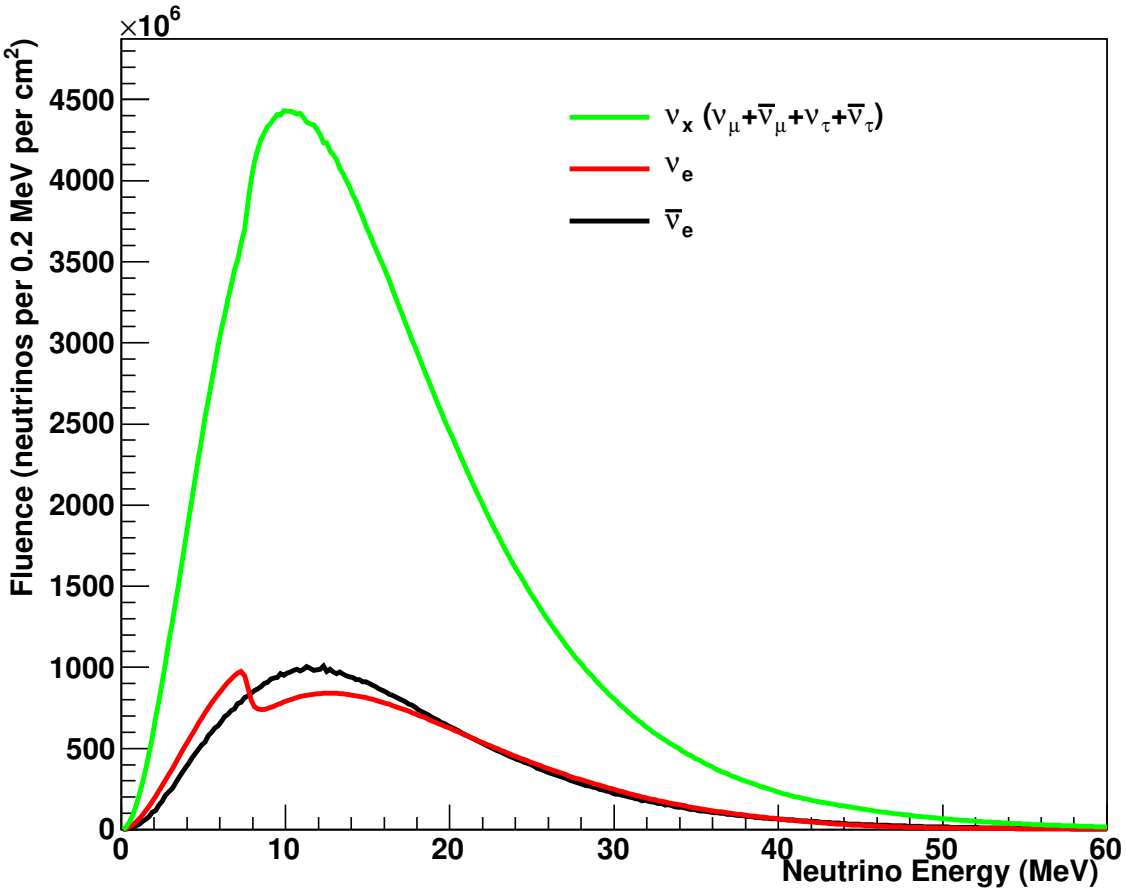


Figure 2.11: Supernova neutrino flux at earth as a function of neutrino energy. The neutrino flux is split into three components, ν_e, $\overline{\nu_e}$, and ν_x, which are plotted in red, black, and green respectively. Figure from [3].

2.5.3 Supernova Neutrino Prospects in DUNE

The current world leading supernova neutrino detector is Super–Kamiokande, which expects to measure around 8000 events for a supernova at 10 kpc[49]. As with Kamiokande–II, these events would primarily be produced via inverse beta decay interactions on the protons within the water. As such, Super–Kamiokande is mostly sensitive to the $\overline{\nu_e}$ component of the supernova neutrino flux. This is common amongst all water Cerenkov and scintillator based detectors, which make up the majority of the current landscape of supernova neutrino detectors. Therefore, DUNE makes a highly complementary detector to the other supernova neutrino detectors, as a result of the large CC cross section for ν_e on argon. In particular, DUNE offers a unique sensitivity to the neutronisation burst, which is primarily made up of ν_e[48].

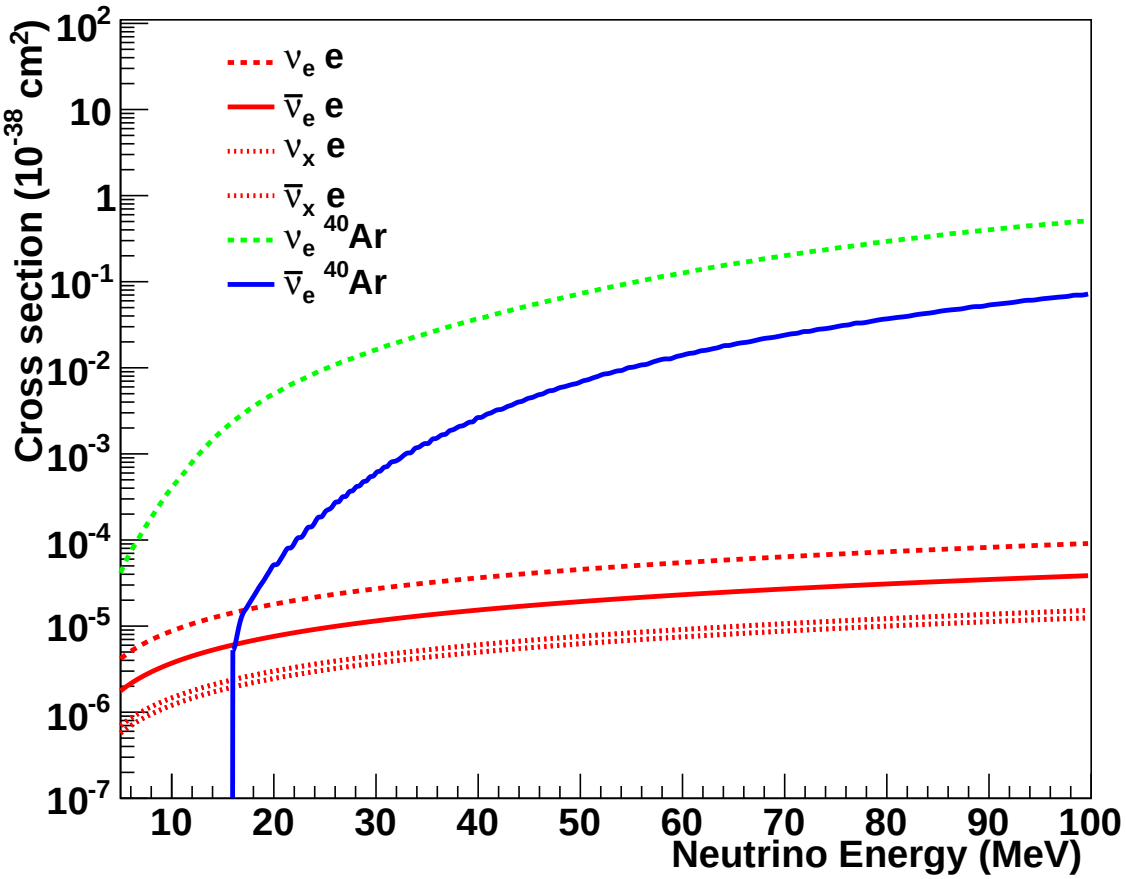


Figure 2.12: Neutrino cross sections in liquid argon as a function of neutrino energy, in the energy regime relevant for supernova neutrinos. The cross section is split into a number of components, depending of the properties of the neutrino and the target, which are labelled in the top left corner of the figure. Figure from [50].

2.5.3.1 Supernova Neutrino Interactions in Liquid Argon

Liquid argon should bring a strong sensitivity to the ν_e component of a supernova neutrino burst, via the charged–current absorption of ν_e on ^{40}Ar,

$$\nu_e + ^{40}\text{Ar} \rightarrow e^- + ^{40}\text{K}^*.$$

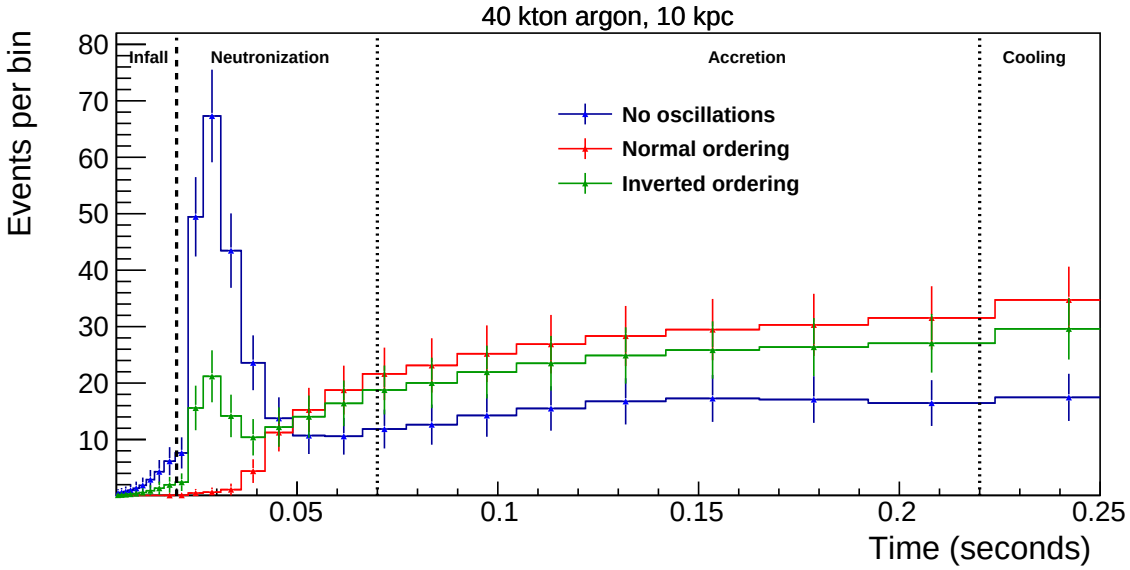
This interaction leaves an e^- and an excited K in the final state. In addition, there are charged–current $\bar{\nu_e}$ interactions with argon, and elastic scattering interactions with electrons, which are available to neutrinos of all flavours. The dominant cross sections in liquid argon as a function of energy are shown in Figure 2.12[50].

2.5.3.2 Supernova Neutrino Events in DUNE

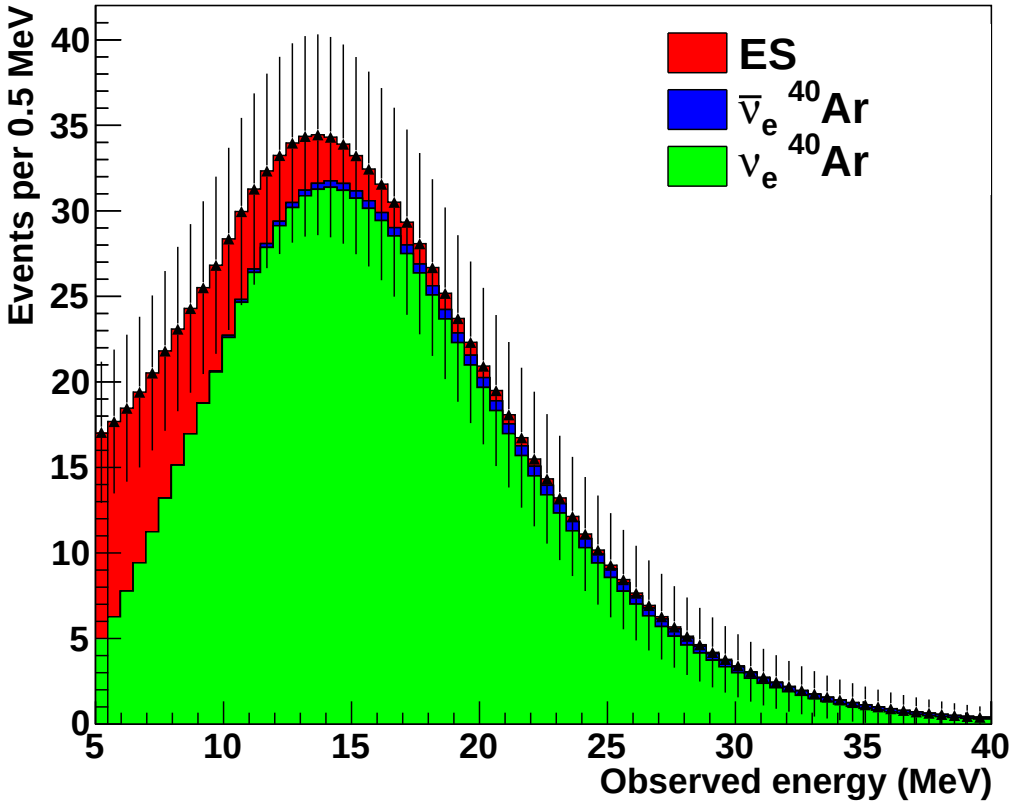
DUNE will be capable of measuring neutrinos from any core–collapse supernova within the Milky Way, and if any should occur during the lifetime of the experiment the measured neutrino events will provide unprecedented data with which to measure the dynamics in the early stages of a core–collapse supernova. The rate of core–collapse supernovae within the Milky Way has been estimated to be 1.9 ± 1.1 per century[51], and therefore there is a reasonable chance that one will occur during the lifetime of the DUNE experiment.

For a 10 kpc supernova, DUNE expects to observe roughly 3000 neutrino events, over a period of around 10 seconds. In DUNE, this will show up as a sudden increase in the rate of low–energy electron events in the detector. Figure 2.13 shows the predicted time structure and energy spectrum for a simulated supernova event in DUNE[50]. The details of the rate and energy of these events as a function of time can hold potential insights into both neutrino physics and the dynamics of supernova bursts.

Detecting and reconstructing electrons with O(10 MeV) of energy is necessary to study supernova neutrinos, as can be seen from the observed energy spectrum for supernova neutrinos in Figure 2.13. There are significant challenges involved in this measurement, as a result of the difficulties involved in the detection of low energy activity and the reconstruction of low–energy electrons. Effectively reconstructing electrons with O(10 MeV) of energy is essential to improving our understanding of supernova neutrinos in DUNE. This is discussed in Chapter 7, using the case of Michel electrons in ProtoDUNE–SP to benchmark the performance of low energy electron reconstruction in liquid argon time projection chambers.



(a) Simulated observed event rate for the first 0.25 seconds of the supernova neutrino burst. The predictions under the assumption of no oscillations, normal neutrino mass ordering, and inverted neutrino mass ordering are shown in blue, red, and green respectively. Figure from [50].



(b) Simulated observed supernova neutrino energy spectra for the full supernova neutrino burst. The contributions from elastic scattering, $\overline{\nu_e}$, and ν_e events are shown in red, blue, and green respectively. Figure from [52].

Figure 2.13: Predicted supernova neutrino event rate and energy spectra for a simulated 10 kpc core collapse supernova measured by DUNE.

3

The ProtoDUNE–SP Experiment

Contents

3.1 ProtoDUNE–SP in the Context of DUNE 36
3.1.1 Liquid Argon Time Projection Chambers 37
3.1.2 ProtoDUNE–SP Goals 39
3.2 The ProtoDUNE–SP Detector 40
3.2.1 The Liquid Argon TPC 41
3.2.2 The Photon Detection System 43
3.3 The H4 beamline 45
3.4 The Cosmic–Ray Tagger 46
3.5 The Data Acquisition System 48
3.6 Simulation 50
3.7 Reconstruction 51
3.8 ProtoDUNE–SP Online Monitoring System 57
3.8.1 Data Processing 58
3.8.2 Monitoring Web Interface 64
3.8.3 System Performance and Possible Improvements 69

ProtoDUNE–SP is one of two prototypes for the DUNE Far Detector modules, which has been operating at the Neutrino Platform at CERN since the summer of 2018. The experiment collected data from a charged particle beam for approximately 3 months before Long Shutdown 2 of the Large Hadron Collider (LHC). Since then, a programme of cosmic–ray data collection has been ongoing.

This chapter will outline the technical details of the ProtoDUNE–SP experiment.

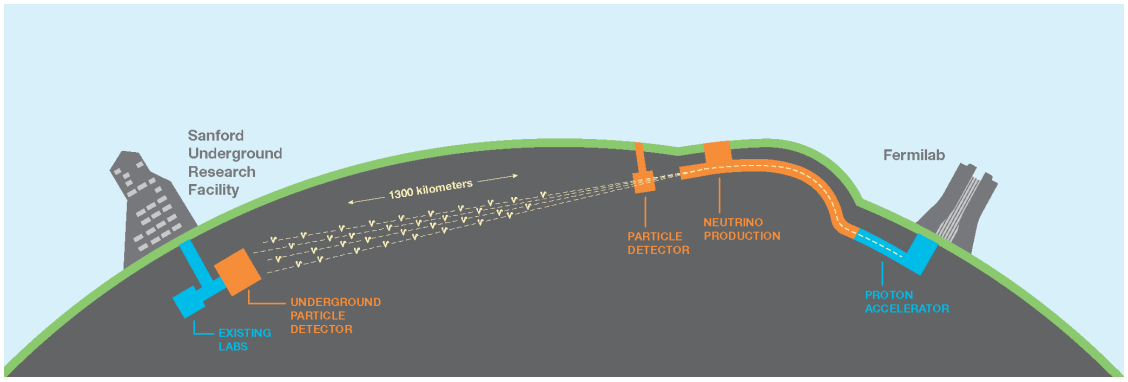


Figure 3.1: The Deep Underground Neutrino Experiment. Figure from [53].

Section 3.1 will outline the role of ProtoDUNE–SP in the context of the DUNE experiment. This will be followed by a discussion of the main elements of the ProtoDUNE–SP detector systems and the H4 beamline, in Sections 3.2 and 3.3 respectively. The tagging of cosmic–ray muons for calibration is done by the Cosmic–Ray Tagger (CRT), which will be discussed in Section 3.4. The data acquisition system responsible for handling the data from each detector component will be discussed in Section 3.5. Sections 3.6 and 3.7 will then discuss the simulation and reconstruction of ProtoDUNE–SP data. Finally, Section 3.8 will cover details of the online monitoring system in ProtoDUNE–SP. During my time at CERN, I was the primary developer and expert on the ProtoDUNE–SP online monitoring system. The development and maintenance of this system represents a significant body of work over a 12 month period, as well as ongoing maintenance work throughout the following 2 years.

3.1 ProtoDUNE–SP in the Context of DUNE

The DUNE experiment will be a next generation neutrino physics and nucleon decay experiment consisting of three principal components; an intense broad band neutrino beam, a precise Near Detector complex, and a Far Detector. The neutrino beam and Near Detector complex will be based at the Fermilab National Accelerator Laboratory (FNAL) near Chicago, and the Far Detector will be based at Sanford Underground Research Facility (SURF) in South Dakota, approximately 1300

km away from the neutrino source. A schematic of the experimental setup is shown in Figure 3.1.

The DUNE experiment identifies three primary scientific goals, which are stated in the DUNE technical design report[50]:

- Perform a comprehensive programme of neutrino oscillation measurements including measurements of δ_{CP}, neutrino mass ordering, and θ_{23}
- Search for proton decay in several decay modes
- Measure neutrinos from a core–collapse supernova, if one occurs within our galaxy during the lifetime of the experiment

In addition, the technical design report also highlights a significant programme of secondary science goals:

- Other accelerator based neutrino physics, such as non–standard interactions, sterile neutrinos, and CPT violation
- Measurements of neutrino properties using atmospheric neutrinos
- Dark matter searches in both the Near and Far Detectors.
- A programme of neutrino interaction physics studies in the DUNE Near Detector.

To achieve these goals, DUNE has opted to base the Near and Far Detector designs on the liquid argon time projection chamber (LArTPC) technology. The DUNE Far Detector will eventually consist of four LArTPC detectors, each with 10 kt of active liquid argon mass. This technology will have never before been used on this scale and, therefore, there has been a significant ongoing programme of LArTPC research and development, to validate and characterise the performance of the technology for DUNE.

3.1.1 Liquid Argon Time Projection Chambers

A LArTPC consists of a large volume of highly–purified liquid argon that is subjected to an electric field. Charged particles traversing the liquid argon produce two primary energy depositions: a trail of ionisation electrons along their path, and prompt ultra–violet scintillation photons. After creation, the ionisation electrons drift in the

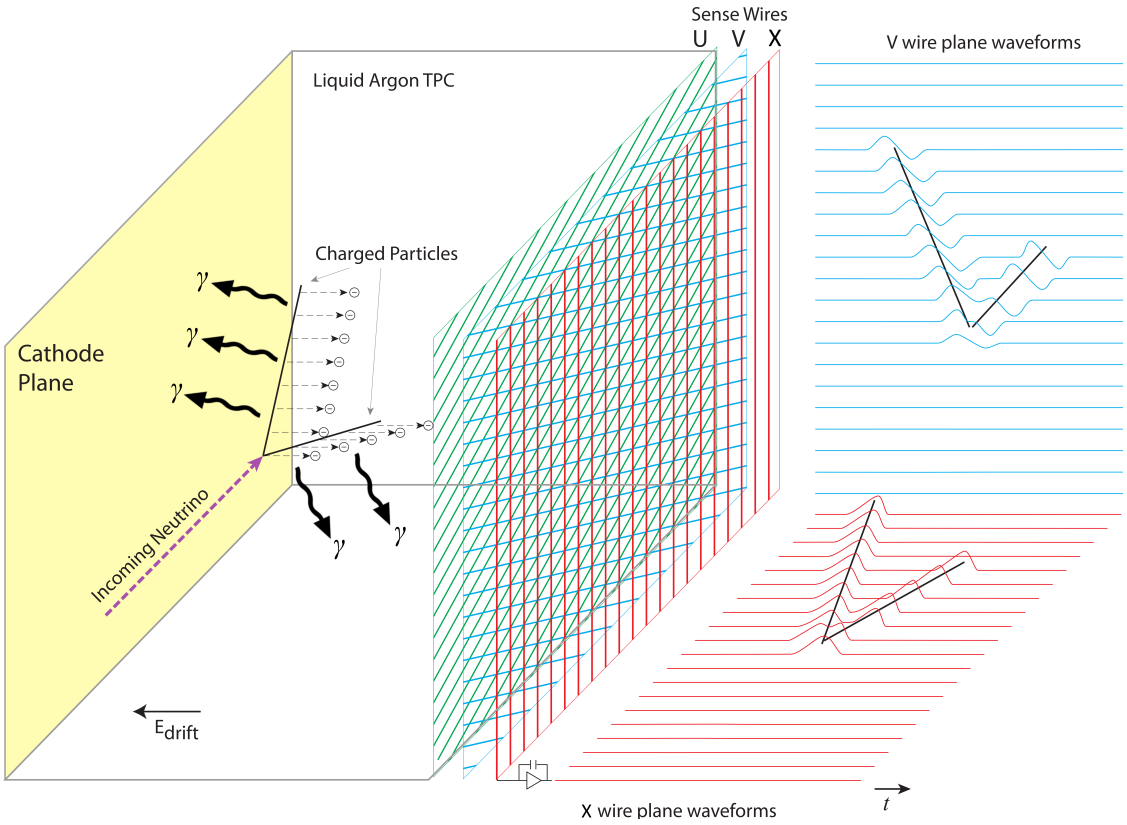


Figure 3.2: A schematic of the LArTPC detection principle. Charged particles, in this case resulting from a neutrino interaction, produce two signals in the liquid argon, ionisation electrons and scintillation light. The ionisation electrons are drifted in the applied electric field and collected on a set of sense wires that record the signal. The scintillation photons are collected by photon detectors, which are not depicted. Figure from [54].

electric field toward the charge readout plane, where they induce electrical signals. Liquid argon is transparent to its own scintillation light, therefore, the scintillation photons can travel through the argon to be collected in a photon detection system. The LArTPC detection principle is illustrated in Figure 3.2.

The Role of Light in LArTPCs

The ionisation signals in a LArTPC are slow, it takes electrons milliseconds to travel from the cathode plane to the anode plane. In contrast, scintillation photons only take nanoseconds to reach the closest anode plane. This scintillation light plays an important role in the accurate 3D reconstruction of interactions in the LArTPC, as it can be used to determine the time of an interaction to within a few nanoseconds.

In a LArTPC, interactions play out much more quickly than the detector is able to record them, and each event actually integrates over a large number of interactions within the readout window. The true time, often referred to as the t_0, of the interactions in the event cannot be accurately reconstructed from the ionisation signals alone; by utilising the much faster scintillation signals, the time of interactions can be calculated with much higher precision. This data can be combined with the arrival time of drifting electrons to reconstruct the true drift distance. Without an accurate t_0, the reconstructed drift distance will be incorrect, and as a result, so too would the predicted charge attenuation.

Whilst the details of the charge readout and photon detection systems are unique in each detector, LArTPC detectors can be split into two main categories: single–phase and dual–phase. In a single–phase detector, the drifting ionisation electrons remain in the liquid argon and the signals are usually read out on three anode wire planes, although pixel readout systems have also been developed[55]. A dual–phase LArTPC contains an additional region of gaseous argon, in which a high electric field, known as the extraction field, is applied to extract the ionisation from the liquid into the gas, before it is amplified and collected on a pair of anode wire planes[53]. The large amplification produced in a dual–phase LArTPC provides an increased signal–to–noise ratio in these detectors when compared to single phase LArTPCs.

3.1.2 ProtoDUNE–SP Goals

ProtoDUNE–SP is one of two large scale prototypes for the DUNE Far Detector modules that focusses on the single–phase LArTPC technology. The DUNE Far Detector is designed such that each module is built up of a number of identical components. ProtoDUNE–SP was designed to prototype the design of many of these components at a 1:1 scale, including the anode planes, cathode plane, and photon detectors. The ProtoDUNE–SP experiment has four primary goals, as outlined in the technical design report[7]:

- Prototype the production and installation procedures for the single–phase Far Detector design.

- Validate the design from the perspective of basic detector performance; this can be achieved with cosmic-ray data
- Accumulate large samples of test-beam data to understand/calibrate the response of the detector to different particle species
- Demonstrate the long-term operational stability of the detector as part of the risk mitigation program ahead of the construction of the first 10 kt Far Detector module

As such, ProtoDUNE–SP represents a significant milestone in the development of the Far Detector for the DUNE experiment. Its successful operation, both in a test beam and with cosmic–rays, provides valuable data with which to understand reconstruction and analysis of the data that will be collected by the DUNE Far Detector.

3.2 The ProtoDUNE–SP Detector

The ProtoDUNE–SP detector is located at the Neutrino Platform at CERN along the H4 beamline. It is a single–phase LArTPC detector with a total liquid argon mass of 0.77 kt, making it the largest monolithic single–phase liquid argon TPC at the time of writing. The TPC comprises the following major components, which are illustrated in Figure 3.3:

- A cathode plane constructed of modular cathode plane assemblies (CPA)
- Two anode planes constructed of modular anode plane assemblies (APA)
- A photon detection system (PDS) that is integrated into the APAs
- A field cage, beam plug, and high voltage systems
- Readout electronics and a data acquisition system (DAQ)

The detector components are designed to be an almost exact replica of the final single–phase Far Detector modules, but the detector has an overall scaling factor of approximately 1 : 20 in terms of total liquid argon mass[7].

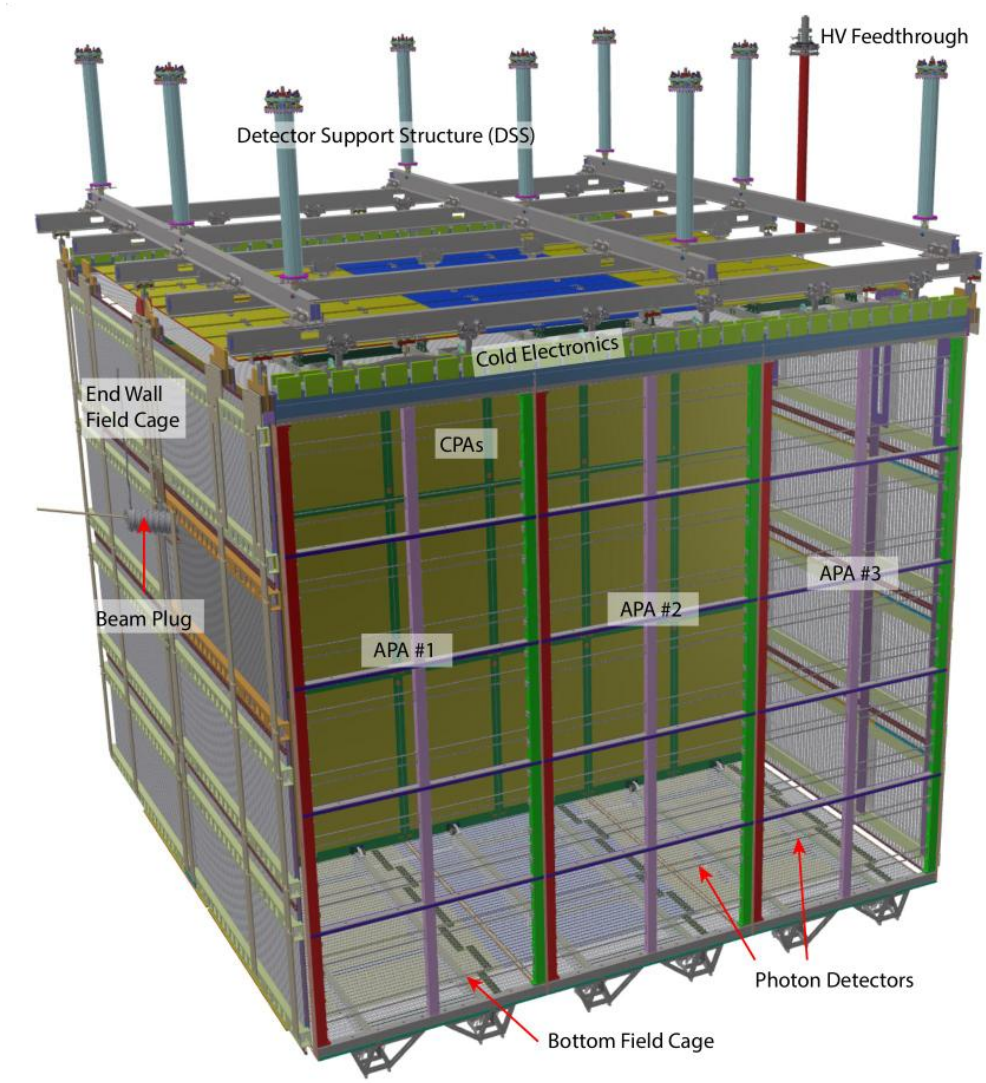


Figure 3.3: A 3D rendering of the main components of the ProtoDUNE–SP TPC, with the main components labelled. Figure from [7].

3.2.1 The Liquid Argon TPC

The ProtoDUNE–SP TPC has an active volume of 6 m (height) × 7.2 m (width, drift direction) × 7 m (length, approximate beam direction). The cathode plane at the center of the active width is flanked by two anode planes that define two 3.6 m drift volumes. The field cage around these two drift volumes helps to ensure a uniform electric field within the drift region.

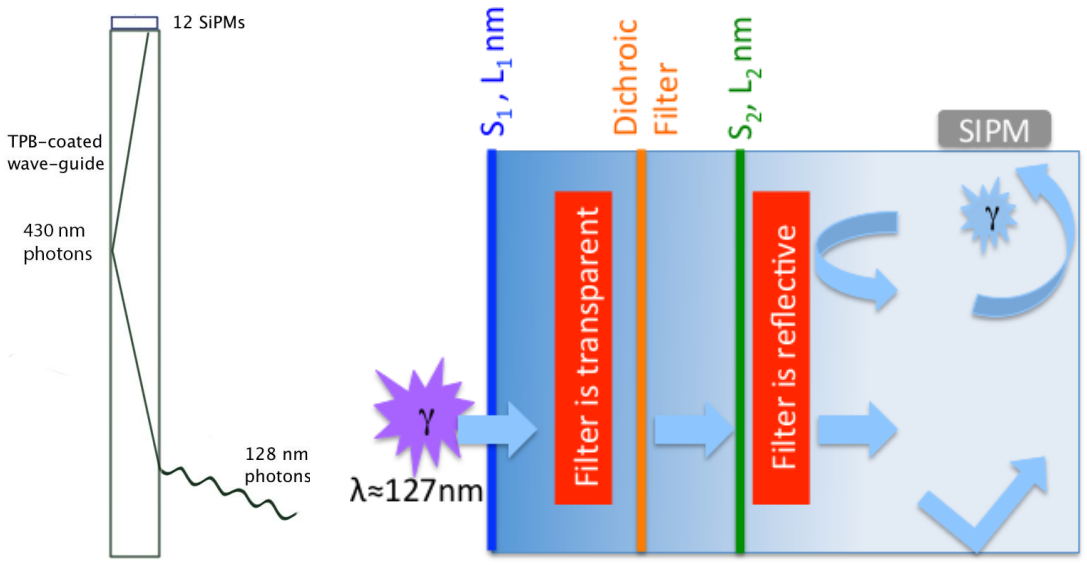
Each anode plane is modularly constructed from three APAs, which have dimensions 6 m (height) × 2.3 m (width). The APA frame holds three sets of parallel wires on the inward and outward facing sides, these are oriented at different

angles to enable 3D reconstruction. The outer two sets of wires are induction wires, these are biased such that they are electrically transparent to the drifting ionisation. Ionisation passing the induction wires causes an induced bi–polar signal. The third set of wires are known as collection wires, which are not biased. When drifting ionisation approaches the collection wires it is absorbed producing a uni–polar signal. In ProtoDUNE–SP, each set of induction wires contains 800 wires at a spacing of 4.67 mm, and each set of collection wires contains 480 wires at a spacing of 4.79 mm.

The wire planes from each APA are read out by electronics boards, known as front–end mother boards (FEMB), which are mounted on the APA frame. The FEMBs are submerged in the liquid argon and, therefore, they are also referred to as cold electronics (CE). A total of 2560 electronics channels are used to read out the data from each APA, with each readout window corresponding to 3 ms, which is 6000 time samples. The CE amplify, shape, and digitise the signals from the wires before transmitting them outside the TPC to the warm interface boards (WIB). The WIBs collate the data from the CE boards and pass the data onto the data acquisition system (DAQ).

The cathode plane in ProtoDUNE–SP consists of an array of 18 CPA modules, 2 m (height) $\times$ 1.2 m (width). The cathode plane is held at -180 kV to provide a 500 V/cm drift field in each of the drift volumes. The field cage surrounding the drift regions ensures that the electric field is uniform across the detector volume, by providing the necessary electrical boundary conditions.

One area in which the design of ProtoDUNE–SP differs from the Far Detector is the inclusion of the beam plug. This is necessary to minimise interactions between the charged particle test beam and the cryostat before the beam enters the active region of the detector. A cylindrical beam plug, which contains nitrogen gas, penetrates from the cryostat wall into the field cage at the location of the incoming test beam.



(a) Waveguide. **(b)** ARAPUCA.

Figure 3.4: The operating principle of the two main types of photon detector module in ProtoDUNE–SP. Figure (a) from [7]. Figure (b) from [56].

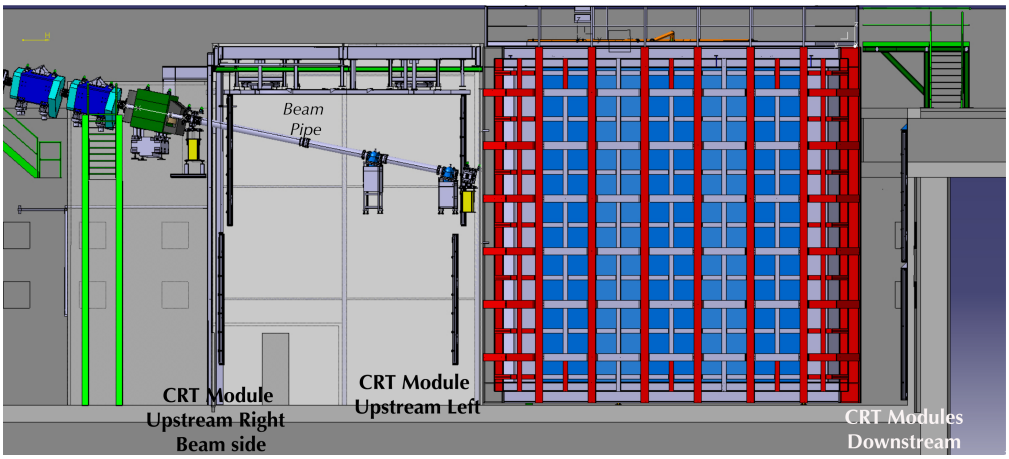
3.2.2 The Photon Detection System

The Photon Detection System (PDS) in ProtoDUNE–SP is integrated into the APAs. Ten photon detector modules are embedded in each APA frame between the layers of wires on either side of the APA, as shown in Figure 3.3. Three types of photon detector module were tested in ProtoDUNE–SP, two very similar module designs based on coupling silicon photomultipliers to wavelength shifting bars, and a third novel design known as the ARAPUCA light trap. The operating principles of these modules are illustrated in Figure 3.4.

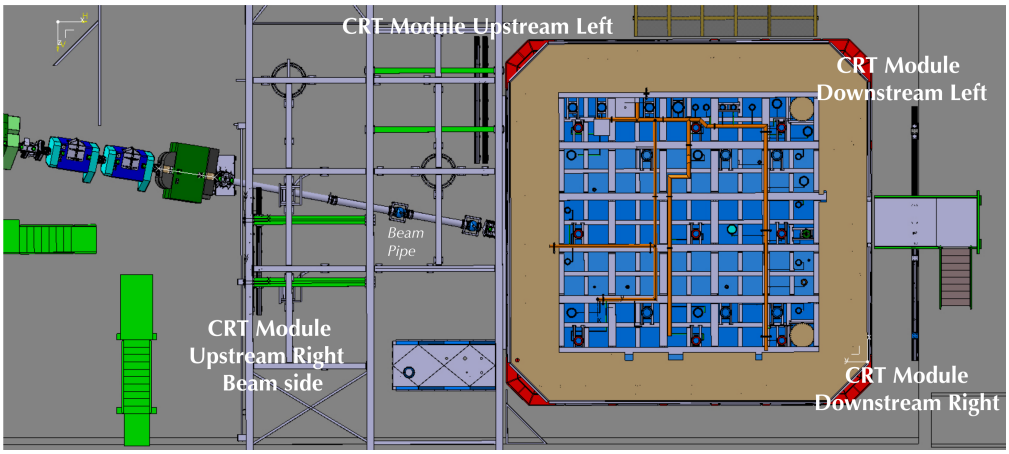
The majority of the photon detector modules in ProtoDUNE–SP consist of wavelength shifting bars coupled to silicon photomultipliers (SiPM). There are two very similar designs of this style that differ only in terms of the wavelength used to transmit the light within the waveguide. In both designs, tetraphenyl-butadiene (TPB) is used to shift the wavelength of the light from ultra–violet to blue before it enters the waveguide. In one case the light is transmitted at the blue wavelength produced by the TPB, whereas in the other, the blue light from the TPB is absorbed in the waveguide, followed by the emission of green light that is transmitted down the waveguide.

The ARAPUCA is a novel design that is only used in a small number of photon detector modules in ProtoDUNE–SP. In this design the photons are trapped in a small box through a sequence of wavelength shifting and optical filtering, thereby significantly increasing the photon detection efficiency[57]. An ARAPUCA light trap consists of a box that is highly reflective on the inside. On one side of the box is the filtering window, and on another a SiPM. Incoming ultra–violet photons are shifted to the blue spectrum before passing through a dichroic filter, which has a tunable wavelength cut–off, after which it transitions from being transparent to reflective. After passing through the filter the photons are shifted again, this time from blue to green, such that the filter becomes reflective to them. Consequently green photons can be trapped within the ARAPUCA until they come into contact with the SiPM, this provides increased photon detection efficiency. Each ARAPUCA photon detector module in ProtoDUNE–SP features an array of these traps arranged in a line across the width of an APA. The ARAPUCA modules contain 12 ARAPUCA light traps, these are constructed from 16 cells, which are 5 cm $\times$ 5 cm $\times$ 1 cm boxes. Half of the light traps are made of one cell and the others are made of two cells joined together. Regardless of the number of cells, each ARAPUCA trap is read out by a single SiPM, making for 12 SiPMs per ARAPUCA module.

Unlike with the TPC electronics, there are no front–end electronics in the LAr volume for the PDS, the unamplified analogue signals are transmitted out of the cryostat before processing and digitisation. The light guide based photon detector modules have 12 SiPMs, which are read out in threes meaning that each module corresponds to 4 readout channels. The ARAPUCA modules are read out by one channel per SiPM, resulting in 12 channels per ARAPUCA module. The readout channels from each APA are processed by four SiPM signal processors (SSP), which are mounted on the top of the cryostat. After processing in the SSPs, the PDS data is passed onto the DAQ along with the TPC data.



(a) Side view.



(b) Top view.

Figure 3.5: A rendering of the location of the H4 beamline and cosmic–ray taggers in relation to the ProtoDUNE–SP detector. The CRT modules are labelled on the figures, in both the side view and the top view. Figures from [58].

3.3 The H4 beamline

The ProtoDUNE–SP experiment is located in the North Area of CERN at the end of the H4 beamline, the location of the beamline with respect to the detector is illustrated in Figure 3.5. The beam has a mixture of particle types including hadrons, muons, and electrons, with momenta in the range 0–7 GeV. The momenta and composition of the beam can be varied, and several runs were taken across a range of particle compositions and beam momenta.

The H4 beamline is a tertiary beamline, which is produced when a secondary beam from the T2 primary target interacts with another target. Particles from

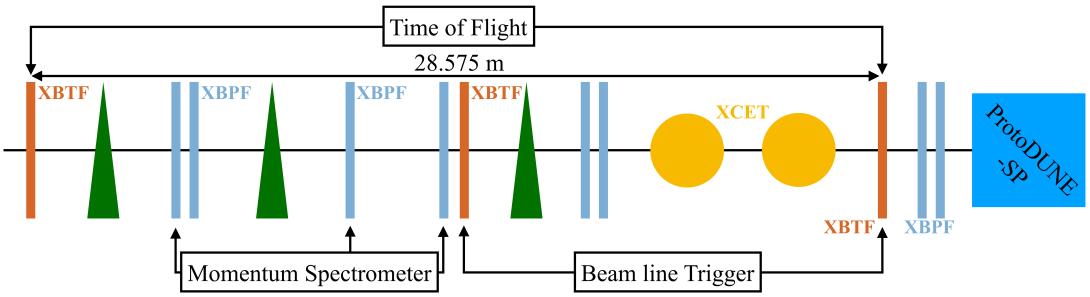


Figure 3.6: A schematic of the H4 beamline magnets and instrumentation. There are eight profile monitors, three trigger counters, and two threshold Cerenkov counters, labelled XBPF, XBTF, and XCET respectively. The bending magnets are depicted as green triangles. The data from these components is combined for triggering, measuring the beam momentum, and measuring the time of flight of beam particles. The components used for each of these tasks are labelled above and below the beamline schematic. Figure from [58].

the secondary target are selected, based on momentum and charge, before travelling down the H4 beamline to ProtoDUNE–SP. A schematic of the beamline instrumentation (BI) and magnets in the H4 beamline is given in Figure 3.6.

By combining momentum measurements from the profile monitors with time of flight (TOF) and Cerenkov measurements, the beam momentum and composition can be measured. The predicted and measured distribution of TOF vs momentum for data, from a number of runs at different momenta, is shown in Figure 3.7. This information is used to trigger the detector during beam running, and is sent to the central trigger board for distribution to the detector components.

During the beam running period, a number of runs were taken with differing beam compositions and momenta. A summary of the collected beam data, as well as predictions of the particle composition from beamline simulations, is given in Table 3.1.

3.4 The Cosmic–Ray Tagger

As a surface level detector with no overburden, ProtoDUNE–SP measures a significant rate, on the order of 20 kHz, of cosmic–ray muons. This corresponds to an average of 60 muons per 3 ms readout window. These muons provide a

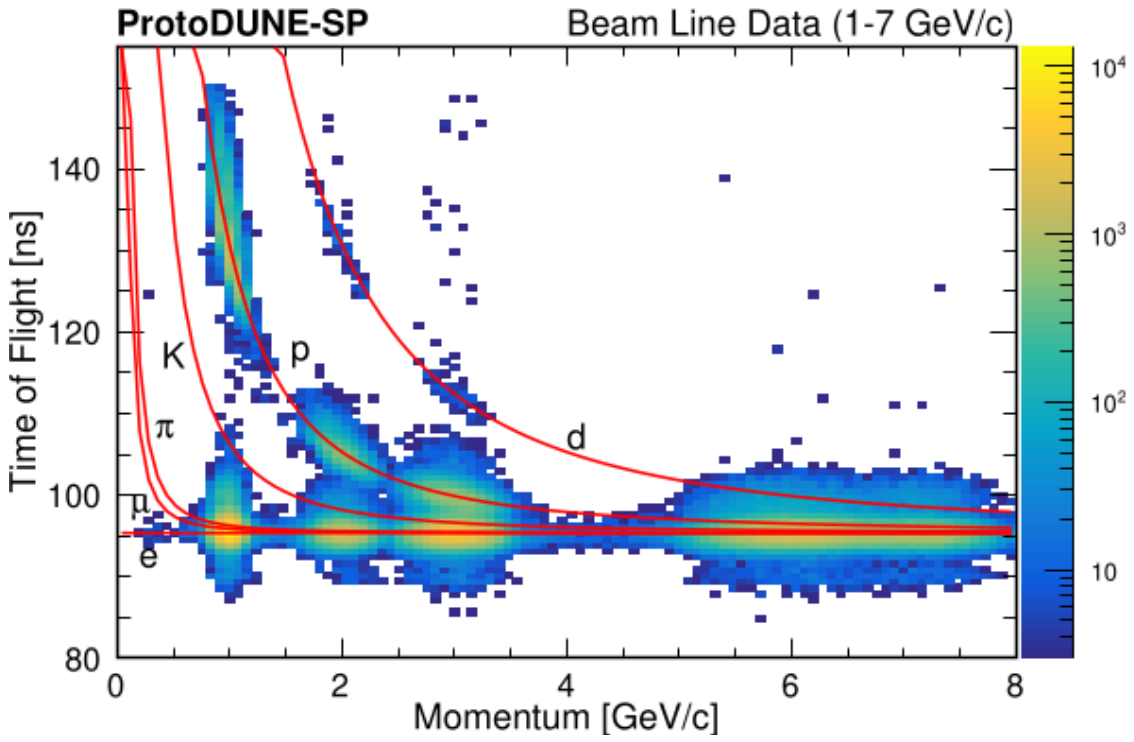


Figure 3.7: Time of flight vs momentum distributions from the H4 beamline instrumentation, for ProtoDUNE–SP runs at average momenta from 1–7 GeV. The theoretical distributions for different beam particle species are labelled on the figure. Figure from [58].

Momentum (GeV/c)	Recorded Triggers (10^3)	Expected Triggers (10^3)	Expected Pions (10^3)	Expected Protons (10^3)	Expected Electrons (10^3)	Expected Kaons (10^3)
0.3	269	242	0	0	242	0
0.5	340	299	1.5	1.5	296	0
1	1089	1064	382	420	262	0
2	728	639	333	128	173	5
3	568	519	284	107	113	15
6	702	689	394	70	197	28
7	477	472	299	51	98	24
All	4173	3924	1693.5	777.5	1381	72

Table 3.1: Summary of expected and recorded beam triggers in ProtoDUNE–SP for all runs taken during the 2018 data taking period. The number of expected triggers represents the number of triggers expected in the initial ProtoDUNE–SP run plan, which are exceeded at every energy. The breakdown of expected triggers by particle species is based on simulations of the H4 beam–line. Data from [59].

useful source of calibration data in the form of long tracks and stopping muons. The cosmic–ray tagger (CRT) in ProtoDUNE–SP was installed on the upstream and downstream faces of the TPC, to trigger the detector for cosmic–ray muons that travel parallel to the anode plane, which are particularly useful for calibration studies. In addition, the CRT can provide accurate timing information for any tracks that are successfully matched between the TPC and CRT.

The CRT consists of four parts, two upstream assemblies and two downstream assemblies, the locations of these assemblies is illustrated in Figure 3.5. Each CRT assembly is constructed from overlapping scintillation counters, which cover an area 6.8 m high and 3.65 m wide. Scintillation strips of length 365 cm and width 5 cm are placed in perpendicular arrays to give 2D reconstruction within each CRT assembly. By combining data from upstream and downstream CRT assemblies, with a time coincidence requirement, the 3D trajectories of tracks can be reconstructed.

3.5 The Data Acquisition System

Data from the TPC, PDS, and CRT in ProtoDUNE–SP are collated by the data acquisition system (DAQ). The DAQ distributes triggers, compresses and packages the data into events, and stores the data ready for future analysis. An overview of the ProtoDUNE–SP DAQ system is seen in Figure 3.8; there are four primary data flows in the system: TPC, PDS, CRT, and BI. Timing and triggering signals are distributed to the detector components by the timing board, which maintains a 50 MHz clock and receives $\sim$ 40Hz of triggers from the central trigger board (CTB), based on data from the BI and CRT[7].

The ProtoDUNE–SP TPC contains a total of 15,360 wires spread across the six APAs. The wires are digitized at a rate of 2 MHz, resulting in an overall data flow of around 480 Gb/s from the TPC electronics. An event in ProtoDUNE–SP corresponds to a continuous readout of the detector for 3 ms, resulting in 6000 samples from each wire; data is buffered such that the readout window can be opened 250 μs before the trigger time. Data from the TPC is received by two systems. One system, which is based on Reconfigurable Computing Elements

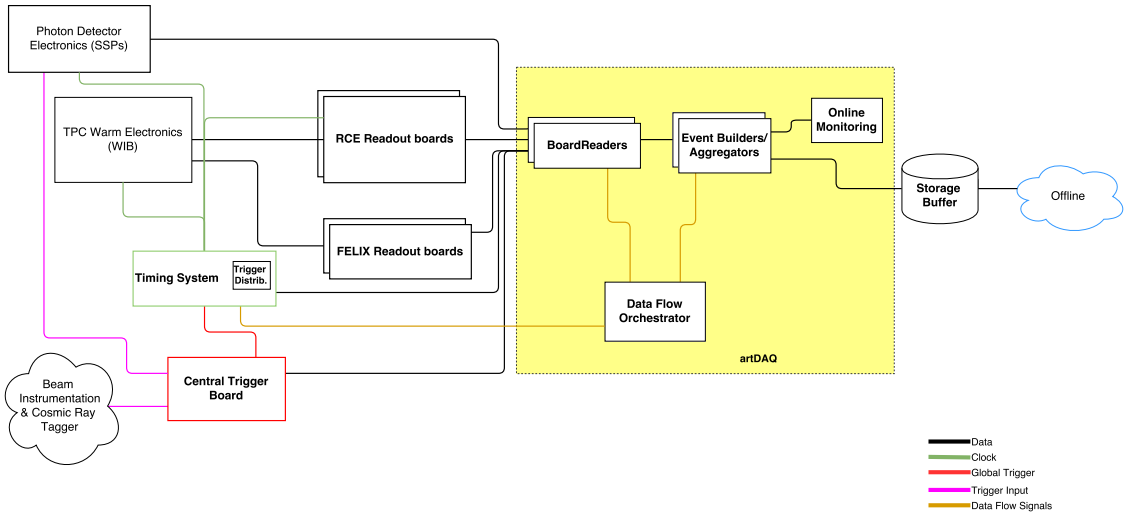


Figure 3.8: Outline of the data acquisition system in ProtoDUNE–SP, showing the data flow from the detector components through to offline storage ready for later analysis. The main elements of the system are shown, including the important software components that are part of the ArtDAQ framework highlighted by the yellow box. Figure from [7].

(RCE)[60], handles the data from five APAs, while the other system handles the data from the sixth APA, and is based on front–end link exchange (FELIX) cards, which have been developed by the ATLAS collaboration[61].

The software layer of the ProtoDUNE–SP DAQ is based on Fermilab's ArtDAQ[62]. This component is primarily responsible for acquiring the data, packaging it, and storing it locally. Triggered events are queued and distributed to the board readers and event builders by the data flow orchestrator. There are multiple board readers that are each responsible for processing the data from specific hardware components. Details vary between specific board readers, but generally these processes are responsible for formatting the data from each component ready for aggregation. Data from the various detector components is aggregated by the event builders, which are responsible for assembling the completed events. After compression, events have an average size of 60 MB. ArtDAQ is also responsible for the real–time monitoring of data quality via the online monitoring system, this system will be discussed in Section 3.8.

3.6 Simulation

Simulation and reconstruction of ProtoDUNE–SP data takes place in the LArSoft framework[63]. LArSoft is a software suite for simulating and reconstructing data collected by LArTPC detectors, based on the Art event framework from Fermilab[64]. LArSoft is under active use and development by a number of participating LArTPC based experiments, with each experiment making use of its core functionality, in addition to experiment–specific code.

Simulation in LArSoft is broken down into three sequential stages: generation, propagation, and detector simulation. The initial state particles are produced in the generation step, the propagation and interaction of these particles is simulated during the propagation step, and finally the transport of energy depositions and simulation of detector effects are handled by the detector simulation phase.

As a surface based detector in a test beam, ProtoDUNE–SP is subject to three main sources of particles: beam particles, cosmic–rays, and radiological decays. The beam particle flux in the vicinity of ProtoDUNE–SP was provided by simulations of the H4 beamline[65] based on two simulation frameworks, G4beamline[66] and FLUKA[67]. The flux of cosmic–ray particles in ProtoDUNE–SP is simulated using CORSIKA[68]. Finally, low energy radiological backgrounds in liquid argon are simulated, these include ^{39}Ar,^{85}Kr, and ^{222}Rn.

After generation, the particles are allowed to propagate and interact in the detector geometry, including the cryostat, external systems, and experimental hall. This is simulated using GEANT4[69], which tracks particles in small steps; the step size, 300 μm, is chosen to be much smaller then the spatial resolution of the detector to allow for small–scale processes like showering to be accurately simulated.

Particles that propagate through the detector, leave energy deposits in the form of ionisation electrons and scintillation photons. A minimum–ionising particle (MIP), which deposits around 2 MeV/cm in liquid argon, will produce tens of thousands of electrons and photons per cm.

The detector simulation is responsible for drifting the ionisation electrons towards the collection wires, propagating the simulated photons to the photon detectors,

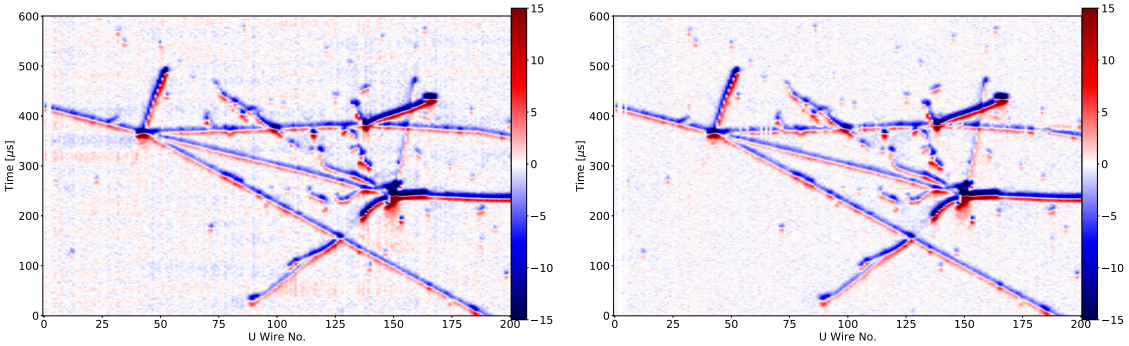
and simulating the response of the wires, SiPMs, and electronics to the signals. The sheer number of ionisation electrons and scintillation photons produced in a ProtoDUNE–SP event makes simulating the propagation of the full set of particles impractical and, therefore, LArSoft employs approximation techniques to predict the observed signals within a reasonable computation time. A number of detector effects are taken into account when simulating the electron drift: electron–ion recombination, transverse and longitudinal diffusion, electric field distortions, and electron capture on impurities.

When the ionisation and scintillation signals arrive at the active detector components, simulated waveforms are produced based on the induced signals produced in the detector components. These signals are converted into electrical signals by electronics simulation, which is completed in LArSoft. The electronics simulation includes simulation of electronics gain, noise, and analogue to digital conversion. The simulated waveforms produced by the ProtoDUNE–SP simulation are then compressed and stored ready for reconstruction.

3.7 Reconstruction

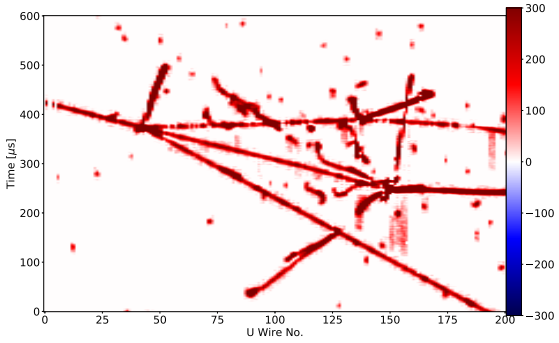
An event consists of a synchronised set of waveforms from all of the TPC channels, with each event corresponding to 3 ms or 6000 samples. The particle interactions within each event are reconstructed into tracks and showers, which can be used for physics analyses. This section provides a brief summary of the steps involved in this reconstruction.

First, the TPC waveforms are passed through a set of filtering algorithms that are designed to reduce the noise on the waveforms, as well as to mitigate effects from the electronics such as sticky–codes and undershooting, which are discussed in [58]. An additional step, known as a 2D deconvolution, is applied to extract the original signal from the measured signal, based on a detector response function for both the wire in question and the neighboring wires. This method has been described in detail by the MicroBooNE collaboration[70]. An important side effect of the 2D deconvolution technique is that, after deconvolution, the bi–polar induction



(a) Before noise filtering.

(b) After noise filtering.



(c) After 2D deconvolution.

Figure 3.9: Example of noise filtering and 2D deconvolution in ProtoDUNE–SP data. Each image shows an example of the wire vs time signals from a region in the U–plane, which is an induction plane. Figures from [58].

signals become unipolar, which simplifies hit reconstruction. An example of the noise filtering and 2D deconvolution techniques applied to a 7 GeV test beam event in ProtoDUNE–SP data is given in Figure 3.9.

Once the deconvoluted waveforms have been calculated, regions of interest (ROI) are defined, these are regions of high amplitude in which charge deposition is likely to have occurred. Reconstructed hits are found within each region of interest by fitting gaussian peaks to the pulses within the region; most hits consist of a single gaussian peak, however in busy regions of the detector, multiple gaussian peaks may be used to fit a single pulse. Each reconstructed hit will have an associated peak time, width, and integral, these objects form the basic input for the subsequent pattern recognition algorithms. An example of the reconstructed hits is shown in Figure 3.10.

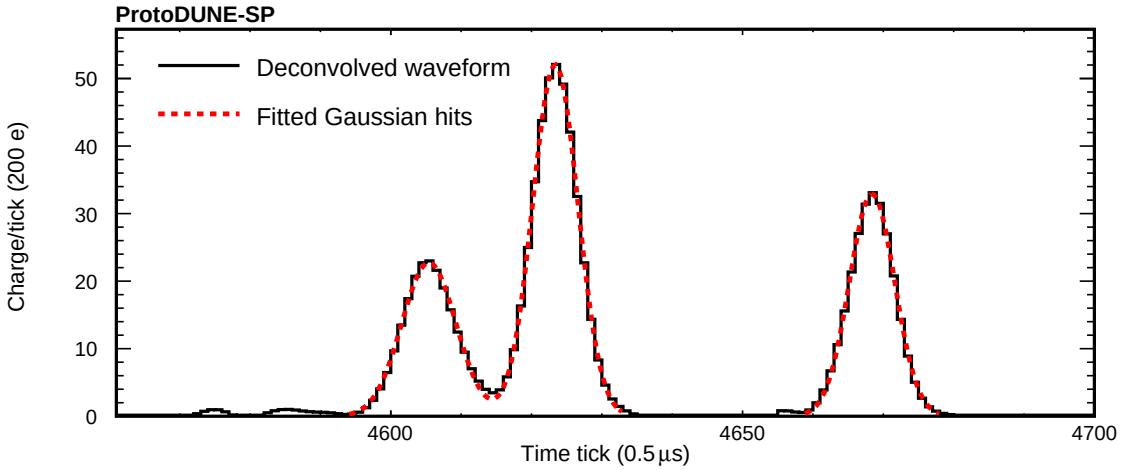


Figure 3.10: An example of reconstructed hits from ProtoDUNE–SP data. Figure from [58].

Pandora[71] is the primary pattern recognition software used in ProtoDUNE–SP. It takes a multi–algorithm approach to reconstructing particle interactions in the detector, and has been successfully used by other LArTPC detectors, such as MicroBooNE[72]. Pandora handles the clustering of hits into 2D clusters, the matching of 2D clusters into 3D clusters, and the reconstruction of 3D objects like tracks and showers. Ultimately, Pandora returns a tree of reconstructed particle flow particles (PFParticles), each corresponding to a distinct track or shower, and connected through parent–daughter relationships, which define the particle flow in the interaction. Pandora reconstruction consists of two stages: a cosmic pass and a beam pass.

First the cosmic pass reconstructs the event with algorithms designed to reconstruct track–like particles. Track stitching algorithms are used to predict the true interaction time for certain tracks during this stage. Any track that crosses either a CPA or APA will produce track segments on either side of the boundary. These segments point in the same direction, but will have been displaced from each other in the drift direction based on the arrival time of the corresponding particle. The true time for that track is equal to half the time shift required to realign the two segments into a continuous track. A visual representation of this algorithm is show in Figure 3.11. After the cosmic pass, any clear cosmic–ray candidates are

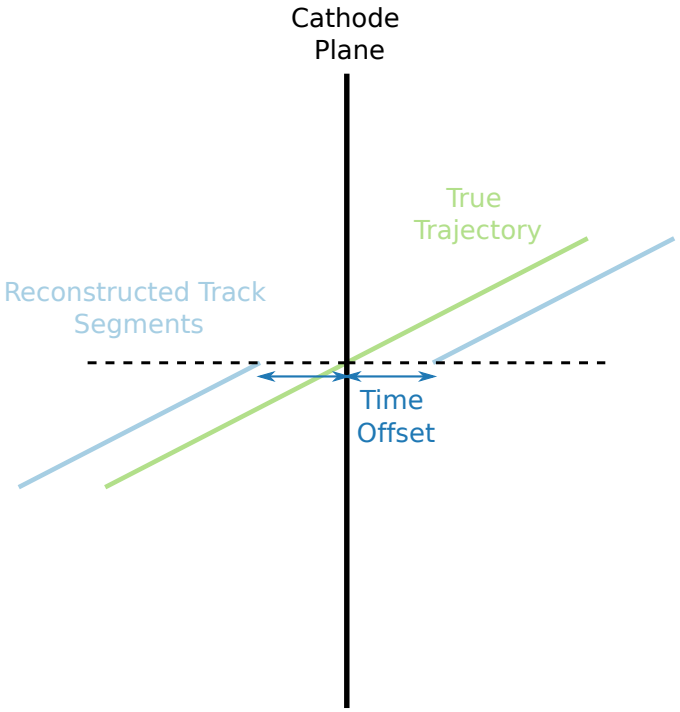


Figure 3.11: Diagram demonstrating the track stitching algorithm in ProtoDUNE–SP. The track is initially reconstructed in two segments on either side of the cathode plane, which are offset by a distance proportional to the arrival time of the track. The two segments can be stitched together by correcting for the time offset, resulting in a single track, which has an associated reconstructed time based on the time offset.

removed if they satisfy any of the following conditions[58]:

- The particle travels through both the top and bottom of the detector
- The assigned t_0 is inconsistent with the beam time
- If no t_0 assigned, try $t_0 = 0$. If any hits are reconstructed with positions outside of the detector this track is inconsistent with the beam time

Beam particle reconstruction considers only the hits that were not removed by being labelled as clear cosmic–rays. These hits are formed into 3D slices that contain all the hits from a single parent particle and its daughters. The slices are reconstructed with both the cosmic–ray and beam particle algorithms, and a boosted decision tree is used to determine whether a given slice is consistent with being a beam particle[58].

The beam particle reconstruction uses a more complex chain of algorithms to reconstruct a given slice. These algorithms are capable of reconstructing the particle hierarchies seen in the complex hadronic interactions from the ProtoDUNE–SP

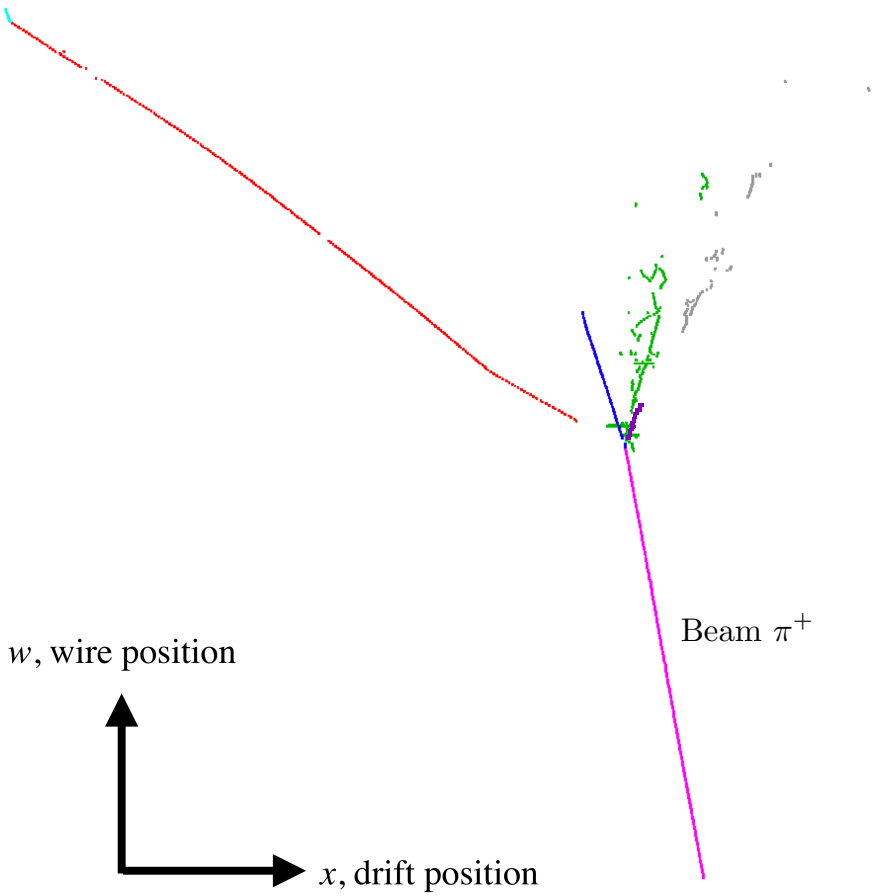


Figure 3.12: An example of a reconstructed particle hierarchy from Pandora, for a beam π^+ event in ProtoDUNE–SP simulation. The 2D representation is in terms of the wire position and drift position in the collection plane. Each reconstructed particle is labelled in a different colour, demonstrating the reconstructed particle hierarchy of the interaction. Figure from [58].

beam. These algorithms return the particle flow of the interactions, in the form of PFParticles, as well as the reconstructed interaction vertex for the primary beam particle. An example of the reconstructed particle hierarchy from a ProtoDUNE–SP beam event is shown in Figure 3.12.

As well as precise spatial reconstruction, LArTPCs provide excellent calorimetric information. Each reconstructed hit has an associated deposited charge, Q, and energy, E, which measure the amount of energy deposited by the charged particle in that small energy deposition. The measured charge in ADC depends on the spatial location of the hit within the TPC, as well as the local ionisation density,

or dQ/dx, in the region of the associated energy deposition. Therefore, to convert the measured charge of each hit into a reconstructed energy, a number of spatial and dQ/dx dependent factors need to be taken into account. The reconstructed charge of each hit is given by the integrated area under the gaussian fit to that hit. The dx for each hit varies based on the direction of the track with respect to the wires, it is equal to the wire spacing divided by the sine of the angle between the wire and the 2D projection of the track onto the anode plane.

A number of factors affect the measured dQ/dx, these effects need to be corrected in order to recover the dQ/dx at the source of the ionisation. These factors are split into two parts in the ProtoDUNE–SP reconstruction:

- Corrections in the drift direction, X corrections. Examples include longitudinal diffusion, attenuation on impurities, and drift velocity variations
- Corrections in the direction of the wire planes, YZ corrections. Examples include wire to wire response variations, and transverse diffusion

Calibration factors are calculated for both sets of corrections by considering a sample of cathode–crossing muons. The aim of these calibration matrices is to normalise the response over the TPC volume based on the median of the measured dQ/dx distribution in each location. The distribution is then normalised to the average value at the anode plane, where the effects of the X corrections are expected to be negligible. The corrected dQ/dx is given by

$$(dQ/dx)_{corr} = N_Q \; C_{yz}(y,z) \; C_x(x) \; \; (dQ/dx)_{reco} \, ,$$

where N_Q normalises the median of the distributions to the median value at the anode, and C_{yz} and C_x are the calibration factors for the YZ and X corrections respectively.

The final step in energy reconstruction is to convert the corrected dQ/dx into a reconstructed dE/dx, this involves accounting for electron–ion recombination at the source. The modified box model is used to simulate the recombination correction,

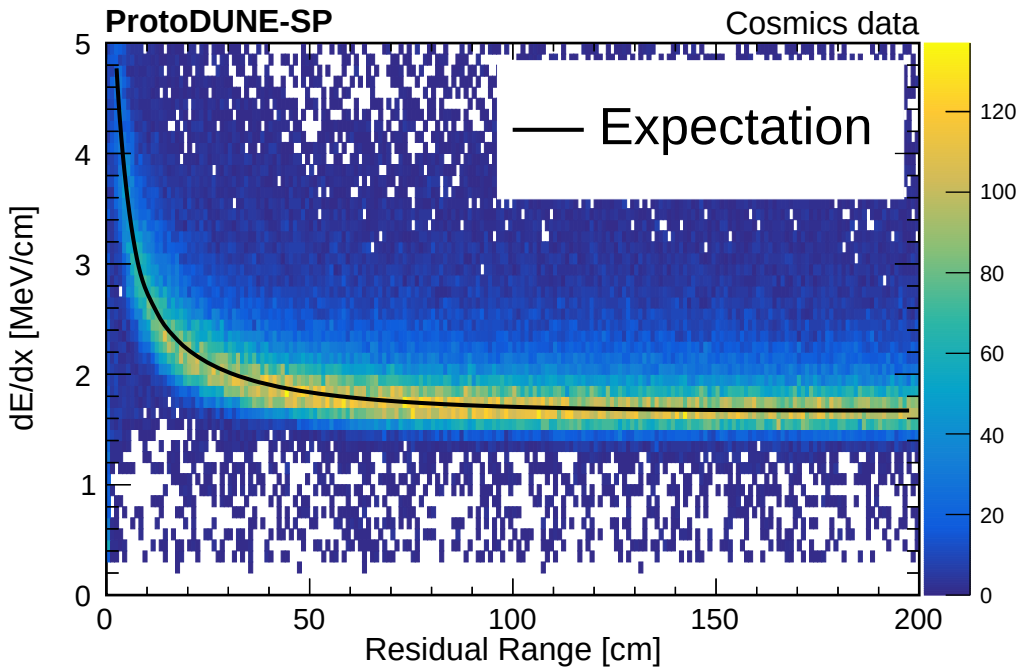


Figure 3.13: dE/dx vs residual range for a stopping muon sample in ProtoDUNE–SP data. The coloured histogram represents the reconstructed dE/dx in ProtoDUNE–SP data, and the black line is the most probable value of the theoretical distribution. Figure from [58].

this model has been studied in a LArTPC by the ArgoNeuT experiment[73]. The reconstructed dE/dx is

$$\frac{dE}{dx} = \left(\exp\left(\frac{\frac{dQ}{dx}}{C_{cal}} \frac{\beta' W_{ion}}{\rho \epsilon} \right) - \alpha \right) \left(\frac{\rho \epsilon}{\beta'} \right)$$

where C_{cal} is a calibration constant used to convert ADC to electrons, W_{ion} is the work function of argon, ϵ is the local electric field, ρ is the liquid argon density, and $\alpha = 0.93 \pm 0.02$ and $\beta' = 0.212 \pm 0.002$ (kV/cm)(g/cm^3)/MeV are the box model parameters as measured by ArgoNeuT[73]. The calibration constant, C_{cal} is calculated by fitting the most probable value of the reconstructed dE/dx distribution as a function of range to the theoretical prediction for dE/dx vs range for a sample of stopping muons, as shown in Figure 3.13.

3.8 ProtoDUNE–SP Online Monitoring System

As well as monitoring the stability of the detector and DAQ systems, the quality of the collected data has to be constantly monitored. This ensures that data

quality issues are quickly identified, and prevents long runs of low quality data being collected. The online monitoring system (OM) is responsible for providing this quality assurance.

As shown in Figure 3.8, the OM is a part of the ProtoDUNE–SP DAQ system. The OM is responsible for processing the data and displaying the results in the control room as soon as possible after the event was triggered. It consists of three main components:

- Analysis processes that decode and analyse the raw data from each detector subsystem
- Merging processes that collate monitoring data from each subsystem
- A web interface that displays monitoring data

An overview of the data flow in the OM is shown in Figure 3.14. Data from the detector are first filtered before being run in any OM processes. The data that is passed to the OM processes is then split up and decoded by the relevant RawDecoder, which reformats the data ready for analysis. A number of Analyser Modules then make use of this data to make plots, which are merged, along with plots from external systems, and sent to a web interface to be displayed in the control room. The following sections will provide a brief summary of these stages, as well as examples of the plots that are produced in the OM.

3.8.1 Data Processing

The data processing for the OM is based on Fermilab's Art and ArtDAQ software frameworks[62, 64], however, the beam instrumentation data is analysed outside of the OM system by the CERN beam group. The beam instrumentation plots are merged with the rest of the OM plots after data processing.

The first step of the data processing is event filtering, the main purpose of which is to control the flow of data into the OM. Two types of filtering take place sequentially. First a random filter is used to cut the data rate into the OM to a manageable level. This is followed by a second filter that aims to increase the likelihood that processed events were triggered by the beam instrumentation or

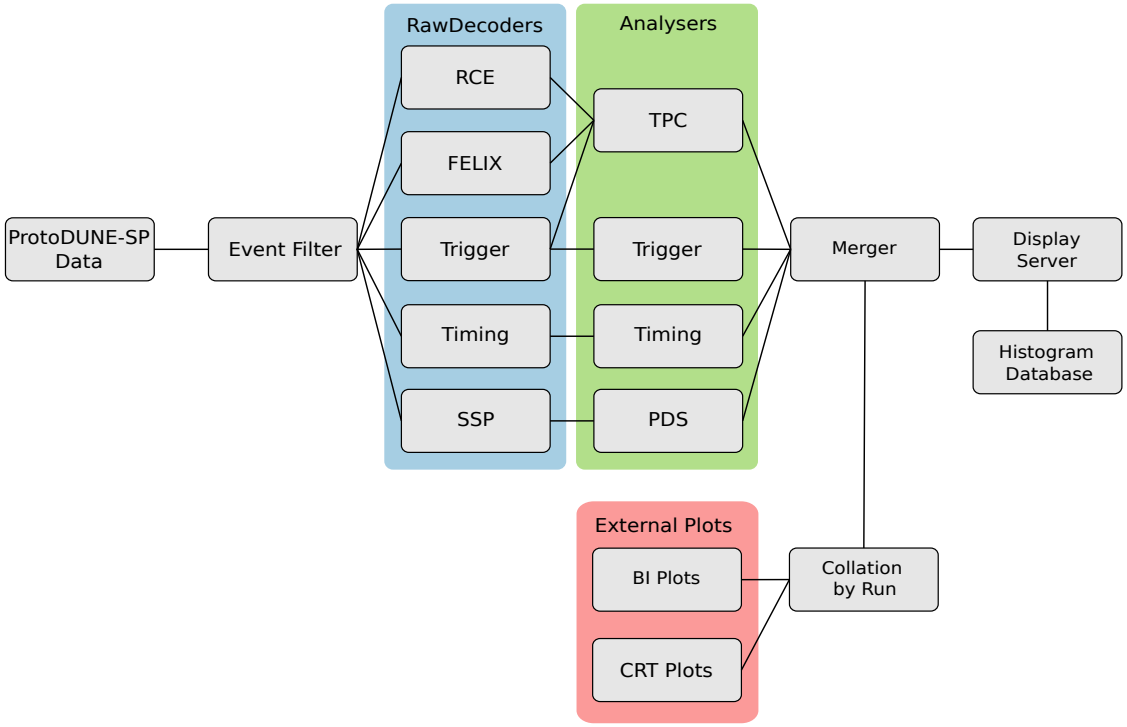


Figure 3.14: A diagram representing the data flow in the ProtoDUNE–SP online monitoring system. The data from each detector subsystem is processed by the RawDecoders before being analysed. The results of the analyses are merged, along with data from external detector systems, before being sent to the display server to be viewed in the control room.

cosmic–ray tagger, as opposed to a random software generated trigger. Different OM processes take different amounts of processing time per event, therefore, multiple filters are used, to control the data rate into each process separately.

After filtering the events are ready to be processed. The data coming into the OM system is in its raw form. It arrives in small pieces known as Fragments that have to be decoded before they can be used by the OM. The RawDecoders are responsible for interpreting the headers and data streams from each detector component, and restructuring the data ready for processing. The details of the decoding vary based on the readout system under consideration. Each readout system defines a class that details the contents of each Fragment. The RawDecoders use the contents of this class to decode the Fragments in order to prepare the data for processing.

After the data has been decoded it is ready to be analysed. This is done by a number of Analyser Modules that analyse the data from different detector

components and produce ROOT[74] plots as output. Details of the processing done for each detector component are given below.

Time Projection Chamber

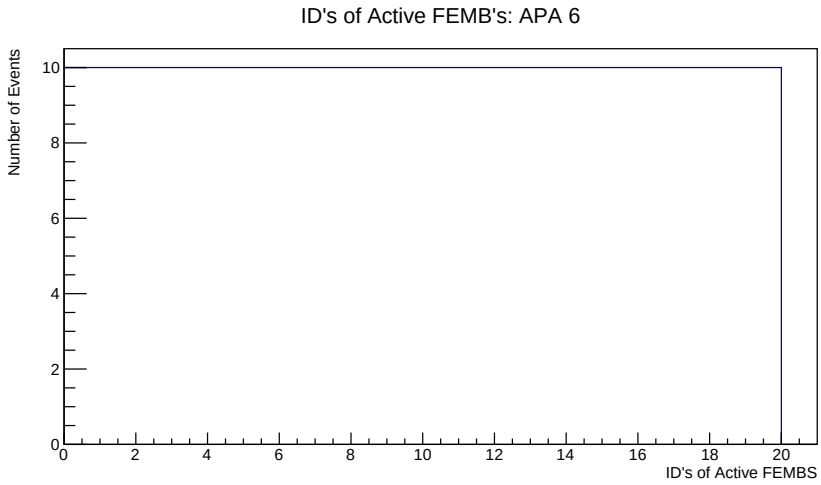
As the largest data source in ProtoDUNE–SP the TPC data is analysed in several small steps: basic data checks, pedestals and noise, Fourier analysis, and event displays. Examples of some of the plots produced by the TPC analysis are detailed below.

The basic data checks are intended to quickly spot any fundamental issues with the incoming data. An example of a basic check is to check which FEMBs are active in the monitored events, Figure 3.15a shows an example of the number of events recorded by each FEMB on APA 1 for a sample of 10 analysed events.

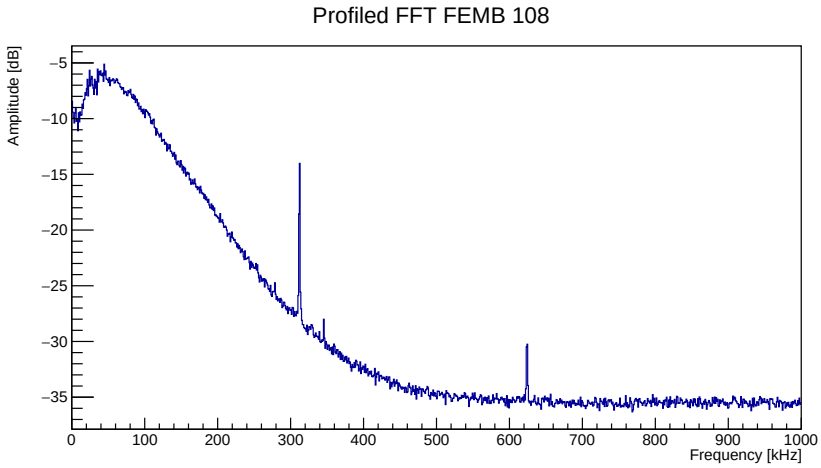
The pedestals and noise are continually evaluated by the online monitoring; these data are displayed in the control room and used later in the monitoring chain to flatten the background in the event displays. Basic hit removal is used to ensure that the pedestals and root mean square deviation (RMS) are only calculated in the regions of the readout corresponding to noise signals. The RMS of all channels can be represented on a single plot by arranging the channels into a 2D grid and displaying the RMS as the colour scale on a 2D histogram, an example of this plot is shown in Figure 3.15c.

To identify noise sources in the APA it is useful to study the Fourier transform of the signal distribution, which allows the noisy frequencies to be identified. Individual noise sources can be identified by studying the Fourier distributions under different operating conditions, for example, the cameras in the TPC were identified as a noise source in this way. The OM, therefore, provides fast fourier transforms (FFT) for each APA, an example of the FFT for a single FEMB is shown in Figure 3.15b.

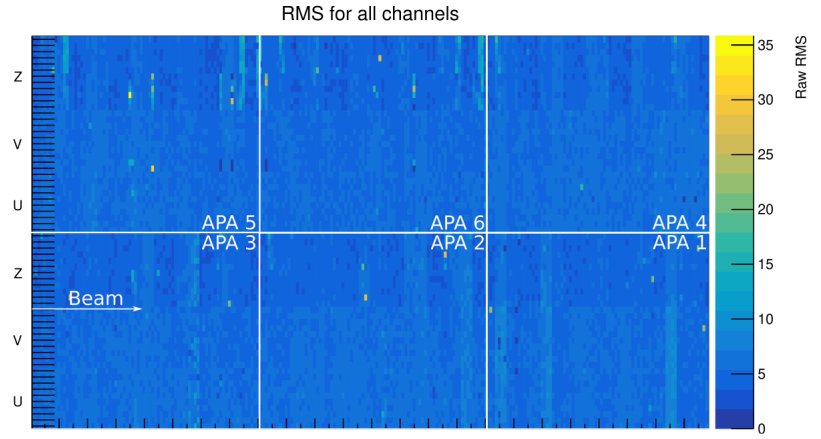
Event displays show all the data from all the TPC readout channels simultaneously and, as such, they provide the most general check of TPC performance in the OM. During the beam run of ProtoDUNE–SP, a number of issues in the data were first identified in the event display, for example, issues in the channel



(a) Number of events recorded by each FEMB on APA 6. 10 events are processed for each OM file, therefore, the histogram should be filled to 10 for every FEMB, if not there has been an issue.



(b) The profiled FFT from all channels from a single FEMB.



(c) RMS for all channels displayed *pseudo–geometrically* on a single histogram. The RMS for every TPC channel is plotted, and the channels are organised by APA and wire plane.

Figure 3.15: Examples of online monitoring plots for the TPC data.

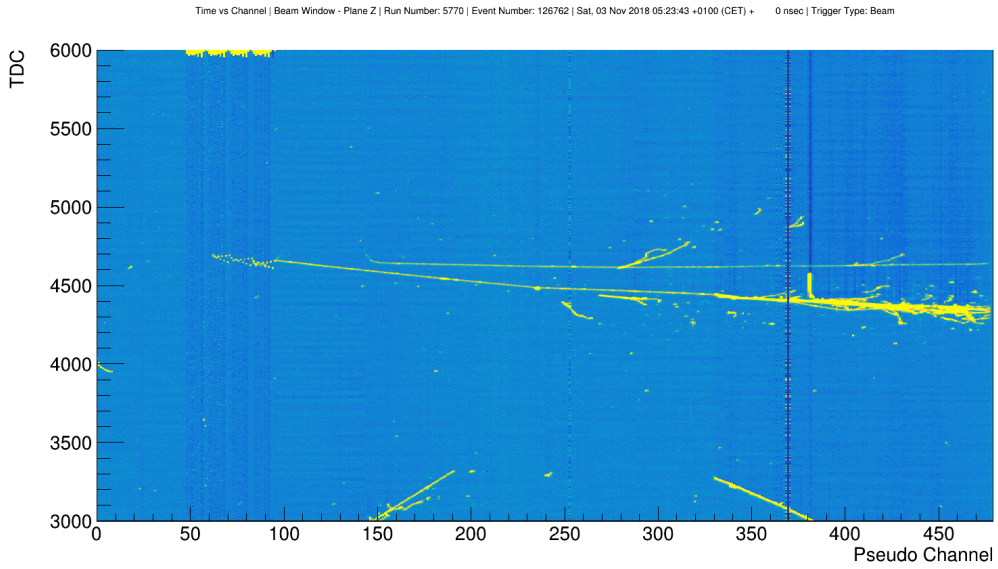
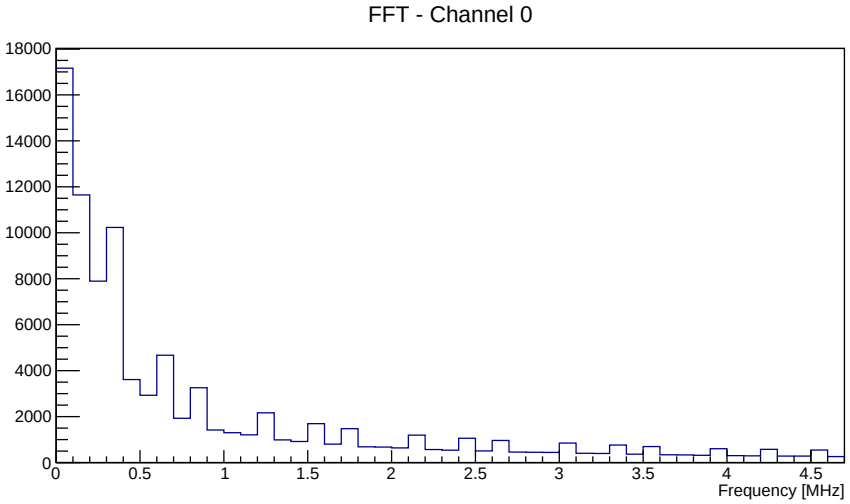


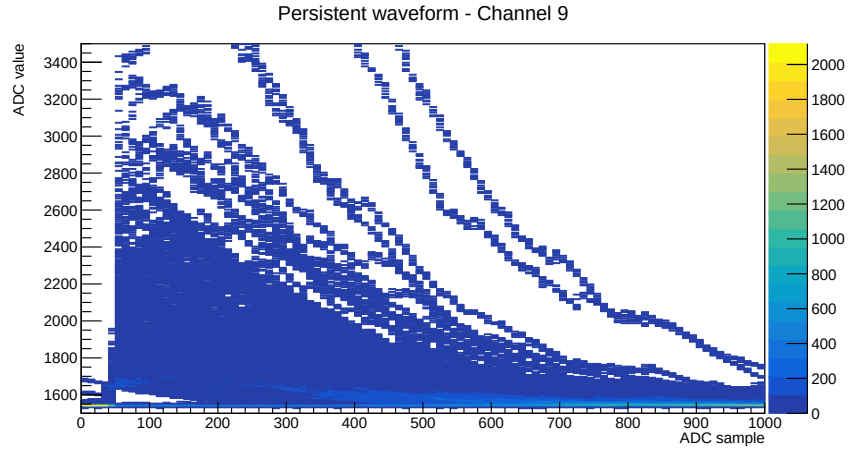
Figure 3.16: An example of a beam window event display from the ProtoDUNE–SP online monitoring system. The measured signal on the collection plane wires in ADC is displayed on wire vs time axes, for the region where the beam enters the TPC. An example of a data defect that was discovered in the OM can be seen on the left side of the plot; this timing issue was mitigated during offline reconstruction.

mapping and timing synchronisation. They are also particularly useful for checking that the beam trigger results in beam particles in the TPC. As a result a significant effort was made to make event displays available in the ProtoDUNE–SP OM system. In particular, changes to the display server were required, these changes are discussed later in this chapter.

Event displays were offered for all views in all APAs. In addition, a special event display, which focussed on the data in the beam window, and event displays that showed the data from all APAs either side of the TPC as a continuous image were produced. The event displays took the longest time to process out of all the plots in the OM, therefore, only one set of event displays was made per OM output file. To maximise the number of beam particles in the event displays, the trigger information was included during processing. An example of a beam–window event display is shown in Figure 3.16; the timing synchronisation issue mentioned previously is visible here, this issue affects a single FEMB and is mitigated during offline reconstruction.



(a) FFT for a single PDS channel.



(b) Waveforms for a single PDS channel that have been accumulated over a number of events. The ADC value as a function of time sample is shown for a number of PDS waveforms from a single channel.

Figure 3.17: Examples of online monitoring plots for the photon detection system data.

Photon Detection System

Similarly to the TPC data, the PDS data is analysed for pedestal, RMS, and FFTs; in addition, the raw waveforms for each channel are accumulated over a number of events and displayed in the online monitoring. Examples of some of the PDS monitoring plots are shown in Figure 3.17.

Trigger and Timing

The trigger and timing systems share much of the same data and, therefore, their monitoring is related. Some examples of useful timing and trigger plots include time–stamp difference distributions, and trigger records. The time–stamp delta plots show the difference in time–stamp between consecutive events, this distribution can be used to monitor the stability of the trigger rate during triggered operation. Trigger records are 2D histograms that detail the time and type of trigger issued by the trigger board. Figure 3.18 contains examples of these two plots.

Beam Instrumentation

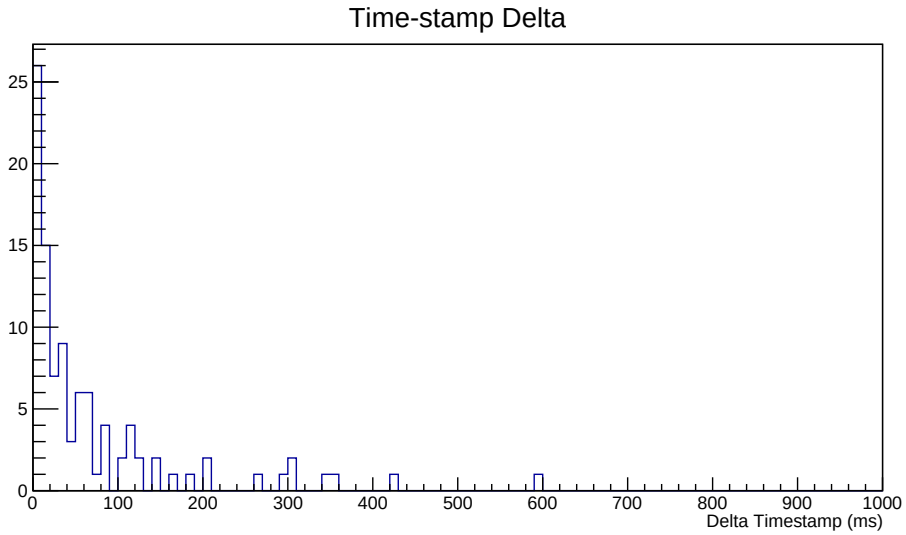
During the ProtoDUNE–SP beam run, the beamline instrumentation was monitored by CERN, who produced beamline monitoring plots[65]. These plots were continually overwritten on a time schedule that was decoupled from the ProtoDUNE–SP running schedule, the OM was responsible for collating the output of these plots in sync with the ProtoDUNE–SP run schedule, so that the collated beam information for each run could be monitored. After collation the beam instrumentation plots, which included momentum and time of flight, examples of which are shown in Figure 3.19, were merged with the rest of the OM plots.

Cosmic–Ray Tagger

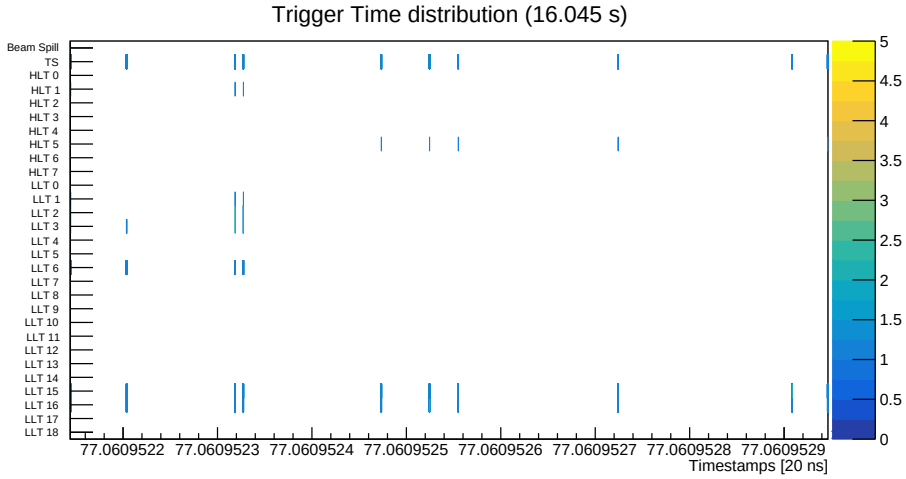
A bug in the DAQ meant that the CRT was run separately during the main beam run. Therefore, the CRT data was analysed separately from the rest of the ProtoDUNE–SP systems and merged in the OM in the same way as the beam instrumentation data. Some examples of CRT monitoring plots include the rate of hits and the mean ADC for each channel, these are shown in Figure 3.20.

3.8.2 Monitoring Web Interface

The plots from the data analysis are stored in a ROOT file on the ProtoDUNE–SP DAQ servers. This data is viewable from anywhere in CERN via a web interface, which was adapted for ProtoDUNE–SP from LHCb's Monet[75]. Monet consists of a



(a) Time–stamp delta distribution. A histogram of the difference in time between subsequent events.



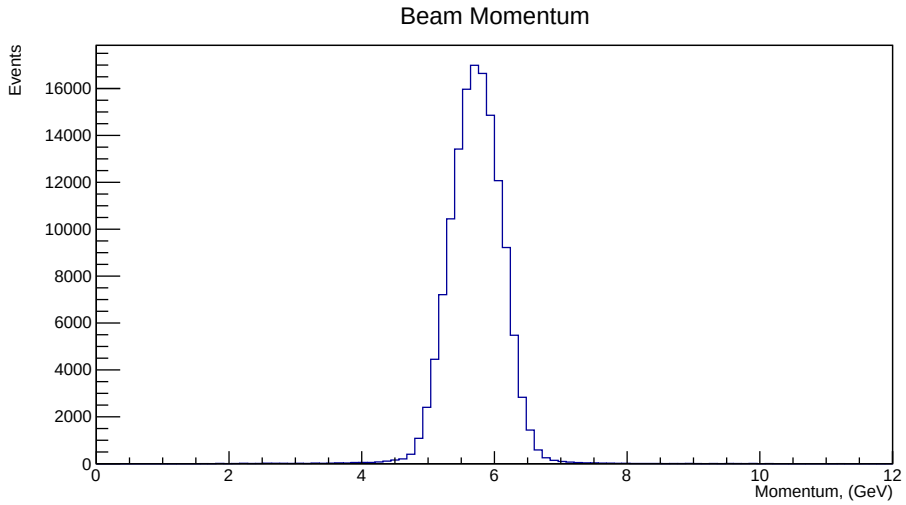
(b) Trigger records. A 2D histogram that shows the active trigger types as a function of the trigger time–stamp for a number of events.

Figure 3.18: Examples of online monitoring plots for the timing and trigger data.

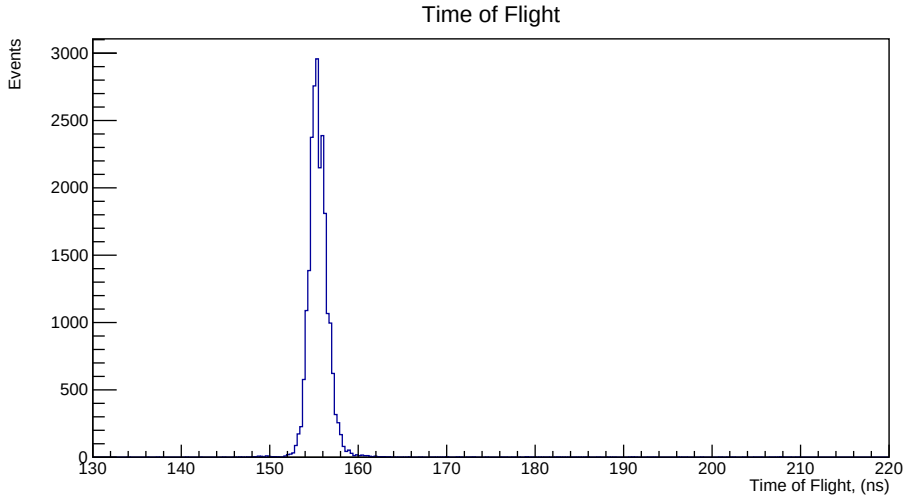
Flask web application[76] that uses Bokeh[77] to render plots for display, ROOT[74] and NumPy[78] are used to process the monitoring data in preparation for display.

Two important modifications were made to Monet for ProtoDUNE–SP: efficiency improvements, which were required to handle the large amounts of data in the event displays, and the addition of a separate rendering server with additional capabilities.

The components of the web interface, and their connections, are outlined in Figure 3.21. When a user makes a request, for example to open a new page of plots



(a) Beam momentum distribution.

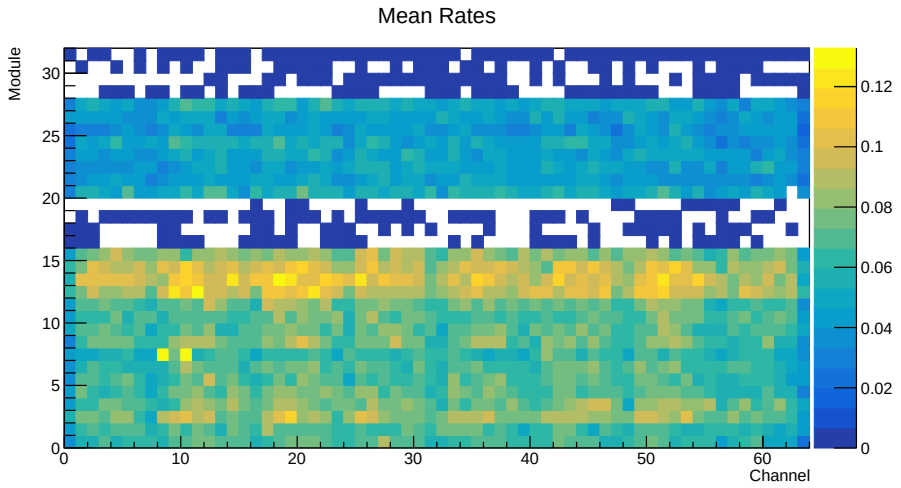


(b) Beam time of flight distribution.

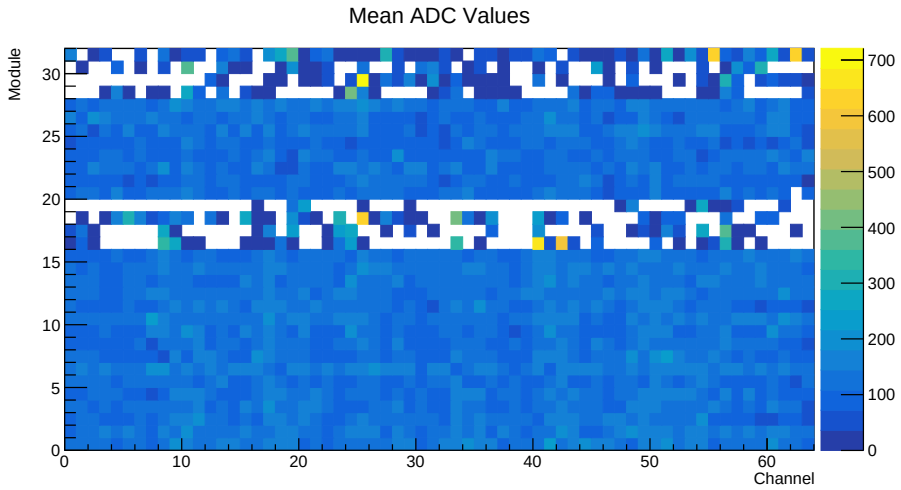
Figure 3.19: Examples of online monitoring plots for the beam instrumentation.

or to interact with a plot, the request is sent to the display server that decides what to do with it. Depending on the nature of the request, the display server will then interact with the histogram database and/or the rendering server, before updating the plots and displaying them to the user.

The plots in the monitoring files are organised into pages in Monet. Each page can contain any number of plots that are organised into an array on the screen, an example of this is shown in Figure 3.22. The histogram database is responsible for the details of each of these pages, this includes the plots on each page, their locations, and the relevant rendering options for each plot. When a new page is



(a) Rate of hits measured by each CRT channel. The channels are arranged into a 2D histogram based on the module number and the channel number within each module. Due to an issue on two of the readout boards, some of the CRT channels had a significantly lower rate of hits.



(b) Average ADC value of hits measured by each CRT channel. The channels are arranged into a 2D histogram based on the module number and the channel number within each module.

Figure 3.20: Examples of online monitoring plots for the cosmic–ray tagger.

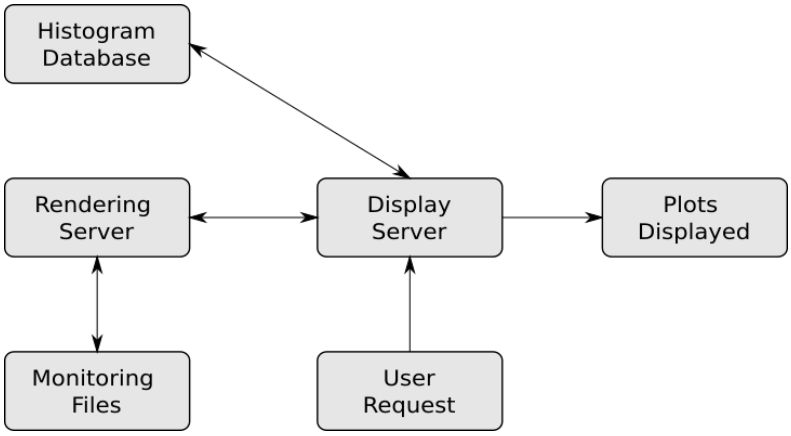


Figure 3.21: Data flow diagram for the web interface in the ProtoDUNE–SP online monitoring system. The display server is responsible for handling user requests by communicating with the histogram database and rendering server, which collect plots from the monitoring files to display to the user.

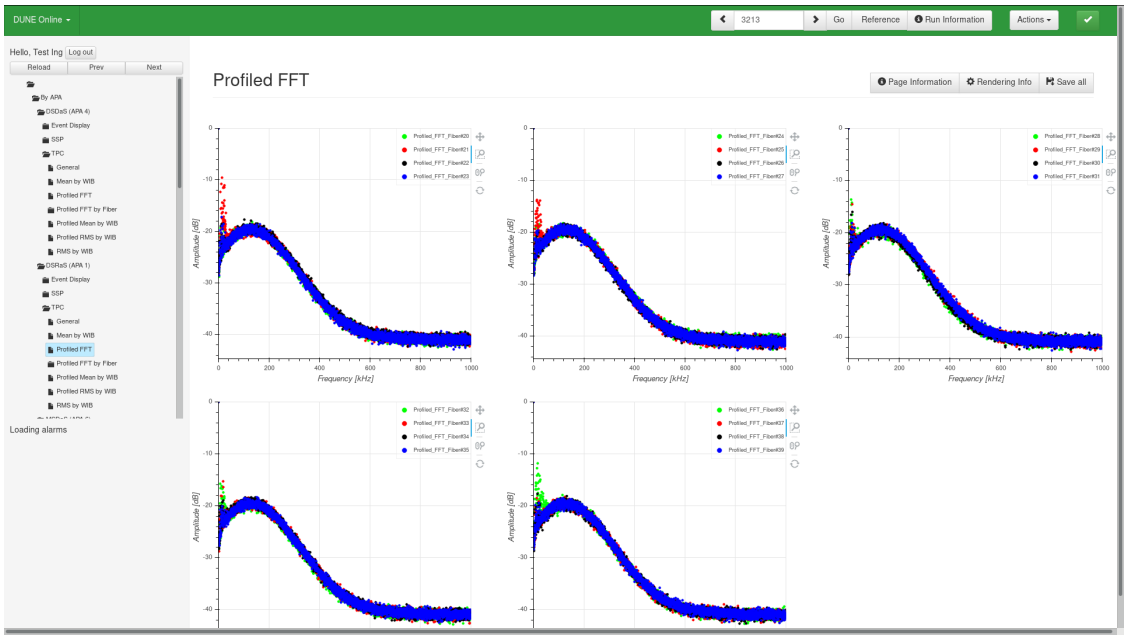


Figure 3.22: An example of an online monitoring page, containing the profiled FFTs for all FEMBs on APA 1. Users are able to select which page to display from the folder list on the left, select a run number in the top bar, and interact with each plot, e.g. zooming in on a specific region.

requested, the display server retrieves these details from the histogram database before requesting the rendering server to render the plots.

The rendering server is responsible for visualising the plots from the monitoring ROOT files, it is capable of rendering all 1D and 2D ROOT plots. In LHCb's Monet, rendering is handled directly on the display server, however in ProtoDUNE–SP an

additional rendering server was added. By separating the rendering processes onto their own server, additional functionality is possible. The main advantage of the rendering server is that it enables the OM to utilise a dynamic binning algorithm to scale the resolution of the displayed images in response to user interaction. This is mainly used for the event displays; the event display is initially loaded in low resolution, but if a user is interested in a specific region of the event display they can zoom in on the region of interest, which is then displayed at a higher resolution by the dynamic binning algorithm as demonstrated in Figure 3.23.

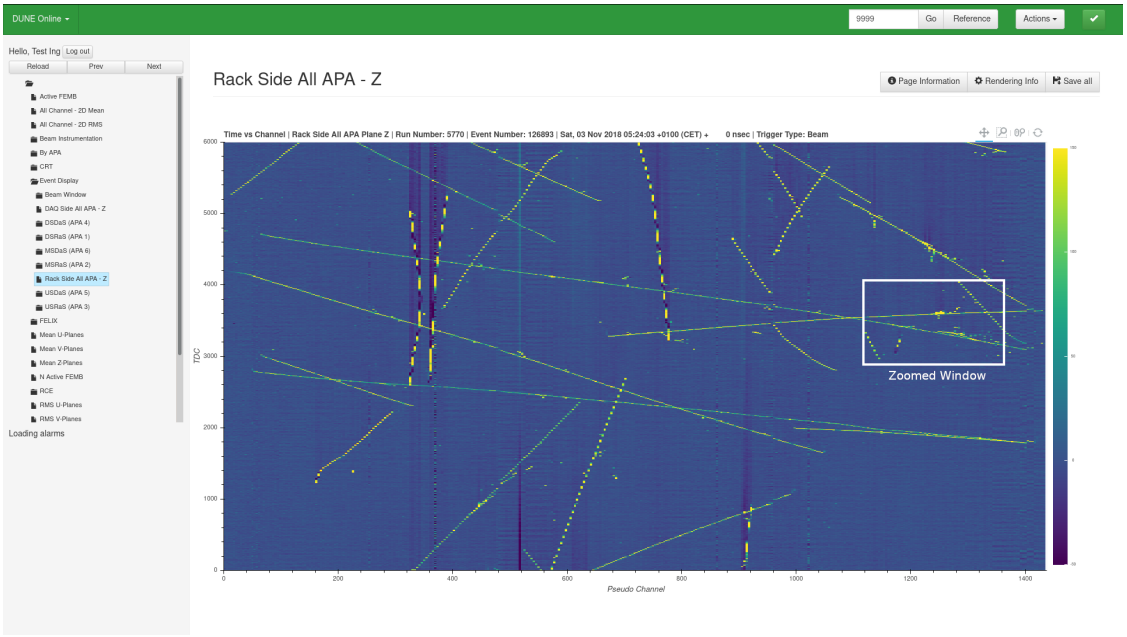
The other significant modification to Monet was to modify the data handling to significantly improve the efficiency of data transfer to the renderer. When Monet was first adapted for ProtoDUNE–SP it was too slow to render some of the larger plots, such as event displays, which would initially take minutes to render. By improving the efficiency of the data transfer from ROOT to Bokeh, the rendering time was reduced to around 4 seconds per event display.

3.8.3 System Performance and Possible Improvements

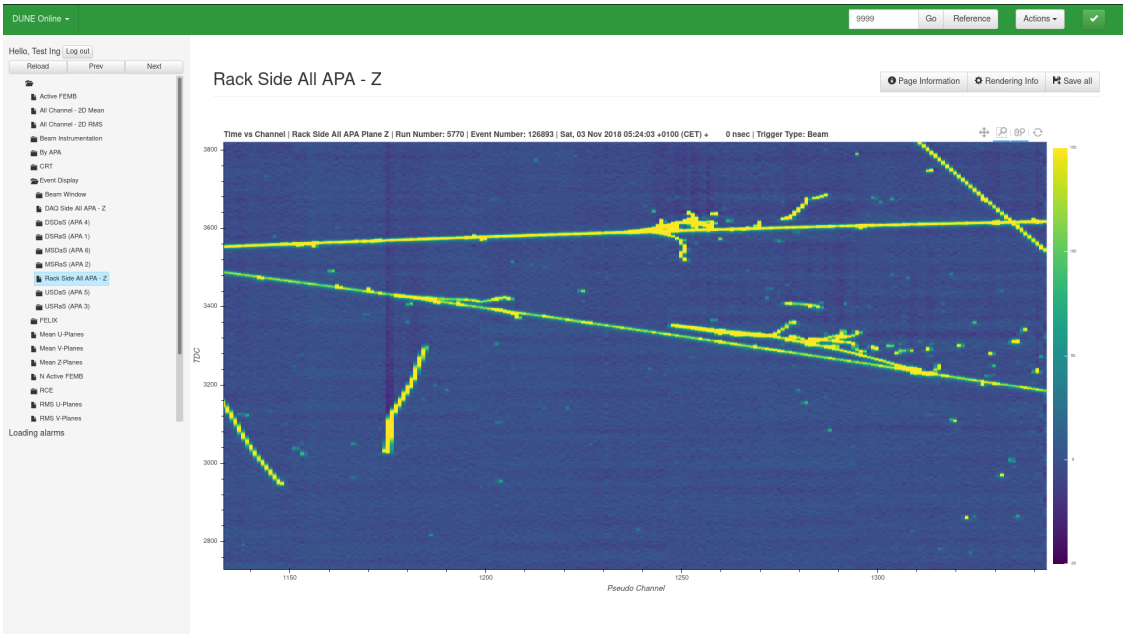
In the ProtoDUNE–SP technical design report[7], the proposed online monitoring system was described as being able to "provide data quality assurance in real–time". In practice, the response time is around one minute. During this time, 10 events will have been processed for each set of plots, however, the event displays and TPC FFTs only contain data from one event.

The performance of the OM was found to be satisfactory for the operation of the ProtoDUNE–SP detector during the beam runs in 2018. However, there is room for improvement, particularly in terms of the turnaround between the data readout and plots being available in the OM. In the current system the latency is made up of four primary components:

- Waiting for the input data file containing the events to be closed
- Processing the data with the Analyser Modules
- Writing the monitoring plots to disk
- Merging TPC monitoring with plots from external systems



(a) The downscaled version of the event display, showing the data from all three APAs on one side of the TPC. The white square represents the zoomed region in Figure 3.23b.



(b) Zoomed region of the event display, which is shown at a higher resolution.

Figure 3.23: Example of the dynamic binning algorithm in the ProtoDUNE–SP online monitoring system.

The latency introduced by each of these components have an approximate ratio of 40:15:30:15.

Input file closure currently makes up the most significant portion of the latency. In the current system, the OM processes have to wait for the event builders to finish processing a batch of 100 events and close the output file before they can access the data. A shared memory system was proposed to tackle this problem, in this system the OM and event builders can simultaneously process the data as it leaves the detector. This would reduce the latency in the OM, by allowing it to analyse the data promptly as it is read out by the detector. Unfortunately, this system wasn't successfully implemented during the first 2018 beam runs of ProtoDUNE–SP, but it remains as a possible improvement to the monitoring system for any future ProtoDUNE–SP runs, and for the DUNE Far Detector monitoring.

Additional improvements could be made to the latency, by using different computing hardware and improving the efficiency of the code. However, at the time of writing there are no specific improvements being investigated in these areas.

Authors Contributions to the ProtoDUNE–SP Online Monitoring

My main contributions to the online monitoring system in ProtoDUNE–SP were made from January 2018 to January 2019, these were:

- Writing code to produce plots for the Analyser Modules
- Developing merger processes to merge the monitoring results from multiple DAQ servers, as well as external plots
- Maintenance and development of the display server including:
  - Significantly improving the efficiency of the plotting algorithms, such that event displays were possible in the OM
  - Integrating a new rendering server into the display server
  - Integrating the histogram database into the display server
- Populating and maintaining the histogram database
- Developing shell scripts for running the online monitoring jobs, and maintaining the online monitoring server
- Ensuring the monitoring software remained up to date with changes to other DAQ software

In addition, I worked as the on–call expert on the online monitoring system during the ProtoDUNE–SP beam run in 2018. I have remained involved with the OM since, mostly for small maintenance tasks.

4

Electromagnetic Energy Loss in Liquid Argon

Contents

4.1 Electromagnetic Energy Loss in Matter 73
4.1.1 Energy Loss for Heavy Charged Particles 74
4.1.2 Energy Loss for Electrons 75
4.1.3 Energy Loss for Photons 78
4.2 Interaction Signatures in ProtoDUNE–SP 79
4.3 Scintillation Light in Liquid Argon 82

The energy loss of particles in liquid argon has important implications for the reconstruction of different particles in a LArTPC, and will be relevant for the reconstruction algorithms developed in Chapters 6 and 7. This chapter will cover in more detail the theory of electromagnetic energy loss in liquid argon, highlighting the important features of the energy loss for muons, electrons, and photons. Much of the discussion in this chapter is summarised from [29].

4.1 Electromagnetic Energy Loss in Matter

In matter, charged particles lose energy through a number of small successive collisions with the electrons in the material, and by radiative processes that produce

additional particles in the material. The relative importance of the collision and radiative stopping power depends on the mass and the energy of the particle. For most particles, which are heavy compared to the electron, radiative energy losses are not important until very high energies, e.g. for a muon they are not important until momenta of around 100 GeV, however, radiative energy loss becomes important for electrons at tens of MeV[29]. As a result, different theories are used to describe the energy loss of heavy particles and electrons in matter.

4.1.1 Energy Loss for Heavy Charged Particles

For heavy particles, such as muons, at moderate energies, the mean rate of energy loss per unit distance is described by the Bethe equation,

$$-\left\langle \frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right]. \tag{4.1}$$

The constants in this equation are detailed in reference [29]. Z and A are the atomic number and mass number of the medium, z is the charge of the scattering particle, W_{max} is the maximum energy transfer possible in a single collision, I is the average ionisation energy, and δ is a density effect correction that is relevant in solids and liquids.

Three important features of the energy loss in the Bethe formula are the minimum ionising region, the relativistic rise, and the Bragg peak, these regions can be seen in Figure 4.1, which shows the dE/dx for muons in argon as a function of momentum.

Delta Rays

Another feature of the electromagnetic energy loss of heavy particles that impacts reconstruction in LArTPCs are delta rays. Delta rays are energetic electrons that are knocked out of their atom when it collides with the heavy particle. In liquid argon detectors, these electrons are seen as small electron tracks that protrude from muon tracks.

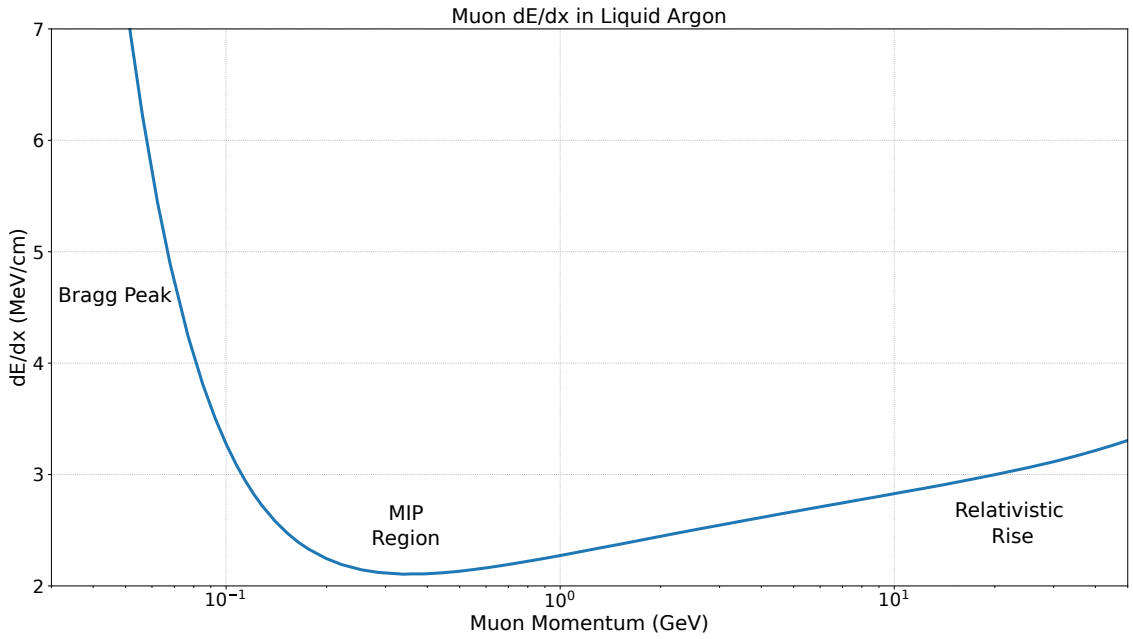


Figure 4.1: Stopping power as a function of energy for muons in liquid argon. Data from [79].

4.1.2 Energy Loss for Electrons

Electrons and positrons undergo different electromagnetic scattering processes in matter, Møller scattering and Bhabha scattering respectively[80, 81]. These processes, which dominate electron and positron energy loss at low energies, have different cross sections, which modify the energy loss in each case. At higher energies, radiative processes such as bremsstrahlung dominate. The two components of the electron stopping power are known as the collision stopping power and the radiative stopping power.

Collision Stopping Power

The collision stopping power of electrons and positrons is calculated with a similar method to the heavy particle stopping power, where individual collisions are considered in succession. The main difference in the calculations is the cross sections used in the calculations, for electrons the Møller scattering cross section is used, and for positrons the Bhabha scattering cross section. Due to the electrons and positrons having the same mass as their targets, the maximum energy transfer in a single collision, W_{max}, is the total kinetic energy. However, this value is halved

for the case of electrons, due to the convention of calculating the stopping power for the final state electron with higher kinetic energy.

The stopping power based on Møller scattering of electrons gives,

$$-\left\langle \frac{dE}{dx} \right\rangle = \frac{1}{2} K \frac{Z}{A} \frac{1}{\beta^2} \left[\ln \frac{m_e c^2 \beta^2 \gamma^2 \left\{ m_e c^2 (\gamma - 1)/2 \right\}}{I^2} + (1 - \beta^2) - \frac{2\gamma - 1}{\gamma} + \frac{1}{8} \left(\frac{\gamma - 1}{\gamma} \right)^2 - \delta \right],$$

while Bhabha scattering, which governs the positron stopping power, gives,

$$-\left\langle \frac{dE}{dx} \right\rangle = \frac{1}{2} K \frac{Z}{A} \frac{1}{\beta^2} \left[\ln \frac{m_e c^2 \beta^2 \gamma^2 \left\{ m_e c^2 (\gamma - 1) \right\}}{2I^2} + 2 \ln 2 - \frac{\beta^2}{12} \left(23 + \frac{14}{\gamma + 1} + \frac{10}{(\gamma + 1)^2} + \frac{4}{(\gamma + 1)^3} \right) - \delta \right],$$

where the terms have the same meanings as in Equation 4.1[29].

Radiative Stopping Power

Above a few tens of MeV, electrons lose most of their energy through the emission of bremsstrahlung photons. Detailed discussion of the energy loss due to bremsstrahlung emission is outside the scope of this thesis, detailed discussions are provided in [29, 82]. A simplified model, which highlights the important factors relevant for the work in this thesis, will be discussed here.

At high energies, where the radiative energy loss is dominant, the energy loss of electrons can be approximated as energy loss over a length scale, X_0, known as the radiation length. In this approximation, the energy loss per unit distance due to bremsstrahlung is,

$$-\left(\frac{dE}{dx} \right)_{brem} = \frac{E}{X_0}. \tag{4.2}$$

The radiation length, X_0, can be parametrised as,

$$\begin{aligned} \frac{1}{X_0} &= 4\alpha r_e^2 \frac{N_A}{A} \left\{ Z^2 \left[L_{rad} - f(Z) \right] + Z L'_{rad} \right\} \\ f(Z) &= \alpha^2 Z^2 \left[\frac{1}{1 + \alpha^2 Z^2} + 0.20206 - 0.0369 \alpha^2 Z^2 + 0.0083 \alpha^4 Z^4 - 0.0002 \alpha^6 Z^6 \right], \end{aligned} \tag{4.3}$$

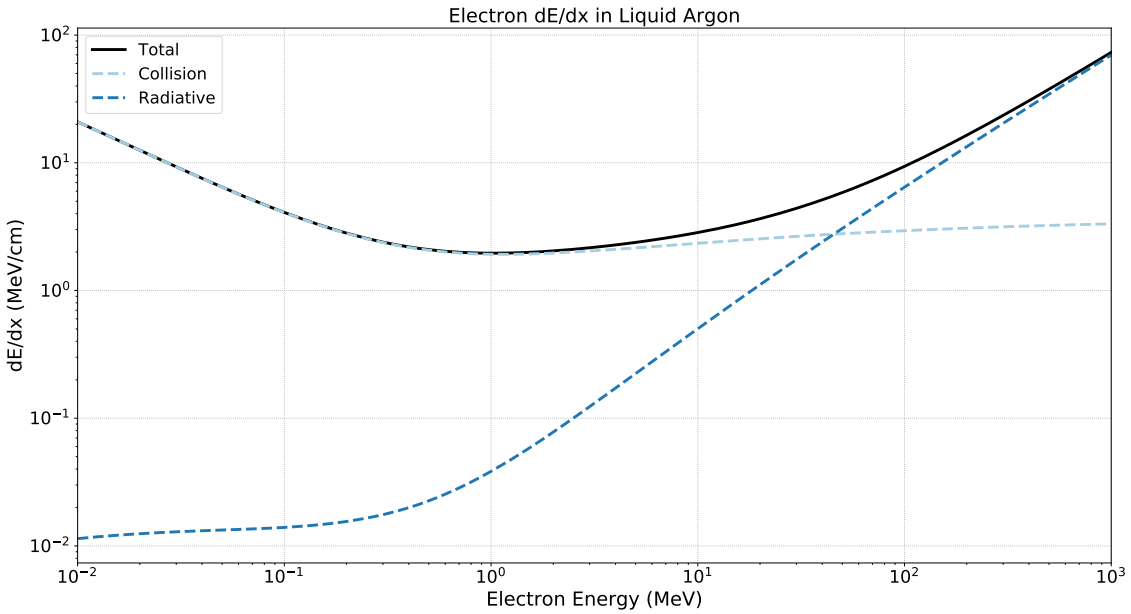


Figure 4.2: Stopping power as a function of energy for electrons in liquid argon. The contributions from collision and radiative stopping power are shown as dashed lines. The total stopping power is shown as a solid line. Data from [84].

where L_{rad} and L'_{rad} are the so–called radiation logarithms, which depend on the atomic number of the material[82].

Critical Energy

Critical energy is often defined as the energy at which the collision and radiative stopping power are equivalent. Other definitions are also used, such as the definition by Rossi, the energy where the ionisation loss per radiation length is equal to the electron energy[83]. Rossi's definition is equivalent to using the approximate dE/dx in Equation 4.2[29]. The value of the critical energy has important implications for reconstruction algorithms, because different approaches are often required above and below the critical energy.

The critical energy is similar but slightly different for electrons and positrons, in liquid argon they are both around 32 MeV based on the Rossi definition[79]. This can be seen in Figure 4.2, which shows the total electron stopping power in liquid argon, in addition to the collision and radiative components that make up the total stopping power.

4.1.3 Energy Loss for Photons

A number of processes contribute to the energy loss of photons in matter, brief descriptions of the main processes are given below.

Photoelectric Effect

The photoelectric effect occurs when a photon collides with an atom, X, the photon is absorbed and an electron is emitted from the atom. As a result, the atom is ionised.

$$\gamma + X \rightarrow e^- + X^+$$

Compton Scattering

Compton scattering occurs when a photon scatters incoherently from an electron within an atom. The electron is typically liberated from the atom, and the photon loses some of its energy.

$$\gamma + e^- \rightarrow \gamma + e^-$$

Pair Production

Pair production is the production of an electron positron pair, in the vicinity of an external electric field. During this process, the photon is destroyed to produce the electron positron pair. In matter, the electric field could be provided by either the electrons in the atom, or the nucleus of the atom.

$$\gamma \rightarrow e^+ + e^-$$

The cross section for these effects vary as a function of photon energy. The cross sections for each process in liquid argon, as well as the total photon cross section, are given in Figure 4.3. The Compton scattering cross section is dominant from around 0.1 MeV to 10 MeV, after which the pair production cross section dominates.

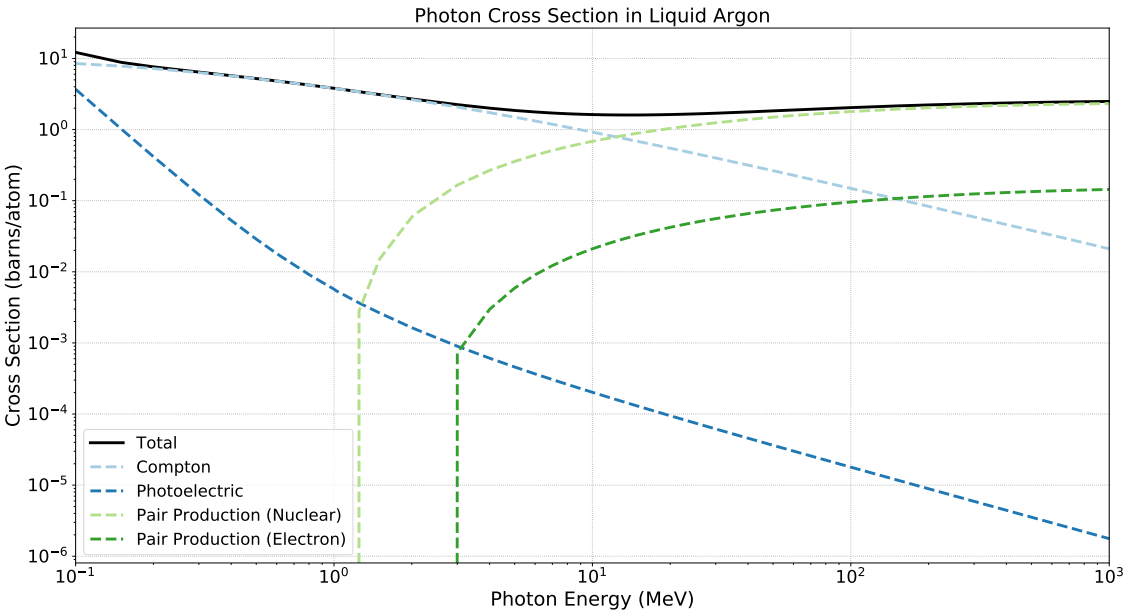


Figure 4.3: Photon interaction cross sections in liquid argon. The contributions from Compton scattering, photo–electric effect, and pair production are shown as dashed lines. The total photon cross section is shown as a solid line. Data from [85].

Photon Mean Free Path

The mean free path of a photon is defined as the average distance travelled by the photon before it interacts with the material. The mean free path for photons has two main components for photons in the MeV range, which are Compton scattering and pair production. The mean free path is given by $\lambda = 1/(n\sigma)$, where n is the number density of targets and σ is the cross section per target. The contribution to the mean free path from pair production is related to the radiation length for electrons, X_0 from Equation 4.3, by $\lambda_{PP} = (9/7)\ X_0$[29]. The mean free path for photons in the MeV range is shown in Figure 4.4, along with the main contributing cross sections.

4.2 Interaction Signatures in ProtoDUNE–SP

The difference in energy loss for each particle species at different energies leads them to leave distinct signatures in the detector. Figure 4.5 shows the typical signatures of four types of interaction in a liquid argon TPC, for events from ProtoDUNE–SP data.

The most common interaction in ProtoDUNE–SP is that of a cosmic–ray muon, seen in Figure 4.5a. Muons, and other heavy particles, leave long tracks in

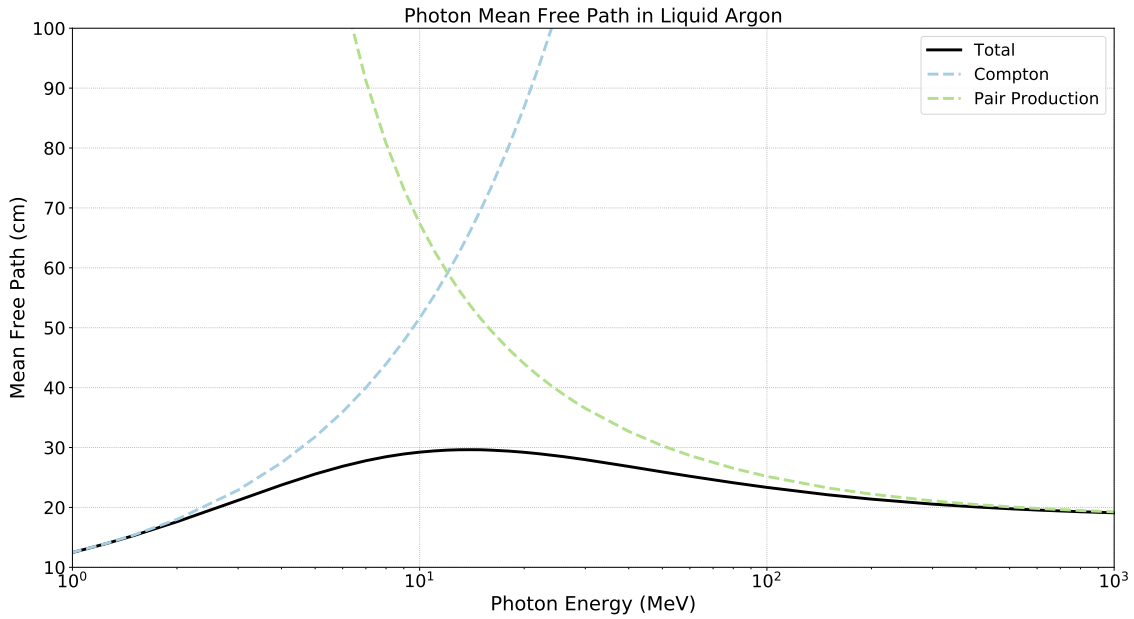
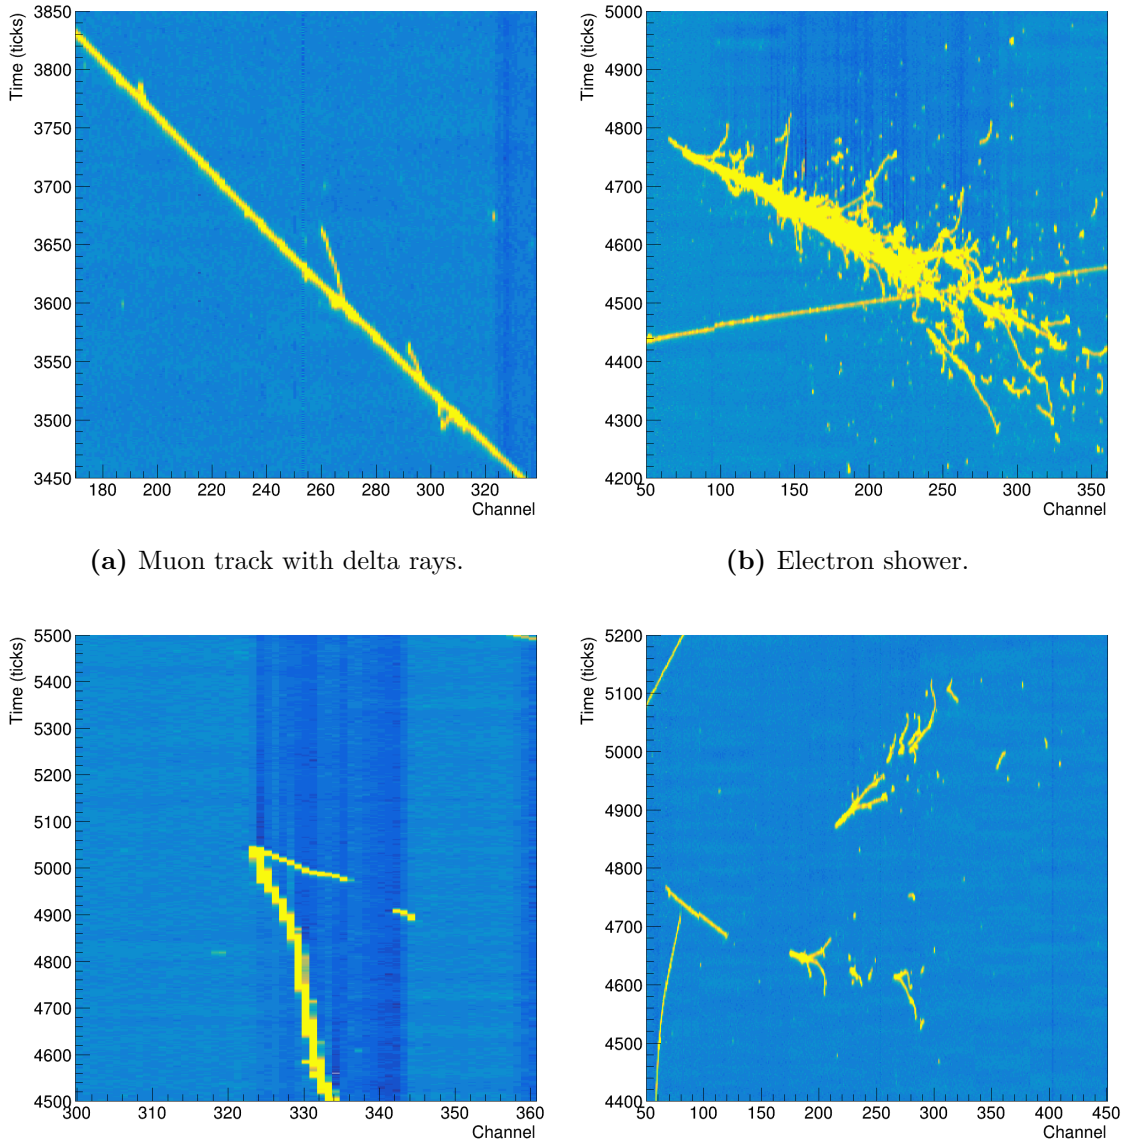


Figure 4.4: Photon mean free path in liquid argon. The mean free path based on the cross sections for Compton scattering, photo–electric effect, and pair production are shown as dashed lines. The total photon mean free path is shown as a solid line. Data from [85].

the detector that deposit around 2.1 MeV/cm in the minimum ionising region. Alongside muon tracks, relatively high energy electrons, known as delta rays, are occasionally produced. Delta rays can be seen as electron activity along the length of the muon track.

Electrons in ProtoDUNE–SP have two distinct signatures, which depend on the energy of the electron. For high energy electrons, with energies above around 300 MeV, radiative energy loss dominates and showers are produced, an example of an electron shower can be seen in Figure 4.5b. These showers are the result of a cascade of electrons, which are produced by a chain of bremsstrahlung photons and electron positron pairs. At energies in the tens of MeV range electrons have a different signature, which consists of a small electron track accompanied by a number of small ionisation energy deposits due to radiated photons.

A common source of tens of MeV electrons in ProtoDUNE–SP are Michel electrons, which are the electrons produced when a muon decays at rest in the detector. An example of a typical Michel electron event is shown in Figure 4.5c. The Michel electron is accompanied by the incoming muon in the image, which



(a) Muon track with delta rays.

(b) Electron shower.

(c) Michel electron.

(d) Photon showers from a π^0 decay.

Figure 4.5: Typical particle signatures in ProtoDUNE–SP data for candidate particles from different particle species. These images are made from the raw detector readout on the collection plane, before noise filtering and deconvolution are applied.

enters from the bottom, because the lifetime of the muon is shorter than the time it takes for the ionisation charge to drift away from the stationary muon.

As with electrons, photon interactions in liquid argon have two distinct signatures at high and low energies. Low energy photons are typically produced by the interactions of low energy electrons in the liquid argon, such as the Michel electron interaction in Figure 4.5c. When these photons interact they produce small isolated energy deposits, which are created by the ionisation from either Compton scattered electrons or an electron positron pair.

At high energies, photons produce similar interactions to high energy electrons, which produce electromagnetic showers in the liquid argon. There is a slight difference in the dE/dx in the first few cm of electron and photon showers, because the start of photon showers contain two electrons whereas the start of electron showers only have one. This difference can allow for electron and photon showers to be distinguished from each other[86]. A common source of photon showers in ProtoDUNE–SP are π^0 decays, shown in Figure 4.5d, which produce a pair of photons.

4.3 Scintillation Light in Liquid Argon

The second source of electromagnetic energy loss in liquid argon is due to the production of scintillation light. As the work in this thesis does not make use of the scintillation energy depositions, the mechanisms involved in the production of scintillation light will only be briefly summarised here for completeness.

Scintillation light in liquid argon is produced by radiative decay of excited argon dimers (or eximers), which are produced when excited argon atoms combine with a ground state argon atom to form an excited Ar_2 molecule. Two excited dimer states contribute to the scintillation light production, a singlet and a triplet, with lifetimes of 6 ns and 1.5 μs respectively[87]. When the dimers decay they produce two ground state argon atoms, and release scintillation light isotropically in the form of 128 nm photons. Argon is transparent to these photons and, therefore, the scintillation light can travel large distances before being collected by a photon detection system.

5
Neural Networks

Contents

5.1 Artificial Neural Networks 84
5.1.1 Convolutional Neural Networks 86
5.1.2 Residual Neural Networks 88
5.2 Supervised Learning 89
5.2.1 The Backpropagation Algorithm 90
5.2.2 Regularisation 93

Machine learning (ML) is a field of research that studies algorithms that can learn to make predictions from data, i.e. predict a set of output variables, given a set of input variables. These algorithms typically take the form of a multivariate function, which is used to predict the output[88].

ML algorithms are often classified into four groups based on two distinctions: the expected output and the level of supervision. The first distinction is based on the nature of the output distribution expected from the network; regression algorithms are designed to predict the outputs of a continuous function, whereas classification algorithms aim to separate data into groups. The second distinction, between supervised and unsupervised algorithms, is based on the inputs used during training. In a supervised algorithm the true output is known, and the network's goal is to predict the true output, while in an unsupervised algorithm the true output is

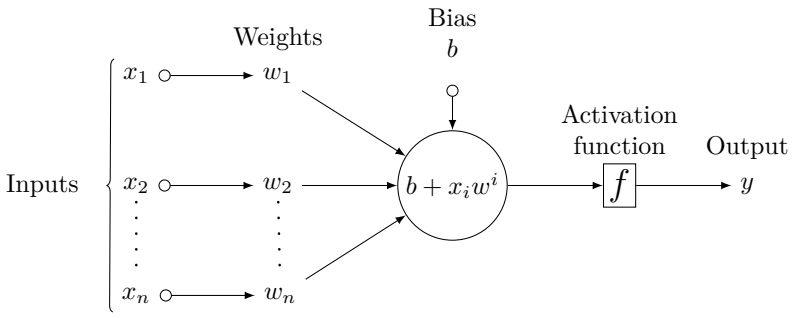


Figure 5.1: Graphical representation of an individual neuron (node) in an artificial neural network.

unknown, and the networks goal is to extract meaningful structure from the data[89].

Chapters 6 and 7 of this thesis describe two examples of the application of Neural networks (NN) for event reconstruction in LArTPC data. These algorithms are classification algorithms based on the supervised learning approach. This chapter will not attempt to provide a full survey of the available ML techniques, which are both numerous and diverse. Instead, this chapter will briefly describe the theory behind NNs, as well as providing details of the techniques used in the subsequent chapters.

5.1 Artificial Neural Networks

Artificial neural networks (ANN) are a class of ML algorithm that draw their inspiration from biological neurons. An ANN consists of a set of nodes, along with connections between those nodes. The set of nodes and connections is often referred to as a graph or architecture. The nodes in the graph take the form of artificial neurons, which pass a number of inputs through an activation function to produce a single output, as depicted in Figure 5.1. The output of each neuron is either distributed to subsequent neurons in the network, or it is part of the output of the network.

One of the most widely used ANNs is the multi–layer perceptron (MLP)[88]; this class of network consists of at least three layers of nodes: an input layer, one or more hidden layers, and an output layer. The layers are connected in a feedforward configuration, such that the graph of nodes contains no cycles. The layers of an MLP are often fully connected or dense, meaning that the output of every node is

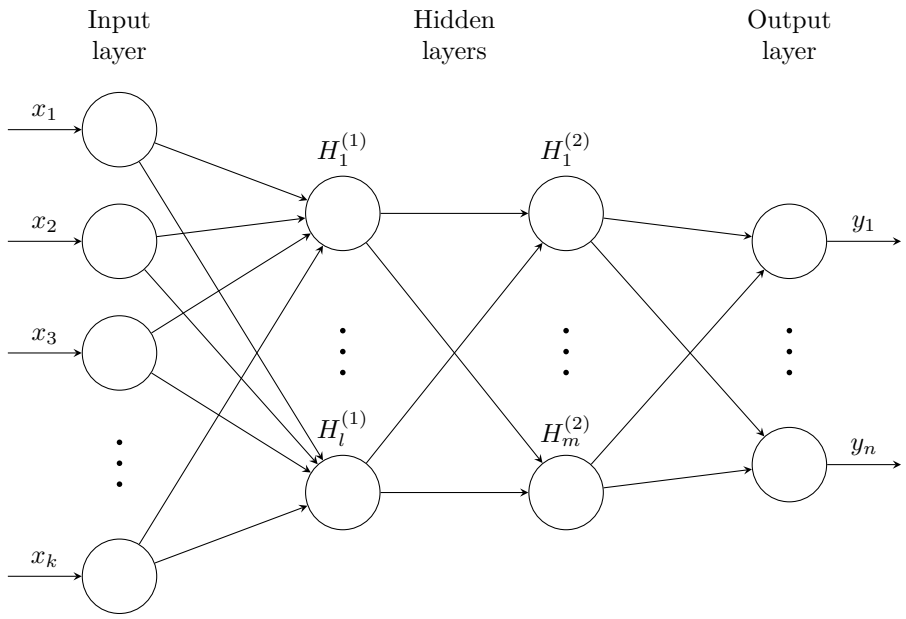


Figure 5.2: A graphical representation of a multi–layer perceptron.

connected to the input of every node in the next layer of the network. An example of a fully connected MLP, with two hidden layers, is shown in Figure 5.2.

Each node receives input from nodes in the previous layer, and uses the inputs to calculate a response function. For the i^{th} node in the j^{th} layer of a network,

$$a^{i,j} = \mathbf{w}^{i,j} \cdot \mathbf{x}^{j-1} + b^{i,j}$$

where $a^{i,j}$ is the response function, $\mathbf{w}^{i,j}$ is the weights vector, $b^{i,j}$ is the bias, and $\mathbf{x}^{j-1}$ is the input vector from the previous layer. This response function is usually passed through a nonlinear activation function, f, to produce the output from the node, y, which is part of the input for the next layer,

$$\left(\mathbf{x}^{j}\right)_i = y^{i,j} = f\left(a^{i,j}\right).$$

Some common choices for the activation function are the sigmoid function, the

hyperbolic tangent, the rectified linear unit (ReLU), and the softmax function[89–91].

$$f(x) = \frac{1}{1+e^{x}} \qquad \text{(Sigmoid)}$$

$$f(x) = \frac{e^{x}-e^{-x}}{e^{x}+e^{-x}} \qquad \text{(Tanh)}$$

$$f(x) = \max(0, x) \qquad \text{(ReLU)}$$

$$f(x_i) = \frac{e^{x_i}}{\sum_j e^{x_j}} \qquad \text{(Softmax)}$$

In the softmax function, the index i indicates the current node, and the index j includes all nodes in the current layer. This construction ensures that the outputs from a softmax layer sum to one, and as such this unit is commonly used in categorisation tasks when the output must belong to one of a given set of categories. The benefits and drawbacks of these common activation functions will be discussed later in this chapter, when the modification of weights with the backpropagation algorithm will be discussed.

The weights and biases of the nodes in an MLP can be adjusted to make accurate predictions of data. For a classification task, there are typically as many output nodes as there are classes, with output values in the range zero to one. The output value of a classification network quantifies how well the input represents each class, based on the networks prediction. In principle, MLPs are able to approximate any function to arbitrary precision with a single hidden layer[92]; however, there is no limit on the number of nodes required in order to achieve a good approximation. In practice, networks with additional hidden layers can attain the required precision with fewer nodes than a network with a single hidden layer, however, there are diminishing returns with more than a few hidden layers[88, 89].

5.1.1 Convolutional Neural Networks

The convolutional neural network (CNN)[91, 93, 94], is an extension of the MLP, which has had considerable success particularly in image classification tasks. These architectures address some of the drawbacks of the traditional MLP, particularly when evaluating data with a high dimensional input. When working with high

dimensional inputs, networks with a fully connected network architecture tend to have a large number of neurons and high computational cost. In addition, for spatially correlated data, such an architecture does not take into account the local spatial structure of the data, instead focussing on all the data at once. A CNN attempts to resolve these issues by exploiting the local spatial structure of the data.

In a CNN the singular input neurons of a traditional MLP are replaced by convolutional kernels. A convolutional kernel is a matrix containing a set of weights, which are multiplied pixel–by–pixel with small regions of the input image. The convolution operator is defined as,

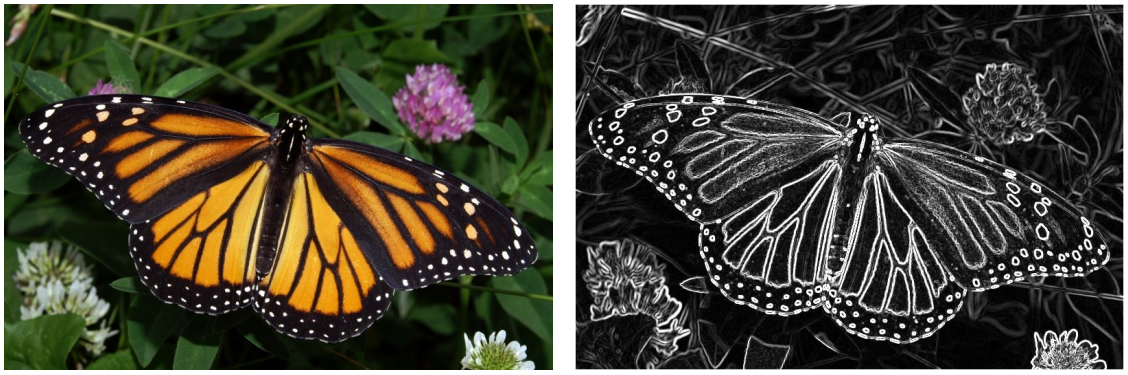
$$(x * y)_i = \sum_{j=-\infty}^{\infty} x_j \, y_{i-j},$$

where x and y are discrete sequences. In the machine learning context, y represents the image, and x the convolutional kernel, which are both finite in extent. Therefore, they are defined to be zero outside of their boundaries, and the sum range can be reduced to $[0, N]$, where N is the number of pixels in the kernel. The convolution operation is applied to all pixels in the input image, which are then replaced by the resulting convolution in the next layer of the network. Many convolutional kernels are used in each layer of the network, each producing an output image that is passed to the next layer.

Convolutional kernels are sometimes referred to as feature detectors, which emphasises the fact that each kernel identifies a given feature in the data, based on the weights in the kernel matrix. The output image from each kernel is known as a feature map, reflecting the fact that they represent the spatial distribution of the features learned by a given feature detector. A common type of feature detector is the edge detector, such as the Sobel edge detector[95]. Figure 5.3 shows an example of the feature map produced when a Sobel edge detector is applied to an image.

Pooling

The use of images as network input often drastically increases the number of input parameters for a network. When combined with a large number of convolutional



(a) Input Image. **(b)** Feature Map.

Figure 5.3: Application of a Sobel edge detector to an image. Images from [96].

filters this can lead to a dramatic increase in the computational cost of training a network. Pooling[94] is a downsampling technique designed to reduce the number of parameters in the network and, therefore, reduce the computational cost of making predictions. In pooling algorithms the data from each $m \times n$ region of the input is downsampled to a single value, the downsampled image is used as the input for the next layer of the network. Two common pooling algorithms are max pooling and average pooling; in max pooling the maximum value from within the downsampling region is used, whereas average pooling uses the average value from the region.

While convolutions are able to extract features from the data, they typically are not able to make predictions in the desired format for classification tasks. Therefore, after convolutions, data is typically flattened and passed through one or more dense layers in order to produce the classification prediction. A graphical representation of a typical CNN architecture is shown in Figure 5.4, it includes convolutional, pooling, and dense layers.

5.1.2 Residual Neural Networks

Adding more layers to a neural network usually improves the performance, typically with each subsequent layer, the network is able to learn more complex features of the image with which to make it's classification. In fact, adding more layers should never decrease the performance of a network, as the new layers could be set to the identity, which would result in the same output from the network. However,

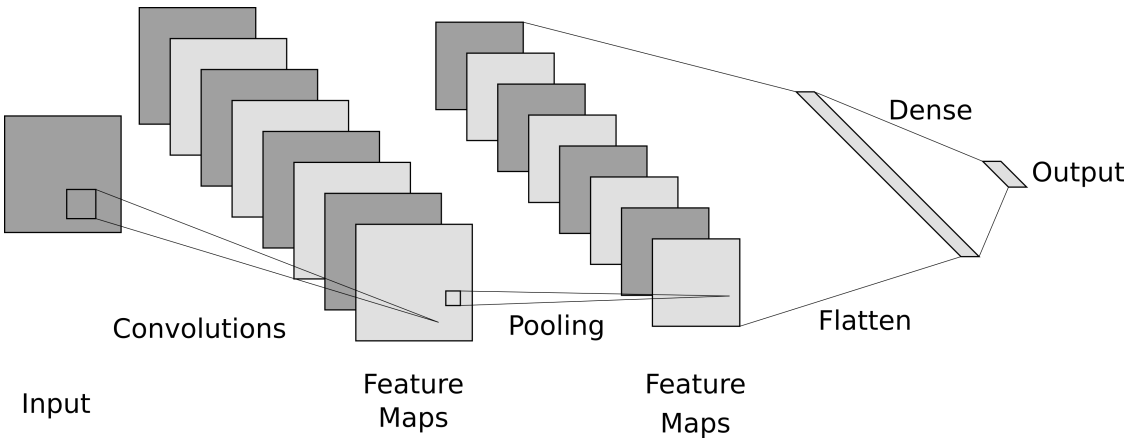


Figure 5.4: A graphical representation of a convolutional neural network. Figure generated with [97].

it has been demonstrated that adding too many layers to a simple neural network will reduce performance, suggesting that the optimisation of larger networks leads to some form of degradation[98].

One method that has been shown to counter the degradation of deep networks, is the use of residual neural networks (ResNet). In a ResNet, the input to each layer is combined with the results of the convolutions as part of the input to a future layer; in effect, the identity transformation has been added to the set of operations in the layer. This results in the weights of each layer remaining as small perturbations around the identity, rather than large corrections that would be required to build up the output from zero. This is a common qualitative intuition for why this method improves optimisation of the network weights. An example of these shortcut connections is shown in Figure 5.5[98].

5.2 Supervised Learning

Supervised learning is the process of teaching a neural network to accurately predict the correct output, or ground truth, for a given input. The network is optimised by minimising the loss function, which quantifies the difference between the network output and the ground truth. The choice of loss function is dependent on the problem being considered and, therefore, specific examples will be discussed in subsequent chapters, which will detail the CNN algorithms developed as part of this thesis.

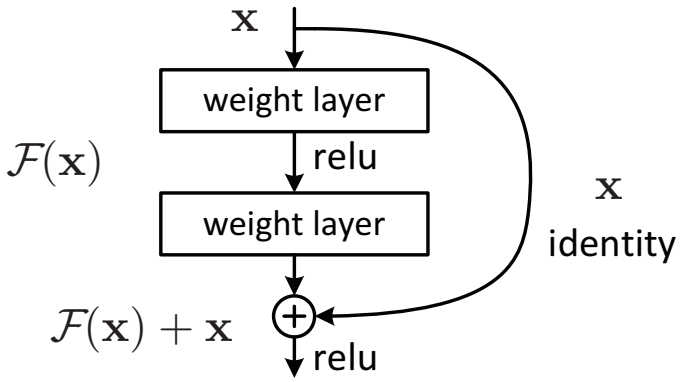


Figure 5.5: An example of a shortcut connection in a convolutional neural network. Figure from [98].

The output of a neural network, $g(x)$, is the result of a chain of function evaluations and matrix multiplication; for a network with n layers,

$$g(x) = f^n(W^n f^{n-1}(W^{n-1} \ldots f^1(W^1 x) \ldots).$$

where g is the output of the network, x is the input, f are the activation functions for each layer, and W are the weights for each layer. The loss of the network is given by the loss function, L, evaluated on the output of the network and the ground truth, y,

$$L(y, g(x)).$$

To make accurate predictions with a NN, quantified by the loss function, the weights and biases of the network need to be optimised. The backpropagation algorithm[99] is currently the most widely used algorithm for minimising the loss function by adjusting the weights and biases in the network.

5.2.1 The Backpropagation Algorithm

Consider the weight for the j^{th} node in the l^{th} layer of a network for an incoming connection from the i^{th} node from the previous layer, W^l_{ij}. The loss function varies with respect to these weights as $\partial L/\partial W^l_{ij}$. If the weights are modified with,

$$\Delta W^l_{ij} = -\eta \frac{\partial L}{\partial W^l_{ij}}, \tag{5.1}$$

for positive learning rates, $\eta > 0$, the loss will be reduced provided it is not already too close to the minimum. The backpropagation algorithm provides a means for calculating the derivatives required to adjust the weights in each layer.

Applying the chain rule to the right hand side of Equation 5.1,

$$\frac{\partial L}{\partial W_{ij}^l} = \frac{\partial L}{\partial a_j^l} \frac{\partial a_j^l}{\partial W_{ij}^l}, \tag{5.2}$$

where a^l are response functions for the nodes in layer l. The derivative of the response function with respect to the weights, is simply the input from the previous layer,

$$\frac{\partial a_j^l}{\partial W_{ij}^l} = \frac{\partial}{\partial W_{ij}^l} \left(\sum_{k=1}^{n_l} W_{kj}^l x_k^{l-1} \right) = x_i^{l-1}, \tag{5.3}$$

where x^{l-1} are the outputs from layer $l-1$, and n_l is the number of nodes in layer l. The other term, which is often denoted as δ_j^l, is known as the error; this term will be discussed shortly. Combining Equations 5.2 and 5.3 gives,

$$\frac{\partial L}{\partial W_{ij}^l} = \delta_j^l \; x_i^{l-1}.$$

The value of the error depends on the loss function being used in each particular model. The error can also be shown to have a recursive relationship via the chain rule,

$$\begin{aligned} \delta_j^l = \frac{\partial L}{\partial a_j^l} &= \sum_{k=1}^{n^{l+1}} \frac{\partial L}{\partial a_k^{l+1}} \frac{\partial a_k^{l+1}}{\partial a_j^l} \\ &= \sum_{k=1}^{n^{l+1}} \delta_k^{l+1} \frac{\partial a_k^{l+1}}{\partial a_j^l}, \end{aligned}$$

where the sum is over all nodes in layer $l+1$. Using the definition of the response function to simplify this further,

$$\begin{aligned} \frac{\partial a_k^{l+1}}{\partial a_j^l} &= \frac{\partial}{\partial a_j^l} \left(\sum_{p=1}^{n^l} W_{pk}^{l+1} f(a_p^l) \right) \\ &= W_{jk}^{l+1} \; f'(a_j^l) \end{aligned}$$

where f is the activation function. The error is, therefore,

$$\delta_j^l = \sum_{k=1}^{n^{l+1}} \delta_k^{l+1} W_{jk}^{l+1} f'(a_j^l),$$

which depends on the value of the error in the following layer.

Based on these manipulations with the chain rule, the backpropagation algorithm works as follows,

1. Calculate the forward pass, while storing the results for $g(x)$, a_j^l, and x_j^l.
2. Starting at the output layer, calculate the derivatives, going backwards through the network.
    - Store the results for δ_j^l from each layer, to use when calculating subsequent layers.
3. Update the weights according to the update rule in Equation 5.1.

The backpropagation algorithm is used to calculate the gradients in order to adjust weights in a neural network. There are a number of optimisation algorithms for adjusting the weights, such as mini–batch stochastic gradient descent (SGD)[100] and adaptive momentum estimation (Adam)[101]. The details of these algorithms are not important for the discussion in this thesis and, therefore, they will only be mentioned in passing here. A review of gradient descent optimisation algorithms can be found in [102]. Some methodologies are common among many optimisation techniques, three important examples are the ideas of mini–batches, training iterations, and training epochs. A mini–batch is a small sample of the input data, on which the backpropagation algorithm is applied simultaneously, this corresponds to one training iteration. A training epoch consists of a number of training iterations, and corresponds to one cycle through the full training dataset, such that the network has been trained on every training example exactly once.

5.2.1.1 The Choice of Activation Function

The rate of learning in the backpropagation algorithm is dependent on the derivative of the activation function, therefore, the choice of activation function has an impact on learning. There are two important factors when considering the impact of this choice on the rate of learning, these are the computational cost of calculating the derivative, and the gradient of the function.

Traditional activation functions, such as sigmoid and tanh, are usually bounded and differentiable, however the presence of exponentials makes computing the gradients expensive. They also typically have vanishing gradients in the tails,

which can lead to slow learning if the value of the activation is consistently in the tail of the distribution.

The ReLU activation function is much cheaper to compute, and while the gradient is undefined at zero, in practice this is not a problem. One issue with ReLU units, is the vanishing gradient in the negative region. This can be resolved by implementing a small negative gradient below zero, which is known as a leaky ReLU or parametric ReLU[90]. In practice, ReLU units have been found to outperform traditional sigmoid units[103, 104].

5.2.2 Regularisation

Regularisation refers to a set of techniques that are used to prevent over–fitting during the training of a NN. Over–fitting is the tendency of models with a high number of parameters to produce overly complex models, which fit the training data well but do not generalise to new data. Regularisation ensures that the trained network can generalise well beyond the training data. A number of regularisation techniques can be used, such as applying penalty terms to the loss function that are proportional to the square of the weights, ensuring that the weights stay relatively small. In this thesis two regularisation techniques are used: dropout[105] and early stopping[106].

The dropout algorithm sets a random subset of weights in each layer to zero for each training iteration. Each weight has a probability, p, of being set to zero for each iteration. The product of the remaining weights is scaled up by a factor of $1/p$, which ensures that the size of the input to the activation function remains reasonably consistent. This reduced network is then trained using the backpropagation algorithm, and the non–zero weights are updated. During the next iteration, the set of non–zero weights is reselected, and a new reduced network is trained. Finally, after training, the final weights of the network are multiplied by p to give the so–called model average of the reduced networks. The dropout algorithm with $p = 0.5$, has proven to provide robust regularisation, and is a popular technique in the field[89].

"Early stopping (is a) beautiful free lunch", according to Geoff Hinton[107]. It refers to the act of monitoring the loss on the validation set during training, in order to stop training before over–fitting has occurred. This technique can be applied alongside other regularisation techniques, and is used throughout this thesis, where the early stopping technique proposed by Lutz Prechelt[106] has been used. This technique involves stopping training after the first checkpoint when the validation loss increases, and using the weights from the previous checkpoint.

6

Hit Classification with Convolutional Neural Networks

Contents

6.1 Hit Classification with Convolutional Neural Networks 96
6.1.1 Data Preparation 97
6.1.2 Network Architecture 99
6.1.3 Training and Validation 102
6.2 Performance on ProtoDUNE–SP Simulation 103
6.2.1 Comparison with Pandora 107
6.3 Validation on ProtoDUNE–SP Data 110
6.3.1 Systematic Uncertainties 119
6.4 Application in ProtoDUNE–SP Analyses 121

The correct categorisation of particle interactions in the detector is a major challenge faced by any particle physics experiment. This typically starts with identifying low level features of the interactions, which can then be used to gradually build up a picture of the full interaction. In DUNE, the classification of neutrino interactions requires identifying the lepton content of the final state, and in ProtoDUNE–SP, cross section analyses rely on accurately identifying the particles in the interaction. Therefore, it is important to be able to distinguish muons and pions from electrons in LArTPCs, or more generally, tracks from showers.

In order to build up a complete picture of an event, it is useful to begin

by identifying the small features of the interaction, which can then be used to gradually build an understanding of the full event. In ProtoDUNE–SP, the smallest reconstructed features are the hits, which correspond to small charge depositions collected on individual wires. Classifying these hits provides useful information for future analyses, and can potentially be used to aid decision making during event reconstruction.

This chapter will describe an approach to hit classification in ProtoDUNE–SP using machine learning techniques. Section 6.1 will detail an approach to hit classification in LArTPCs based on identifying the source of energy depositions with a convolutional neural network. The performance of this approach will then be analysed with ProtoDUNE–SP simulation and data in sections 6.2 and 6.3 respectively. Finally, Section 6.4 will briefly highlight some of the current applications of the network in ProtoDUNE–SP analyses.

6.1 Hit Classification with Convolutional Neural Networks

Effective tracks–shower separation forms the basis of many reconstruction challenges in DUNE and ProtoDUNE–SP, for example, it is used to define pure calibration samples, such as minimum ionising muons and π^0 decays, and it is an important part of neutrino event reconstruction. Each particle species leaves a characteristic signature in the detector, but the first step in reconstructing them is the same: reconstructing tracks and showers, which can be combined to build the final state.

In a LArTPC, tracks and showers are built from collections of hits; these hits then have to be clustered and identified as track or shower objects. In this section, a method for identifying the source of hits in the ProtoDUNE–SP LArTPC will be described. The classification of each hit is stored as part of the reconstructed output in LArSoft, and can be used by subsequent reconstruction and analysis algorithms.

In addition to track and shower objects, Michel electrons are a useful calibration sample in LArTPCs. Michel electrons are electrons produced when a muon decays at rest, and have an energy spectrum in the range of 0–50 MeV. As discussed in

Chapter 4, the critical energy for electrons in liquid argon is around 30 MeV, at this point the collision energy deposition and the radiative energy deposition are equal. Therefore, Michel electrons have a distinctive signature in LArTPCs, including a significant fraction of both track–like and shower–like energy deposition, and they are included as a separate category in the classification algorithm.

6.1.1 Data Preparation

A CNN was designed to classify hits by using images of the detector readout around the reconstructed hits. The network was trained to predict,

$$[p_t, p_s, p_e] \text{ and } [p_m]$$

where p_t, p_s, p_e, and p_m, are the probabilities for track, shower, empty, and Michel electron classifications respectively. The empty category is included to ensure that the network doesn't learn to assign track–like or shower–like classifications to empty or noisy regions of the data. In addition, because the Michel electron category has an overlap with the shower and track categories, the Michel electron probability is decoupled from the other probabilities. The track, shower, and empty (TSE) probabilities are constrained to sum to one, such that every hit is classified into one and only one of these categories.

The input data for the CNN was made using a small–patch approach, in which a small input image is made for each reconstructed hit. This approach was chosen based on the CPU time and memory requirements for the CNN, when run during the ProtoDUNE–SP reconstruction chain. Another approach could be to create a single large input image for each wire plane and perform semantic segmentation, however, this would trade a higher memory usage for a lower CPU time.

An input image was produced for every reconstructed hit, with the hit being classified at the centre of the image. These images were produced from the wire readout data, after the noise removal and 2D deconvolution steps described in Chapter 3. Each input image is 48×48 pixels, with each pixel being filled with the ADC value from the wire readout data. The width of the image corresponds

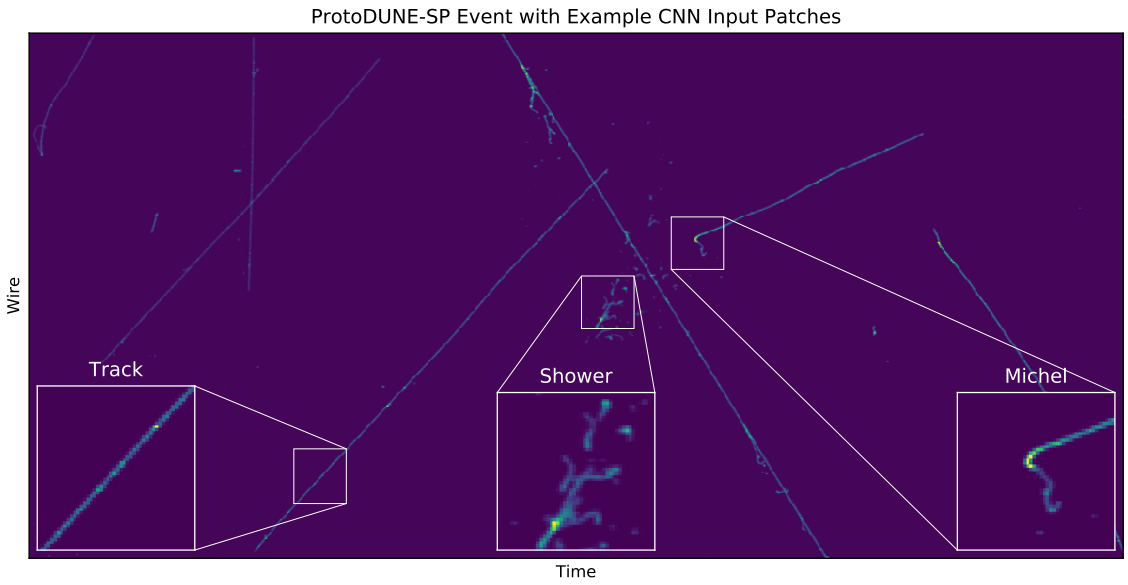


Figure 6.1: Example CNN input images for the track, shower, and Michel classes. The input images are 48×48 pixel images, which are taken from regions of the full readout from a single APA. In the diagram, one example of each non–empty image class is displayed. The white outlines demonstrate the location of each input image in the APA readout.

to one wire per pixel, and the height is the time coordinate. The time data is downsampled, using an average over time samples, such that the spatial dimensions of the image are the same in both directions. Therefore, each image represents around 24×24 cm^2 of wire data. Examples of an input image from each non–empty class are shown in Figure 6.1, which also shows the relationship between the images and the detector readout.

The true classification for each sample is obtained from the simulation, by associating the measured ionisation energy depositions to the corresponding simulated particle. If the reconstructed hit is not associated to a true particle, then this hit is due to noise and the true classification is empty. Additional images were produced in empty regions of the detector, in the vicinity of charged particles, to increase the training sample for the empty category. The truth vectors for different true particle sources are detailed in Table 6.1. These values were chosen based on the typical interaction type in ProtoDUNE–SP for these particle species, which is based on the electromagnetic energy loss in liquid argon for energies in the MeV–GeV range.

The training data for the CNN was built using simulations of the ProtoDUNE–

Ionisation Source	Track–Shower–Empty	Michel
Muons	[1,0,0]	[0]
Hadrons	[1,0,0]	[0]
Michel Electrons	[0,1,0]	[1]
Other Electrons	[0,1,0]	[0]
Noise	[0,0,1]	[0]

Table 6.1: Truth labels for the CNN output for different particle species.

Dataset	Shower	Track	Empty	Michel	Total
Training	13,493,982	9,727,604	2,517,882	731,456	26,470,925
Validation	734,673	562,038	141,388	42,727	1,480,826
Test	764,659	518,805	139,987	39,674	1,463,125
Total	14,993,314	10,808,447	2,799,257	813,857	29,414,876

Table 6.2: Number of input images with each truth label in the training, validation, and test sets.

SP detector in the LArSoft framework; the simulations used were under beam operating conditions and, therefore, included simulations of cosmic–rays, and test beam particles in the range of 1–7 GeV. Around 29 million input images were produced in total. This sample was split into three datasets, the training, validation, and test sets. The training set was used to train the CNN, the validation set was used to monitor the performance of the CNN during training, and the test set was used as an initial verification of the performance of the network after training. Details of the number of patches of each type in these three datasets are detailed in Table 6.2, in total 26 million images were used to train the network.

6.1.2 Network Architecture

The network architecture for this CNN gave the best performance in terms of false–positive rates of the architectures considered given constraints on running time and memory usage during network evaluation. This CNN is run on CPUs as part of the low level reconstruction chain for ProtoDUNE–SP, and its run time is

required to be no more than 10 seconds per event. In addition, the CNN should not increase the maximum memory usage during reconstruction beyond around 4 GB. There is currently ongoing work looking into using GPUs during the network evaluation, which would decrease the evaluation time, and allow more complex architectures to be used. Therefore, there is potential to increase performance of this method in the future.

The network architecture used is shown in Figure 6.2. A simple network architecture was chosen, which was necessary in order to meet the current running requirements. The images are first processed by a convolutional layer with fourty–eight 5×5 filters, this layer extracts feature maps from the data. The responses from the convolutional layer are passed through the leaky ReLU activation function, which is discussed in Chapter 5. The feature maps are then processed by a pair of dense layers, with 128 and 32 nodes respectively. These layers also use the leaky ReLU activation function. After the second dense layer the network is split into two branches in order to make its prediction. The first branch returns the prediction for the TSE categories, and the second branch for the Michel electron category. The first branch uses a softmax activation function, which ensures that the scores for the TSE categories sum to one. The second branch uses a sigmoid activation function, which ensures that the Michel electron scores are bounded between zero and one. The choice of activation functions in the output layers allows for a pseudo–probabilistic interpretation of the scores from the CNN, due to the constraint of the scores summing to one.

Two dropout layers are used in the network, these layers are used as part of the regularisation of the network. A dropout probability of 0.5 is used in both of the dropout layers. The dropout algorithm is discussed in Chapter 5.

The loss of the network is the weighted sum of the losses for the two output branches,

$$L = 0.1 \cdot L_{TSE} + L_{M},$$

where the Michel electron loss is given higher precedence due to the smaller training dataset available for the Michel electron output. The loss function for the TSE

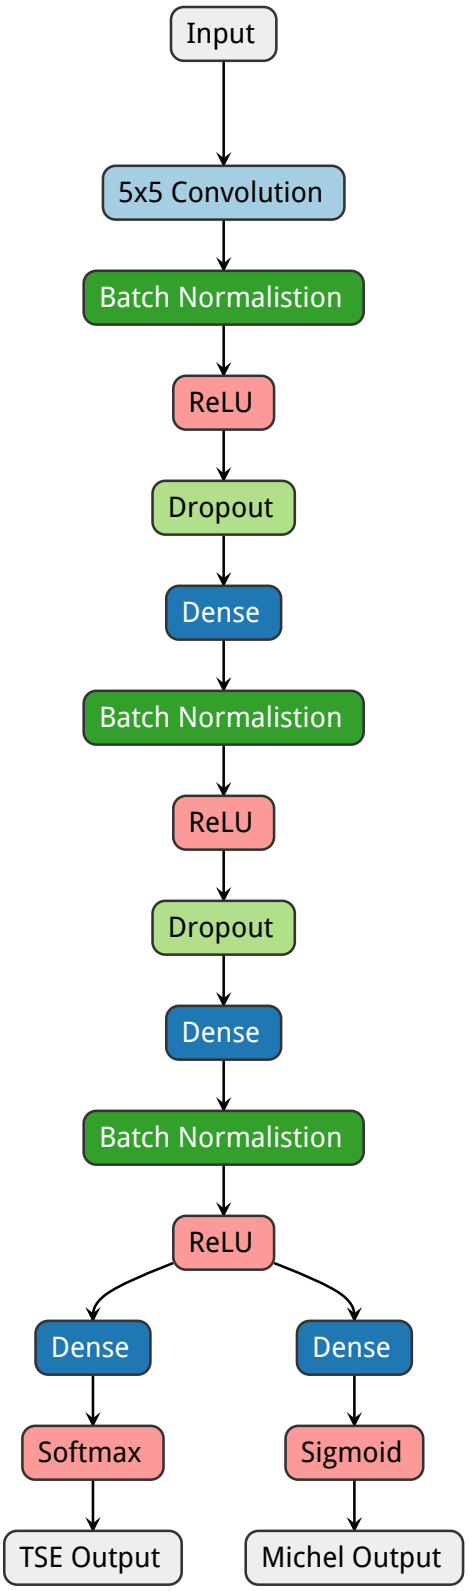


Figure 6.2: Network architecture used for hit classification. Visualisation adapted from the Netron neural network viewer[108].

branch is the categorical cross entropy loss[109], and for the Michel electron branch it is the mean squared error[110],

$$L_{TSE} = -\frac{1}{N}\sum_{j=1}^{N}\sum_{i=0}^{2}(t_j)_i \log(p_j)_i,$$

$$L_M = \frac{1}{N}\sum_{j=1}^{N}(t_j - p_j)^2$$

where t_j and p_j are the truth and the prediction for the j^{th} sample in the training batch, and i sums over all outputs in the TSE branch. These losses are commonly used for classification tasks within the field of machine learning, the categorical cross entropy is related to the negative log–likelihood between the true distribution and the predictions from the CNN[111].

6.1.3 Training and Validation

The TensorFlow[112] library and the Keras[113] application programming interface (API) were used to design and train the CNN. The training was completed on a dedicated ProtoDUNE–SP server at CERN, with an NVIDIA GTX 1080 GPU. Training was monitored using the TensorBoard visualisation toolkit[114], which is part of the TensorFlow library.

The CNN was trained using the mini–batch stochastic gradient descent (SGD) algorithm, including both the momentum and decay algorithms[88]. The momentum algorithm reduces the oscillations of the weights during learning, while the decay of the learning rate allows for rapid learning during early stages of SGD, and increased precision as the model converges.

During training the learning metrics were monitored with TensorBoard. The losses for each branch and the total loss were monitored for the training and validation datasets. The validation loss was calculated once per training epoch, and the training loss was averaged at the end of each epoch. The training and validation losses as a function of training epoch are shown in Figure 6.3. The weights of the network were saved at the end of each epoch, and the final weights were chosen based on the early stopping algorithm discussed in Chapter 5, focussing on the TSE branch of the network, this is highlighted in Figure 6.3.

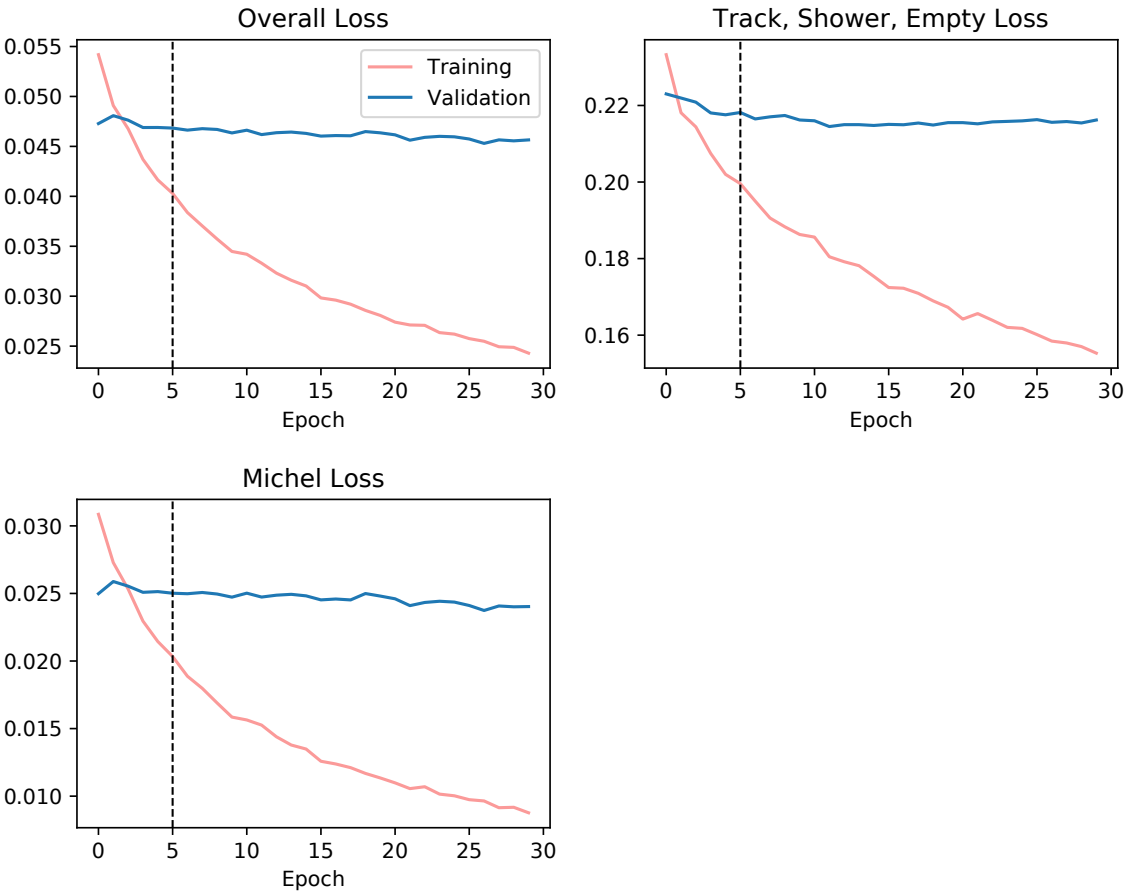


Figure 6.3: Evolution of the validation and training set losses during the training process for the hit tagging CNN. The losses for the TSE branch and the Michel electron branch are shown, as well as the total loss. The weights were frozen based on an early stopping algorithm, focussing on the loss in the TSE branch. The losses with the weights chosen by the early stopping algorithm are indicated by a vertical dashed line on each diagram.

During training, the validation loss remained stable over a number of epochs, which suggests that the dropout algorithm was successful in preventing over–fitting. After training, the network's performance was verified against the test set, which found that the test set losses were comparable with the validation set loss.

6.2 Performance on ProtoDUNE–SP Simulation

The performance of the hit tagging algorithm was evaluated with reconstructed events from ProtoDUNE–SP simulation, the dataset used for this performance analysis was distinct from the training, validation, and test sets.

The distributions of the shower score from the TSE classifier for true shower

hits and all other hits is given in Figure 6.4, where there is a strong separation between the distributions for the shower and track hits, showing that the network has strong discriminating power. In practice, the empty score of the TSE classifier was found to be on the order of 10^{-9} or smaller for all hits, as such,

$$\text{TrackScore} \approx 1 - \text{ShowerScore}.$$

Since the track and shower samples are of similar size, and the scores are related in this way, the results of the analysis of track and shower scores give similar results. Therefore, the following analysis will only focus on the shower score.

The shower classification threshold for subsequent algorithms should be tuned on a case by case basis, however, for this study a simple optimisation strategy is presented in order to quantify the basic network performance. This is based on the F_1 metric, a specific case of the F_β metric[115] that places equal importance on purity and efficiency. The F_1 metric is given by,

$$\frac{1}{F_1} = \frac{1}{2}\left(\frac{1}{\text{purity}} + \frac{1}{\text{efficiency}}\right),$$

where the purity is defined as the fraction of correctly classified shower hits in the sample of all selected shower hits, and the efficiency as the fraction of all true shower hits that were selected as shower hits. The F_1 score was calculated across a range of selection thresholds, this is shown in Figure 6.4. This score peaks at a classifier threshold of 0.72 where,

$$F_1 \approx \text{purity} \approx \text{efficiency} \approx 0.863.$$

The overall performance of the TSE classifier can also be evaluated with a receiver operating characteristic (ROC) curve[116]. The ROC curve shows the true–positive rate vs the false–positive rate for the classifier, as the selection threshold is varied. Figure 6.5 shows two ROC curves for the TSE classifier, one is evaluated in simulation including the space charge effect (SCE), and the other excludes the SCE. Both curves demonstrate that the network is capable of achieving high true–positive rates, while maintaining low false–positive rates. In addition, there is a

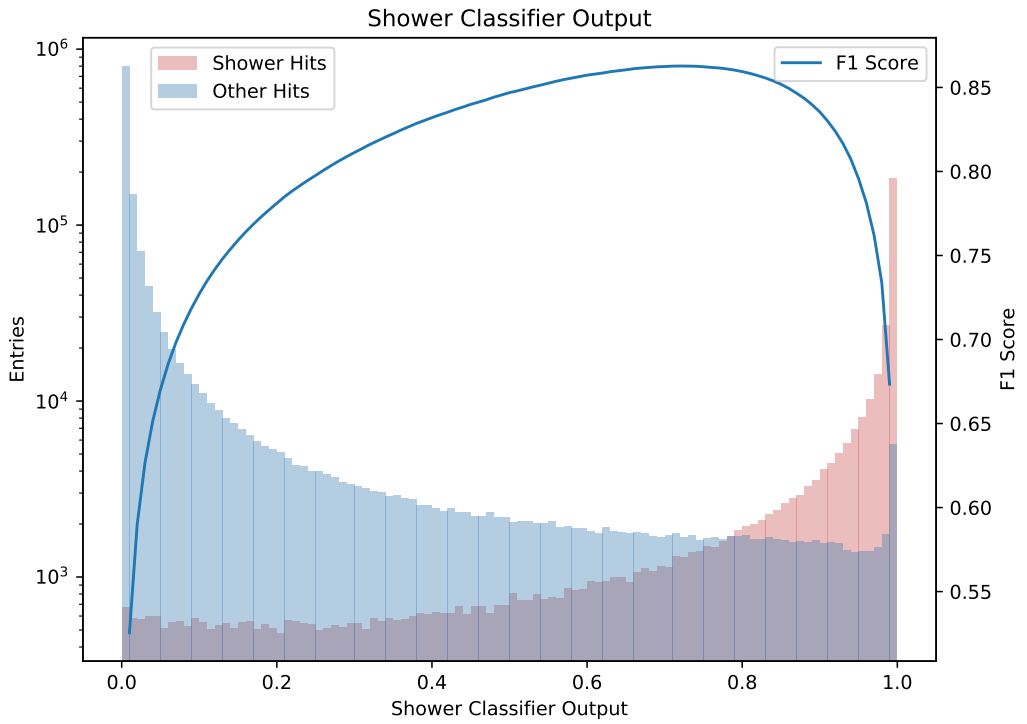


Figure 6.4: Shower score output distributions for the TSE classifier. The score distributions for true shower hits and all other hits are shown in red and blue respectively. Threshold optimisation was done using the F_1 score metric, which is also plotted as a function of the classification threshold.

very close agreement between the two curves, which suggests that the TSE classifier is robust to changes in the SCE model.

The performance of the Michel electron classifier was analysed with the same methods as the shower classifier. The Michel electron score distribution for true Michel electron hits and all other hits is shown in Figure 6.6. In this case, the large discrepancy in sample size between Michel electron hits and other hits leads to a low F_1 score of around 0.2 when considering the performance on a hit–by–hit basis. However, in Chapter 7, it will be shown that despite the low performance of the classifier for individual hits, a pure sample of Michel electron events can be selected by searching for clusters of hits with high Michel electron scores.

The ROC curve is independent of the relevant size of each true sample and, therefore, it provides a more instructive evaluation of the performance of the Michel electron classifier than the score distribution and F1 score. The ROC curves for the Michel electron classifier are shown in Figure 6.7, where both SCE and no SCE

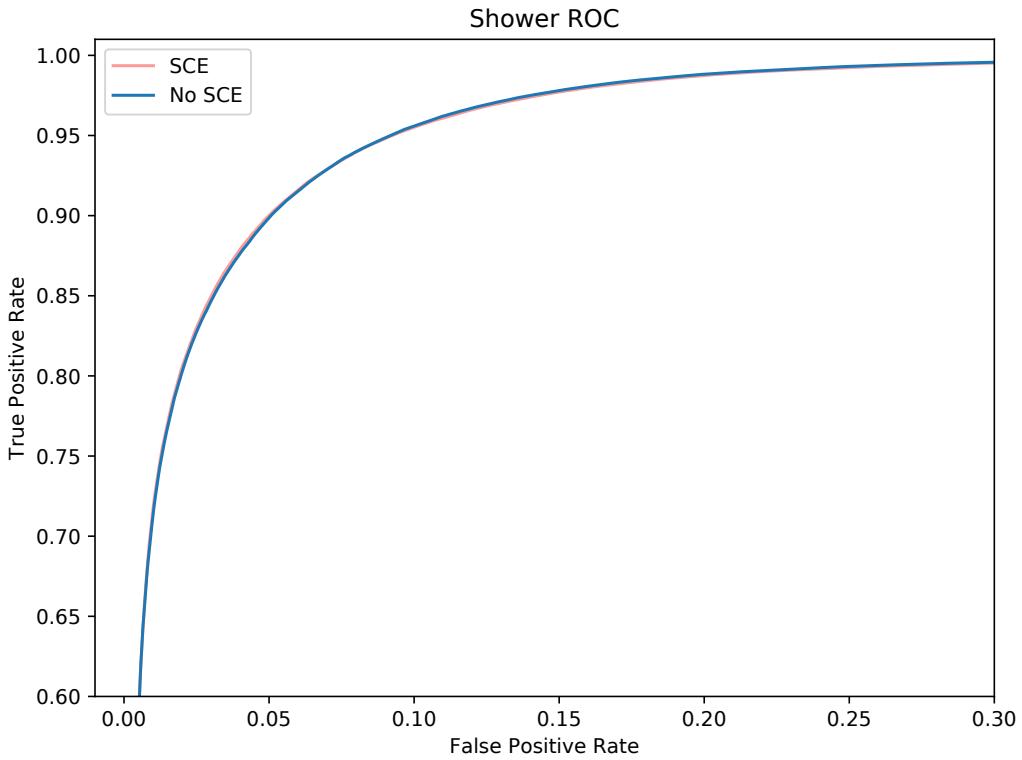


Figure 6.5: ROC curves for the shower score from the TSE classifier. These curves show the true–positive rate vs the false–positive rate for varying classification thresholds. The ROC curves for the ProtoDUNE–SP simulation including SCE and excluding SCE are shown in red and blue respectively. The ROC curve in the presence of SCE is concealed by the curve without SCE.

samples are shown. The jitter in the lines is due to the smaller sample size in this case. In both cases the Michel electron classifier is able to achieve a true–positive rate of over 90%, while maintaining a false–positive rate of less than 2.5%. There is a difference between the SCE and no SCE curves for the Michel electron sample, however, the difference is no bigger than 4% and is not currently understood. It is important, therefore, to validate the performance of the Michel electron classifier in data, where space charge distortions are present. Possible causes of the discrepancy in the ROC curves include, distortion of images due to space charge effects, which could lead to energy depositions near the boundary of images to be included in the SCE sample but not in the no SCE sample, or differences in the recombination correction between the SCE and no SCE samples.

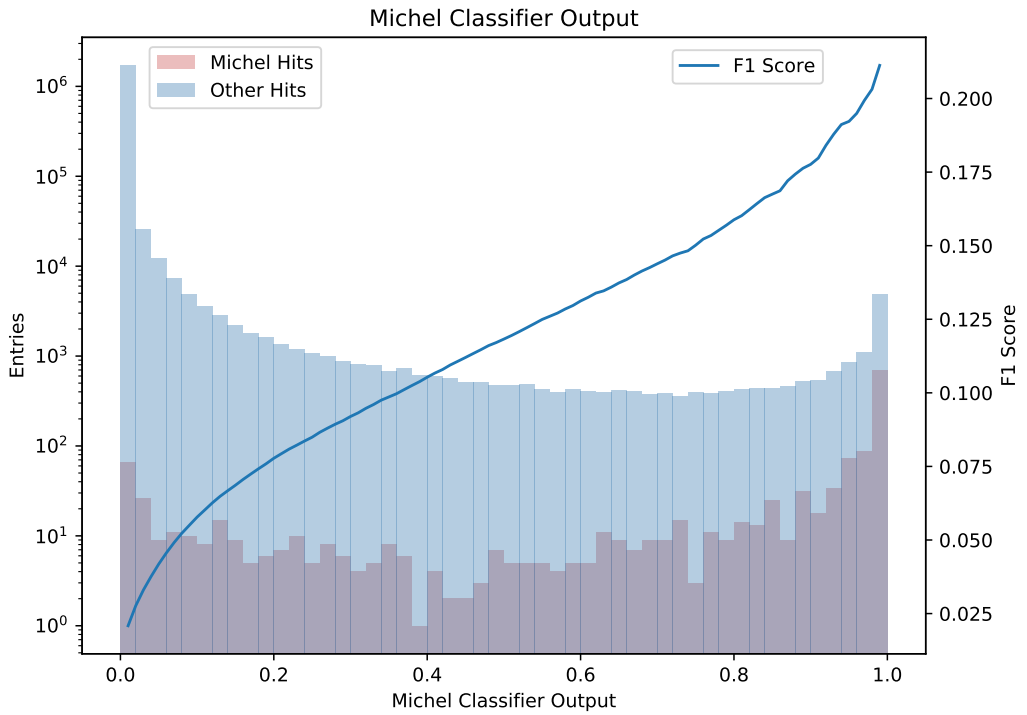


Figure 6.6: Michel electron score distributions for the Michel electron classifier. The score distributions for true Michel electron hits and all other hits are shown in red and blue respectively. The F1 score metric was used to assist threshold selection, and is also plotted as a function of selection threshold. The final threshold was modified when combined with a clustering algorithm, see the analysis in Chapter 7.

6.2.1 Comparison with Pandora

Pandora is the primary reconstruction framework used in ProtoDUNE–SP, it is discussed in Chapter 3. One of the goals of the CNN is to provide supplementary information to Pandora, to assist analysers in defining pure event samples. This is possible because Pandora and the CNN have slightly different goals: Pandora aims to cluster hits in the most appropriate way based on their spatial distribution, while the CNN aims to classify the hits based on the particle that caused the energy deposition.

The comparison of the particle classification to Pandora was done with a beam particle sample in ProtoDUNE–SP simulation, including the primary particle and any daughters. The true particle type was obtained from the simulation, and the track–shower categorisation was compared between Pandora and the CNN. For Pandora, the type of reconstructed cluster was taken as the classification. For the CNN the average shower score for the hits in the cluster was calculated and compared

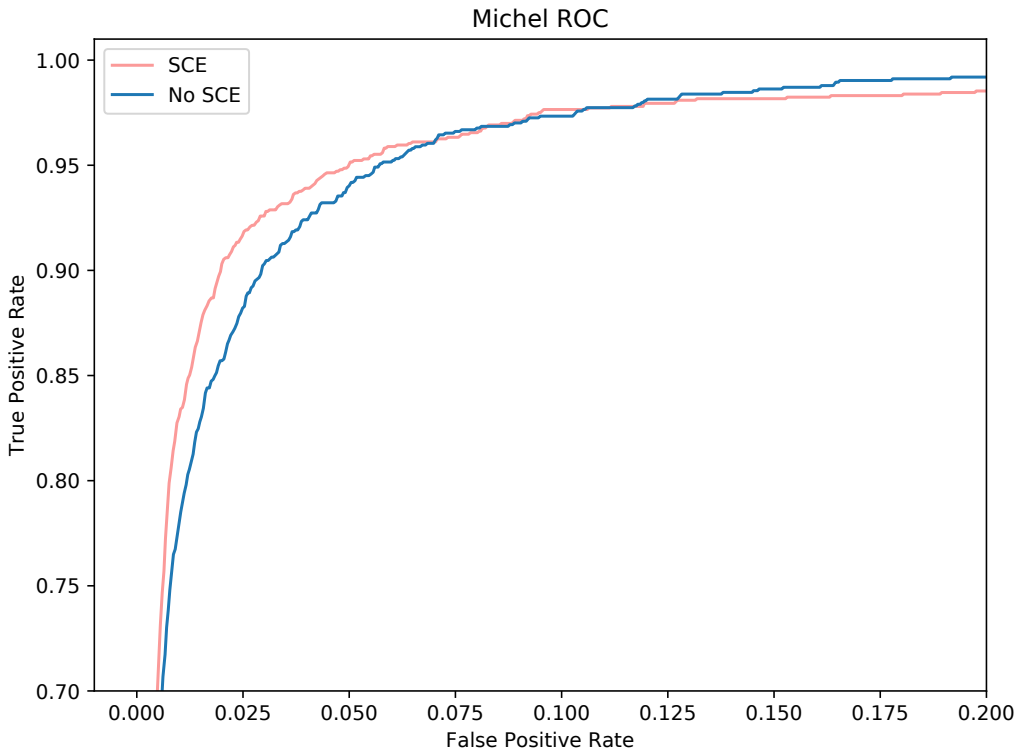


Figure 6.7: ROC curves for the Michel electron classifier. These curves show the true–positive rate vs the false–positive rate for varying classification thresholds. The ROC curves for the ProtoDUNE–SP simulation including SCE and excluding SCE are shown in red and blue respectively.

to a classification threshold of 0.72, based on the peak of the F1 score distribution in Figure 6.4. Any particles with an average shower score above the threshold were classified as showers, and any with score below the threshold were classified as tracks.

The track–shower classification by particle species based on the Pandora method and the CNN method are compared in Figure 6.8. The figure shows the fraction of each particle species classified as tracks and showers by both Pandora and the CNN, and it can be seen that CNN gives a stronger classification than Pandora for all particle species. Electrons and photons are more often classified as showers, and pions, muons, and protons are more often classified as tracks; this is as was hoped.

The false–positive rates (FPR) for Pandora and the CNN for each particle species are given in Table 6.3, as well as the ratio of Pandora to the CNN. The CNN gives an improved false–positive rate for all species, with a modest improvement over Pandora for pions and muons, and a significant improvement for electrons,

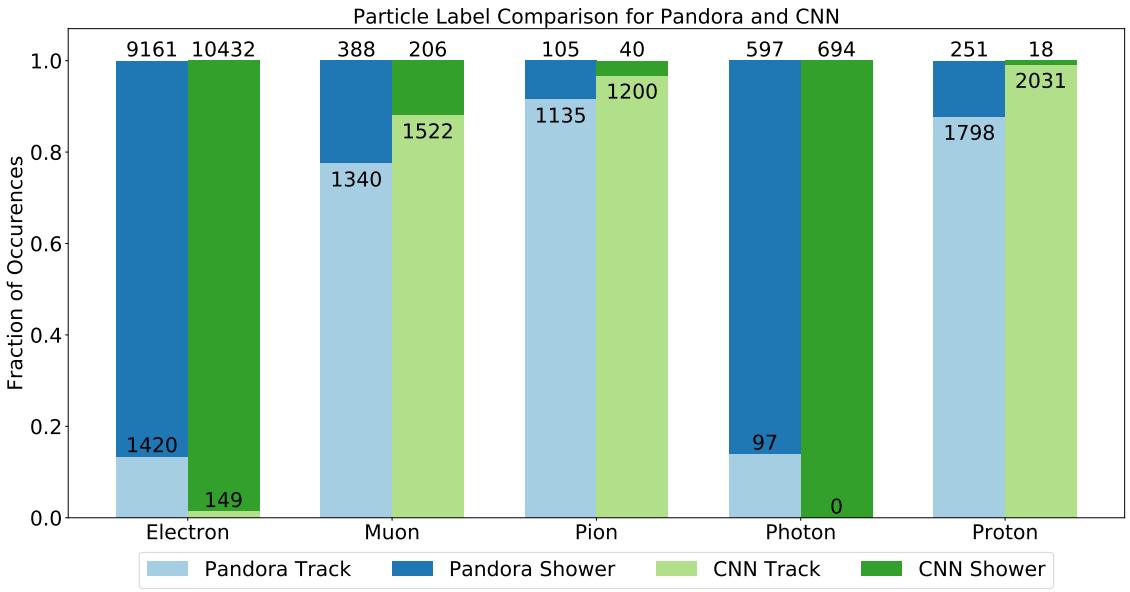


Figure 6.8: Comparison of track and shower classifications by Pandora and the CNN. For each true particle species the fraction of events classified as tracks and showers is shown, for Pandora and the CNN. Each bar is labelled with the number of events, for example, the CNN classified 149 electron events as tracks and 10432 electron events as showers.

Particle Species	Electron	Photon	Muon	Pion	Proton
Sample Size	10,581	694	1,728	1,240	2,049
Pandora FPR (%)	13.4	13.9	22.5	8.5	12.2
CNN FPR (%)	1.4	0.0	11.9	3.2	0.9
Pandora / CNN	9.53	–	1.88	2.62	13.9

Table 6.3: False–positive rates for track shower classification with Pandora and the CNN. The false–positive rates for both the CNN and Pandora are given for different particle species. In addition, the sample size for each particle species, and the ratio of the false positive rate of Pandora to the CNN are given whenever possible.

photons, and protons. The biggest improvement is for photons where, for a sample of 694 photons, the CNN has a false–positive rate of zero.

The stronger classification from the CNN gives supplementary information to Pandora, which can be used in analyses to improve event selection or background rejection. This information is being utilised in a number of ProtoDUNE–SP analyses, some of these are highlighted briefly in Section 6.4.

As well as providing supplementary information to Pandora for analysis, the CNN

scores could, in future, be used during Pandora reconstruction as an additional guide for the reconstruction algorithms. This approach is being developed by Pandora for the DUNE Far Detector, based on a similar neural network for discriminating tracks and showers[117].

6.3 Validation on ProtoDUNE–SP Data

For validation on real ProtoDUNE–SP data three approaches were used: visual validation with event scans, comparisons of the overall score distributions, and the comparison of score distributions for different particle species. Data from ProtoDUNE–SP runs 5387 and 5809 were used for these validations; the data for these runs were taken under stable operating conditions, with an average beam energy of 1 GeV. The beam composition in run 5387 was tuned to contain mostly hadrons, while in run 5809 it was tuned to have an enhanced electron component. The same operating conditions were used for the simulated data, in order to give the best possible comparison between data and simulation.

As discussed in Section 6.2, the sample of Michel electron hits is smaller than the other hits by orders of magnitude. In order to make a meaningful validation of the performance of the Michel electron classifier on data, the fraction of Michel electron hits in the sample needed to be increased. Michel electron event selection and reconstruction will be discussed in Chapter 7, where the validation of the Michel electron classifier will also be covered.

Hand scans of the events show qualitatively that the performance on the data is good. Figure 6.9 shows an example of the track score for hits in a reconstructed event from run 5387. The hits in this image are from the collection plane near the beam entry point, an electron shower tagged by the beam instrumentation (BI) can be seen near the centre of the image. In the event it can be seen that for hits in the tracks the CNN produces a large output score, and for the shower like activity in the event the score is low, as expected. In addition, the classifier assigns a high shower score to the hits adjacent to the tracks, which are from delta–rays and low–energy photons.

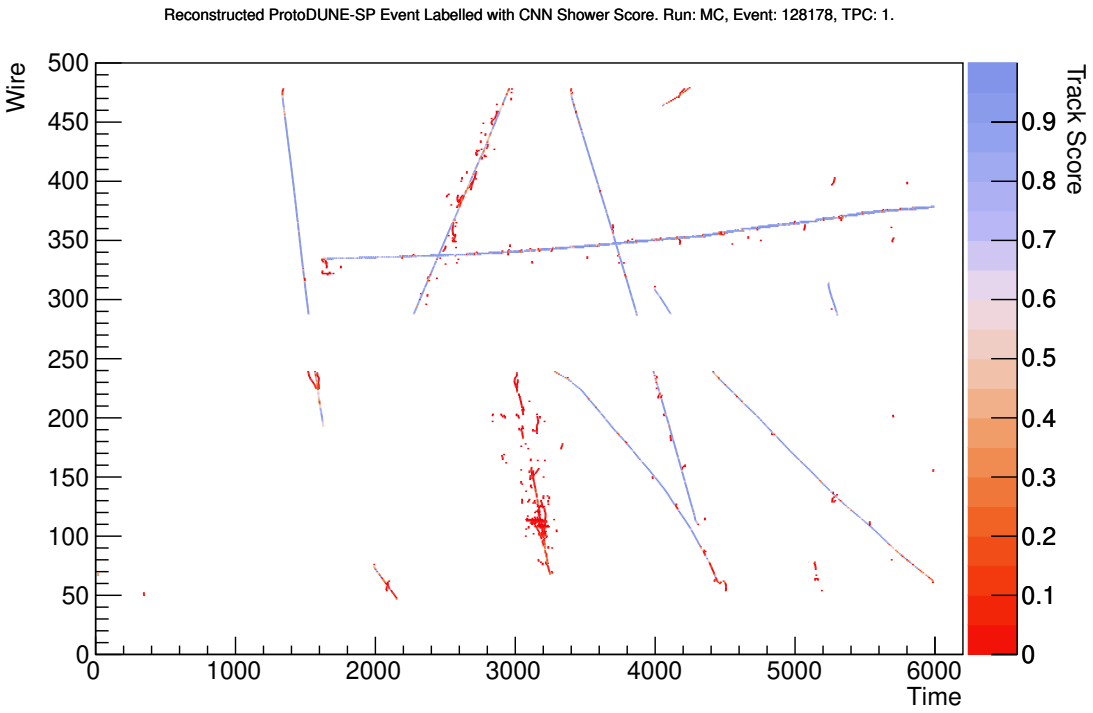
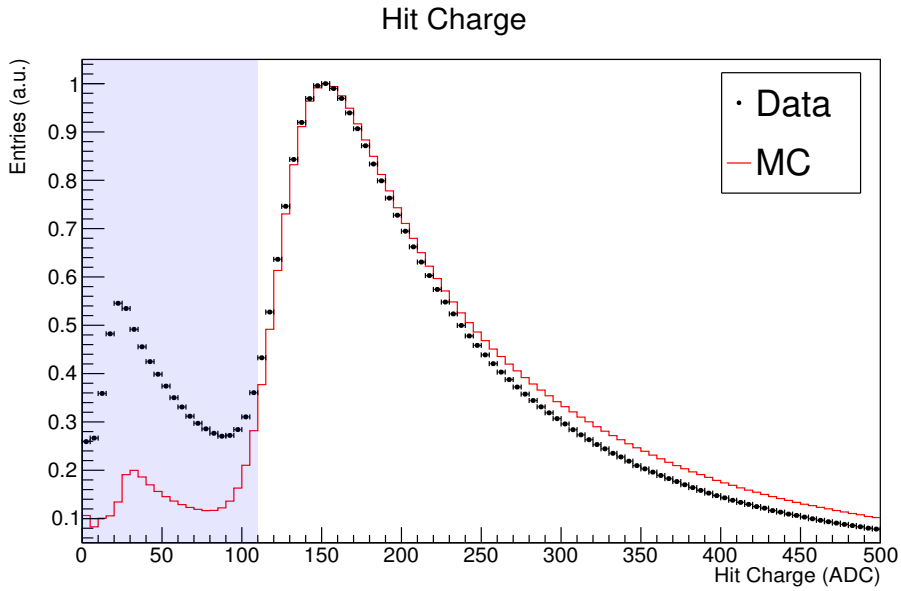


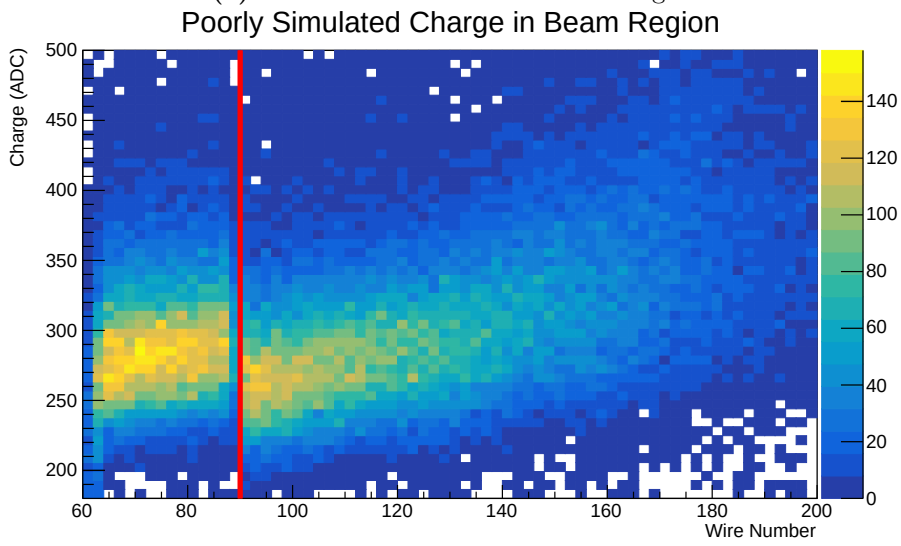
Figure 6.9: Track classifier scores for a reconstructed event in ProtoDUNE–SP data. Each reconstructed hit is labelled with its track classifier score, and located at the reconstructed wire and time coordinates.

The shower score distribution for all hits gives an overall validation of the network performance between data and simulation. This is the only quantitative validation method for the CNN that remains decoupled from the Pandora reconstruction algorithm. Therefore, it is independent of the differences in Pandora between data and simulation. The comparison is still dependent on the noise removal algorithm, the hit tagging algorithm, and the particle flux, which impact the input images to the CNN.

In this comparison, and the following beam particle comparisons, a cut is made on the reconstructed charge of each hit. This removes the excess of low charge hits in data, as illustrated in Figure 6.10a. In addition, all hits in the first 45 cm of the APA closest to the beam entry point were removed. This region has a known issue in charge simulation that can be seen as a discontinuity in the reconstructed charge. This charge discontinuity can be seen in in Figure 6.10b, which shows the discontinuity for a simulated proton beam sample. The cause of the discontinuity is not currently understood, possible causes include poor simulation of the energy



(a) Cut on the reconstructed hit charge.



(b) Cut on the reconstructed wire, due to a charge discontinuity in simulation.

Figure 6.10: Quality cuts on reconstructed hits for CNN comparison.

loss of particles in the wall of the TPC and poor modelling of the space charge effect in the region of the beam window.

The overall shower score distribution for data and simulation is shown in Figure 6.11, as well as the fractional residuals in each bin. These distributions are normalised by the number of hits, which allows the shape of the score distribution in data and simulation to be compared. Overall there is a good agreement between the data and the simulation, however, the distribution in data has a slightly different

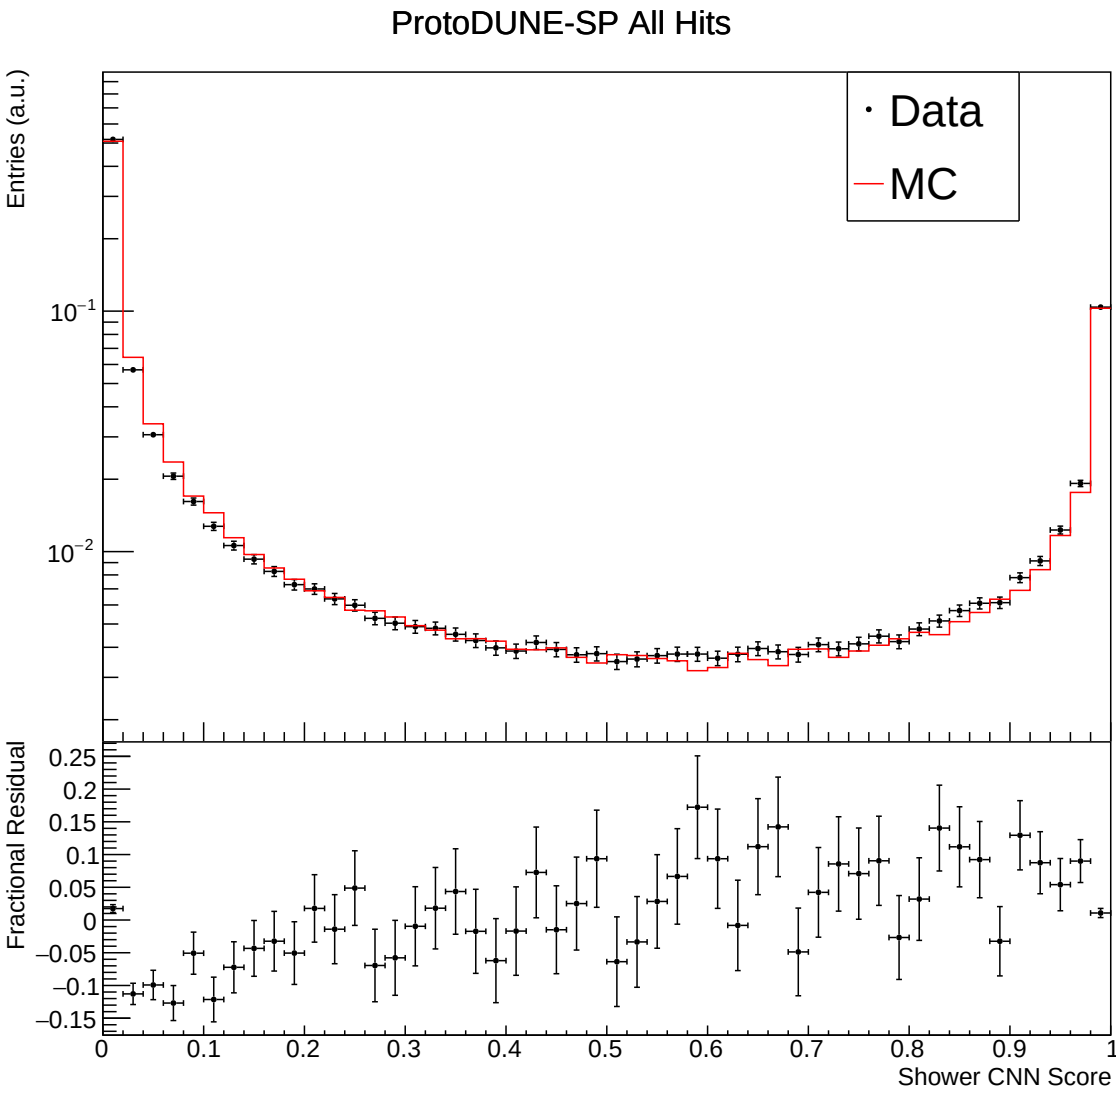


Figure 6.11: Overall shower score distribution from the CNN for all hits in data and simulation. The top panel shows the shower score distribution for all hits in data and simulation, normalised by the number of hits, and including statistical error bars. The fractional residuals in each bin are shown in the bottom plot.

shape at low CNN scores. This can be seen in the residuals, which are consistently negative for hits with shower score less than around 0.2.

It is also interesting to consider the comparison of the CNN between data and simulation for different particle species, which can be achieved with the beam instrumentation (BI). The BI in ProtoDUNE–SP provides particle identification (PID) for beam particles with a very high purity. For particles originating from the charged particle beam, the BI can be used to provide an effective truth source. This allows the results of the CNN to be compared between data and simulation

for different particle species. Any particles that arrive out–of–time with the beam can be assumed to be cosmic–rays and, therefore, form an additional truth source. The shower score distributions for electrons, pions, protons, and cosmic–rays are produced based on these samples.

It is important to note that the results for these true particle samples rely on both Pandora and the CNN scores, because the clustering from Pandora is used to select the appropriate hits within each particle. In order to minimise the impact of Pandora on the observed CNN score distributions, a set of phase space cuts were made. These cuts selected regions of the detector where Pandora had a good agreement between data and simulation, and included cuts on the start and end points of the reconstructed tracks. The same cuts were applied to both the start and end point of each track. These cuts are illustrated for the end point of the tracks in Figure 6.12a. The reconstructed angular distributions for tracks after these cuts are shown in Figure 6.12b.

The cuts made were as follows:

- X (drift) direction:
  - $|x| > 330$ cm, to remove tracks with end–points close to or outside of the detector volume.
- Y (vertical) direction:
  - $|y| < 50$ cm, to remove tracks with end–points close to the bottom of the detector volume.
- Z (beam) direction:
  - $|z| < 50$ cm and $|z| > 650$ cm, to remove tracks with end–points close to the edge of the detector volume.
  - $|z-233| < 10$ cm and $|z-466| < 10$ cm, to remove tracks with end–points close to an APA boundary.

After the cuts, there is one region in which there is a large discrepancy between data and simulation, near the center of the X direction. This discrepancy is due to

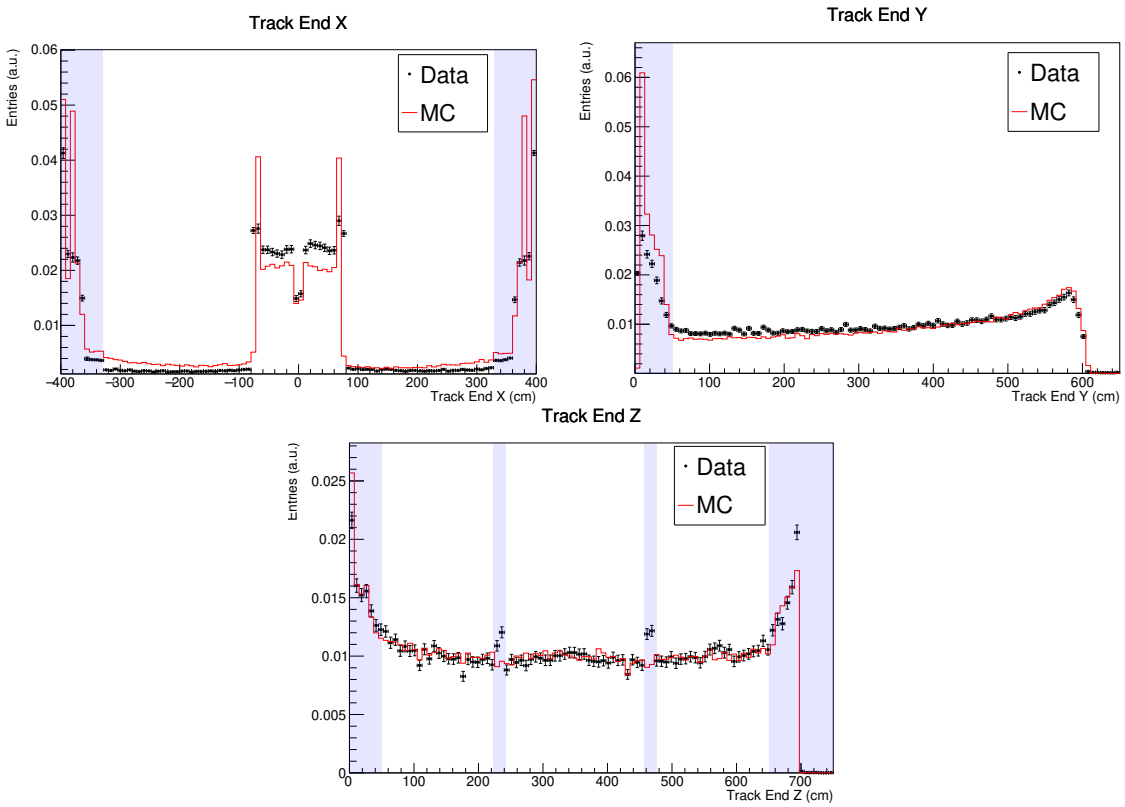
the fact that the track stitching algorithm that is used to assign the true time to tracks is more efficient in simulation than data, and the distribution of true times impacts the end point distribution of tracks in the drift direction (X).

The shower classifier scores for the pion, proton, and electron test beam samples are shown in Figure 6.13. The data in the pion and proton distributions were taken from ProtoDUNE–SP run 5387, and the data in the electron distribution were from run 5809. The data in all of these beam particle distributions are normalised by the number of triggered beam particles of the given flavour.

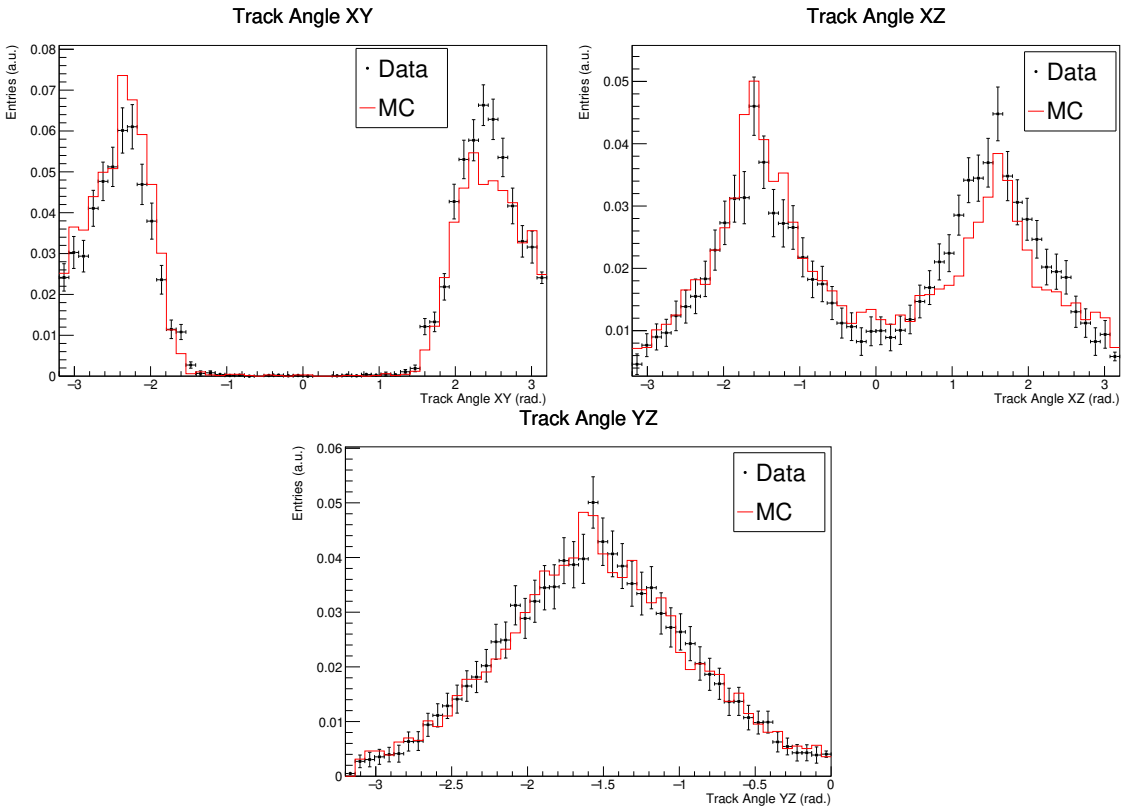
Overall, there is a good agreement between the data and simulation in terms of the shower score distributions for each particle species. However, there are still some discrepancies, which highlight the differences between the data and the simulation. The difference is most pronounced for the electron sample, which gets consistently more entries at low CNN scores. One consistent difference between data and simulation for all particle species, is the tendency for Pandora to cluster more hits into each reconstructed particle in data than it does in simulation. This can be seen in the bins at the two ends of the distributions, which contain the majority of the hits, and consistently have more entries in data than they do in simulation. This difference is the result of the Pandora reconstruction algorithms, but it is not yet fully understood.

The cosmic–ray sample was selected based on the cosmic–ray part of the Pandora reconstruction chain, discussed in Chapter 3. Any reconstructed particles that have been labelled as both clear cosmic–rays and tracks by Pandora are included in the sample. The data in this sample was taken from run 5387; while a cosmic–ray only run would be preferable, all of the ProtoDUNE–SP simulations include test beam particles. Therefore, a test beam run was chosen in order to match as closely as possible to the simulation.

The shower score distribution for cosmic–ray tracks is shown in Figure 6.14. These distributions are normalised by the total number of hits in each distribution, because there is a large difference in the average number of reconstructed hits in clear cosmic–rays from Pandora between data and simulation. Again a good



(a) Spatial cuts applied to the X (drift), Y (vertical), and Z (beam) coordinates of the track end.



(b) Track angular distributions after spatial cuts in (a) applied.

Figure 6.12: Data quality cuts for CNN comparison, and the resulting muon angular distributions after the quality cuts.

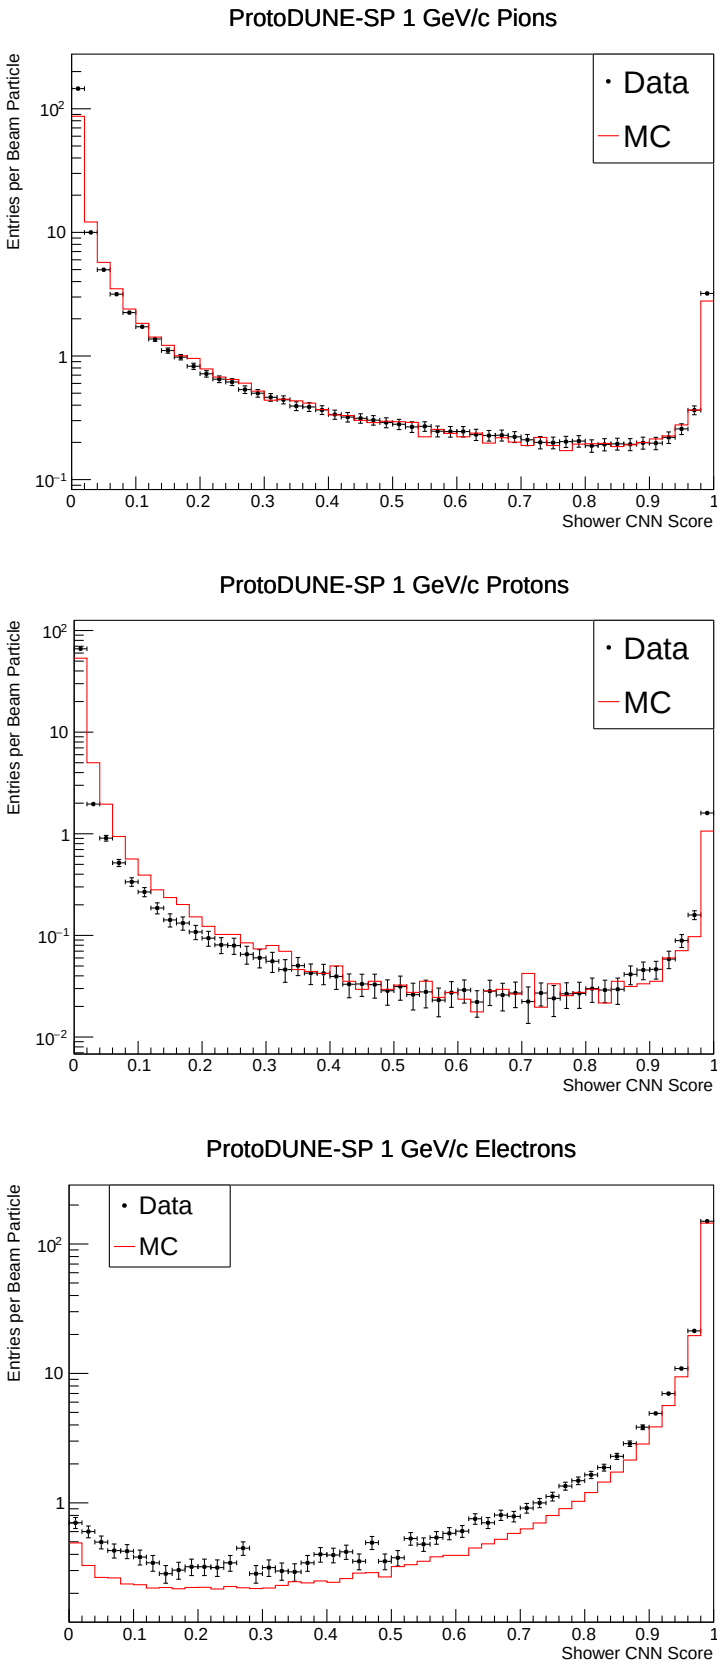


Figure 6.13: Shower classifier scores for different particle species from the ProtoDUNE–SP beam, in data and simulation, with statistical error bars for the data.

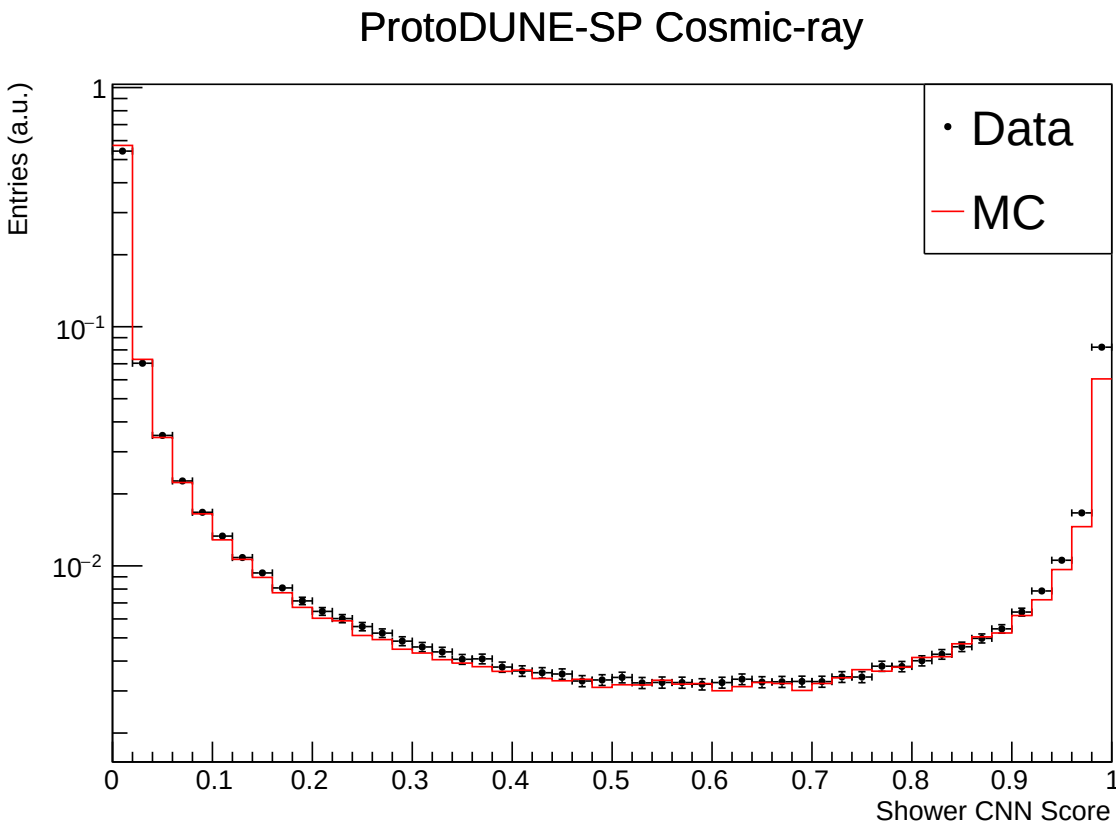


Figure 6.14: Shower classifier score for hits from reconstructed cosmic–rays, in data and simulation, with statistical error bars for the data.

agreement is seen over the majority of the range, but in data there is a prominent excess of hits with a high shower score. The cause of this excess is not currently understood, one possibility is that delta–rays are more often reconstructed as part of cosmic–ray tracks in data than in simulation.

In practice, the CNN is used to select events based on choosing hits with scores above a threshold. This could be based on an average over a number of hits, on a hit–by–hit basis, or a more complicated algorithm. For the purposes of understanding the approximate errors involved in selecting hits with the CNN, a hit–by–hit algorithm was used. The uncertainties involved with selecting hits in this way were evaluated by considering the fraction of hits selected into the sample given a selection threshold.

Table 6.4 lists the fraction of individual hits selected into the appropriate category for each sample in data and simulation. The fractional difference between

Hit Source	Class	MC Fraction (%)	Data Fraction (%)	Data / MC
All	Track	77.1	76.2	0.99
Pion	Track	93.8	95.4	1.02
Proton	Track	97.1	96.6	0.99
Electron	Shower	96.9	95.8	0.99
Cosmic–ray	Track	92.9	90.3	0.97

Table 6.4: Fraction of hits in correct class for different samples in ProtoDUNE–SP data and simulation.

the numbers in each case is an estimate of the percentage uncertainty associated with this simple selection algorithm. For the sample of all hits, the track category was used, and a fractional difference of 1% was seen. The fractional difference varies depending on the particle species, with the largest difference being 3% for the cosmic–ray sample.

Despite some small discrepancies between data and simulation, overall, the CNN results have been shown to give a good agreement between data and simulation across a range of particle species. The discrepancies seen are also influenced by the difference between data and simulation for the Pandora reconstruction framework, which makes the cause of the discrepancies difficult to disentangle. These differences will impact each analysis differently, therefore, the uncertainties involved with using the CNN classifier should be evaluated on a case–by–case basis; for hit selection with a simple hit–by–hit algorithm, the uncertainties are on the order of 1-3% depending on the particle species. The typical error in the analyses that make use of these hits is around 15-20%[118], therefore, the error due to this hit tagging algorithm is a small contribution to the overall errors.

6.3.1 Systematic Uncertainties

While no extensive study has been made into the magnitude of the systematic uncertainties associated with this hit tagging algorithm, this section will aim to discuss the possible sources of systematic uncertainty in the network, for the benefit of those working on analyses with the algorithm and those who would like to

develop the algorithm further in the future. Possible mitigation strategies will also be discussed for the likely uncertainties.

Two possible sources of systematic uncertainty in the hit tagging algorithm relate to the training of the network and the images used during the training process. There are systematic uncertainties due to differences in the noise model between data and simulation, and due to the differences in the space charge model. Each of these effects impact the images generated during the training and evaluation of the network, which the network uses in order to make its prediction of scores for each hit. Differences in the noise model could lead to a change in the number of hits being misidentified if, for example, the noise model caused an increase in small hits that mimic shower like energy deposits. In the case of a difference in the space charge model between data and simulation, the images could be distorted differently in data which could lead to some hits being included or excluded from images. Additionally, the changes in the local electric field could lead to a different recombination correction, which could also lead to misidentified hits. While it has not been confirmed, it is possible that differences in the recombination effect are responsible for the difference in the ROC curve for the Michel electron events, which can be seen in Figure 6.7.

In order to mitigate these effects in the image production, I recommend that any further developments to this network consider producing a more diverse training sample including a variety of simulated detector effects. For example, the network could be trained on images that include a wide range of noise models so that the network learns to make accurate predictions irrespective of the noise present in the images. This training data set could be produced directly from simulation by varying the parameters in the noise model for each batch of images, but a particularly promising technique for building such a data set for training at lower computational cost is that of a generative adversarial network (GAN). An in–depth discussion of GANs is beyond the scope of this thesis; they are mentioned here as a possible extension of this work and have been discussed in detail in [119].

Other possible sources of systematic uncertainty are the cuts made during this analysis. These include the fiducial volume cuts in Figures 6.10b and 6.12a, the cut on the reconstructed hit charge in Figure 6.10a, and the threshold used to select hits as shower–like. These cuts may exclude certain interactions to a greater or lesser extent in data than in simulation, which could lead to a systematic effect. Since the cuts used in each analysis involving these CNN scores are likely to be different, these effects were not investigated here. However, they should be considered for analyses that make use of the output of this network for event selection or reconstruction.

6.4 Application in ProtoDUNE–SP Analyses

The general agreement between data and simulation is a sign that the CNN results are sensible, and the additional classification strength of the CNN over Pandora makes it a useful tool in analyses of the ProtoDUNE–SP data. The output of the CNN classifier has been used in a number of ProtoDUNE–SP analyses since it was developed. Chapter 7 details one use of the network in a Michel electron analysis. In addition, the scores from the network are being incorporated into analyses by other members of the ProtoDUNE–SP experiment, including:

- Selecting shower candidates for neutral–pion event selection[118]
- Identifying charge exchange candidates in charged–pion cross section analyses[120]
- Identifying Michel electron contaminated tracks for calibration studies with stopping muons[121]

In the neutral pion analysis, the CNN score is being used as part of the sample selection algorithm, which is used to select a sample of events with pairs of secondary photons. An example of the CNN shower score distribution for secondary particles in this sample is shown in Figure 6.15, demonstrating the breakdown of the CNN score by particle species for candidates in this sample. An example of a neutral pion candidate selected with this algorithm is shown in Figure 4.5d in Chapter 4. These plots are courtesy of Milo Vermeulen, who developed the neutral pion event selection algorithm in ProtoDUNE–SP[118].

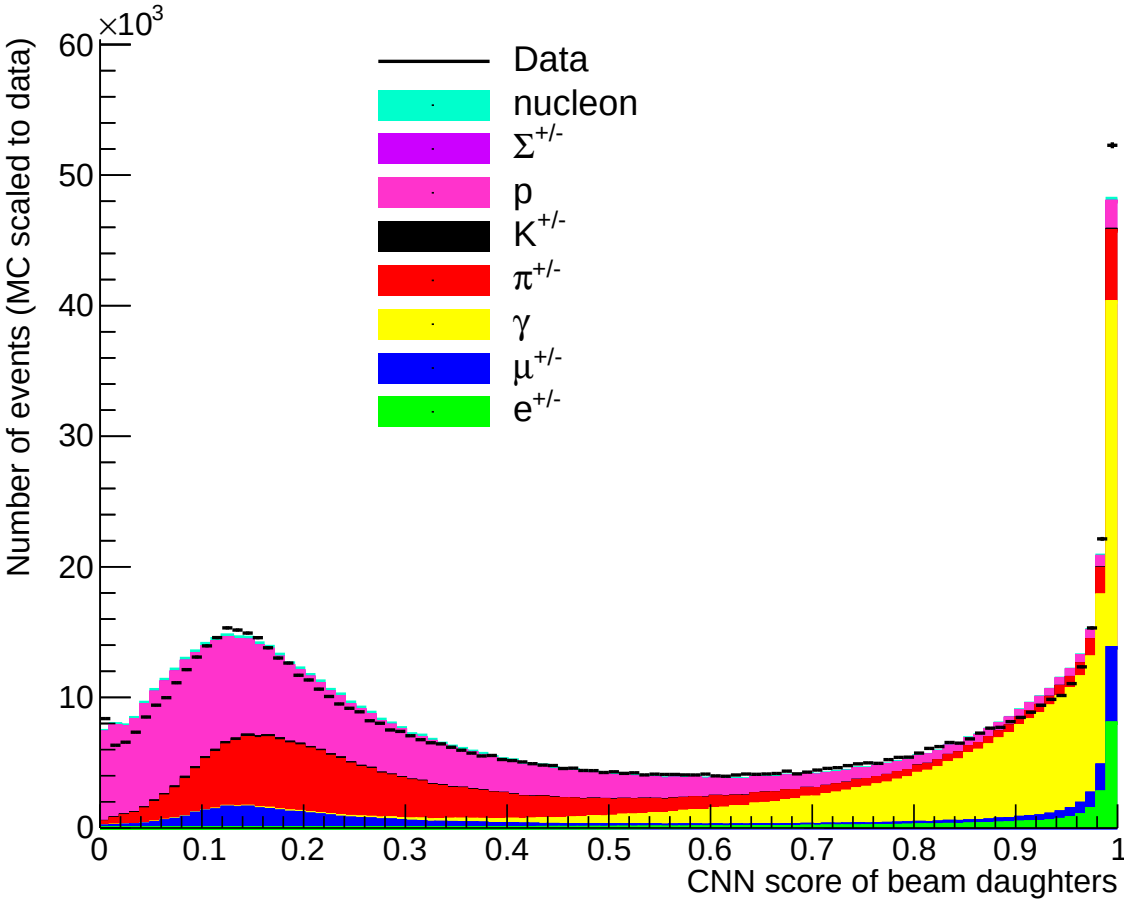


Figure 6.15: CNN shower score distribution for secondary particles in a neutral pion event selection sample. The coloured distributions represent different true particle species in ProtoDUNE–SP simulation, and the black points represent the distribution in ProtoDUNE–SP data. Figure courtesy of Milo Vermeulen[118].

7

Study of Michel Electron Reconstruction in ProtoDUNE–SP

Contents

7.1 Michel Electrons in Liquid Argon 124
7.2 Michel Electron Event Selection 131
7.3 Michel Electron Energy Reconstruction 134
7.3.1 Michel Electron Hit Tagging with a U–ResNet 134
7.3.2 Michel Electron Reconstruction 139
7.4 Conclusion 153

Studying electrons in the tens of MeV energy range can provide valuable input into reconstruction techniques and energy uncertainty for astrophysical neutrinos from supernova bursts. Understanding the response of LArTPC detectors to electrons in this range will be important for any large scale LArTPC experiment wishing to study supernova bursts. At these energies, electron interactions have large contributions from both ionisation energy loss and radiative energy loss and, therefore, they have a distinctive signature that is neither track–like nor shower–like. As a result, low–energy electrons require dedicated reconstruction algorithms to maximise the overall reconstruction performance.

This chapter will discuss an approach to low–energy electron reconstruction in LArTPC detectors, based on the use of convolutional neural networks and semantic

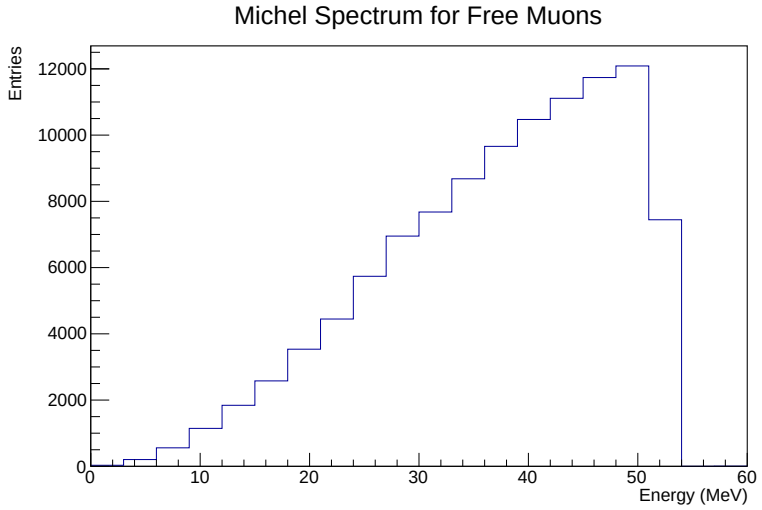
segmentation. Michel electron events from ProtoDUNE–SP will be used to test the performance of this technique, and to provide an estimate of the energy uncertainty of LArTPC detectors for low–energy electrons.

In this Chapter, Section 7.1 will discuss the signature left by Michel electrons in liquid argon, and the implications of this signature on reconstructing Michel electron events in ProtoDUNE–SP. This will be followed by a discussion of the algorithm used to select Michel electron events in Section 7.2. The Michel electron event reconstruction will be discussed in Section 7.3, along with a discussion of the reconstruction performance and a comparison to the results of other LArTPC experiments. Finally, the results presented in this chapter and their implications for the DUNE Far Detector will be summarised in Section 7.4.

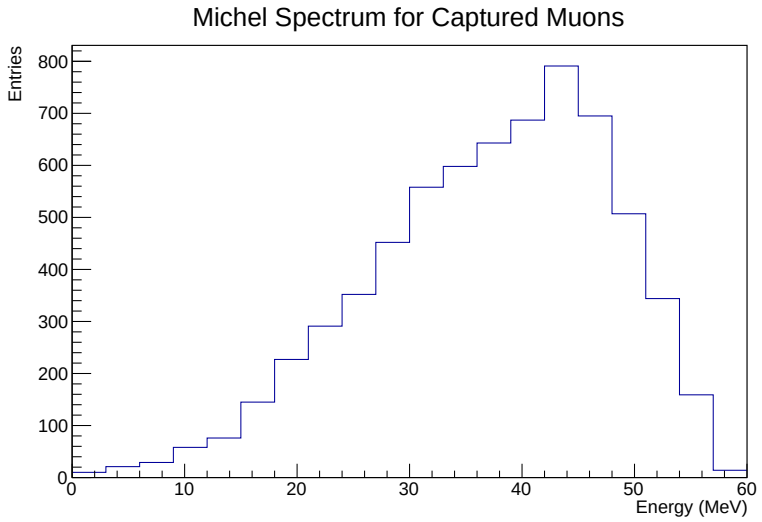
7.1 Michel Electrons in Liquid Argon

Michel electrons are produced when a muon decays at rest. In vacuum, this decay gives rise to a characteristic energy spectrum that has a sharp cut–off at around 50 MeV, corresponding to half the muon mass. In matter, it is also possible for μ^- to be captured on nuclei before they decay, this causes a broadening of the Michel electron spectrum for these events. A comparison of the Michel electron energy spectrum for free $\mu^\pm$ and captured μ^- is given in Figure 7.1. The capture process occurs roughly 70% of the time for negative muons in liquid argon, therefore, in ProtoDUNE–SP, the observed Michel electron energy spectrum is a combination of the two processes in similar quantities.

As discussed in Chapter 4, the energy loss for electrons in liquid argon passes from an ionisation dominated regime to a radiation dominated regime in the tens of MeV region. The crossover point for this transition occurs at around 30 MeV, close to the peak of the Michel electron spectrum. This leads to a distinctive signature for Michel electrons in liquid argon detectors: a short ($\sim$ 5 cm) track segment is surrounded by a number of small radiated energy deposits. Figure 7.2 shows an example of a Michel electron candidate from ProtoDUNE–SP data, along with labels of the key features.



(a) Free Muons.



(b) Captured Muons.

Figure 7.1: Michel electron energy spectra for free muons and captured muons in liquid argon from ProtoDUNE–SP simulation.

One of the main challenges for Michel electron reconstruction in liquid argon is to successfully associate the radiated energy depositions back to the initial Michel electron once they have produced ionisation in the detector. Photons have a radiation length of around 20–30 cm in liquid argon, which is many times larger than the size of the typical track–like part of the event. Figure 7.3 shows the spectrum of radiated photons from Michel electron events in ProtoDUNE–SP simulation, alongside the photon multiplicity as a function of Michel electron energy. While most of the radiated photons only carry a small fraction of the Michel electron's

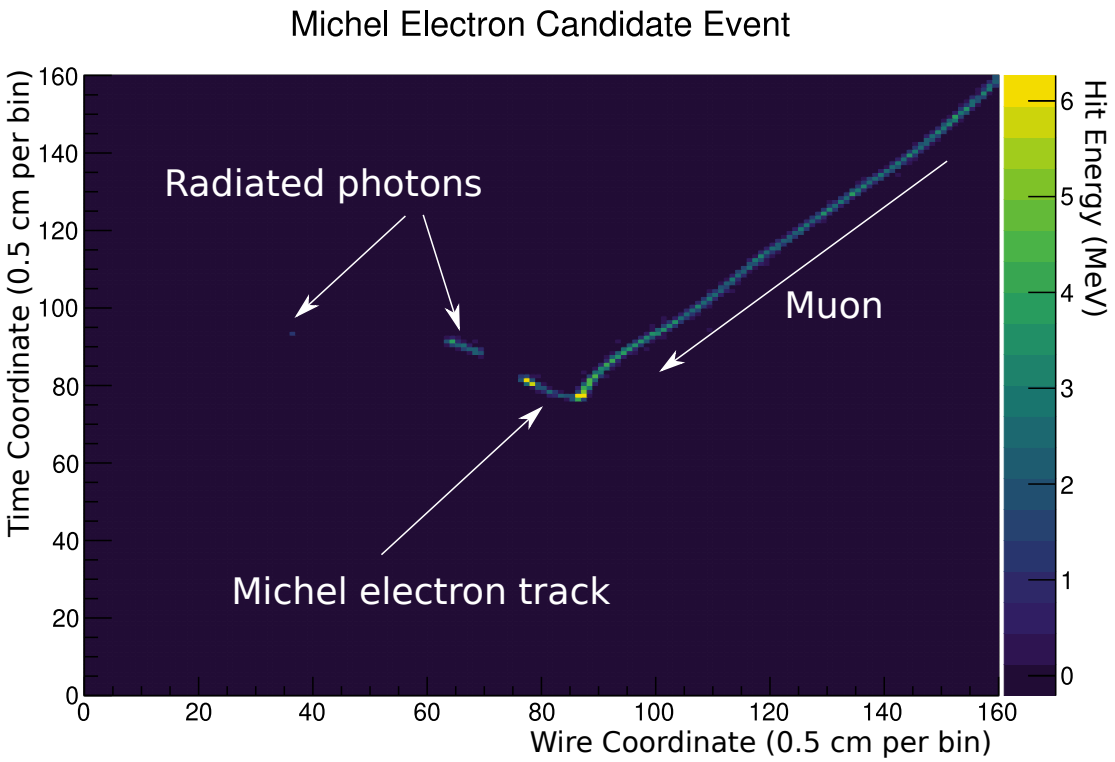
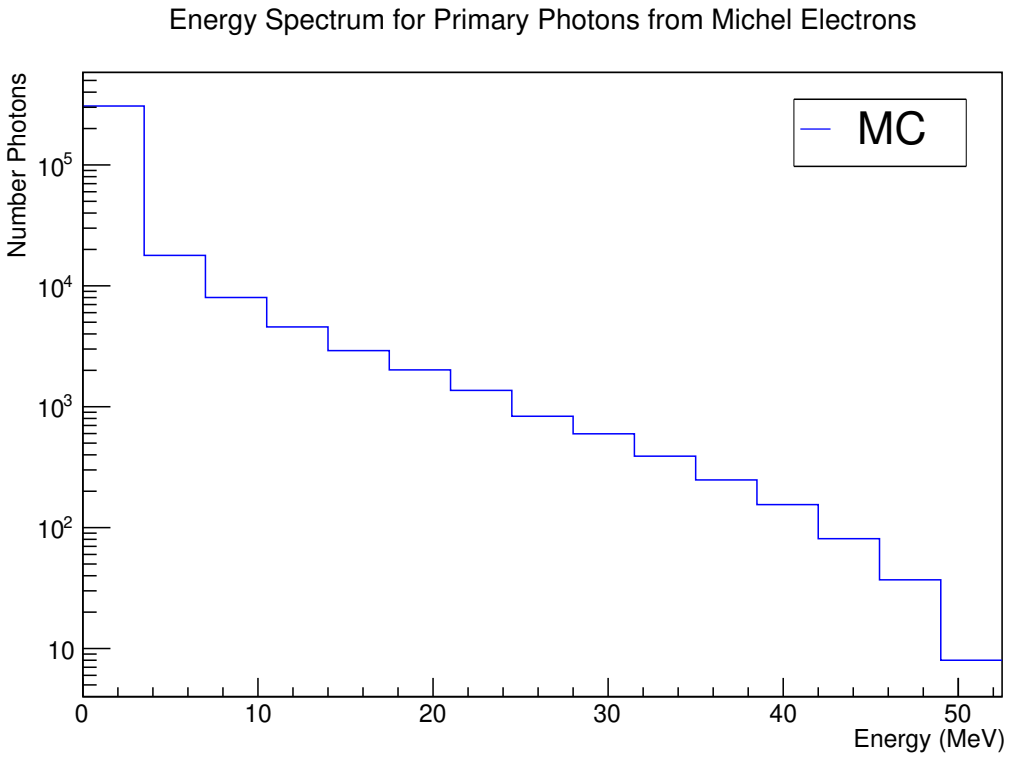


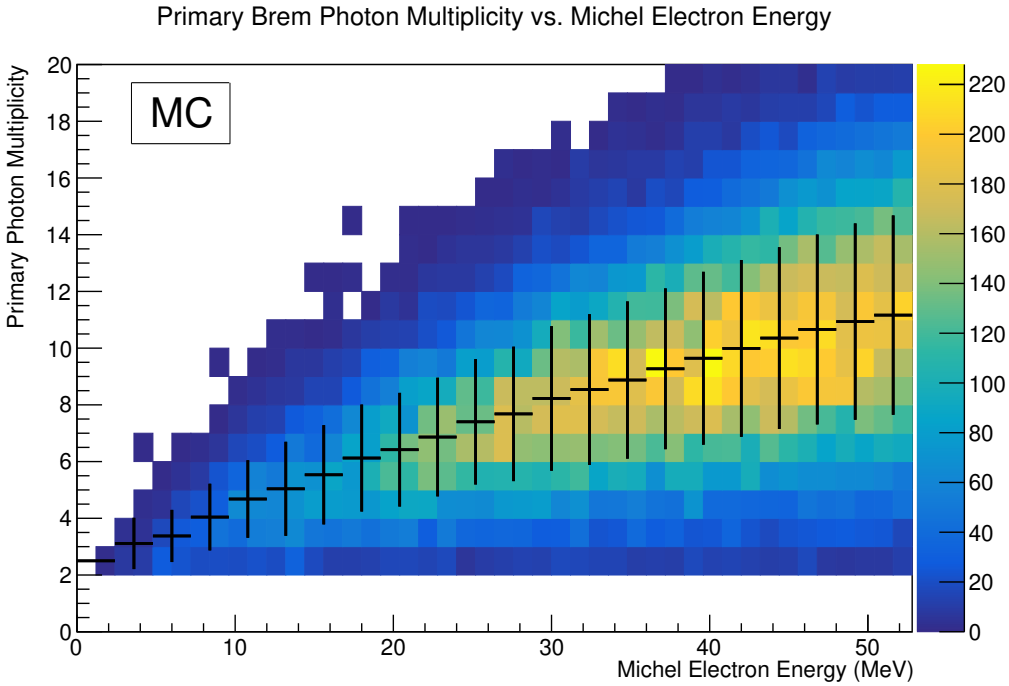
Figure 7.2: Michel electron candidate event from ProtoDUNE–SP data. The wire vs time data in the region of the Michel electron event is shown. The incoming muon, initial Michel electron track, and energy depositions from radiated photons have been labelled.

energy, in some cases a single radiated photon can carry a significant fraction of the electron energy. In addition, around the peak of the Michel electron spectrum ($\sim$ 45 MeV) there is a high photon multiplicity and a large spread in the multiplicity distribution. The combination of these effects leads to a significant spread in the fraction of radiated energy for Michel electron events.

The energy that is lost into radiated photons is only visible once the photons interact in the argon to produce secondary electrons, which then ionise the argon. These secondary electrons are produced over large angles and distances in the detector when compared to the short Michel electron track, the spatial distribution of secondary electrons is shown in Figure 7.4. This shows that the radiated energy deposits are spread over a large area, when compared to the size of the initial Michel electron track, which is typically around 5–10 cm. Therefore, any reconstruction algorithm hoping to recover the radiated energy, will need to use data from a relatively large volume in order to maximise energy recovery.



(a) Energy spectrum of radiated photons from Michel electron events.



(b) Photon multiplicity as a function of true Michel electron energy. The distribution is overlaid with a profile of the mean of the distribution, the error bars on the profile represent the RMS of a gaussian fit to the distribution in each bin.

Figure 7.3: Properties of radiated energy deposits from Michel electrons in ProtoDUNE–SP simulation.

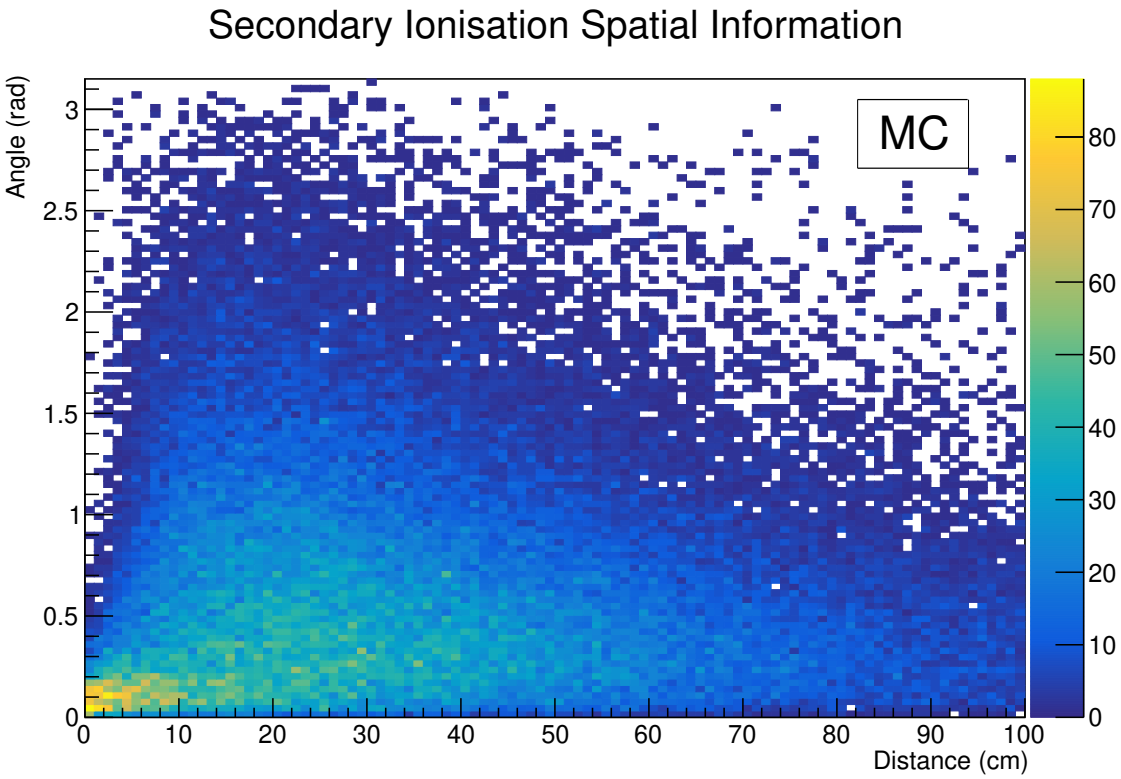


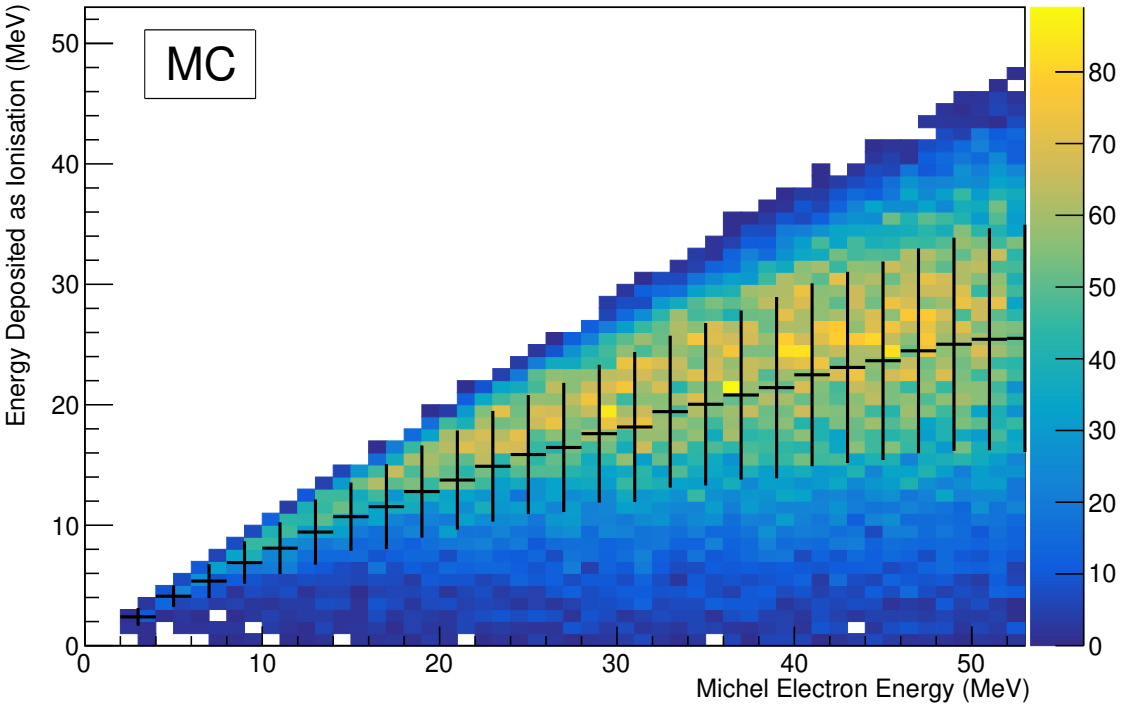
Figure 7.4: Spatial distribution of radiated ionisation deposits from Michel electrons. This plot shows a 2D histogram of the distance and angle of each radiated energy deposition from the primary Michel electron track.

To highlight the impact of the radiated energy deposits, the results of perfect energy reconstruction were considered in two cases:

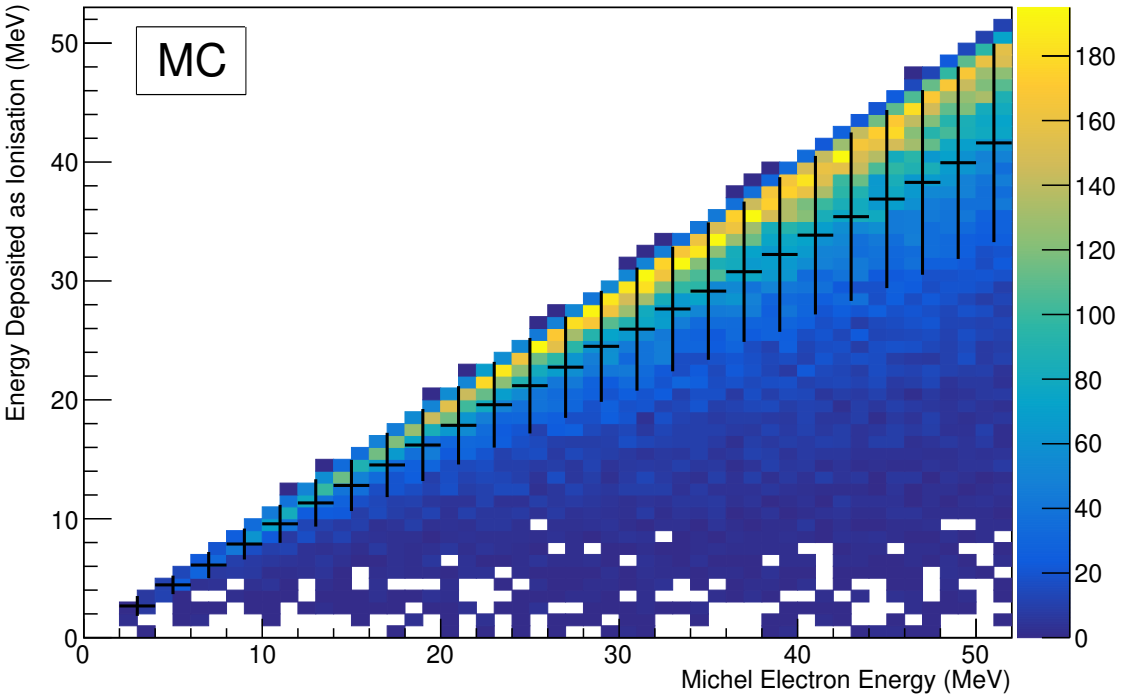
- Only considering the Michel electron track
- Considering all ionisation energy within some radius and angle of the Michel electron track

Figure 7.5 illustrates the considerable increase in energy collected if radiated energy is considered. The distribution is significantly narrower, and much more energy is recovered when considering the energy deposited within a spherical–section of height 40 cm and angle 30° of the Michel electron vertex. In addition, the fraction of energy recovered as a function of Michel electron energy has a more linear distribution for a collection radius of 40 cm.

The average fractional energy recovery as a function of collection radius is shown in Figure 7.6, where the error bars represent the RMS of the distribution. By increasing the collection radius from 0 cm to 40cm, the average energy recovered is



(a) True collected ionisation energy from the initial Michel electron track only.



(b) True collected ionisation energy from the initial Michel electron track, and any radiated ionisation within a 40 cm spherical–section at a 30° opening angle of the initial Michel electron track.

Figure 7.5: 2D histograms of the true collected ionisation energy vs the true Michel electron energy in ProtoDUNE–SP simulation based on two different collection radii.

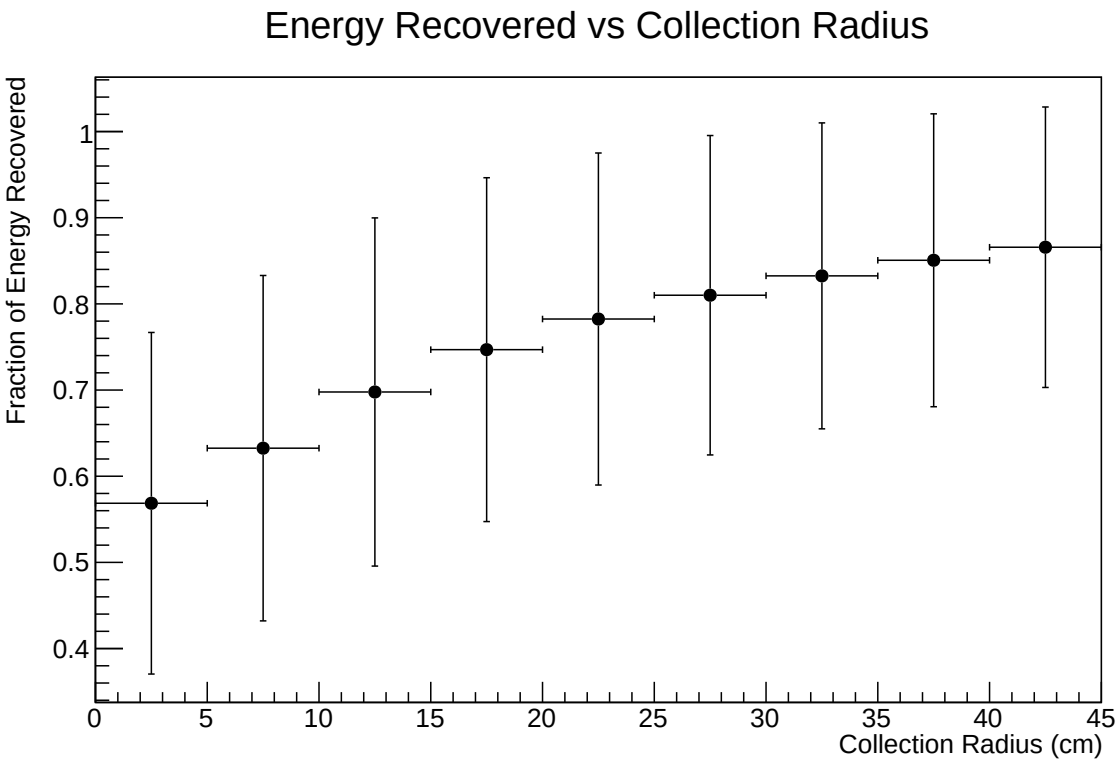


Figure 7.6: Fraction of Michel electron energy collected as ionisation vs collection radius. The error bars here represent the RMS spread of the distribution.

increased from 57% to 87%. The RMS of the fractional energy recovered reduces slightly from 20% to 16% across this range. Therefore, the ratio of the RMS to the mean is reduced from 35% to 18%.

Increasing the collection radius beyond around 30–40 cm gives minimal improvement in the fractional energy recovery from ionisation. This can been seen in Figure 7.7, which shows the total true ionisation energy collected for the case of a 40 cm collection radius and an 80 cm collection radius. While increasing the collection radius does not significantly increase the collected ionisation, it was found to impact the purity and efficiency of reconstruction algorithms. Therefore, the energy reconstruction algorithm discussed in this chapter considers a 40 cm collection radius when collecting radiated energy deposits.

The study of ProtoDUNE–SP simulation presented in this section highlights the importance of radiated energy deposits in Michel electron and other low–energy electron events. Based on these results, it is clear that to reduce energy uncertainties for these events it is important to increase the amount of energy

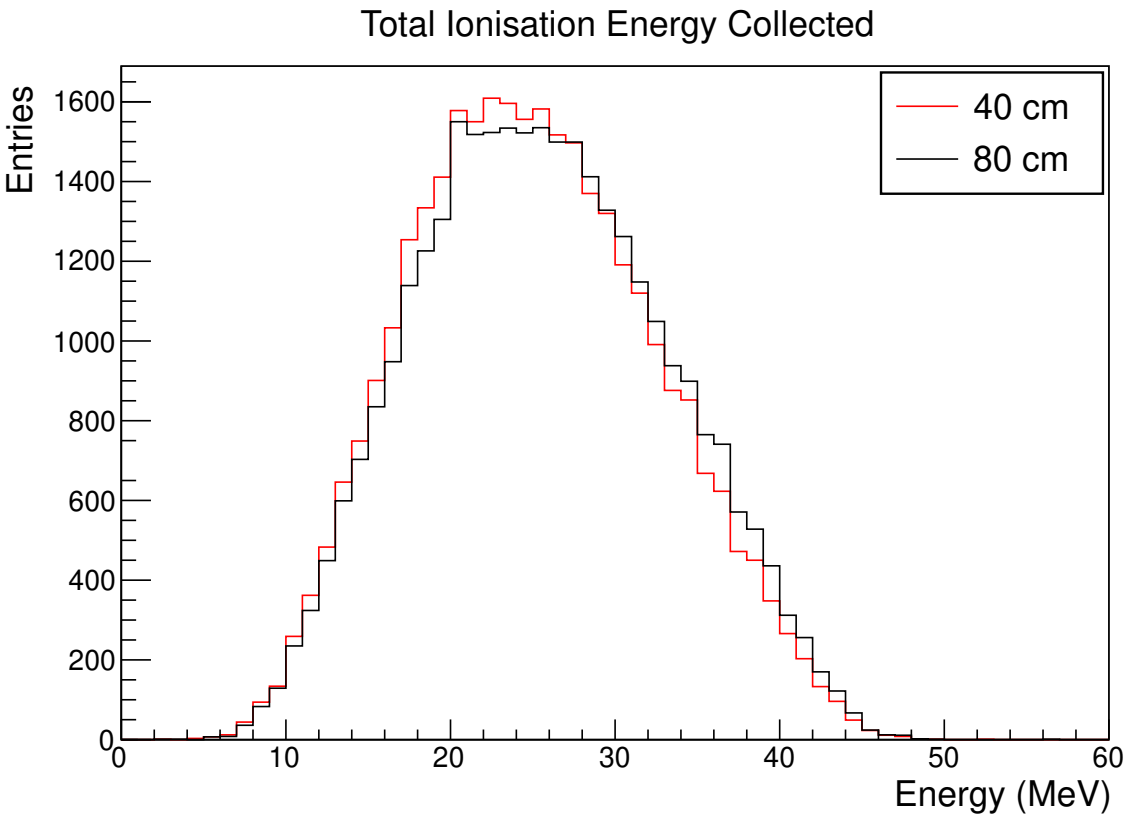


Figure 7.7: Total true ionisation energy recovered for a 40 cm collection radius, and an 80 cm collection radius, in red and black respectively.

collected from radiated photons. The rest of this chapter will discuss an algorithm that was developed to tackle this problem, and its application on Michel electron events in ProtoDUNE–SP data.

7.2 Michel Electron Event Selection

In order to select Michel electrons in ProtoDUNE–SP data, an event selection algorithm was developed based on combining the results from the hit tagging CNN from Chapter 6 with clustering performed by the main ProtoDUNE–SP reconstruction framework, Pandora. This algorithm has four steps, which are detailed below.

1. Select Primary Tracks

This step defines the initial sample of muon candidates, which consists of all primary tracks from the Pandora reconstruction chain.

2. Define Michel Electron Candidates

Based on the sample of muon candidates, a set of Michel electron candidates is defined. This sample includes any daughter of the muon candidate that satisfies the following conditions:

- It has a start point within 5 cm of the primary track endpoint
- It contains a minimum of 5 reconstructed hits on the collection plane

3. Define Best Michel Electron Candidate

The sample of Michel electron candidates is analysed to select the best Michel electron candidate. This is the Michel electron candidate with the largest fraction of Michel–like hits, based on the output of the Michel electron score from the hit tagging CNN, with a threshold of 0.9. In the case of a tie, the candidate with the most hits is chosen.

4. Event Selection Cuts

Events are selected based on the fraction of Michel–like hits in the best Michel electron candidate. Figure 7.8 shows a comparison of the fraction of Michel–like hits in Michel electron candidates for ProtoDUNE–SP data and simulation, normalised by the number of muon candidates. There is a less good modelling in the fraction of Michel–like hits in data for low scores, but a good agreement at higher scores, above around 0.6. Events are selected if the Michel electron candidate contains more than 80% Michel–like hits, this threshold was selected by setting a requirement of greater than 97.5% purity on the event selection algorithm in simulation. This results in 7.3% of the best candidates being selected in simulation, and 6.8% in data.

Based on this algorithm, Michel electron events are selected with a purity of 98% and an overall efficiency of 6% in ProtoDUNE–SP simulation. Figure 7.9 shows the event selection efficiency as a function on Michel electron energy. The efficiency is reasonably consistent across the energy range, however, there is a significant reduction for energies below 5 MeV, which is likely a result of the cut on the minimum number of collection plane hits. The efficiency is highest for Michel electrons in the

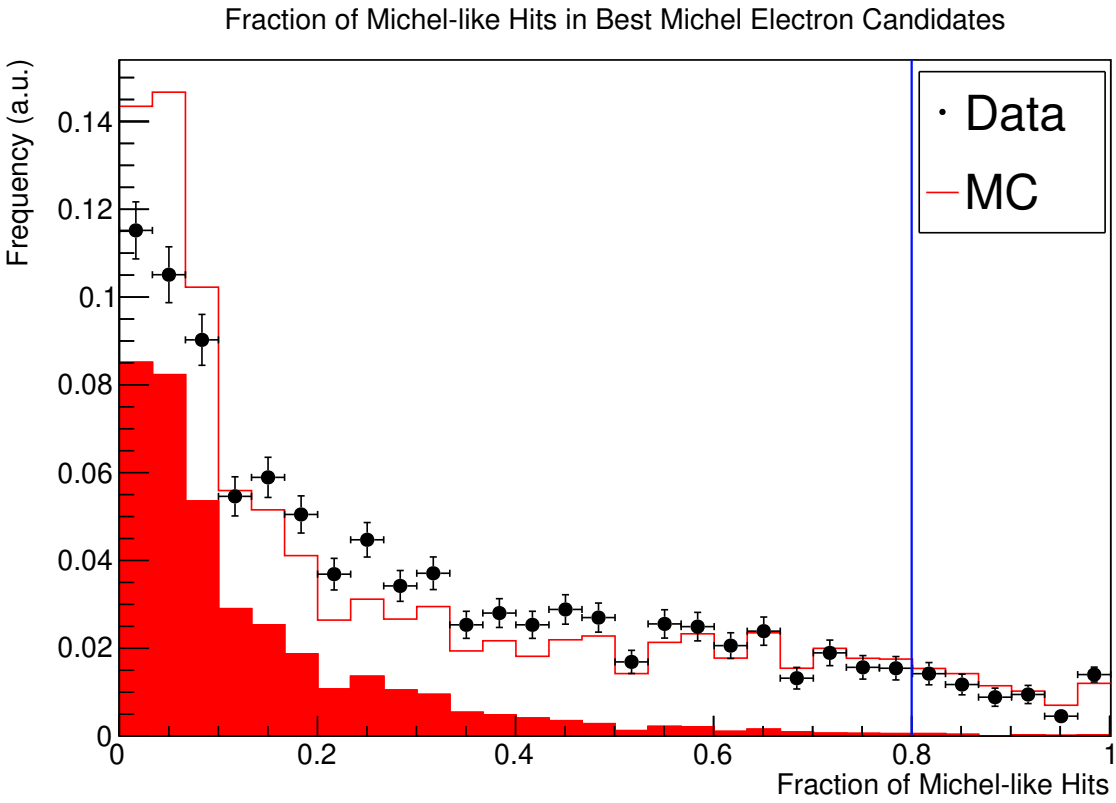


Figure 7.8: Fraction of Michel–like hits in the best Michel electron candidate in data and simulation. The red line is the full distribution in simulation, and the solid red region represents the background in simulation. The black points are the results in data, with statistical error bars only. The vertical blue line represents the cut used to select Michel electron candidates as Michel electron events, with any candidates containing more than 80% of Michel–like hits being selected.

range of 10–40 MeV, and reduces slightly at high energies. At 40–50 MeV, around the peak of the Michel electron energy spectrum, the efficiency is around 5%.

To accurately estimate the energy of hits during energy reconstruction, the true time of the Michel electron needs to be known. If not, then the drift time of the charge is unknown, and the charge attenuation cannot be estimated. Therefore, only tracks with reconstructed true time were considered for the Michel electron analysis. The spatial and angular distributions of the primary muons, normalised by the number of selected muons, are shown in Figure 7.10. There is a reasonable agreement between data and simulation for both the spatial and angular distributions.

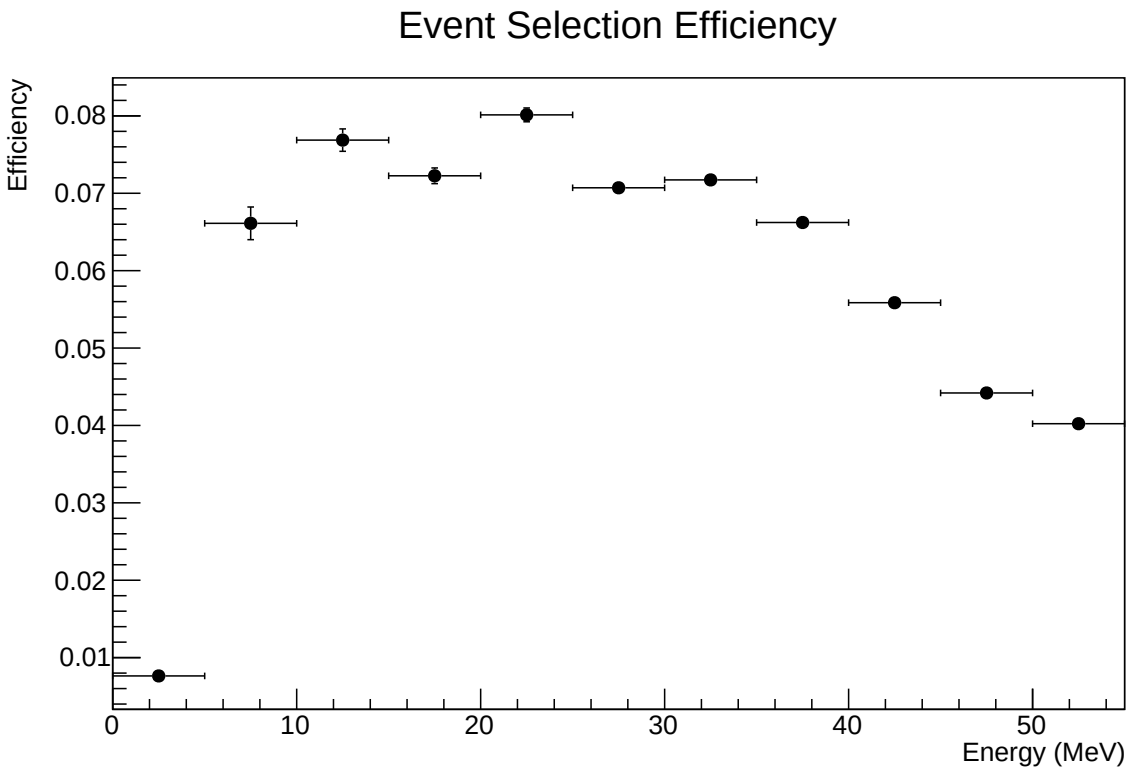


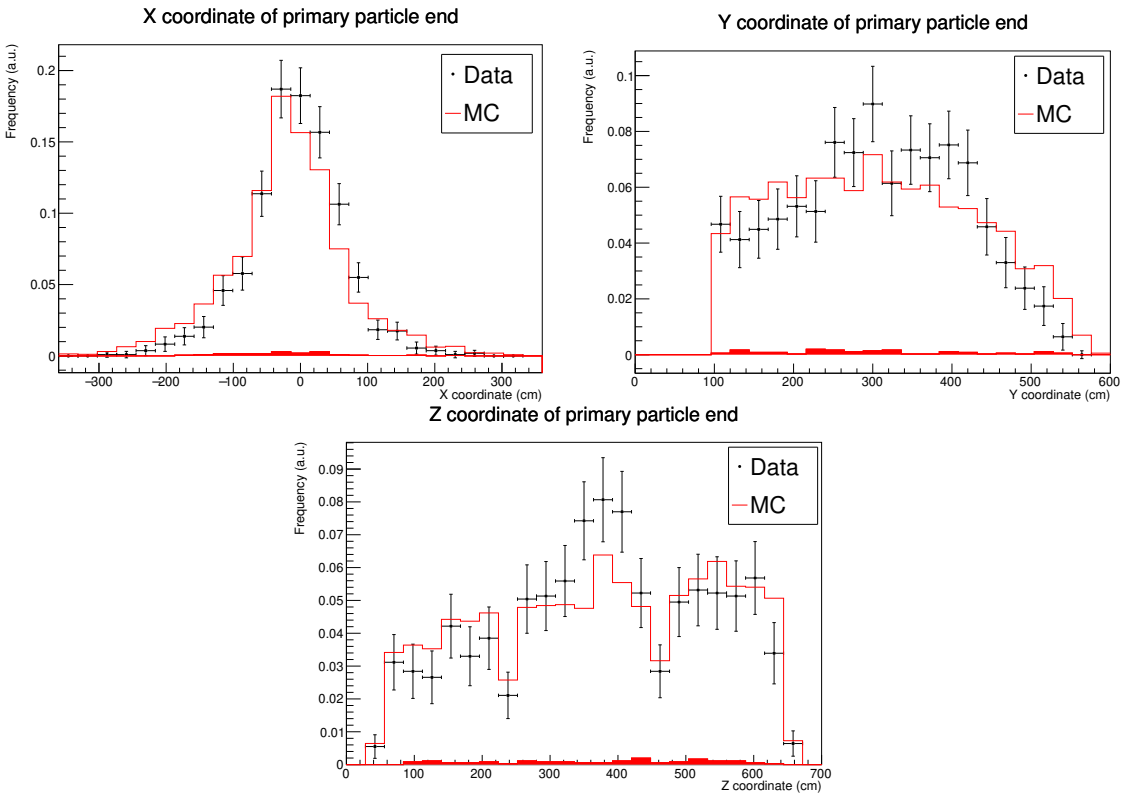
Figure 7.9: Efficiency of Michel electron event selection as a function of true Michel electron energy in ProtoDUNE–SP simulation, with statistical error bars.

7.3 Michel Electron Energy Reconstruction

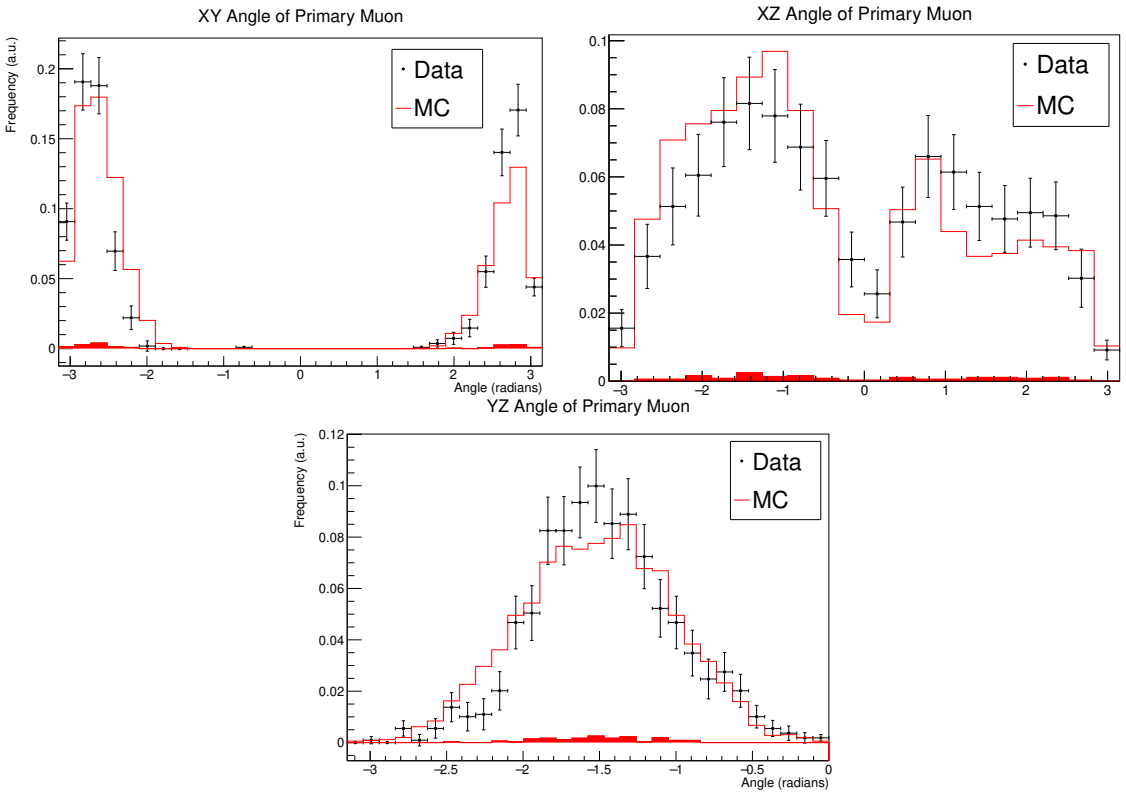
To reconstruct the energy of Michel electrons in liquid argon the relevant hits must first be selected. The reconstructed ionisation energy for the Michel electron is the sum of the reconstructed ionisation energy for all of the selected hits. This section will detail a hit selection algorithm based on a type of convolutional neural network called a U–ResNet, which returns hit selection maps for the Michel electron energy reconstruction. This algorithm is used to select Michel electron hits with a high purity and efficiency, and the resulting reconstructed energy spectrum is used to estimate the energy resolution of ProtoDUNE–SP for electrons in the tens of MeV range.

7.3.1 Michel Electron Hit Tagging with a U–ResNet

A U–ResNet is a type of convolutional neural network that is designed to perform semantic segmentation of images. The U–ResNet network architecture was first



(a) Primary muon end–points: X (drift), Y (vertical), Z (beam).



(b) Primary muon angular distributions.

Figure 7.10: Spatial and angular distributions for primary muons associated with selected Michel electrons, with statistical error bars.

developed for biomedical image segmentation[122]. In semantic segmentation the aim is to return a map of pixels that correspond to the areas of interest; the output of the network is the same dimension as the input with a one–to–one correspondence between input pixels and output pixels. The architecture used for the hit selection algorithm is shown in Figure 7.11. During the first half of the network architecture the resolution of the output is decreased, this is analogous to many conventional CNNs, and it is during this phase the network learns about the content of the image. The second phase of the architecture allows the U–ResNet to rebuild the locations of different learned features within the initial image. This is achieved by passing the details of previous layers to the network, through residual connections, allowing the resolution of the output map to be slowly increased back to the original resolution.

The network architecture used for the Michel electron reconstruction is shown in Figure 7.11. The network consists of a repeating structure that contains the following key components:

- Convolutional blocks that contain multiple convolutional layers
- Pooling layers for downscaling in the first half of the network
- Residual connections and up–sampling in the second half of the network

The convolutional blocks each contain three convolutional layers, which is shown in Figure 7.11b. These take the form of inception units[91] with leaky ReLU activation functions; a schematic of the inception units is shown in Figure 7.11c. During the first half of the network, maximum pooling is used to downsample the size of the feature maps, which reduce by a factor of two in each max pooling layer. Then, during the second half of the network, the up–sampling layers increase the resolution of the feature maps by a factor of two, such that the output of the network has the same resolution as the input. The residual connections from early layers, are combined with the up–sampled feature maps, to allow the network to reconstruct the location of the learned features. As with the CNN from the previous chapter, both dropout and early–stopping are implemented to prevent over–fitting.

In the Michel electron case, the aim of the network is to return a map of all ionisation energy deposits from the Michel electron, including the initial track and

any secondary deposits from radiated photons. The inputs and outputs are two dimensional images of the location of reconstructed hits, centered on the selected Michel electron candidate. The amplitude of each input pixel is given by the integrated charge of any reconstructed hits within the pixel. For the true output images the pixels have an amplitude of 1 if they contain a Michel electron hit, and 0 otherwise. Only data from the collection plane is used, because there is a higher signal to noise ratio on these wires.

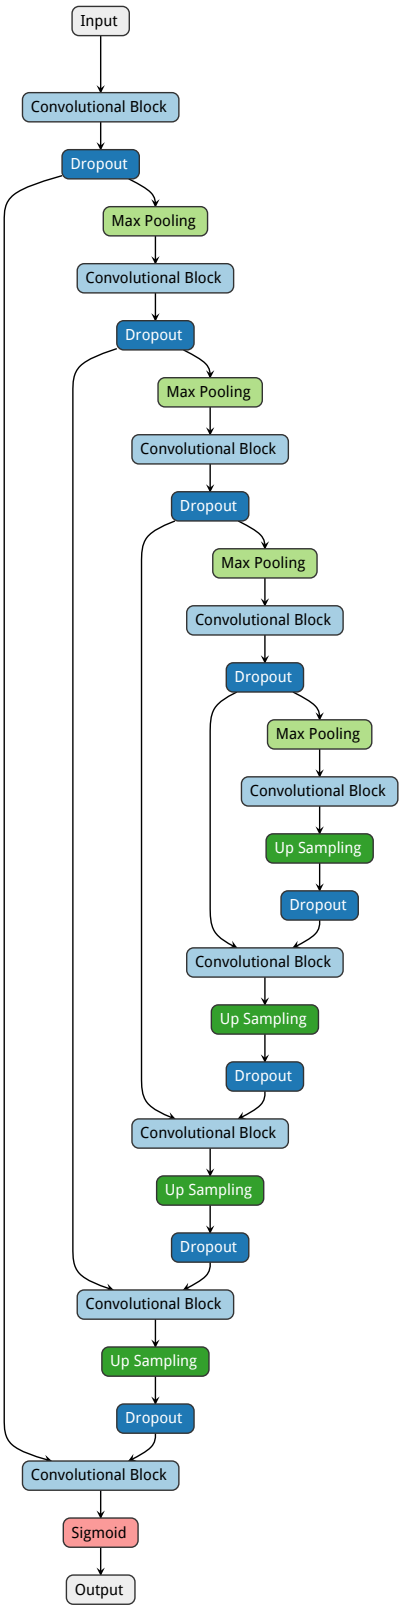
The F_β metric[115] was used as the loss function for the U–ResNet. This loss is a generalised version of the F_1 metric discussed in Chapter 6, which allows for the relative importance of purity and efficiency to be tuned,

$$F_\beta = \left(1 + \beta^2\right) \frac{\text{purity} \cdot \text{efficiency}}{(\beta^2 \cdot \text{purity}) + \text{efficiency}}.$$

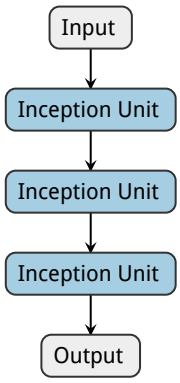
This loss rewards the network for selecting as many correct hits as possible (high efficiency), while penalising it for selecting more hits that necessary (low purity). The F_β score lies between 0 and 1, with a score of 1 corresponding to a perfect match. The β parameter was chosen to be 2, such that efficiency was considered to be more important than purity. This was found to improve the performance of the Michel electron reconstruction in terms of energy resolution and bias, however, no extensive optimisation of β was performed.

The datasets for the training process were generated from a full simulation of the ProtoDUNE–SP detector under beam operation, including both cosmic–ray and beam particles. The images produced contain the location and integrated charge for each hit within the image window. The training data was split into training, test, and validation sets in the ratio 80:10:10. In total around 15,000 images were used in the combined training, test, and validation datasets.

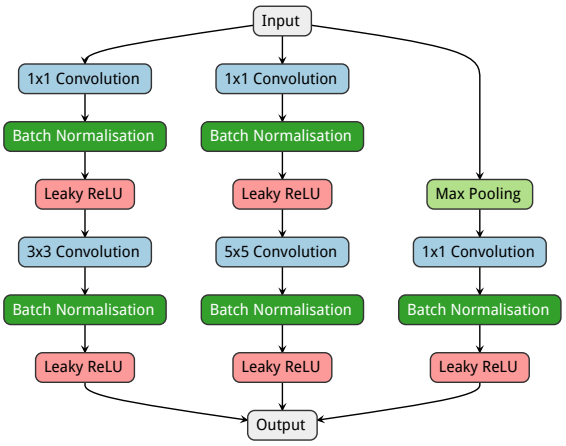
As with the hit tagging CNN discussed in Chapter 6, the training and validation scores were monitored throughout training using TensorBoard. The weights of the network were saved after each epoch, and the final weights were those from one epoch before the validation score first increased. Figure 7.12 shows the evolution



(a) Overall structure.



(b) Convolutional block.



(c) Inception unit.

Figure 7.11: U–ResNet CNN architecture used to select Michel electron ionisation energy deposits. Visualisation adapted from the output of the Netron neural network viewer[108].

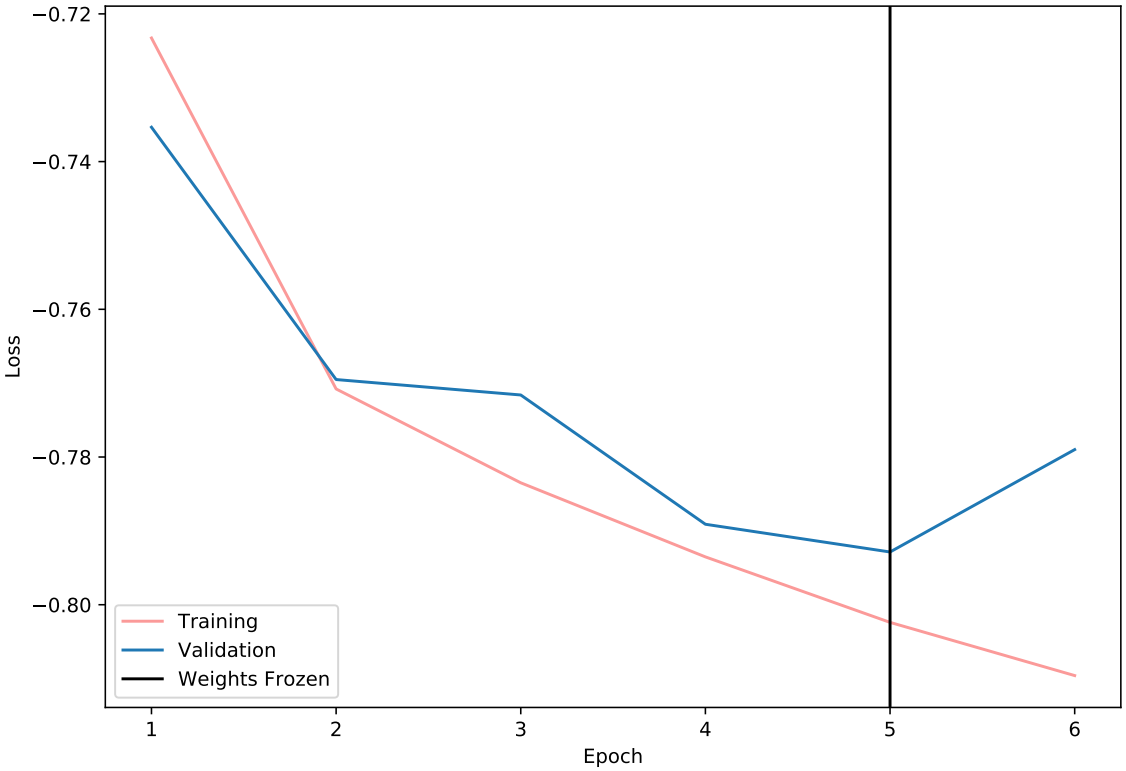


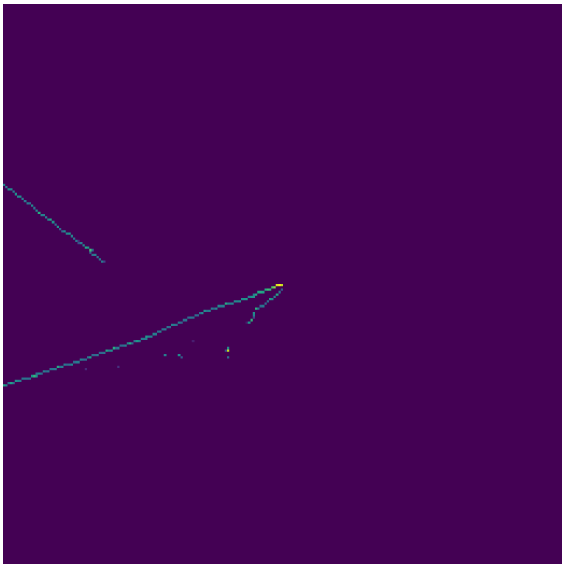
Figure 7.12: U–ResNet training and validation loss as a function of epoch. The weights were frozen based on an early stopping algorithm, which is demonstrated by the vertical black line.

of the loss over time, along with a vertical line representing the loss at which the weights were chosen.

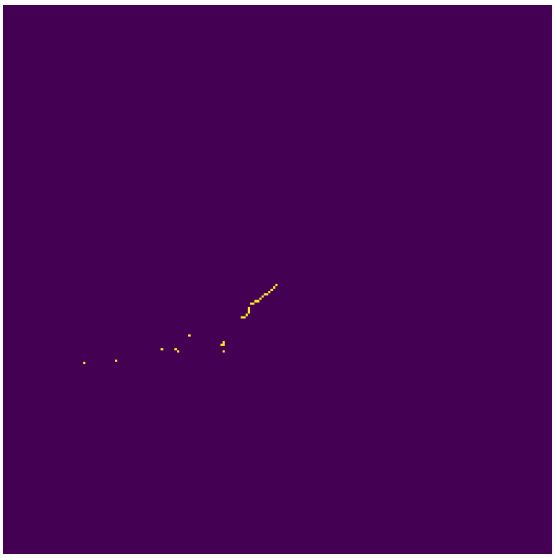
A demonstration of the output of the U–ResNet is given in Figure 7.13, which shows the input, output, and truth images for an event from ProtoDUNE–SP simulation.

7.3.2 Michel Electron Reconstruction

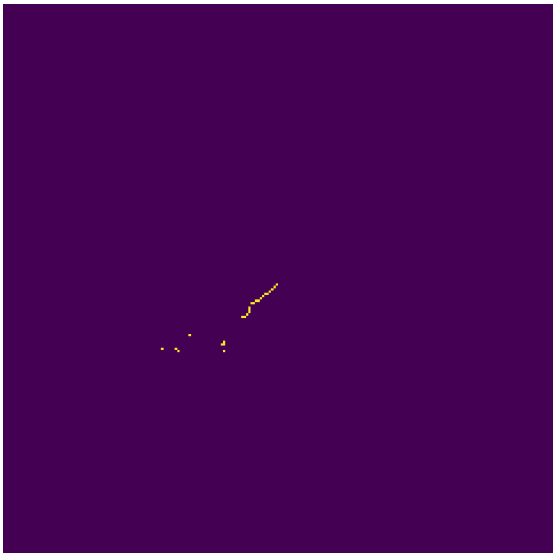
An additional dataset was prepared to evaluate the performance of the Michel electron reconstruction in simulation; this dataset was part of the same batch of simulation as the training, test, and validation data. For the ProtoDUNE–SP data sample, all data from runs 5387, 5809, and 5770 were used; at the time of writing, these runs are the only runs for which the energy calibration constants have been calculated.



(a) Input.



(b) Truth.



(c) Output.

Figure 7.13: Example input, truth, and predicted images for U–ResNet.

7.3.2.1 Hit Selection

The U–ResNet produces a sharply peaked output distribution in both data and simulation, as seen in Figure 7.14. The histograms here are normalised by the number of entries. Both distributions show sharp peaks at zero and one, but the distribution has slightly sharper peaks in simulation. Hits from the input images are selected as Michel electron hits if their score exceeds a selection threshold of

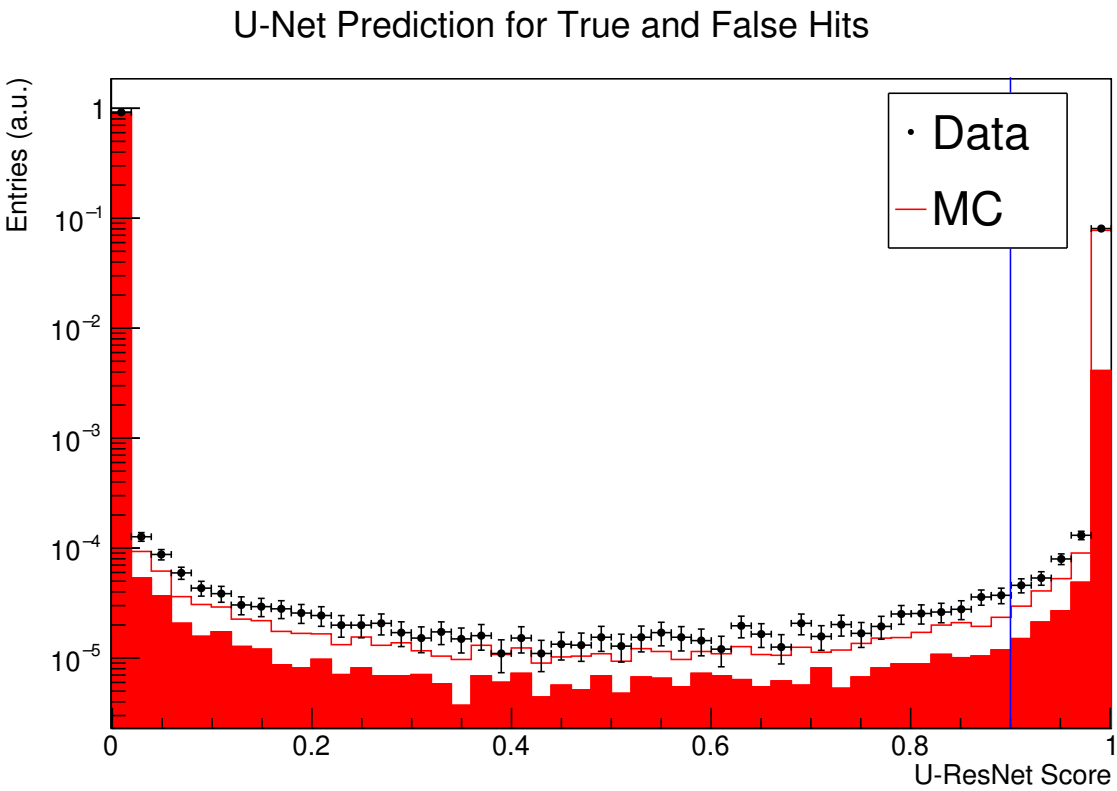


Figure 7.14: U–ResNet predicted score distribution for hits in Michel electron ionisation images. The simulated results are shown by the red line, with the backgrounds, including track hits and hits due to noise, as a shaded red region. The data is shown in black, with statistical error bars. The vertical blue line represent the selection cut for hits that were included in the energy reconstruction, all hits with a score above 0.9 where included in the energy reconstruction.

0.9. In simulation 7.8% of hits are selected, while in data the fraction selected is 8.0%. Therefore, there is a fractional difference of 2.5% in the fraction of selected hits between data and simulation.

The performance of the hit tagging algorithm was analysed with the simulated sample. Based on the score distributions for true and false hits, the purity and efficiency of the hit tagging algorithm can be evaluated. The purity and efficiency are defined as

$$\begin{aligned} \text{Purity} &= \frac{N_{TP}}{N_{TP} + N_{FP}} \\ \text{Efficiency} &= \frac{N_{TP}}{N_{TP} + N_{FN}} \end{aligned}$$

where N_{TP}, N_{FP}, and N_{FN} are the number of true–positives, false–positives, and false–negatives respectively. These parameters give a quantitative evaluation of

the performance of the hit tagging algorithm.

The purity and efficiency of the hit tagging was calculated for selection thresholds in the range $[10^{-7}, 1-10^{-7}]$, which returned purities in the range 0.91–0.93, and efficiency in the range 0.94–0.96. For thresholds outside of this range the purity and efficiency remained constant, except for at thresholds of 0 and 1, this is due to the underlying precision in Keras, which sets the smallest non–zero float to be 10^{-7}. The hit tagging algorithm produces a high purity and efficiency throughout the range of thresholds, however, the combination of the steep score distribution with the underlying precision in Keras means that little can be done to optimise the performance by varying the threshold in this case.

The number of hits selected per event for data and simulation is shown in Figure 7.15. Around 10 hits are selected on average per event, with a tail extending up to around 30 hits. The distribution of the backgrounds peaks at a lower energy than the true events, but is only a small contribution to the overall distribution. A good agreement is seen between the data and simulation for the number of selected hits; these hits are then used to reconstruct the deposited ionisation energy of the Michel electron.

7.3.2.2 Ionisation Energy Reconstruction

The total ionisation energy is reconstructed by summing the hit–by–hit ionisation energy for hits selected by the U–ResNet. The ionisation energy for each hit is reconstructed from the hit integral in ADC as

$$E_{hit} = \frac{I_{hit} \times C_X \times C_{YZ} \times N \times W_{ion}}{C \times R},$$

where E_{hit} is the reconstructed hit energy in MeV, I_{hit} is the integrated hit charge in ADC, C_X is an X–correction factor that is dependent on the X (drift) coordinate of the hit within the TPC, C_{YZ} is a YZ–correction factor that is dependent on the Y (vertical) and Z (beam) coordinates of the hit within the TPC, N is a dimensionless normalisation factor that normalises the data and MC distributions to give the same magnitude, W_{ion} is the ionisation energy of argon in MeV per

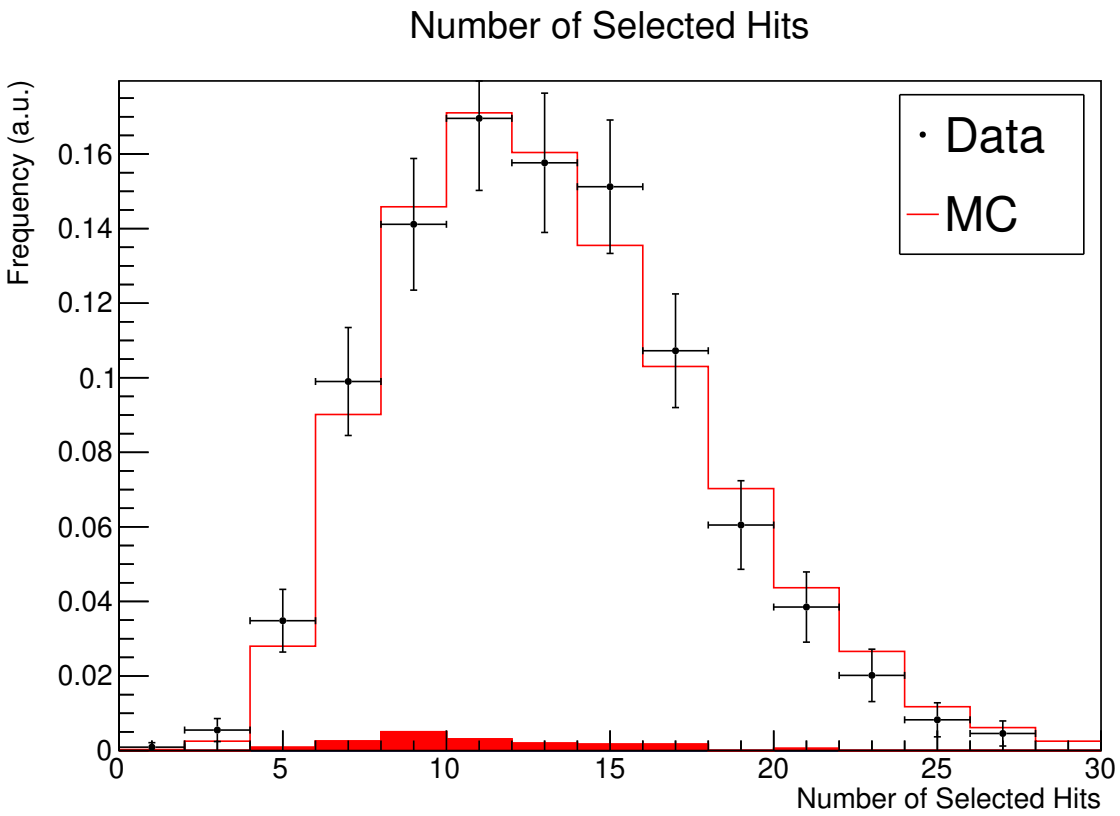


Figure 7.15: Number of hits selected by the U–ResNet for Michel electron candidate events in data and simulation. The distribution in data includes statistical error bars.

electron, C is a constant conversion factor that has units ADC per electron, and R is the recombination factor. The distribution of reconstructed hit energies in ProtoDUNE–SP data and simulation is shown in Figure 7.16, which shows a good agreement between data and simulation.

The position dependent calibration matrices correct for any non-uniformity in the detector response across the TPC. In the X direction, which is parallel to the drift direction, the main contributing factors are attenuation due to electron absorption, and variations in the electron drift velocity due to space charge effects. The main contributing factor for the YZ–correction factor is wire–to–wire response variations.

As discussed in Chapter 4, the recombination factor is a dQ/dx dependent factor that depends on the conditions in the liquid argon. Due to the shortness of Michel electron tracks and the other charge deposits, it is challenging to assign dQ/dx on a hit–by–hit basis for this sample. Therefore, an average recombination factor is

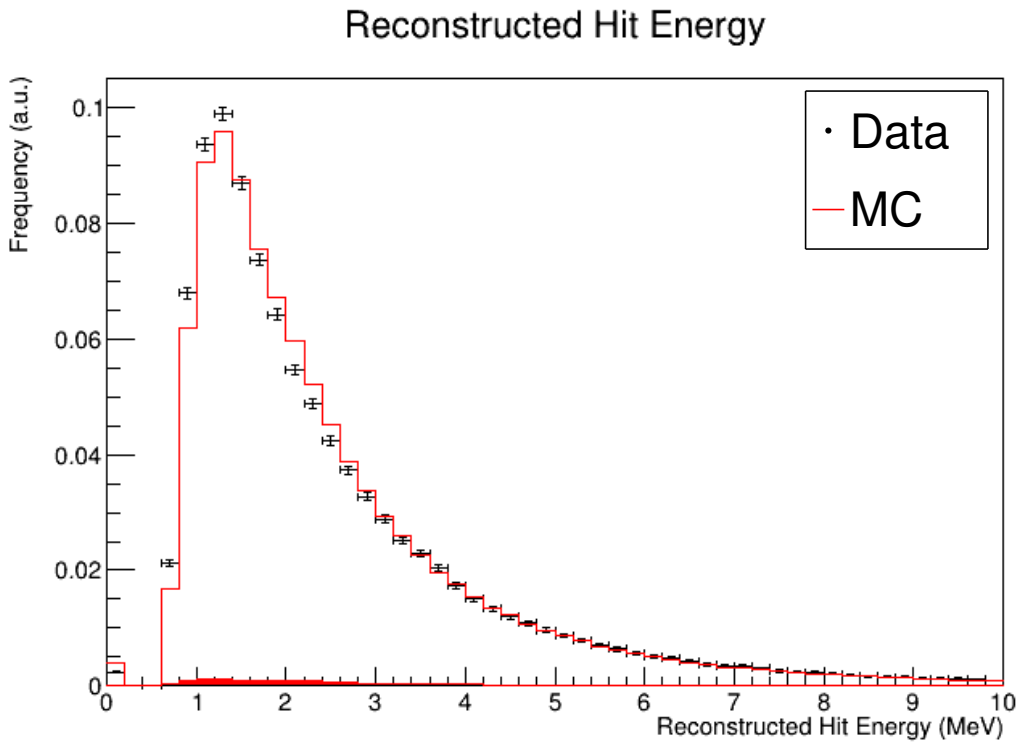


Figure 7.16: Reconstructed ionisation energy for hits in U–ResNet input images in data and simulation. The distribution in data includes statistical error bars.

used for all hits. The recombination factor was calculated using the box model[73] under ProtoDUNE–SP operating conditions, which gives an average value of 0.69.

As a result of the spatially dependent calibration factors, the energy of hits can only be accurately reconstructed for events with a reconstructed true time, which means that the track needs to have been stitched by the track stitching algorithm described in Section 3.7. This algorithm has a low efficiency due to the fact that only cathode crossing tracks can be used, and has a different efficiency in data and simulation due to differences in the alignment properties of track segments. The overall efficiency of the track stitching algorithm is around 15% in simulation, and only 4% in data.

The reconstructed Michel electron energy is the sum of the reconstructed ionisation energy for all selected Michel electron hits. The reconstructed ionisation energy spectrum for Michel electrons in ProtoDUNE–SP is shown in Figure 7.17. The distribution peaks at around 20 MeV and has a tail up to just under 50 MeV, which

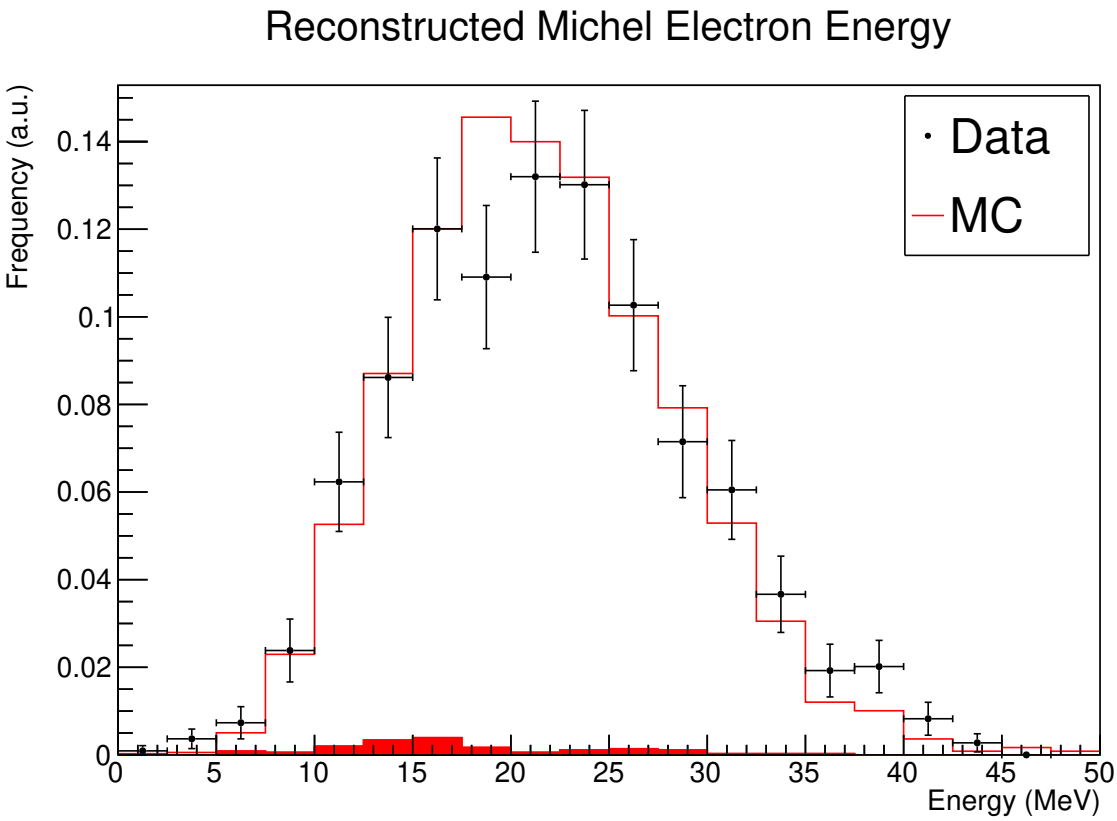


Figure 7.17: Reconstructed ionisation energy for Michel electron candidates in ProtoDUNE–SP data and simulation. The distribution in data includes statistical error bars.

is consistent with the true ionisation energy deposited by Michel electrons within a 40 cm radius, shown in Figure 7.7. Overall, there is a good agreement between the data and simulation in terms of reconstructed energy. The main difference between the data and simulation is a slightly longer tail at high energy in data.

The performance of the ionisation energy reconstruction was evaluated with the simulated Michel electron sample. First, the reconstructed ionisation energy was compared to the true ionisation energy, and a linear scaling factor was calculated to correct for any energy offset. The correction factor was calculated based on a linear fit to the distribution of reconstructed energy against true energy, which is shown in Figure 7.18. The reconstructed energy was corrected based on this linear fit through the origin,

$$E_{corr} = E_{reco} \,/\, 0.9823,$$

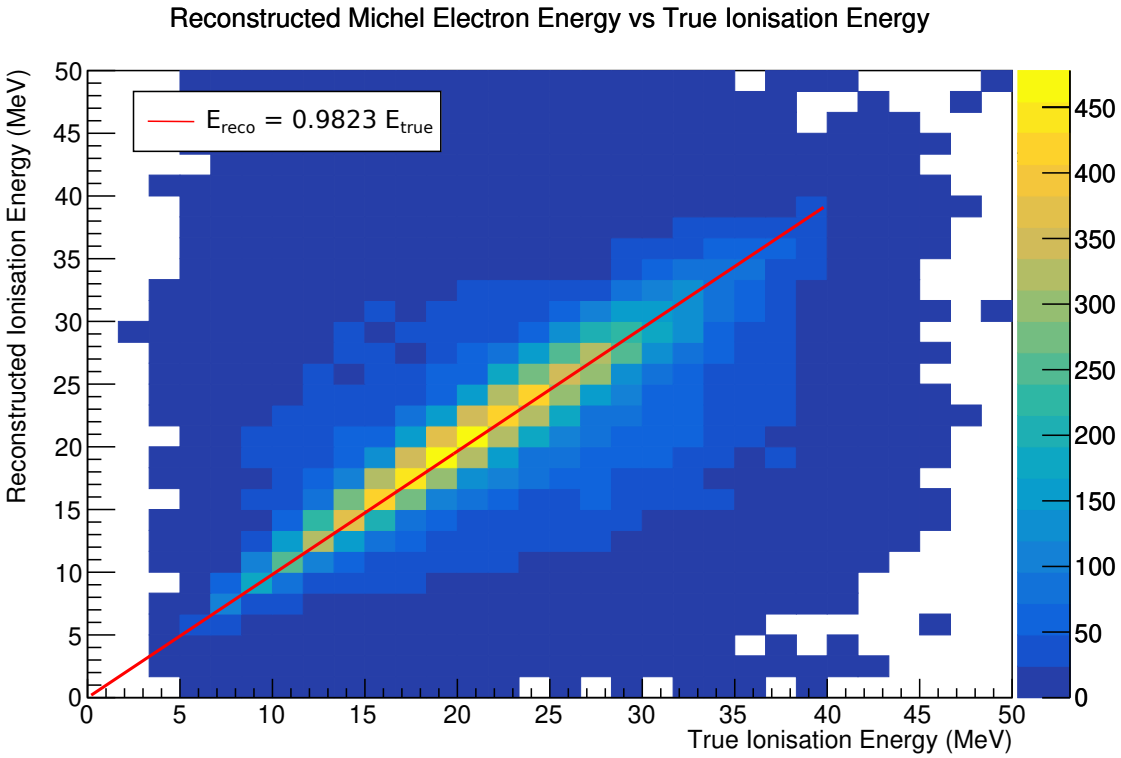


Figure 7.18: Reconstructed ionisation energy vs total deposited ionisation energy for Michel electron events in ProtoDUNE–SP simulation.

and the corrected energy was used to estimate the energy resolution and bias of the Michel electron energy reconstruction.

To estimate the resolution and bias of the ionisation energy reconstruction, the fractional difference between the reconstructed and true ionisation energy was considered. The fractional difference is defined as,

$$\text{Fractional Difference} = \frac{E_{corr} - E_{true}}{E_{true}},$$

such that the spread of the distribution is a measure of the fractional uncertainty of the reconstructed energy. The energy resolution and bias were estimated by fitting a gaussian to the fractional difference distribution; the standard deviation of the gaussian was taken to be an estimate of the energy resolution, and the mean as an estimate of the bias. The distribution of the fractional difference for all events, and the associated gaussian fit, are shown in Figure 7.19. Based on this fit the energy resolution was estimated to be around 10% and the bias around 1%. However, it is worth noting that the fractional difference distribution has a

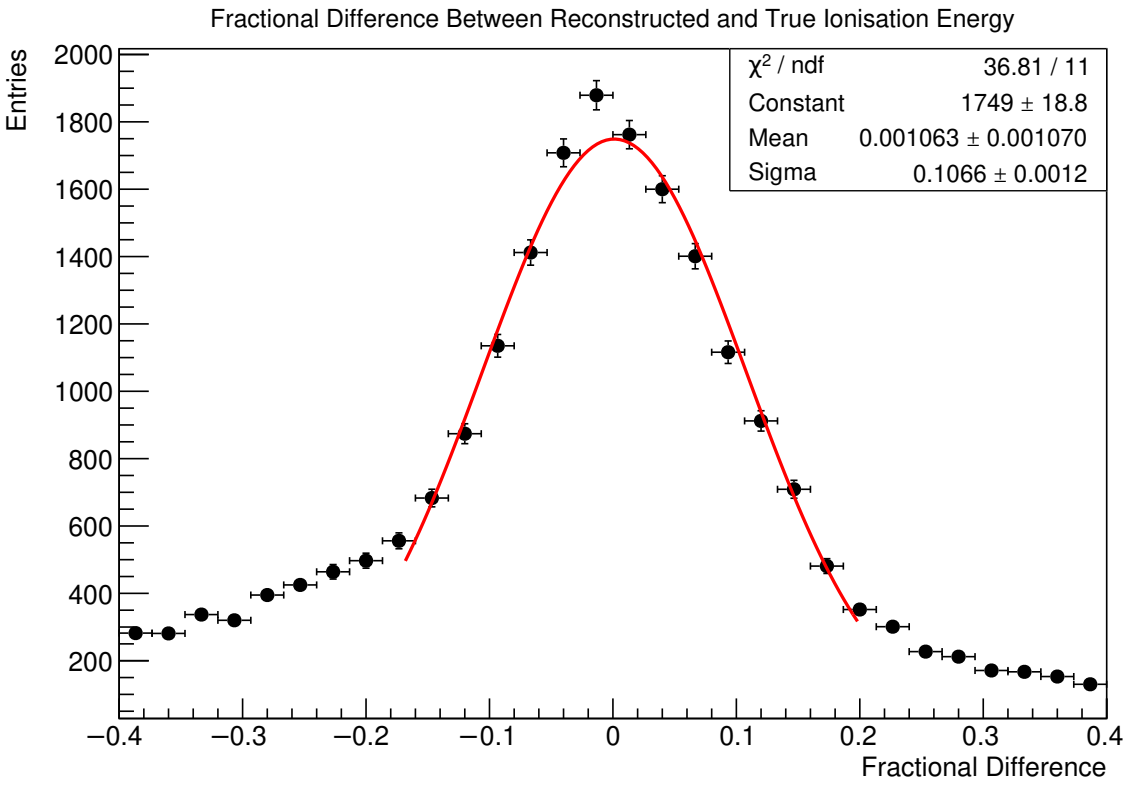
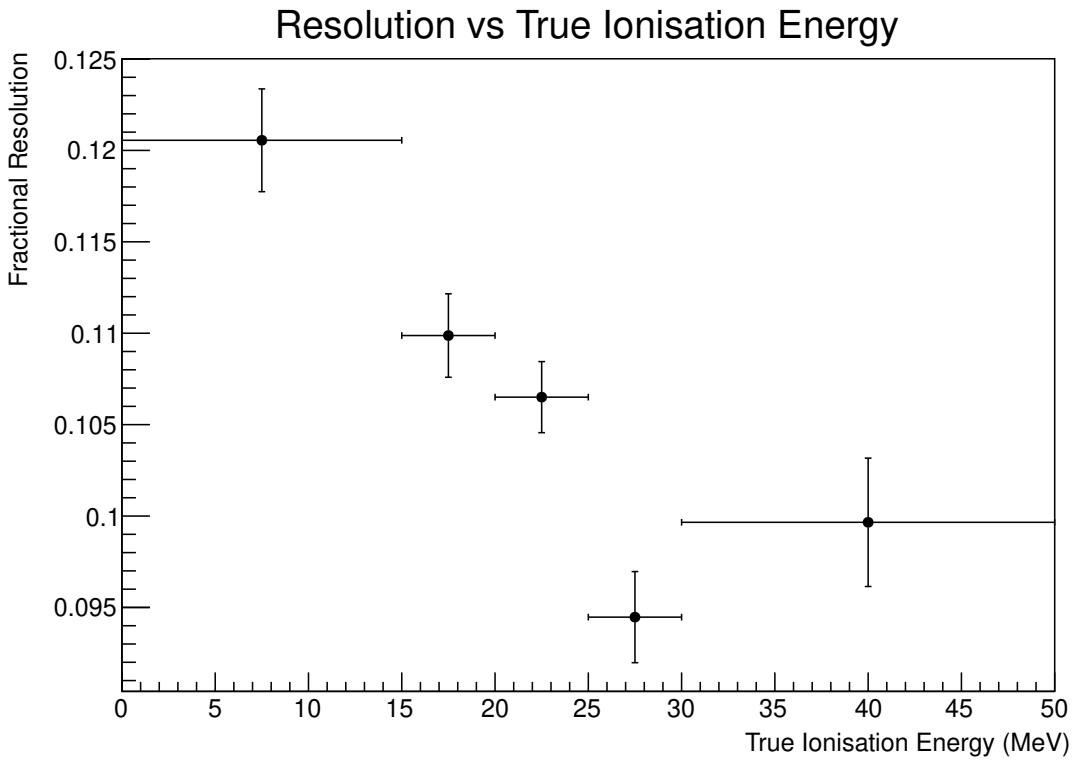


Figure 7.19: Fractional energy difference between reconstructed and true Michel electron ionisation energy in ProtoDUNE–SP simulation, with statistical error bars.

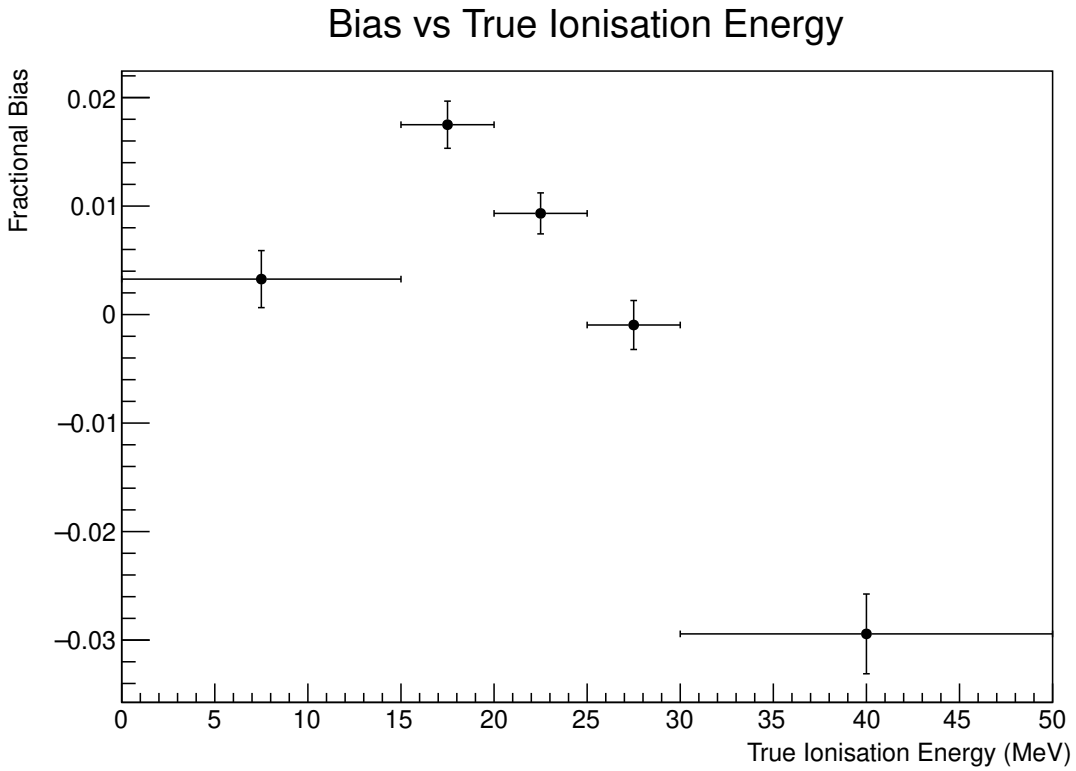
fairly significant tail for negative fractional differences and, therefore, the fit was restricted to the central region of the distribution.

Studies into the energy resolution and bias as a function of energy gave insights into the source of the tail in Figure 7.19. These were investigated by binning the fractional energy difference distribution in terms of true ionisation energy, and performing the same gaussian fit to the fractional difference distribution for each energy bin. The measured energy resolution and bias as a function of true ionisation energy are shown in Figure 7.20, the associated fits are given in Appendix A.

The energy resolution and bias as a function of energy demonstrate that the energy reconstruction algorithm is capable of good energy reconstruction performance. However, it is important to note that above around 25 MeV there is significant tail for negative fractional differences. This is due to a reduction in the hit tagging efficiency in the high energy region, which causes less of the initial Michel electron energy to be collected by the CNN. No attempt was made to fit this tail



(a) Resolution.



(b) Bias.

Figure 7.20: Energy resolution and bias for Michel electrons as a function of true ionisation energy deposition in ProtoDUNE–SP simulation. The error bars represent the statistical uncertainty only.

Energy Bin	Gaussian Mean (%)	Gaussian Sigma (%)	Mean (%)	RMS (%)
0–15 MeV	$+0.3 \pm 0.3$	12.1 ± 0.3	$+0.8 \pm 0.2$	15.0 ± 0.2
15–20 MeV	$+1.8 \pm 0.2$	11.0 ± 0.2	$+1.8 \pm 0.2$	14.7 ± 0.2
20–25 MeV	$+0.9 \pm 0.1$	10.6 ± 0.2	$+0.0 \pm 0.2$	14.5 ± 0.1
25–30 MeV	-0.1 ± 0.2	9.5 ± 0.2	-3.4 ± 0.2	14.9 ± 0.2
30–50 MeV	-2.9 ± 0.4	10.0 ± 0.4	-8.6 ± 0.3	16.0 ± 0.2

Table 7.1: Comparison of energy resolution estimates for Michel electron ionisation energy from ProtoDUNE–SP simulation.

and, therefore, the estimated energy resolution and bias based on the gaussian fits overstate the performance in the high energy region. A more conservative estimate can be made by considering the mean and RMS of the distribution. Table 7.1 gives a comparison between the estimated bias and resolution based on these two methods.

Contributions to Ionisation Energy Resolution

The ionisation energy resolution is made up of a combination of the hit–by–hit energy resolution, and the spread introduced due to imperfect hit tagging. The contribution from the hit–by–hit energy resolution can be estimated by considering the fractional energy difference in the case of perfect hit tagging, which sums the reconstructed energy of all the true Michel electron hits in the image. The sum over the hits in each image is considered because this allows the error in the hit–by–hit recombination factor, which arises due to the use of an average recombination factor during reconstruction, to be averaged over the hits in the image. The fractional energy difference in this case is fit with a gaussian distribution, which is shown in Figure 7.21.

The contribution to the ionisation energy resolution due to the hit energy reconstruction was estimated based on the parameters from the gaussian fit in Figure 7.21. The bias in the hit energy reconstruction was estimated to be around 3%, and the resolution was estimated to be approximately 8%. The remaining energy resolution and bias seen in the Michel electron ionisation energy reconstruction is a result of the spread due to imperfect hit tagging.

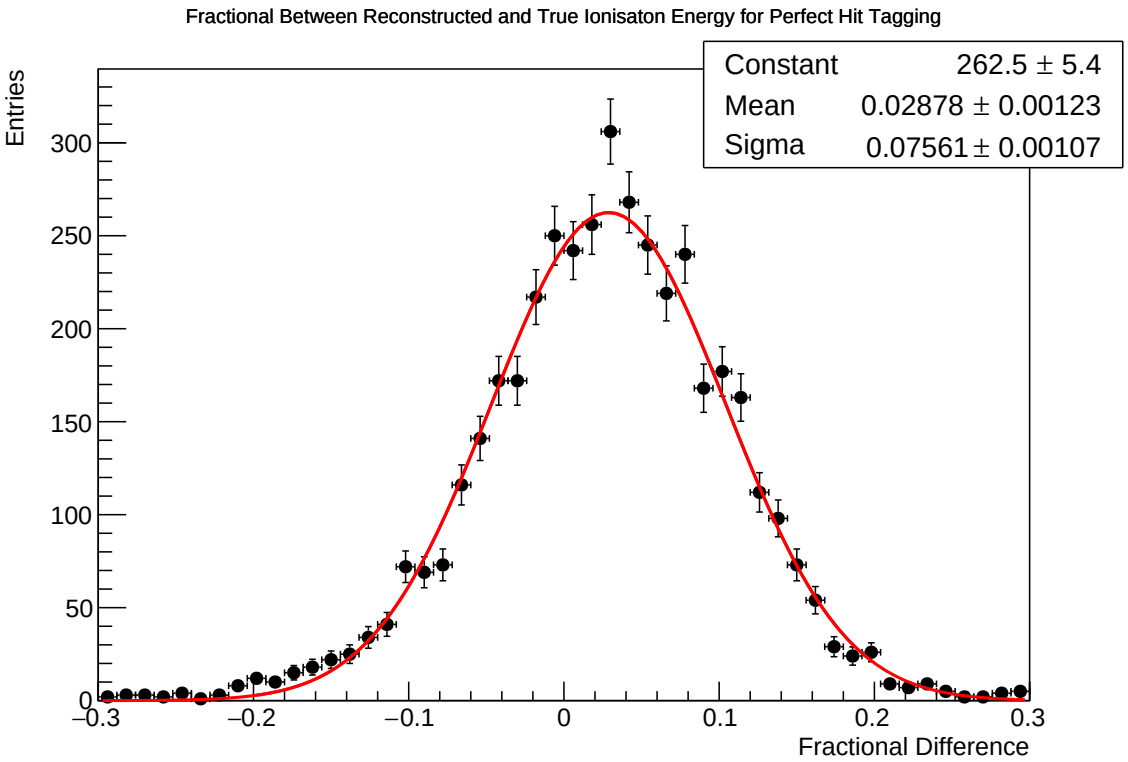


Figure 7.21: The fractional difference between the true ionisation energy deposition and the reconstructed ionisation energy deposition for Michel electron reconstructed with perfect hit tagging, with statistical error bars. The distribution is fit with a gaussian, which is shown in red, and the fit parameters are shown in the top right corner.

Systematic Uncertainties

While no study has been made into the magnitude of the systematic uncertainties associated with this reconstruction algorithm, this section will aim to discuss the possible sources of systematic uncertainty in the network, and to suggest possible mitigation strategies for any similar algorithms that may be developed for supernova neutrino energy reconstruction in the future.

As with the hit tagging CNN in Chapter 6, there are likely to be systematic uncertainties associated with the effects that impact the generation of images for the training and evaluation of the network. These include the noise model and space charge model in the simulation, which will impact the U–ResNet scores due to the effects discussed in Section 6.3.1. These effects could be mitigated with a network that has been trained on images produced with a variety of noise and SCE models, which could be generated with a generative adversarial network,

as discussed in Section 6.3.1.

Other possible sources of systematic uncertainty for the Michel electron energy reconstruction include:

- Fiducial volume cuts
- The requirement of a true time (t_0), and the different efficiency for t_0 tagging in data and simulation
- The calibration constants used during energy reconstruction
- The use of an average recombination correction
- The calculated linear energy correction factor

A study of the size and implications of these uncertainties for low energy electron reconstruction in ProtoDUNE–SP and DUNE remains as a possible future study. In addition, it is worth noting that the analysis of these uncertainties for Michel electrons in ProtoDUNE–SP should not be considered as an estimate of systematic uncertainties for supernova neutrino events in DUNE. Some of the reasons for this will be outlined briefly in Section 7.4.

Comparison to Other LArTPC Experiments

In order to put the performance of the Michel electron energy reconstruction into context, it is useful to compare the results presented here with those from other LArTPC experiments. The comparisons here are restricted to the 0–30 MeV region, recognising that improvements are still required in order to reduce the tail for higher ionisation energies. The closest comparison is against the MicroBooNE experiment, which is a surface level LArTPC that used ionisation energy for the reconstruction[123]. LArIAT and ICARUS have also performed Michel electron reconstruction[124, 125], however, LArIAT used the scintillation light for reconstruction, and ICARUS is deep underground and, therefore, has significantly less backgrounds.

The MicroBooNE experiment gives the closest comparison for Michel electron reconstruction, as both MicroBooNE and ProtoDUNE–SP are surface level TPCs,

and they both use ionisation energy for Michel electron reconstruction. In MicroBooNE, the ionisation energy resolution for Michel electrons was found to be between 15–25% in the 0–30 MeV region[123]. In this region the algorithm developed here gives an improved energy resolution of around 10–12%, which highlights the potential of the semantic segmentation approach if the bias at higher energies can be improved such that the approach is applicable across the full energy range.

ICARUS also used ionisation energy for Michel electron reconstruction, however, the ICARUS detector was located deep underground and, therefore, they were subject to much lower backgrounds. In ICARUS, the ionisation energy resolution was measured to be in the range of 3–8%[124], however, these results did not take into account any radiated energy deposits, which form an important part of the Michel electron signature. Therefore, the results from ICARUS are not directly comparable to the measured energy resolution presented here.

LArIAT is a significantly smaller TPC than ProtoDUNE–SP and, therefore, despite being at surface level, there is a sufficiently low rate of cosmic–rays that the photon detection system could be used to tag and measure Michel electron events. LArIAT measured an energy resolution on the order of 20% for Michel electrons based on the scintillation light signal[125].

In DUNE, both the ionisation energy and scintillation light will be used to reconstruct supernova neutrinos. The results in the DUNE technical design report are based on energy resolution estimates from the standard LArSoft reconstruction framework, which give an ionisation energy resolution of around 20%, and a scintillation light energy resolution in the range of 10-20%[50]. Therefore, if the tail at high energy can be improved, a semantic segmentation algorithm, similar to the one presented here, may be able to improve the energy reconstruction performance for supernova neutrinos in the DUNE Far Detector. This is an important factor in many supernova neutrino physics studies, in which the mean neutrino energy is an important parameter[50].

7.4 Conclusion

In the energy regime of Michel electrons, and supernova neutrinos, the electromagnetic energy deposition of electrons moves from an ionisation–dominated to a radiation–dominated region. As such, these events have a signature that contains both track–like and a shower–like components. Therefore, they require a unique reconstruction algorithm in order to maximise the energy collected from the events.

This chapter has presented a novel Michel electron reconstruction algorithm based on machine learning techniques. A sample of Michel electrons was selected in ProtoDUNE–SP data with a purity of over 98% and an overall efficiency of around 6%. The low efficiency is acceptable due to the very high rate of Michel electron events in ProtoDUNE–SP, which will lead to a large high purity sample once more calibrated data has been collected and processed. The ionisation energy of these events was reconstructed based on a semantic segmentation algorithm, which uses a U–ResNet CNN architecture. The performance of this algorithm showed an improvement in energy resolution over similar experiments, with an ionisation energy resolution of around 10-15%. However, there is still room for improvement, particularly in the high energy region where a large tail is seen in the fractional difference distribution.

The promising results in the low energy region based on this method suggest that, with further study, the semantic segmentation approach could prove to be an effective algorithm for reconstructing low energy particles in LArTPCs. In order to increase the performance in the high energy region, it may be useful to consider increasing the amount of training data in this region. At the time of writing, there is insufficient ProtoDUNE–SP simulation to produce a flat training sample with enough data. Therefore, this remains as a possible future study, which has the potential to significantly increase the performance of this algorithm by increasing the generalisability across the energy range.

While Michel electrons provide an interesting sample for studying low energy electron reconstruction, it is important to note the differences between Michel electron and supernova neutrino events in LArTPCs, as well as the different

conditions between ProtoDUNE–SP and DUNE, when considering the applicability of these results to supernova neutrino reconstruction in DUNE. A number of relevant differences will now be detailed.

Firstly, Michel electron events are accompanied by a muon track in LArTPC events. This aids event selection but also masks some of the Michel electron energy due to the presence of the Bragg peak of the muon.

Secondly, the peak of the Michel electron energy spectrum is around 40 MeV as opposed to around 15 MeV in the case of supernova neutrinos. As a result, Michel electron events contain a larger fraction of radiated energy deposition than supernova neutrinos. Therefore, the energy reconstruction results in the lower energy region are likely to give a more representative prediction of the expected energy resolution for supernova neutrinos in DUNE.

Finally, DUNE will have a significantly lower background of cosmic–ray events than ProtoDUNE–SP, which will give two benefits: the hit tagging algorithm will have fewer unnecessary hits to filter due to cosmic–ray activity, and it will be significantly easier to match signals from the TPC to those in the photon detectors. This will allow for both ionisation and scintillation energy deposits to be utilised during energy reconstruction in DUNE, which was not possible in ProtoDUNE–SP.

As a result of these differences, it will be important to conduct separate studies of supernova neutrino reconstruction under the operating conditions expected in the DUNE far detector. This study has demonstrated the considerable promise of a semantic segmentation approach to low energy electron reconstruction in LArTPCs, one that outperforms the current supernova neutrino energy resolution estimates based on ionisation energy deposits in the region of interest for supernova neutrinos. This would be an interesting approach to consider when studying the capabilities of DUNE for supernova neutrino reconstruction, and may lead to improved energy resolution in the associated analyses when compared to the benchmark set in the technical design report.

8 Conclusion

The Deep Underground Neutrino Experiment (DUNE) is a next generation neutrino oscillation experiment that aims to make precision measurements of neutrino oscillation probabilities, in order to search for CP–violation in the lepton sector. In addition, DUNE provides an opportunity to study the mechanisms involved in supernova bursts through the measurement of supernova neutrinos. To achieve these goals, the charged particles in the DUNE detectors will have to be accurately reconstructed, in order to allow the details of the neutrino interactions to be determined. This thesis presents the results of the development of two reconstruction algorithms in the ProtoDUNE–SP detector, which is a surface level proto–type for the single phase DUNE Far Detector modules.

The data used in this thesis was collected using the ProtoDUNE–SP detector during a beam run in late 2018. The online monitoring system used during this time was developed and maintained by the author, who was also the on call expert for the system. Many components of the OM system, particularly the data analysis processes will be adopted and developed for the online data quality monitoring of the DUNE Far Detector.

In this thesis, the development of a hit tagging algorithm for track, shower, and Michel electron hit classification was discussed in Chapter 6. This algorithm is based on convolutional neural networks, with a small patch approach. In terms of

track–shower discrimination, the algorithm was shown to reduce the false–positive rate for all particle species, with a reduction of more than a factor of 10 in some cases, when compared to the previous algorithm. In addition, this algorithm was shown to give a good agreement, of around 1–3%, between data and simulation for hit–by–hit event selection. This algorithm is now seeing significant use within ongoing ProtoDUNE–SP data analyses, including pion cross–section analyses and detector calibrations. Additionally, this network was used as part of the Michel electron event selection algorithm, which was discussed in Chapter 7 of this thesis.

A sample of Michel electrons from the ProtoDUNE–SP detector was used to study Michel electron reconstruction in LArTPCs, and to estimate the energy resolution and bias of ProtoDUNE–SP for electrons in the tens of MeV range. Michel electron events were selected based on a new algorithm, which combined the clustering of Pandora with the output of the hit classification CNN to select clusters with a high fraction of Michel–like hits. This algorithm was shown to agree well between data and simulation, with around a 2% difference between the selected fraction in data and simulation. The simulation was used to estimate the purity and efficiency of the event selection algorithm, which were found to be around 98% and 6% respectively. This represents an improvement in both purity and efficiency over similar LArTPC Michel electron studies from other experiments, such as MicroBooNE.

The Michel electron sample was used to investigate a novel reconstruction technique, which was based on the semantic segmentation of images with a convolutional neural network. A U–ResNet architecture was used to perform the segmentation, and the results were promising, however, more work is required to improve the performance of this algorithm across the whole energy range.

In the 0–30 MeV region, where the performance of the U–ResNet algorithm was good, the ionisation energy resolution for Michel electrons based on this algorithm was found to be around 10–12%, and the bias was less than 5%. This represents an improvement over the measured Michel electron ionisation energy resolution of MicroBooNE, which is the closest comparison to ProtoDUNE–SP in terms of Michel electron reconstruction. Above 30 MeV, the reconstruction algorithm

began to develop a significant degradation in energy resolution. This represents the tendency of the algorithm to under–reconstruct the Michel electron ionisation energy, and is a result of a drop in hit selection efficiency above 30 MeV. Some additional work is required to improve the performance of this algorithm, but the excellent performance in the low energy region shows that this is a promising approach. A similar approach could be developed, in the future, for other low–energy reconstruction tasks in LArTPCs, such as the reconstruction of supernova neutrinos in the DUNE far detector.

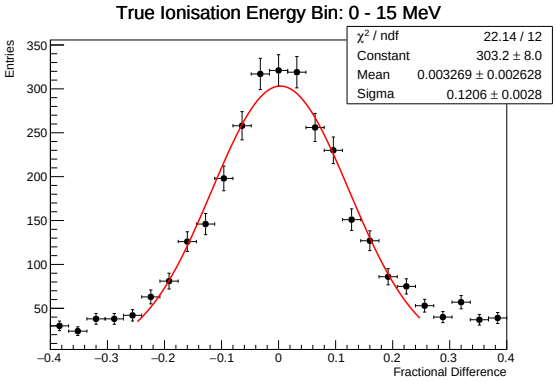
The DUNE experiment will use a liquid argon time projection chamber to make precision studies of neutrino oscillations, supernova neutrinos, and beyond the standard model phenomena. The development of effective reconstruction algorithms and understanding the detector response are both crucial components of understanding particle interactions in the DUNE detector. This thesis has presented the results of various studies into electromagnetic interactions in the ProtoDUNE–SP detector, focussing on the application of convolutional neural networks to event reconstruction in LArTPCs. The results of these analyses are currently being used to assist in physics measurements based on ProtoDUNE–SP data, and can also be used to guide further development of reconstruction strategies for LArTPC detectors.

Appendices

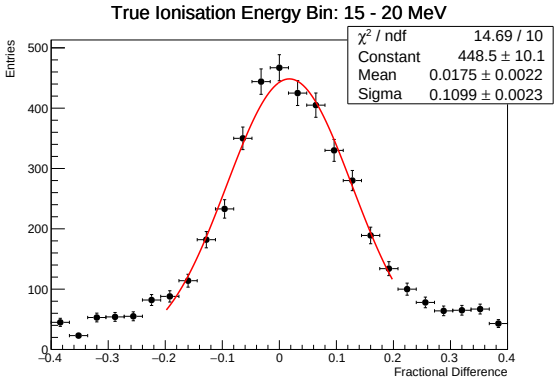
A

Additional Details for Michel Electron Reconstruction

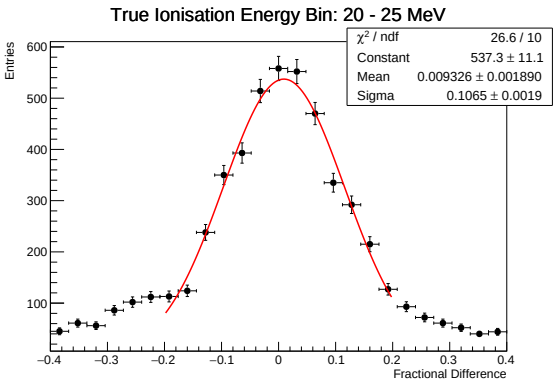
Chapter 7 provided a study of the fractional energy resolution of the developed Michel electron reconstruction algorithm, in terms of ionisation only deposited energy. The bias and resolution estimates were based on gaussian fits to the fractional difference between the reconstructed ionisation energy and the true ionisation energy. The fits used to populate Figure 7.20 are shown in Figure A.1. The data is binned in terms of the total true ionisation energy deposited by the Michel electron, and each plot contains the data from a single bin. No attempt was made to fit the tails of the distribution, which is particularly noteworthy in Figure A.1e, where there is a significant tail for negative fractional differences.



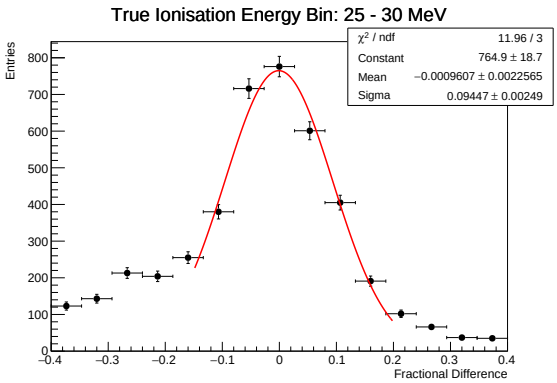
(a) 0 - 15 MeV



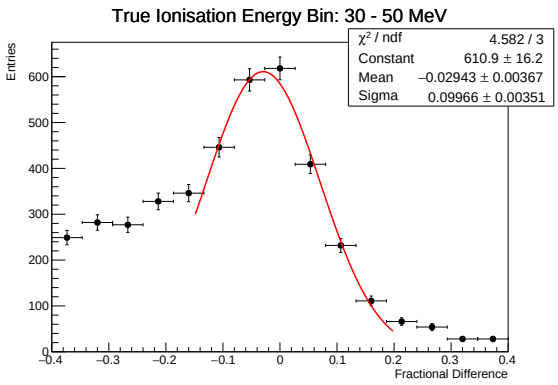
(b) 15 - 20 MeV



(c) 20 - 25 MeV



(d) 25 - 30 MeV



(e) 30 - 50 MeV

Figure A.1: Gaussian fits to the fractional energy difference between reconstructed ionisation energy and true ionisation energy, used to estimate the energy resolution and bias as a function of energy in Figure 7.20. No Attempt was made to fit the tails of the distribution.

References

[1] B. P. Abbott et al. "GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral". In: *Physical Review Letters* 119 (2017), p. 161101. arXiv: 1710.05832 [gr-qc].

[2] M. G. Aartsen et al. "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A". In: *Science* 361.6398 (2018), eaat1378. arXiv: 1807.08816 [astro-ph.HE].

[3] Kate Scholberg. "Supernova Neutrino Detection". In: *Annual Review of Nuclear and Particle Science* 62 (2012), pp. 81–103. arXiv: 1205.6003 [astro-ph.IM].

[4] R. Acciarri et al. "Design and Construction of the MicroBooNE Detector". In: *Journal of Instrumentation* 12.02 (2017), P02017. arXiv: 1612.05824 [physics.ins-det].

[5] F. Cavanna et al. *LArIAT: Liquid Argon In A Testbeam*. 2014. arXiv: 1406.5560 [physics.ins-det].

[6] M. Antonello et al. *A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam*. 2015. arXiv: 1503.01520 [physics.ins-det].

[7] B. Abi et al. *The Single-Phase ProtoDUNE Technical Design Report*. 2017. arXiv: 1706.07081 [physics.ins-det].

[8] James Chadwick. "Intensitätsverteilung im magnetischen Spectrum der β-Strahlen von radium B+ C". In: *Verhandlungen der Deutschen Physikalischen Gesellschaft* 16 (1914), p. 383.

[9] Wolfgang Pauli. "Pauli letter collection: letter to Lise Meitner". Typed copy. URL: https://cds.cern.ch/record/83282.

[10] E. Fermi. "Versuch einer Theorie der β-Strahlen. I". In: *Zeitschrift für Physik* 88.3-4 (1934), pp. 161–177.

[11] H. Bethe and R. Peierls. "The Neutrino". In: *Nature* 133.532 (1934).

[12] F. Reines and C. L. Cowan. "Detection of the free neutrino". In: *Physical Review* 92.3 (1953), pp. 830–831.

[13] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Physical Review Letters* 9 (1962), pp. 36–44.

[14] F. J. Hasert et al. "Observation of Neutrino Like Interactions without Muon or Electron in the Gargamelle Neutrino Experiment". In: *Physics Letters B* 46.1 (1973), pp. 138–140.

[15] K. Kodama et al. "Observation of tau neutrino interactions". In: *Physics Letters B* 504 (2001), pp. 218–224. arXiv: hep-ex/0012035.

[16] G. Alexander et al. “Electroweak parameters of the Z^0 resonance and the Standard Model: the LEP Collaborations”. In: *Physics Letters B* 276 (1992), pp. 247–253.

[17] R. Davis, D. S. Harmer, and K. C. Hoffman. “Search for neutrinos from the sun”. In: *Physical Review Letters* 20.21 (1968), pp. 1205–1209.

[18] J. N. Bahcall, N. A. Bahcall, and G. Shaviv. “Present status of the theoretical predictions for the Cl36 solar-neutrino experiment”. In: *Physical Review Letters* 20.21 (1968), pp. 1209–1212.

[19] K. S. Hirata et al. “Experimental study of the atmospheric neutrino flux”. In: *Physics Letters B* 205.2-3 (1988), pp. 416–420.

[20] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Physical Review Letters* 81.8 (1998), pp. 1562–1567. arXiv: hep-ex/9807003.

[21] Q. R. Ahmad et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Physical Review Letters* 89.1 (2002), pp. 6–11. arXiv: 0204008 [nucl-ex].

[22] A. Yu. Smirnov. *Solar neutrinos: Oscillations or No-oscillations?* 2016. arXiv: 1609.02386 [hep-ph].

[23] K. Eguchi et al. “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance”. In: *Physical Review Letters* 90.2 (2003), p. 6. arXiv: 0212021 [hep-ex].

[24] T. Araki et al. “Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion”. In: *Physical Review Letters* 94.8 (2005), pp. 1–5. arXiv: 0406035 [hep-ex].

[25] N. Agafonova et al. “Observation of a first ν_τ candidate event in the OPERA experiment in the CNGS beam”. In: *Physics Letters B* 691.3 (2010), pp. 138–145. arXiv: 1006.1623 [hep-ex].

[26] N. Agafonova et al. “Final Results of the OPERA Experiment on ν_τ Appearance in the CNGS Neutrino Beam”. In: *Physical Review Letters* 120.21 (2018), p. 211801. arXiv: 1804.04912 [hep-ex].

[27] M. Aker et al. “First operation of the KATRIN experiment with tritium”. In: *European Physical Journal C* 80.3 (2020), p. 264. arXiv: 1909.06069 [physics.ins-det].

[28] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.

[29] M. Tanabashi et al. “Review of Particle Physics”. In: *Physical Review D* 98 (3 2018), p. 030001.

[30] S. King. “Neutrino Mass and Mixing in the Seesaw Playground”. In: *Nuclear Physics B* 908 (2015). arXiv: 1511.03831 [hep-ph].

[31] C. Arnaboldi et al. “CUORE: A cryogenic underground observatory for rare events”. In: *Nuclear Instruments and Methods in Physics Research A*. 518.3 (2004), pp. 775–798. arXiv: 0212053 [hep-ex].

[32] V. Álvarez et al. "NEXT-100 technical design report (TDR). Executive summary". In: *Journal of Instrumentation* 7.6 (2012). arXiv: 1202.0721 [physics.ins-det].

[33] S. Andringa et al. "Current Status and Future Prospects of the SNO+ Experiment". In: *Advances in High Energy Physics* 2016 (2016). arXiv: 1508.05759 [physics.ins-det].

[34] I. Esteban et al. "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23}, δ_{CP}, and the mass ordering". In: *Journal of High Energy Physics* 01 (2019), p. 106. arXiv: 1811.05487 [hep-ph].

[35] S. Fukuda et al. "Solar ^8B and hep Neutrino Measurements from 1258 Days of Super-Kamiokande Data". In: *Physical Review Letters* 86 (25 2001), pp. 5651–5655. arXiv: hep-ex/0103032.

[36] K. Abe et al. "Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector". In: *Physical Review D* 96 (9 2017), p. 092006. arXiv: 1707.01048 [hep-ex].

[37] P. Adamson et al. "Combined Analysis of ν_μ Disappearance and $\nu_\mu \to \nu_e$ Appearance in MINOS Using Accelerator and Atmospheric Neutrinos". In: *Physical Review Letters* 112 (19 2014), p. 191801. arXiv: 1403.0867 [hep-ex].

[38] M. A. Acero et al. "First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA". In: *Physical Review Letters* 123 (15 2019), p. 151803. arXiv: 1906.04907 [hep-ex].

[39] F. P. An et al. "Observation of electron-antineutrino disappearance at Daya Bay". In: *Physical Review Letters* 108 (2012), p. 171803. arXiv: 1203.1669 [hep-ex].

[40] Y. Abe et al. "First Measurement of θ_{13} from Delayed Neutron Capture on Hydrogen in the Double Chooz Experiment". In: *Physics Letters B* 723 (2013), pp. 66–70. arXiv: 1301.2948 [hep-ex].

[41] J. K. Ahn et al. "Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment". In: *Physical Review Letters* 108 (2012), p. 191802. arXiv: 1204.0626 [hep-ex].

[42] K. Abe et al. "Constraint on the Matter-Antimatter Symmetry-Violating Phase in Neutrino Oscillations". In: *Nature* 580 (2020), pp. 339–344. arXiv: 1910.03887 [hep-ex].

[43] J. A. Formaggio and G. P. Zeller. "From eV to EeV: Neutrino Cross Sections Across Energy Scales". In: *Reviews of Modern Physics* 84 (2012), pp. 1307–1341. arXiv: 1305.7513 [hep-ex].

[44] I. Tamborra et al. "Self-sustained asymmetry of lepton-number emission: A new phenomenon during the supernova shock-accretion phase in three dimensions". In: *The Astrophysical Journal* 792.2 (2014), p. 96. arXiv: 1402.5418 [astro-ph.SR].

[45] K. Hirata et al. "Observation of a Neutrino Burst from the Supernova SN 1987a". In: *Physical Review Letters* 58 (1987), pp. 1490–1493.

[46] R. M. Bionta et al. "Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud". In: *Physical Review Letters* 58 (14 1987), pp. 1494–1496.

[47] T. J. Loredo and D. Q. Lamb. "Bayesian analysis of neutrinos observed from supernova SN-1987A". In: *Physical Review D* 65 (2002), p. 063002. arXiv: astro-ph/0107260.

[48] A. Mirizzi et al. "Supernova Neutrinos: Production, Oscillations and Detection". In: *Rivista de Nuovo Cimento* 39.1-2 (2016), pp. 1–112. arXiv: 1508.00785 [astro-ph.HE].

[49] K. Abe et al. "Real-Time Supernova Neutrino Burst Monitor at Super-Kamiokande". In: *Astroparticle Physics* 81 (2016), pp. 39–48. arXiv: 1601.04778 [astro-ph.HE].

[50] B. Abi et al. *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II DUNE Physics*. 2020. arXiv: 2002.03005 [hep-ex].

[51] Roland Diehl et al. "Radioactive Al-26 and massive stars in the galaxy". In: *Nature* 439 (2006), pp. 45–47. arXiv: astro-ph/0601015.

[52] R. Acciarri et al. *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF*. 2015. arXiv: 1512.06148 [physics.ins-det].

[53] B. Abi et al. *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume I Introduction to DUNE*. 2020. arXiv: 2002.02967 [physics.ins-det].

[54] B. Abi et al. *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV Far Detector Single-phase Technology*. 2020. arXiv: 2002.03010 [physics.ins-det].

[55] M. Auger, A. Ereditato, and J. R. Sinclair. *ArgonCube: a Modular Approach for Liquid Argon TPC Neutrino Detectors for Near Detector Environments*. URL: http://cds.cern.ch/record/2268439/.

[56] A. A. Machado and E. Segreto. "ARAPUCA a new device for liquid argon scintillation light detection". In: *Journal of Instrumentation* 11.02 (2016), p. C02004.

[57] E. Segreto et al. "Liquid Argon test of the ARAPUCA device". In: *Journal of Instrumentation* 13.08 (2018), P08021. arXiv: 1805.00382 [physics.ins-det].

[58] B. Abi et al. "First results on ProtoDUNE-SP LArTPC performance from a test beam run at the CERN Neutrino Platform". In: *Manuscript in preparation for submission to Journal of Instrumentation* (2020). arXiv: 2007.06722 [physics.ins-det].

[59] M. Spanu. "The status and results from ProtoDUNE Single Phase". In: *Journal of Physics: Conference Series* 1312 (2019).

[60] R. Herbst et al. "Design of the SLAC RCE Platform: A general purpose ATCA based data acquisition system". In: *Nuclear Science Symposium and Medical Imaging Conference*. 2014, pp. 1–4.

[61] J. Anderson et al. "FELIX: a PCIe based high-throughput approach for interfacing front-end and trigger electronics in the ATLAS Upgrade framework". In: *Journal of Instrumentation* 11.12 (2016), pp. C12023–C12023.

[62] K. Biery et al. “Artdaq: An Event-Building, Filtering, and Processing Framework”. In: *Transactions of Nuclear Science* 60.5 (2013), pp. 3764–3771.

[63] E. L. Snider and G. Petrillo. “LArSoft: Toolkit for simulation, reconstruction and analysis of liquid argon TPC neutrino detectors”. In: *Journal of Physics: Conference Series* 898.4 (2017).

[64] C. Green et al. “The Art Framework”. In: *Journal of Physics: Conference Series* 396 (2012), p. 022020.

[65] A. C. Booth et al. “Particle production, transport, and identification in the regime of 1−7 GeV/c”. In: *Physical Review Accelerators and Beams* 22.6 (2019), p. 061003.

[66] T.J. Roberts et al. “G4Beamline Particle Tracking in Matter-dominated Beam Lines”. In: *Particle Accelerator Conference* 0806233 (2008), WEPP120.

[67] T. T. Böhlen et al. “The FLUKA Code: Developments and Challenges for High Energy and Medical Applications”. In: *Nuclear Data Sheets* 120 (2014), pp. 211–214.

[68] D. Heck et al. *CORSIKA: A Monte Carlo code to simulate extensive air showers.* 1998.

[69] S. Agostinelli et al. “GEANT4: A Simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research A* 506 (2003), pp. 250–303.

[70] C. Adams et al. “Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation”. In: *Journal of Instrumentation* 13.07 (2018), P07006. arXiv: 1802.08709 [physics.ins-det].

[71] J. S. Marshall and M. A. Thomson. “The Pandora software development kit for pattern recognition”. In: *European Physical Journal C* 75.9 (2015), pp. 1–16. arXiv: 1506.05348 [physics.data-an].

[72] R. Acciarri et al. “The Pandora multi–algorithm approach to automated pattern recognition of cosmic-ray muon and neutrino events in the MicroBooNE detector”. In: *European Physical Journal C* 78.1 (2018), p. 82. arXiv: 1708.03135 [hep-ex].

[73] R. Acciarri et al. “A study of electron recombination using highly ionizing particles in the ArgoNeuT Liquid Argon TPC”. In: *Journal of Instrumentation* 8.8 (2013). arXiv: 1306.1712 [physics.ins-det].

[74] I. Antcheva et al. “ROOT – A C++ framework for petabyte data storage, statistical analysis and visualization”. In: *Computer Physics Communications* 180.12 (2009), pp. 2499–2512. arXiv: 1508.07749 [physics.data-an].

[75] M. Adinolfi et al. “LHCb data quality monitoring”. In: *Journal of Physics: Conference Series* 898 (2017), p. 092027.

[76] M. Grinberg. *Flask Web Development: Developing Web Applications with Python.* O’Reilly Media, Inc., 2014. URL: https://flask.palletsprojects.com.

[77] Bokeh Development Team. *Bokeh: Python library for interactive visualization.* 2018. URL: https://bokeh.pydata.org/en/latest/.

[78] T. Oliphant. *NumPy: A guide to NumPy.* Trelgol Publishing. 2006. URL: http://www.numpy.org/.

[79] *Atomic and Nuclear Properties of Matierials*. URL: http://pdg.lbl.gov/2020/AtomicNuclearProperties/.

[80] C. Møller. "On the Radiative Collision between Fast Charged Particles". In: *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 152.876 (1935), pp. 481–496.

[81] H. J. Bhabha. "The Scattering of Positrons by Electrons with Exchange on Dirac's Theory of the Positron". In: *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 154.881 (1936), pp. 195–206.

[82] Y. Tsai. "Pair Production and Bremsstrahlung of Charged Leptons". In: *Reviews of Modern Physics* 46 (1974), p. 815.

[83] B. Rossi. *High Energy Particles*. 1952.

[84] *ESTAR: Stopping-power and range tables for electrons*. URL: https://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html.

[85] *XCOM: Photon Cross Sections Database*. URL: https://www.nist.gov/pml/xcom-photon-cross-sections-database.

[86] R. Acciarri et al. "First Observation of Low Energy Electron Neutrinos in a Liquid Argon Time Projection Chamber". In: *Physical Review D* 95.7 (2017), p. 072005. arXiv: 1610.04102 [hep-ex].

[87] W.H. Lippincott et al. "Scintillation time dependence and pulse shape discrimination in liquid argon". In: *Physical Review C* 78 (2008), p. 035801. arXiv: 0801.1531 [nucl-ex].

[88] R. Reed and R. J. Marks. *Neural smithing: supervised learning in feedforward artificial neural networks*. MIT Press, 1999.

[89] Y. Lecun, Y. Bengio, and G. Hinton. "Deep learning". In: *Nature* 521.7553 (2015), pp. 436–444.

[90] K. He et al. "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification". In: *Proceedings of the International Conference on Computer Vision and Pattern Recognition* (2015), pp. 1026–1034. arXiv: 1502.01852 [cs.CV].

[91] C. Szegedy et al. "Going deeper with convolutions". In: *Proceedings of the International Conference on Computer Vision and Pattern Recognition*. 2015. arXiv: 1409.4842 [cs.CV].

[92] G. Cybenko. "Approximation by superpositions of a sigmoidal function". In: *Mathematics of Control, Signals and Systems* 2 (1989), pp. 303–314.

[93] L. D. Jackel et al. "Backpropagation Applied to Handwritten Zip Code Recognition". In: *Neural Computation* 1.4 (1989), pp. 541–551.

[94] Y. LeCun, K. Kavukcuoglu, and C. Farabet. "Convolutional networks and applications in vision". In: *Proceedings of the International Symposium on Circuits and Systems*. 2010, pp. 253–256.

[95] N. Kanopoulos, N. Vasanthavada, and R. L. Baker. "Design of an image edge detection filter using the Sobel operator". In: *Journal of Solid-state Circuits* 23.2 (1988), pp. 358–367.

[96] D. Esterhuizen. *C# How to: Image Edge Detection*. URL: https://softwarebydefault.com/2013/05/11/image-edge-detection.

[97] A. Le Nail. *NN–SVG: Publication ready NN–architecture schematics*. URL: http://alexlenail.me/NN-SVG.

[98] K. He et al. "Deep Residual Learning for Image Recognition". In: *Proceedings of the International Conference on Computer Vision and Pattern Recognition*. 2016. arXiv: 1512.03385 [cs.CV].

[99] D. E Rumelhart, G. E. Hinton, and R. J. Williams. "Learning representations by back-propagating errors". In: *Nature* 323.6088 (1986), pp. 533–536.

[100] M. Li et al. "Efficient Mini-Batch Training for Stochastic Optimization". In: *Proceedings of the International Conference on Knowledge Discovery and Data Mining*. 2014, pp. 661–670.

[101] D. P. Kingma and L. J. Ba. "Adam: A Method for Stochastic Optimization". In: *Proceedings of the International Conference on Learning Representations*. 2015. arXiv: 1412.6980 [cs.LG].

[102] S. Ruder. *An overview of gradient descent optimization algorithms*. 2016. arXiv: 1609.04747 [cs.LG].

[103] A. L. Maas, A. Y. Hannun, and A. Y. Ng. "Rectifier nonlinearities improve neural network acoustic models". In: *Proceedings of the International Conference on Machine Learning*. 2013.

[104] V. Nair and G. E. Hinton. "Rectified Linear Units Improve Restricted Boltzmann Machines". In: *Proceedings of the International Conference on Machine Learning*. 2010, pp. 807–814.

[105] N. Srivastava et al. "Dropout: a simple way to prevent neural networks from overfitting". In: *Journal of Machine Learning Research* 15 (2014), pp. 1929–1958.

[106] L. Prechelt. "Early Stopping – But When?" In: *Neural Networks: Tricks of the Trade*. Springer, 1998, pp. 55–69.

[107] G. Hinton, Y. Bengio, and Y. Lecun. "Deep Learning". In: *Neural Information Processing Systems Tutorial* (2015).

[108] L. Roeder. *Viewer for neural network models*. URL: https://www.electronjs.org/apps/netron.

[109] P. de Boer et al. "A Tutorial on the Cross-Entropy Method". In: *Annals of Operations Research* 134.1 (2005), pp. 19–67.

[110] C. Sammut and G. I. Webb. "Mean Squared Error". In: *Encyclopedia of Machine Learning*. Springer, 2010, pp. 653–653.

[111] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. Adaptive computation and machine learning. MIT Press, 2016. URL: http://www.deeplearningbook.org.

[112] Abadi. M. et al. "TensorFlow: A system for large-scale machine learning". In: *Proceedings of the Symposium on Operating Systems Design and Implementation*. 2016, pp. 265–283.

[113] F. Chollet et al. *Keras*. https://keras.io.

[114] *TensorBoard: TensorFlow's visualization toolkit.* https://www.tensorflow.org/tensorboard.

[115] C. J. Van Rijsbergen. *Information retrieval.* Butterworths, 1975.

[116] T. Fawcett. "An Introduction to ROC Analysis". In: *Pattern Recognition Letters* 27.8 (2006), pp. 861–874.

[117] A. Chappell. *A deep neural network to direct the Pandora multi–algorithm LArTPC event reconstruction.* Poster presented at Neutrino 2020. URL: https://indico.fnal.gov/event/19348/contributions/186187.

[118] M. Vermeulen. *Neutral pion event selection in ProtoDUNE–SP.* Unpublished, analysis ongoing.

[119] Ian J. Goodfellow et al. "Generative Adversarial Nets". In: *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2.* NIPS'14. MIT Press, 2014, pp. 2672–2680.

[120] J. Calcutt, F. Stocker, and L. Jiang. *Pion charge–exchange analysis in ProtoDUNE–SP.* Unpublished, analysis ongoing.

[121] F. Spagliardi. *Contaminated track removal for stopping muon calibration.* Unpublished, analysis ongoing.

[122] O. Ronneberger, P. Fischer, and T. Brox. "U–net: Convolutional networks for biomedical image segmentation". In: *Proceedings of the International Conference on Medical Image Computing and Computer-assisted Intervention.* Springer. 2015, pp. 234–241.

[123] R. Acciarri et al. "Michel Electron Reconstruction Using Cosmic-Ray Data from the MicroBooNE LArTPC". In: *Journal of Instrumentation* 12.09 (2017), P09014. arXiv: 1704.02927 [physics.ins-det].

[124] S. Amoruso et al. "Measurement of the μ decay spectrum with the ICARUS liquid argon TPC". In: *European Physical Journal C* 33 (2004), pp. 233–241. arXiv: hep-ex/0311040.

[125] W. Foreman. "Light-based triggering and reconstruction of Michel electrons in LArIAT". In: *Journal of Instrumentation* 11.01 (2016), pp. C01037–C01037.