# A Data-Driven Extrapolation Technique for the DUNE-PRISM Oscillation Analysis

# DOCTORAL THESIS

University College London

Department of Physics and Astronomy

Author:

Supervisor:

Alex Wilkinson

Prof. Ryan Nichol

Submitted to UCL in partial fulfilment of the requirements for the award of the degree of Doctor of Philosophy.

# Declaration

I, Alex Wilkinson, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

## Abstract

The Deep Underground Neutrino Experiment (DUNE) is a next-generation long-baseline neutrino oscillation experiment that will make precision measurements of the parameters governing neutrino oscillations. A core objective of DUNE's design is to minimise the impact of neutrino-nucleus interaction uncertainties on these precision measurements. The capability of DUNE's near detector to move transverse to the neutrino beam to sample many distinct off-axis fluxes, known as the Precision Reaction Independent Spectrum Measurement (PRISM), is key to fulfilling this objective. PRISM is used to construct a unique oscillation analysis where near detector data is extrapolated directly to the far detector via linear combination of measurements at different off-axis positions. If constructed carefully, such an oscillation analysis incorporates the near detector constraint with very little reliance on potentially fraught neutrino-nucleus interaction models.

This thesis presents a novel method for extrapolating detector effects from the near to the far detector for the PRISM oscillation analysis. Two distinct implementations are developed that use machine learning to map event-level information between the detectors to minimise reliance on simulation of the neutrino-nucleus interaction. Pairs of identical neutrino interactions at the two detectors are generated to permit supervised training. One implementation uses a generative model to map between the reconstructed quantities. The other uses computer vision methods to perform a style-transfer on the detector response. Both implementations are demonstrated to enhance the PRISM oscillation analysis by reducing the impact of systematic uncertainties associated with the interaction model.

# Impact Statement

The Deep Underground Neutrino Experiment is poised to significantly advance our understanding of nature's most abundant matter particle via precision measurements of neutrino oscillations. This thesis contributes to this programme of precision measurements by introducing a novel technique that reduces the impact of a key class of systematic uncertainties. The technique is developed and demonstrated to be effective, establishing a path towards its use in oscillation analyses once detector data becomes available.

Within the wider field of neutrino physics, the specific application of deep learning used in this thesis is of benefit. Both the use of a generative model to overcome ill-defined mappings between detectors and the application of style-transfer to detector readout are novel and contribute to advancing research methodologies in the field.

Outside of academia, measurements of neutrino oscillations enrich society by advancing our fundamental understanding of nature and offering a promising solution to the mystery of matter-antimatter asymmetry in the Universe. While there are no immediate practical applications for neutrinos, future technologies may one day be influenced, directly or indirectly, by their properties. Making precise measurements in the present ensures that future breakthroughs are not severely delayed by gaps in our foundational knowledge.

# Acknowledgements

Firstly, I would like to thank my supervisor Prof. Ryan Nichol for providing valuable guidance, exercising patience, and being rather good humoured. I am sincerely grateful for his support during one particularly challenging time.

I would like to thank the fellow DUNE-PRISM people I had the pleasure of working with for their enthusiasm and curiosity: Mike, Wei, Ioana, Radi, Ciaran, Cris, and Colin. This thesis would have been *considerably* harder to write without their work and support.

I am thankful for all the collaborators that became friends over the years. To name a few, Francisco has been an excellent companion for terrible coffee and great beer. Gajendra has been responsible for much laughter and many happy memories. Radi has been an invaluable source of both good and questionable ideas and introduced me to the ever true credo "We do these things not because they are easy, but because we thought they would be easy".

During my PhD, I was fortunate enough to spend an enjoyable eight months at Fermilab. Thank you to Sungbin, Tingjun, Marco, Bill, and Nicola for making me feel very welcome as I joined the SBND collaboration shortly after arriving. I met many people at Fermilab across different experiments and disciplines that I consider to be friends and would therefore like to acknowledge the warmth of the Fermilab community.

The UCL HEP group and CDT has been an engaging and friendly place to work. I would like to thank Andrius, Emily, Matthijs, Thandi, Jackson, Margot, and Mustafa for their friendship and formalisation of lunchtime procrastination.

I am deeply grateful for my Mum, Dad, Grandma, and Sister whose encouragement, interest, and support throughout all my years has put me here today. Lastly, a special thank you to Jade who shared the burdens and joys of the last four years and beyond. This thesis is yours also.

# Contents

D	eclara	ation		1
$\mathbf{A}$	bstra	$\operatorname{\mathbf{ct}}$		2
In	ıpact	t Statement         3           wledgements         4           roduction         9           atrino Physics         11           The Neutrino Chronology         11           Neutrinos in the Standard Model         15           2.2.1 The Standard Model         15           2.2.2 Electroweak Symmetry Breaking         17           2.2.3 Mechanisms for Neutrino Mass         20           2.2.4 Lepton Mixing         21           Neutrino Oscillations         22           2.3.1 Three Flavour Oscillations         22           2.3.2 Matter Effects         23           2.3.3 Physics From Oscillation Measurements         26           Neutrino Interactions         29           Neutrino Oscillation Measurements         32           2.5.1 Solar and Long Baseline Reactor Neutrinos         32           2.5.2 Short Baseline Reactor Neutrinos         33           2.5.3 Atmospheric Neutrinos         35           2.5.5 Current Oscillation Parameter Constraints         38           Deep Underground Neutrino Experiment         39           Overview         39           Physics Goals         40           3.2.1 Standard Model Neutrino Oscillations         40		
A	cknov	wledge	ments	4
1	Intr	oducti	on	9
2	Neu	itrino l	Physics	11
	2.1	The N	eutrino Chronology	11
	2.2	Neutri	nos in the Standard Model	15
		2.2.1	The Standard Model	15
		2.2.2	Electroweak Symmetry Breaking	17
		2.2.3	Mechanisms for Neutrino Mass	20
		2.2.4	Lepton Mixing	21
	2.3	Neutri	no Oscillations	22
		2.3.1	Three Flavour Oscillations	22
		2.3.2	Matter Effects	23
		2.3.3	Physics From Oscillation Measurements	26
	2.4	Neutri	no Interactions	29
	2.5	Neutri	no Oscillation Measurements	32
		2.5.1	Solar and Long Baseline Reactor Neutrinos	32
		2.5.2	Short Baseline Reactor Neutrinos	33
		2.5.3	Atmospheric Neutrinos	33
		2.5.4	Accelerator Neutrinos	35
		2.5.5	Current Oscillation Parameter Constraints	38
3	The	Deep	Underground Neutrino Experiment	39
	3.1	Overvi	iew	39
	3.2	Physic	es Goals	40
		3.2.1	Standard Model Neutrino Oscillations	40
		3.2.2	Supernova Neutrino Bursts	41
	3.3	Liquid	Argon Time Projection Chamber Technology	43
		3.3.1	Detection Principle	43

		3.3.2	Particle Propagation in Liquid Argon	:5
	3.4	The L	BNF Neutrino Beam	:7
	3.5	The N	ear Detector	8
		3.5.1	ND-LAr	.9
		3.5.2	TMS and ND-GAr	0
		3.5.3	PRISM	1
		3.5.4	SAND	3
		3.5.5	Role in Oscillation Measurements	3
	3.6	The Fa	ar Detector	5
		3.6.1	Horizontal Drift	6
		3.6.2	Vertical Drift	7
4	DU	NE Sin	nulation and Reconstruction 6	0
	4.1	Event	Generation in LAr	0
	4.2	The Fa	ar Detector	1
		4.2.1	Detector Response Simulation	1
		4.2.2	Signal Processing and Hit Finding 6	2
		4.2.3	Track and Shower Reconstruction	4
		4.2.4	Neutrino Signal Selection	5
		4.2.5	Neutrino Energy Estimation	6
	4.3	The N	ear Detector	7
		4.3.1	Detector Response Simulation	9
		4.3.2	Parametrised Reconstruction	′1
5	PRI	ISM O	scillation Analysis 7	3
	5.1	Overvi	iew	3
	5.2	Monte	Carlo Data	'4
		5.2.1	Near Detector	4
		5.2.2	Far Detector	7
	5.3	Extrap	polation of Off-Axis Measurements	80
		5.3.1	Near Detector Backgrounds	0
		5.3.2	Near Detector Effects	1
		5.3.3	$\nu_e/\nu_\mu$ Cross Section Ratio	3
		5.3.4	Far Detector Effects	34
	5.4	Genera	ating a PRISM Prediction	86
		5.4.1	Linear Combination	86
		5.4.2	Far Detector Backgrounds	9
		5.4.3	Flux-Matching Correction	0
		5.4.4	Predictions	0
	5.5	Sensiti	ivities	0

		5.5.1	Four-Channel Oscillation Analysis
		5.5.2	Missing Proton Energy Fake Data
	5.6	Impro	vements to Extrapolation Method
		5.6.1	Geometric Efficiency Correction
		5.6.2	Data-Driven Near-to-Far Translation
6	A P	Paired 1	Near to Far Dataset 99
	6.1	Gener	ating Near-Far Pairs
		6.1.1	Shared Neutrino Interaction
		6.1.2	Detector Simulation and Reconstruction
		6.1.3	Near-Far Pairing
		6.1.4	Oscillation Analysis Compatible Pairs
		6.1.5	Steps to a Complete Four-Channel Sample
	6.2	Detect	tor Response Pairs
	6.3	Recon	struction Pairs
7	Rec	onstru	ction Translation 113
	7.1	Traini	ng Data
	7.2	Autor	egressive Transformer for Translation
		7.2.1	Architecture
		7.2.2	Training and Sampling
		7.2.3	Results
	7.3	Imple	mentation in PRISM Analysis
		7.3.1	Monte Carlo Data
		7.3.2	PRISM Prediction
		7.3.3	Sensitivities
8	Res	ponse	Translation 130
	8.1	Propo	sed Method
	8.2	ND-L	Ar Inactive Volume Infill
		8.2.1	Sparse Tensor Networks
		8.2.2	Data Preparation
		8.2.3	Architecture
		8.2.4	Training and Sampling
		8.2.5	Results
	8.3	Near-l	Far Image Translation
		8.3.1	Data Preparation
		8.3.2	Architecture
		8.3.3	Training
		8.3.4	Results

References						
9	Conclusions and Outlook					
		8.5.3	Effect of Cross Section Systematics	. 156		
		8.5.2	Extrapolating to the Far Detector	. 154		
		8.5.1	Monte Carlo Data	. 153		
	8.5	Cross	Section Systematics	. 152		
		8.4.2	Results	. 149		
		8.4.1	Generating Samples	. 149		
	8.4	Valida	ation with Reconstruction	. 149		

Introduction 1

The Standard Model (SM) of particle physics is the leading theory of the fundamental particles and their interactions via the electromagnetic, weak, and strong forces. It provides a mathematical framework for making predictions from a small number of free parameters and has proved remarkably consistent with experimental data at all accessible energy scales thus far. However, many observed phenomena cannot be explained by the SM, suggesting the necessity of extensions and reformulations to yield a deeper understanding of nature. This motivates a programme of precision measurements to probe the least well understood areas. One area that warrants particular scrutiny is neutrino oscillations due to compelling experimental evidence for the phenomenon of violating charge-parity (CP) symmetry.

The Deep Underground Neutrino Experiment (DUNE) is a next-generation long-baseline neutrino oscillation experiment that will make precision measurements of leptonic CP violation, the neutrino mass ordering, and the oscillation parameters  $\theta_{13}$ ,  $\theta_{23}$ ,  $\Delta m_{32}^2$ , and  $\delta_{\text{CP}}$ . To accomplish this, DUNE must minimise the impact of systematic uncertainties. Of particular importance are the systematic uncertainties associated with the neutrino-nucleus interaction model. Such a model is necessary to infer the incident neutrino energy and flavour from final state charged particles observed in a detector. The near detector (ND) of DUNE is an integral part of minimising these uncertainties. Placed close to the source of the neutrino beam prior to the effect of oscillations, the ND will strongly constrain the interaction model via measurements of neutrino-nucleus interactions. However, since no complete neutrino-nucleus interaction model exists, empirical corrections will need to be applied to force agreement between ND data and the model. Fitting the empirically corrected simulation to oscillated data at the far detector (FD) can significantly bias oscillation parameter measurements. Overcoming this challenge is key to DUNE's programme of precision oscillation parameter measurements.

The aspect of DUNE's design that aims to overcome this challenge is the Precision Reaction Independent Spectrum Measurement (PRISM). PRISM is the capability of DUNE's ND to move transverse to the neutrino beam to sample different neutrino fluxes. The set of neutrino fluxes accessible through this transverse movement can be linearly combined into a flux comparable to the oscillated flux expected at the FD. This facilitates a unique oscillation analysis methodology where ND data can be *extrapolated* to the FD for a data-

driven prediction that is compared directly to FD data. This largely removes the necessity of constraining any particular interaction model and so reduces the risk of bias oscillation parameter measurements.

An important challenge in the PRISM oscillation analysis is to correct for distinct detector effects so that a prediction composed of ND measurements may be compared with FD data. Performing this correction with minimal reliance on an interaction model is a key component in realising the full potential of the PRISM oscillation analysis methodology. This thesis introduces two novel techniques for making this correction. These techniques are shown to be robust to changes in the interaction model owing to their utilisation of the full phase space of individual ND data events rather than a single binned variable. They are implemented using modern deep learning methods trained on carefully manipulated simulation.

This thesis begins with an overview of neutrino physics in Chapter 2. The chapter covers the discovery of the neutrino and the first indications of its mass from a historical perspective, the formulation of a massive neutrino in the SM, the various phenomena relevant to neutrino oscillations, and the experimental status of neutrino oscillations. The aim of the chapter is to provide context and motivation for DUNE's oscillation physics programme.

Chapter 3 covers the design of DUNE's detectors and neutrino beam and is followed by a discussion of the simulation and reconstruction in Chapter 4. These chapters describe aspects of DUNE necessary for context and understanding of later chapters focusing on developments within the PRISM oscillation analysis.

In Chapter 5 the PRISM oscillation analysis is explained in detail. This includes the current method for the detector effects correction that is reliant on simulation of an interaction model. At the end of the chapter, potential improvements to the analysis are outlined. In this discussion the techniques that comprise the subject of this thesis are first introduced and motivated.

Chapter 6 describes the creation of a paired dataset of identical neutrino interactions at the ND and FD. Such a dataset is required for the deep learning approach that drives the proposed techniques of an event-level detector effects correction.

The two proposed detector effects correction techniques, one operating on reconstructed quantities and the other directly on detector response, are presented in Chapters 7 and 8 respectively. For each technique, there is a discussion of the methodology and of the results of a study intended to probe the technique's interaction model dependence.

Chapter 9 concludes this thesis by summarising the key results and identifying important avenues for future work.

Neutrino Physics

This chapter provides an overview of the development in understanding of the neutrino from its initial proposal through to observations of flavour oscillation. The formulation of the neutrino sector in the Standard Model is reviewed with particular emphasis on the phenomenology of flavour mixing. An overview of neutrino-nucleus interaction modelling is given. Lastly, relevant experimental results on the measurement of neutrino oscillations are highlighted.

# 2.1 The Neutrino Chronology

The experimental necessity of the neutrino originates in the early 20th century as experiments looked to probe matter beyond the atomic level. Armed with an understanding of the atom based on the Bohr model [1], Rutherford measured the scattering of alpha particles on air. The findings of this experiment, presented in 1920 [2], led to the discovery that the hydrogen nucleus is present in all other nuclei. To emphasis the importance of this seemingly fundamental constituent, the positively charged hydrogen nucleus was coined the *proton*. This discovery prompted the formulation of an atomic model where the nucleus consists of protons and a smaller number of *nuclear* electrons. The model explained  $\beta$ -decay as the emission of one of the nuclear electrons.

A key deficit in this model was the prediction of a monochromatic  $\beta$ -radiation spectrum, as is observed for  $\alpha$ -radiation. In 1927, Ellis and Wooster studied the energy of  $\beta$ -radiation from radium decay [3]. They found conclusively that  $\beta$ -radiation forms a continuous energy spectrum. In Rutherford's description of the atom this implied a violation of energy conservation. This revelation, along with the failure of the model to predict the correct spin for multiple nuclei, prompted Pauli's famous letter [4] in which the addition to the nucleus of an unobserved spin-1/2 neutral particle with very small mass is proposed. This new particle, named the *neutron*, would be emitted alongside the electron in a  $\beta$ -decay such that if its energy and the energy of the electron were measured the total energy spectrum would be monochromatic.

Experiments starting with Bothe and Becker in 1930 [5] discovered that a highly penetrating radiation made up of neutral particles can be generated through the absorption of alpha particles by some specific light nuclei. In 1932 Chadwick made a precise measure-

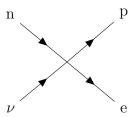


Figure 2.1: Four-point interaction of Fermi's theory of weak interaction

ment of the mass of these neutral particles and postulated the existence of the neutron that we know today [6]. Just a few months later, atomic models with a proton-neutron nucleus are proposed that resolve many outstanding experimental problems. However, the origin of the  $\beta$ -decay process remained unaccounted for, as Chadwick's neutron has a much larger mass than that proposed by Pauli. This was remedied in 1934 with Fermi's theory on  $\beta$ -decay [7] that incorporated Pauli's neutron (now known as the neutrino) in the decay process. The theory predicts a four-point interaction, meaning there is no mediator particle, originating from a term in the Hamiltonian that is a product of the yet unobserved neutrino field along with the other fields shown in Figure 2.1. The interaction has a matrix element

$$M = \frac{G_F}{\sqrt{2}} \left( \bar{u}_p \gamma^\mu u_n \right) \left( \bar{u}_e \gamma^\nu u_\nu \right), \tag{2.1.1}$$

where  $G_F$  is the Fermi Constant and  $\gamma^{\mu}$  are Dirac matrices. In analogy to electromagnetism, the bilinear terms are only of vector type. This theory of weak interaction correctly predicted the observed continuous  $\beta$ -decay spectrum. Although, it would later be found that having only a vector current is not supported by experiment. One notable discrepancy was the lack of a mechanism for parity violation, a necessity for any good theory of the weak interaction following Wu's measurement of maximal parity violation in the  $\beta$ -decay of <sup>60</sup>Co nuclei [8].

In the years following Fermi's theory, interest in making an indirect measurement of the neutrino grew. The detection came in the mid 1950s when Cowan and Reines observed the inverse  $\beta$  decay process,

$$\bar{\nu}_e + \mathbf{p} \to e^+ + \mathbf{n},\tag{2.1.2}$$

a prediction of Fermi's theory, using an intense flux of antineutrinos from the Savannah River nuclear reactor [9]. The experiment used a water-based detector doped with cadmium. The observed signal was a pair of prompt photons from the positron annihilation followed by a delayed light signal from the neutron capture on the cadmium.

Shortly after this discovery, a table-top experiment by Goldhaber, Grodzins, and Sunyar measured the neutrino to be exclusively left-handed [10]. This measurement fit well with the major theoretical development of the time, the refinement of Fermi's theory to a vector minus axial (V-A) theory of the weak interaction developed by both Feynman and Gell-Mann [11] and Marshak and Sudarshan [12]. The V-A theory takes its name

from the combination of vector and axial currents which make up interactions of the form

$$j^{\mu} \propto \bar{\psi} \gamma^{\mu} \frac{1}{2} \left( 1 - \gamma^5 \right) \phi, \tag{2.1.3}$$

where  $\psi$  and  $\phi$  are Dirac fields. It can be shown that the product of axial and vector currents, present in the cross-term of scattering amplitude calculations, violates parity, as demanded by the results of Wu's experiment. Using the properties of the projection operators, we can also show that the weak interaction only couples left-handed particles or right-handed antiparticles,

$$\bar{\psi}\gamma^{\mu}\frac{1}{2}\left(1-\gamma^{5}\right)\phi = \left(\bar{\psi}_{L}+\bar{\psi}_{R}\right)\gamma^{\mu}\phi_{L}$$

$$=\bar{\psi}_{L}\gamma^{\mu}\phi_{L}.$$
(2.1.4)

This means the neutrino, which only interacts via the weak force, can only ever be produced as a left-handed particle.

Thus far, there was experimental evidence for only the electron neutrino. This changed in 1962 when the first accelerator neutrino experiment at Brookhaven National Laboratory discovered the existence of the muon neutrino [13]. By impacting a 15 GeV proton beam on a beryllium target, a flux of kaons and pions were produced which provided a neutrino beam through the decays  $K/\pi \to \mu\nu$ . The interaction of the beam neutrinos on aluminium nuclei was measured using a spark chamber. The experiment observed a large excess of muon events over electron events. This provided strong evidence for the neutrino that couples to the muon being distinct from the neutrino that couples to the electron. Fast forward to 2001, the DONUT accelerator neutrino experiment at Fermilab measured the tau neutrino [14]. This experiment followed much of the same principles as the earlier experiment at Brookhaven but used neutrinos from the decay of the charm-strange D meson,  $D_S \to \tau\nu$  and subsequent  $\tau \to \nu X$ , and with a detector optimised to observe a track with a kink that is characteristic of the very short-lived tau decaying.

The V-A theory was very successful in its experimental predictions but had some significant theoretical shortcomings. It was not known why it possessed no conserved current as in quantum electrodynamics, how to come to terms with the unitarity violation caused by the neutrino cross section being linear with the neutrino energy, and why the interaction is very short ranged. In the 1960s work on developing a unified electroweak theory to address these problems was progressing rapidly. Work in the early part of the decade by Glashow [15] and Salam and Ward [16] explored the weak force being mediated by massive bosons and unification under the symmetry group  $SU(2) \times U(1)$  but failed to produce a consistent theory. The path to a unified theory seemed to be through spontaneous symmetry breaking (SSB) of exact symmetries to explain the approximate conserved currents of the weak interaction. A key challenge was in Goldstone's theorem, that states every symmetry group generator broken via SSB gives rise to a massless boson

[17, 18]. Any such boson would have been easily observed experimentally and so their existence is ruled out. The breakthrough came when, independently of each other, Higgs [19], Brout and Englert [20], and Guralnik, Hagen, and Kibble [21] showed that if the broken symmetry was a gauge symmetry rather than a global symmetry, the resulting boson would be a massive spin-1 vector particle. The consequence of this was understood by both Salam [22] and Weinberg [23] who constructed a renormalisable, as proved by 't Hooft a few years later [24], electroweak theory based on an exact symmetry group that undergoes the SSB SU(2) × U(1)<sub>Y</sub>  $\rightarrow$  U(1)<sub>e</sub>. The resulting theory contains the photon and the three massive vector bosons  $W^+$ ,  $W^-$ , and Z that mediate the weak force. The mass of these bosons appears in the denominator of the propagator which solved the problem of unitarity and explains the short range of the weak force. The neutrino fits neatly into this theory as part of the left-handed lepton field SU(2) doublets

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix}, \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}, \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}.$$
 (2.1.5)

At the time of the theory's formulation no mass term for the neutrino was included due to an absence of any experimental support for neutrino mass. This electroweak theory is the basis of the modern SM and so will be discussed in more detail in § 2.2.

The next significant development in the understanding of the neutrino's nature came with Ray Davis' experiment in Homestake Mine to measure the solar neutrino flux starting in the 1960s [25]. The Homestake experiment employed a tank of  $C_2Cl_4$  to detect electron neutrinos from the Sun through the inverse  $\beta$ -decay process

$$\nu_e + ^{37} \text{Cl} \to ^{37} \text{Ar} + e^-.$$
 (2.1.6)

The Argon atoms were periodically counted to give a solar neutrino event rate. Over decades of operation, the experiment consistently measured only  $\sim 1/3$  of the predicted solar neutrino rate [26]. This deficit became known as the 'solar neutrino problem'. It was confirmed by the later experiments GALLEX [27], Kamiokande [28], and SAGE [29] which when combined also found it to be an energy dependent effect. Since these results were only for solar neutrinos and only sensitive to electron neutrinos, they could be explained by an incorrect model of the solar neutrino flux or as an incorrect description of the neutrino in the SM.

The theory of neutrino oscillations offered an explanation for the solar neutrino problem. Developed from work by Maki, Nakagawa, and Sakata [30] and Pontecorvo and Gribov [31, 32], the theory proposed neutrinos having small masses which is demonstrated to induce oscillations between flavours. In the early 2000s the SNO [33], KamLAND [34], and Super-Kamiokande [35] experiments were able to confirm the oscillation hypothesis by measuring the flux from multiple flavours and from solar, reactor, and atmospheric sources. These results confirmed that the solar model was correct and so resolved the solar neutrino problem as a consequence of neutrinos changing flavour. This was evidence of a non-zero neutrino mass. The mechanism of this oscillation will be discussed more in § 2.3.

### 2.2 Neutrinos in the Standard Model

### 2.2.1 The Standard Model

The SM Lagrangian is determined by  $SU(3) \times SU(2) \times U(1)_Y$  gauge symmetry, the Poincaré group spacetime symmetry, the requirement it produces a renormalisable quantum field theory (QFT), and the choice of group representations for the particle fields. The gauge symmetry gives rise to the the gauge fields  $G^g_\mu$ , where  $g \in \{1, ..., 8\}$ ,  $A^a_\mu$ , where  $a \in \{1, 2, 3\}$ , and  $B_\mu$  for the groups SU(3), SU(2), and  $U(1)_Y$  respectively. To allow for parity violation, the fermion fields are treated as separate left- and right-handed two-component Weyl spinors rather than as a single four-component Dirac spinor. Using the notation  $(\mathbf{p}, \mathbf{q})_n$ , where  $\mathbf{p}$  is the dimension of the SU(3) representation,  $\mathbf{q}$  is the dimension of the SU(2) representation, and n is the  $U(1)_Y$  hypercharge, the fermions of the SM are listed in Table 2.1. In addition to the fermions, there is single a Higgs doublet  $\phi$ , which is a scalar field in the  $(\mathbf{1}, \mathbf{2})_{1/2}$  representation.

The SM Lagrangian is then constructed by writing the most general gauge invariant, Poincaré invariant, and renormalisable Lagrangian possible with the available fields. Ignoring the theta terms, collections of terms in the SM Lagrangian can be grouped as

$$\mathcal{L}_{SM} = \mathcal{L}_{gauge} + \mathcal{L}_{spinor} + \mathcal{L}_{higgs} + \mathcal{L}_{yukawa}.$$
 (2.2.1)

In an attempt at brevity, we will only consider the leptonic sector when discussing these terms, i.e. ignore the quarks and SU(3) gauge field. The Minkowski metric convention

**Table 2.1:** Fermion fields of the Standard Model and their  $SU(3) \times SU(2) \times U(1)_Y$  representations. Colour indices are omitted.

Representation	Fields
Leptons	
$({f 1},{f 2})_{-1/2}$	$l_L^1 = \begin{pmatrix}  u_{eL} \\ e_L \end{pmatrix}, \ l_L^2 = \begin{pmatrix}  u_{\mu L} \\ e_L \end{pmatrix}, \ l_L^3 = \begin{pmatrix}  u_{ au L} \\  au_L \end{pmatrix}$
$({f 1},{f 1})_{-1}$	$e_R^1 = e_R, \ e_R^2 = \mu_R, \ e_R^3 = \tau_R$
Quarks	
$({f 3},{f 2})_{1/6}$	$q_L^1 = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \ q_L^2 = \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \ q_L^3 = \begin{pmatrix} t_L \\ b_L \end{pmatrix}$ $u_R^1 = u_R, \ u_R^2 = c_R, \ u_R^3 = t_R$
$({f 3},{f 1})_{2/3}$	$u_R^1 = u_R, \ u_R^2 = c_R, \ u_R^3 = t_R$
$(3,1)_{-1/3}$	$d_R^1 = d_R, \ d_R^2 = s_R, \ d_R^3 = b_R$

 $\eta_{\mu\nu} = \text{diag}(1, -1, -1, -1)$  is used and Einstein notation is implied except for with generation indices.

The gauge term consists of contractions of the field strength tensors,

$$\mathcal{L}_{\text{gauge}} = \frac{1}{4} A^{a}_{\mu\nu} A^{a\mu\nu} + \frac{1}{4} B_{\mu\nu} B^{\mu\nu}, \qquad (2.2.2)$$

where,

$$A^{a}_{\mu\nu} = \partial_{\mu}A^{a}_{\nu} - \partial_{\nu}A^{a}_{\mu} + g_{2}\epsilon^{ijk}A^{j}_{\mu}A^{k}_{\nu}, \qquad (2.2.3)$$

$$B_{\mu\nu} = \partial_{\mu}B_{\nu} - \partial_{\nu}B_{\mu}, \tag{2.2.4}$$

where  $g_2$  is a coupling constant and  $\epsilon^{ijk}$  are the SU(2) structure constants. This term includes cubic and quartic interactions for the SU(2) gauge field which, after SSB, describe self-interactions of the W and Z bosons.

The spinor term for the lepton fields is

$$\mathcal{L}_{\text{spinor}} = \sum_{m=1}^{3} \left( \bar{l}_{L}^{m} i \not \!\!\!D l_{L}^{m} + \bar{e}_{R}^{m} i \not \!\!\!D e_{R}^{m} \right). \tag{2.2.5}$$

where D is the contraction of the gamma matrices with the covariant derivative. The covariant derivative is defined by the representation of the field it acts on as

$$D_{\mu} = \partial_{\mu} + ig_2 A_{\mu}^a T_{SU(2)}^a + i\frac{g_1}{2} Y B_{\mu} \mathbb{1}, \qquad (2.2.6)$$

where  $T_{SU(2)}^a$  are the generators of the field's SU(2) representation and Y is the field's U(1) hypercharge. From this, we see that the spinor term (2.2.5) gives rise to coupling of the fermion fields to the gauge fields which in turn is responsible for the interactions between the fermion fields we see in nature.

The Higgs term consists of all allowed couplings of the scalar field with itself,

$$\mathcal{L}_{\text{higgs}} = (D_{\mu}\phi)^{\dagger} D_{\mu}\phi + \mu^{2}\phi^{\dagger}\phi - \lambda(\phi^{\dagger}\phi)^{2}, \tag{2.2.7}$$

where  $\mu^2$  and  $\lambda$  are the mass parameter and self-coupling constant of the Higgs potential respectively. So far we have not been able to write a mass term for the fermion fields since a Dirac mass made up of left and right-handed fields,

$$\mathcal{L} \supset m \left( \bar{e}_L e_R + \bar{e}_R e_L \right), \tag{2.2.8}$$

would break SU(2) symmetry. But, using gauge invariant couplings to the Higgs field we

can write a Yukawa term,

$$\mathcal{L}_{\text{yukawa}} = -\sum_{m=1}^{3} \sum_{n=1}^{3} y_l^{mn} \bar{l}_L^m \phi e_R^n + h.c., \qquad (2.2.9)$$

where  $y_l^{mn}$  are the Yukawa coupling constants of a complex  $3 \times 3$  Yukawa matrix  $Y_l$ . This appears to result in a complicated mixing of the lepton generations via coupling to the Higgs and controlled by the 18 real parameters of  $Y_l$ . But, many of these parameters can be removed through a redefinition of the lepton fields. For any complex matrix Y there exists a biunitary transformation using two unitary matrices  $V_L$  and  $V_R$ , such that  $V_L^{\dagger}YV_R$  is real and diagonal. We can use these unitary matrices for the Yukawa matrix to redefine the lepton fields,

$$l_L^m \to V_L^{mn} l_L^n, \ e_R^m \to V_R^{mn} e_R^n.$$
 (2.2.10)

Since the redefinition matrices are unitary, it is easy to see the spinor terms for the leptons (2.2.5) are left unchanged. The field redefinitions also diagonalise the Yukawa term in generation index to give

$$\mathcal{L}_{\text{yukawa}} = -\sum_{m=1}^{3} y_l^m \bar{l}_L^m \phi e_R^m + h.c..$$
 (2.2.11)

This procedure of diagonalising the Yukawa matrix can be repeated for the quark fields. However, due to the presence of an additional right-handed field (see Table 2.1) there are two possibilities for left-right Yukawa couplings that results in two Yukawa matrices to be diagonalised. This requires that the two components of  $q_L^m$  are rotated differently which results in the kinetic  $\mathcal{L}_{\text{spinor}}$  term changing following the diagonalisation of the Yukawa matrices. Couplings of the quarks to the gauge fields pick up factors of the field redefinition matrices which creates couplings that mix generation indices. This behaviour is described by the Cabibbo-Kobayashi-Maskawa (CKM) matrix and is responsible for interactions mediated by the W boson changing quark families. We will see in § 2.2.4 this mixing behaviour can also occur for the lepton fields.

### 2.2.2 Electroweak Symmetry Breaking

The key component that relates the SM Lagrangian to our observable world is electroweak symmetry breaking. This is responsible for giving mass to fermions and generating massive gauge bosons through the Higgs mechanism. Spontaneous symmetry breaking arises when there is a symmetry of the Lagrangian's equations but not of the solution to these equations. For the Higgs Lagrangian (2.2.7), the potential term has a set of minima defined by  $\phi^{\dagger}\phi = \mu^2/2\lambda = v^2/2$ . Without loss of generality, we may choose the minimum to be

$$\phi_0 = \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix}, \tag{2.2.12}$$

this is known as the vacuum expectation value (VEV) of this Higgs field. When the Higgs field assumes this VEV it is clear that much of the  $SU(2) \times U(1)_Y$  symmetry will not be present in the equations. To test this, we consider how the generators of the group act on the VEV. If  $T^a\phi_0 = 0$ , the generator does not change the VEV and is *unbroken*, if  $T^a\phi_0 \neq 0$ , the generator is *broken*. We find that the original  $SU(2) \times U(1)_Y$  generators are broken. However, the linear combination of generators  $T^3 + Y$  is unbroken,

$$(T^{3} + Y) \phi_{0} = \left(\frac{1}{2}\sigma^{3} + \frac{1}{2}\mathbb{1}\right)\phi_{0}$$

$$= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix}$$

$$= 0.$$
(2.2.13)

This single unbroken generator generates the residual symmetry group  $U(1)_{em}$  which is the symmetry associated with charge conservation. The electroweak symmetry breaking pattern is  $SU(2) \times U(1)_Y \to U(1)_{em}$ .

To examine the effect of the Higgs field's VEV on the rest of the Lagrangian, we expand the Higgs field around  $\phi_0$  in the parametrised form

$$\phi(x) = \exp\left(i\frac{\sqrt{2}}{v}\xi(x)^a\hat{T}^a\right) \begin{pmatrix} 0\\ \frac{1}{\sqrt{2}}(v+\varphi(x)) \end{pmatrix}, \tag{2.2.14}$$

where  $\hat{T}^a$  are the three broken generators we obtain through linear combinations of the original generators such that one of the generators is the unbroken  $T^3 + Y$  and the set of new generators still forms a vector space basis for the Lie algebra. When this parametrisation is substituted into the Higgs potential term (2.2.7), the fields  $\xi(x)^a$  appear in the particle spectrum as massless scalar fields known as *Goldstone bosons*. This was the observation of Goldstone's theorem mentioned in § 2.1. However, since we built our Lagrangian with local symmetries, we have the redundancy to fix the gauge. Making a gauge transformation such that the aspirant massless scalar fields  $\xi(x)^a$  are set to zero is known as going to the unitary gauge,

$$\phi_{\text{unitary}}(x) = \begin{pmatrix} 0\\ \frac{1}{\sqrt{2}}(v + \varphi(x)) \end{pmatrix}. \tag{2.2.15}$$

Using the Higgs field in the unitary gauge, we may examine the particle spectrum after SSB.

Starting with the Higgs potential (2.2.7) we find

$$\mathcal{L}_{\text{higgs}} = \ldots + \lambda v^2 \varphi + \lambda v \varphi^3 + \frac{\lambda}{4} \varphi^4. \tag{2.2.16}$$

The scalar field  $\varphi$  is the Higgs boson with a mass  $m_{\varphi} = \sqrt{2\lambda v^2} = \sqrt{2}\mu$  and self-interaction terms. The new vector boson fields and their masses can be found be substituting  $\phi_{\text{unitary}}$  into the kinetic term for the Higgs field (2.2.7). After some manipulations to reveal the particle spectrum, the resulting fields and their masses are,

$$\mathcal{A}_{\mu} = \cos \theta_W B_{\mu} + \sin \theta_W A_{\mu}^3, \ m_{\gamma} = 0, \tag{2.2.17}$$

$$W_{\mu}^{\pm} = \frac{1}{\sqrt{2}} (A_{\mu}^{1} \mp iA_{\mu}^{2}), m_{W} = \frac{g_{2}v}{2},$$
 (2.2.18)

$$Z_{\mu} = \cos \theta_W A_{\mu}^3 - \sin \theta_W B_{\mu}, \ m_Z = \frac{v}{2} \sqrt{g_1^2 + g_2^2},$$
 (2.2.19)

where we have introduced the Weinberg angle

$$\sin \theta_W = \frac{g_1}{\sqrt{g_1^2 + g_2^2}}, \ \cos \theta_W = \frac{g_2}{\sqrt{g_1^2 + g_2^2}}.$$
 (2.2.20)

The gauge covariant derivative (2.2.6) is then written in terms of these new gauge fields,

$$D_{\mu} = \partial_{\mu} + ig_2 \sin \theta_W (T^3 + Y) \mathcal{A}_{\mu} + i(g_2 \cos \theta_W T^3 - g_1 \sin \theta_W Y) Z_{\mu} + i \frac{g_2}{\sqrt{2}} (W_{\mu}^+ T^+ + W_{\mu}^- T^-),$$
(2.2.21)

where  $T^{\pm} = T^1 \pm i T^2$ . We identify  $W_{\mu}^{\pm}$  and  $Z_{\mu}$  as the gauge bosons of the weak interaction and  $\mathcal{A}_{\mu}$  as the photon field which corresponds to the unbroken residual symmetry  $\mathrm{U}(1)_{\mathrm{em}}$ . The residual symmetry's coupling constant is the electric charge which we find is a mixture of the coupling constants from the original  $\mathrm{SU}(2) \times \mathrm{U}(1)_Y$  gauge group,  $e = g_2 \sin \theta_W$ .

We have seen how SSB generates the massive vector bosons we observe in nature. Another effect of the Higgs field assuming its VEV is on the Yukawa couplings (2.2.11). Substituting  $\phi_{\text{unitary}}$  in for these terms we find

$$\mathcal{L}_{\text{yukawa}} = -\sum_{m=1}^{3} y_l^m \left( \bar{\nu}_{eL}^m \quad \bar{e}_L^m \right) \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix} e_R^m + h.c.$$

$$= -\sum_{m=1}^{3} \frac{y_l^m v}{\sqrt{2}} \bar{e}_L^m e_R^m + h.c..$$
(2.2.22)

This has the form of a mass term for each generation with mass  $m_l^m = y_l^m v / \sqrt{2}$ . SSB has resulted in two separate massless Weyl spinors,  $e_L^m$  and  $e_R^m$ , forming into a single massive Dirac spinor which correspond to the leptons we observe experimentally.

With this discussion of SSB, we now have a complete description of the electroweak SM. It describes the particle content of nature and all possible tree-level interactions with a total of 19 free parameters to be measured experimentally. A notable feature is that there is no mass term for the neutrino, it is described as a left-handed massless particle. As we have discussed in §2.1, there is strong experimental evidence that the neutrino does have mass. The reason for neutrinos being massless in the SM is both a historical

artefact and because it is not clear which mechanism for generating a massive neutrino is best supported experimentally. We will discuss how we might incorporate neutrino mass into the SM in the next section.

### 2.2.3 Mechanisms for Neutrino Mass

Multiple experiments in the late 1990s and early 2000s confirmed oscillations of neutrinos between flavours, a phenomenon implying a neutrino mass that differs by generation. So we must look for a theoretical description of the neutrino mass that, ideally, explains why their mass is so low,  $< 1\,\mathrm{eV}$  compared to the next lightest fermion with a mass of  $m_e = 511\,\mathrm{keV}$ .

Without extending the particle content of the SM there is no way to write a mass term for the left-handed neutrino that is both gauge invariant and satisfies renormalisability. The simplest extension to the SM would be to include a right-handed neutrino field  $\nu_R^m$  which would be *sterile*, meaning in the trivial representation of the gauge group, i.e.  $(1,1)_0$ . With these extra fields, the most general allowed Lagrangian picks up extra terms,

$$\mathcal{L} \supset -\sum_{m=1}^{3} \sum_{n=1}^{3} p^{mn} \bar{l}_{L}^{m} \tilde{\phi} \nu_{R}^{n} - \sum_{m=1}^{3} M^{m} \bar{\nu}_{R}^{m} (\nu_{R}^{m})_{c} + h.c., \qquad (2.2.23)$$

where  $p^{mn}$  are Yukawa couplings,  $M^m$  are the Majorana mass couplings,  $\tilde{\phi} = i\sigma_2\phi^*$  such that the Yukawa term is  $U(1)_Y$  invariant, and  $(\nu_{Rm})_c = c(\bar{\nu}_{Rm})^T$  are the charge-conjugated fields. The first term is a familiar Yukawa coupling and the second term is known as the Majorana mass term. This Majorana mass term can only be written for the new right-handed fields since they are fully neutral under the SM gauge group.

With this addition of a right-handed neutrino, there are two ways we can reproduce a small neutrino mass. One is to have the Yukawa couplings  $p^{mn}$  be of order unity and the Majorana mass  $M^m$  be very large at  $\sim 10^{14}\,\text{GeV}$ . After diagonalising the Yukawa couplings and performing electroweak symmetry breaking, the mass potential of the neutrinos becomes

$$\mathcal{L} \supset -\sum_{m=1}^{3} \frac{p^{m} v}{\sqrt{2}} \bar{\nu}_{L}^{m} \nu_{R}^{m} - \sum_{m=1}^{3} M^{m} \bar{\nu}_{R}^{m} (\nu_{R}^{m})_{c} + h.c..$$
 (2.2.24)

Then diagonalising the resulting mass matrix, we find, for each generation, one light eigenstate with mass  $\sim 0.1\,\mathrm{eV}$  and one very heavy eigenstate at  $\sim 10^{14}\,\mathrm{GeV}$ . This mechanism of generating a light neutrino is called the seesaw mechanism. The heavy eigenstate would be well beyond the energies we can probe experimentally and the light eigenstate is the familiar neutrino. An interesting consequence of the Majorana mass term is that it would require the neutrino spinor field to follow the Majorana equations of motion which enforce  $\psi_c = \psi$ , i.e. the neutrino is its own antineutrino. There would only be the neutrino, which when produced by the decay of a  $W^-$  is left-handed and when produced by the decay of a  $W^+$  is right-handed. In this scenario we say the neutrino is Majorana. The other

way to generate small neutrino mass is to forbid the introduction of the Majorana mass term. This can be done by enforcing lepton number conservation in the Lagrangian which so far has only been an accidental symmetry of the SM. The Yukawa couplings for the neutrinos would then be set much smaller than those of the other leptons to generate a small neutrino mass after SSB. The neutrino spinor field would then follow the Dirac equation only. There would be the left-handed neutrino and right-handed antineutrino we observe through the weak interaction and the right-handed neutrino and left-handed antineutrino that do not couple to the weak force. In this scenario we say the neutrino is *Dirac*.

### 2.2.4 Lepton Mixing

With the addition of the right-handed neutrino field required for generating a mass term, the full post-SSB Yukawa term for the leptons becomes

$$\mathcal{L}_{\text{yukawa}} = -\sum_{m=1}^{3} \sum_{n=1}^{3} \left( \frac{y_l^{mn} v}{\sqrt{2}} \bar{e}_L^m e_R^n + \frac{p_{\nu}^{mn} v}{\sqrt{2}} \bar{\nu}_L^m \nu_R^n \right) + h.c..$$
 (2.2.25)

To diagonalise this term the field redefinitions in (2.2.10) are no longer sufficient. We must now individually transform the fields of the left-handed doublet,

$$e_{L}^{m} \to V_{eL}^{mn} e_{L}^{n}, \ e_{R}^{m} \to V_{eR}^{mn} e_{R}^{n},$$

$$\nu_{L}^{m} \to V_{\nu L}^{mn} \nu_{L}^{n}, \ \nu_{R}^{m} \to V_{\nu R}^{mn} \nu_{R}^{n},$$
(2.2.26)

where the U(3) transformations are chosen such that they diagonalise the Yukawa couplings,  $V_{eL}^{\dagger}Y_lV_{eR}$  and  $V_{\nu L}^{\dagger}P_{\nu}V_{\nu R}$ . These field redefinitions diagonalise the Yukawa couplings (2.2.25) but, since we were not free to choose the  $\nu_L^m$  transformation to be the same as for  $e_L^m$ , the leptonic charged current couplings are affected,

$$\mathcal{L}_{\text{spinor}} \supset \sum_{m=1}^{3} i \frac{g_{2}}{\sqrt{2}} W_{\mu}^{+} \bar{\nu}_{L}^{m} \gamma^{\mu} e_{L}^{m} + h.c.$$

$$\to \sum_{m=1}^{3} \sum_{m=1}^{3} i \frac{g_{2}}{\sqrt{2}} W_{\mu}^{+} \left( V_{\nu L}^{np\dagger} V_{eL}^{pm} \right) \bar{\nu}_{L}^{n} \gamma^{\mu} e_{L}^{m} + h.c..$$
(2.2.27)

These field redefinitions have revealed charge current interactions that mix lepton generation indices. The mixing is controlled by the matrix  $U_{mn}^{\text{PMNS}} = V_{\nu L}^{np\dagger} V_{eL}^{pm}$ , called the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix.

The PMNS matrix is responsible for the phenomenon of neutrino oscillations and we have seen that it is a consequence of requiring the neutrino to have mass. It will be useful for the next section to parametrise the PMNS matrix here. The PMNS matrix is a  $3 \times 3$  unitary matrix so it appears to depend on  $3^2 = 9$  real parameters. However, it is possible to make non-symmetry phase transformations of the fermion fields that simplify

 $U^{\mathrm{PMNS}}$  but leave the kinetic and mass terms in the Lagrangian invariant. If we assume the existence of a Majorana mass term, we only have phase changes of  $e_L^m$  alongside identical phase changes of  $e_R^m$  that leave the Lagrangian invariant. This means the PMNS matrix has 9-3=6 physical parameters and may be parametrised by 3 orthogonal transformation parameters and 3 phases in the form  $U^{\mathrm{PMNS}} = U(\theta_{12}, \theta_{23}, \theta_{13}, \delta_{\mathrm{CP}}) \times \mathrm{diag}(e^{i\alpha_1}, e^{i\alpha_2}, 1)$ . The two phases  $\alpha_{1,2}$  are called Majorana phases. If we could consider the neutrino as Dirac this allows phase transformations of  $\nu_L^m$  and  $\nu_R^m$  which reduces the number of physical parameters to 9-3-(3-1)=4. The -1 coming from the transformation of all lepton flavours by the same phase having no effect on  $U^{\mathrm{PMNS}}$ . The standard parametrisation of the PMNS matrix without the Majorana phases is,

$$U^{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{\text{CP}}} \\ -c_{23}s_{12} - s_{23}c_{12}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{12} - s_{23}s_{12}s_{13}e^{i\delta_{\text{CP}}} & s_{23}c_{13} \\ s_{23}s_{12} - c_{23}c_{12}s_{13}e^{i\delta_{\text{CP}}} & -s_{23}c_{12} - c_{23}s_{12}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{13} \end{pmatrix},$$

$$(2.2.28)$$

where  $c_{ij} = \cos \theta_{ij}$  and  $s_{ij} = \sin \theta_{ij}$ . The parameters  $\theta_{ij}$  are referred to as the mixing angles and the parameter  $\delta_{CP}$  is referred to as the CP violating phase.

# 2.3 Neutrino Oscillations

The neutrino eigenstates associated with the weak interaction are called the *flavour* eigenstates. We define the flavour eigenstate by observing the accompanying charged lepton produced in interactions. The eigenstates associated with the mass terms are called, surprisingly, the *mass* eigenstates. The mass eigenstates are those of the free Hamiltonian and so are the eigenstates that are most convenient to use when considering time evolution. The experimental phenomenon of neutrino oscillations is that the flavour eigenstate can differ between creation and interaction. We will see how this phenomenon arises through the PMNS matrix.

### 2.3.1 Three Flavour Oscillations

The PMNS matrix acts as a mixing matrix between the flavour eigenstates,  $|\nu_{\alpha}\rangle$  where  $\alpha \in \{e, \mu, \tau\}$ , and the mass eigenstates,  $|\nu_{i}\rangle$  where  $i \in \{1, 2, 3\}$ ,

$$|\nu_{\alpha}\rangle = \sum_{i} U_{\alpha i} |\nu_{i}\rangle.$$
 (2.3.1)

The mass eigenstates have the standard quantum mechanical time evolution,

$$|\nu_{\alpha}(t)\rangle = \sum_{i} U_{\alpha i} e^{-i\hat{H}t} |\nu_{i}\rangle = \sum_{i} U_{\alpha i} e^{-i\sqrt{p^{2} + m_{i}^{2}t}} |\nu_{i}\rangle \approx \sum_{i} U_{\alpha i} e^{-iEt} e^{-i\frac{m_{i}^{2}t}{2E}} |\nu_{i}\rangle, \quad (2.3.2)$$

where, assuming the neutrinos to be ultrarelativistic  $p = |\vec{p}| \gg m_i$ , we have made the approximation  $E_i = \sqrt{p^2 + m_i^2} \approx E + \frac{m_i^2}{2E}$ . We have also made the simplification that the mass eigenstates are created with the same energies,  $E_i = E$ . The probability of measuring a neutrino in the flavour state  $\beta$  after it was created in a flavour state  $\alpha$  and travelled for some time T is given by the amplitude of the overlap of the states,

$$P_{\alpha \to \beta} = \left| \langle \nu_{\beta} | \nu_{\alpha}(T) \rangle \right|^2 = \left| \sum_{i} U_{\alpha i} U_{\beta i}^* e^{-i \frac{m_i^2 T}{2E}} \right|^2. \tag{2.3.3}$$

It is more common to describe oscillation probabilities in terms of distance travelled L rather than T. Since the neutrinos are ultrarelativistic, the two are interchangeable. Making the change to L and performing some manipulations, we can rewrite the oscillation probability as

$$P_{\alpha \to \beta} = \delta_{\alpha\beta} - 4 \sum_{i>j} \operatorname{Re} \left( U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j} \right) \sin^2 \left( \frac{\Delta m_{ij}^2 L}{4E} \right) + 2 \sum_{i>j} \operatorname{Im} \left( U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j} \right) \sin \left( \frac{\Delta m_{ij}^2 L}{2E} \right),$$

$$(2.3.4)$$

where  $\Delta m_{ij}^2 = m_i^2 - m_j^2$  and are called the mass splitting terms.

From the form of the oscillation probability (2.3.3) it is clear that the Majorana phases, which appear as a factor of diag( $e^{i\alpha_1}$ ,  $e^{i\alpha_2}$ , 1), have no effect on oscillations. The equation (2.3.4) then tells us that neutrino oscillations depend on the 3 mixing angles and single phase of the PMNS mixing matrix (2.2.28) along with 2 independent mass splitting terms which by convention are  $\Delta m_{21}^2$  and  $\Delta m_{32}^2$ . It is also a function of L/E which can be controlled experimentally. When we consider antineutrino oscillations  $\bar{\nu}_{\alpha} \to \bar{\nu}_{\beta}$ , the term involving the imaginary part of the PMNS matrix elements flips sign. If this term is non-zero, neutrino oscillations would be a CP violating process since  $P_{\alpha \to \beta} \neq P_{\bar{\alpha} \to \bar{\beta}}$ . The magnitude of this term is controlled by the aptly named CP violating phase parameter of the PMNS matrix.

### 2.3.2 Matter Effects

Thus far we have only considered oscillations in a vacuum. Experimentally, neutrinos will often be propagating through a non-negligible amount of matter and so be subject to a potential caused by coherent interactions with the medium. This results in the matter effect. We must formulate oscillation probabilities with matter effects taken into

account in order to properly fit oscillations parameters to experimental data. This was first done by Wolfenstein [36] and later expanded on by Mikheyev and Smirnov [37] and so is also known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect. It is because the potential experienced by different neutrino flavours is not equal that a phase difference arises which causes oscillations through matter effects. For accelerator neutrinos, where the propagation medium is the Earth, all flavours will interact via the Z boson, called a neutral current (NC) interaction, with the same amplitude. However, the electron neutrinos can also undergo scattering with electrons via the W boson, called a charge current (CC) interaction. These interaction channels, shown in Figure 2.2, result in a different potential for the electron neutrino which alters the dynamics of oscillations.

For simplicity, we will examine the matter effect in the two flavour paradigm where there is only one mixing angle and the PMNS matrix is given by the  $2 \times 2$  submatrix of the matrix found on the right of the product in (2.2.28). In a vacuum, the mass eigenstates obey the Schrödinger equation,

$$i\frac{\mathrm{d}}{\mathrm{d}t} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = H_0 \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}, \tag{2.3.5}$$

where  $H_0$  is the vacuum Hamiltonian . Using the same approximations as in the previous section and ignoring any constant phase factors since they do not affect probabilities, the Hamiltonian is

$$H_0 = \frac{1}{2E} \begin{pmatrix} m_1^2 & 0\\ 0 & m_2^2 \end{pmatrix}. \tag{2.3.6}$$

We can then use  $2 \times 2$  mixing matrix U to write (2.3.5) in terms of the flavour basis,

$$i\frac{\mathrm{d}}{\mathrm{d}t} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = UH_0U^{\dagger} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}$$

$$= \left[ \frac{m_1^2 + m_2^2}{4E} \mathbb{1} + \frac{\Delta m^2}{4E} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} \right] \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}$$

$$= \frac{\Delta m^2}{4E} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix},$$
(2.3.7)

where in the last equality the constant phase factor is omitted.

We can now introduce matter effects through the additional potential for the electron neutrinos, diag  $(V_e, 0)$ . Considering matter to be a homogeneous and isotropic gas of electrons, this potential is given by  $V_e = \pm \sqrt{2}G_F N_e$  where  $G_F$  is the Fermi constant,  $N_e$  is the electron density of matter, and the  $\pm$  is for neutrinos and antineutrinos respectively. A derivation of this potential term can be found in Section 9.1 of [38]. Using the freedom

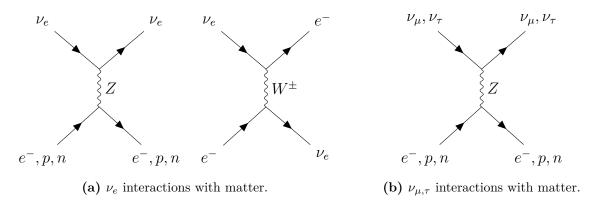


Figure 2.2: NC and CC neutrino interactions with matter.

to add a constant phase factor, the effective Hamiltonian of neutrinos in matter is

$$H_{M} = \frac{\Delta m^{2}}{4E} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} + \begin{pmatrix} \frac{V_{e}}{2} & 0 \\ 0 & -\frac{V_{e}}{2} \end{pmatrix}$$

$$= \frac{\Delta m^{2}}{4E} \begin{pmatrix} -\cos 2\theta + A & \sin 2\theta \\ \sin 2\theta & \cos 2\theta - A \end{pmatrix},$$
(2.3.8)

where

$$A = \pm \frac{2\sqrt{2}G_F N_e E}{\Delta m^2}. (2.3.9)$$

Defining an effective mixing angle  $\theta_M$  and an effective mass splitting  $\Delta m_M$  the matter Hamiltonian takes the same form as the vacuum Hamiltonian in (2.3.7),

$$H_M = \frac{\Delta m_M^2}{4E} \begin{pmatrix} -\cos 2\theta_M & \sin 2\theta_M \\ \sin 2\theta_M & \cos 2\theta_M \end{pmatrix}.$$
 (2.3.10)

Which means the oscillation probabilities will follow the same functional form as in the previous section with  $\theta \to \theta_M$  and  $\Delta m \to \Delta m_M$ . For two flavour oscillations this is

$$P_{e \to \mu} = \sin^2 \theta_M \sin^2 \left( \frac{\Delta m_M^2 L}{4E} \right). \tag{2.3.11}$$

The effective parameters are found by re-expressing the matter Hamiltonian (2.3.8) in the same form as the vacuum Hamiltonian (2.3.7). This gives

$$\Delta m_M^2 = C\Delta m^2, \tag{2.3.12}$$

$$\sin 2\theta_M = \frac{\sin 2\theta}{C},\tag{2.3.13}$$

where

$$C = \sqrt{(\cos 2\theta - A)^2 + \sin^2 2\theta}.$$
 (2.3.14)

The treatment of matter effects for three flavours follows similarly but is far more

tedious, the results can be seen in [39]. Examining the two flavour oscillation probability (2.3.11) is enough to provide insight into the consequences of the matter effect. One such consequence is a resonance condition known as the MSW resonance that occurs when  $\cos 2\theta = A$ . Even with a very small vacuum mixing angle, oscillations will be significantly enhanced when this condition is met. We also see that the oscillation probabilities will be different for neutrinos and antineutrinos due to the  $\pm$  in A (2.3.9), meaning experimentally observing  $P_{\alpha \to \beta} \neq P_{\overline{\alpha} \to \overline{\beta}}$  in the presence of matter effects is not enough to indicate CP violation, precision measurements of the oscillations probabilities are required. Lastly, we see that the MSW resonance depends on the sign of the mass splitting,  $\Delta m^2$ , as it appears in the denominator of A (2.3.9). This means that through the matter effect we can have strong experimental sensitivity to not just the difference in mass between the mass eigenstates, but also their ordering.

### 2.3.3 Physics From Oscillation Measurements

The neutrino mass and in turn the phenomenon of flavour oscillations is a divergence with the SM strongly supported by experimental data. It is important that precision measurements of the parameters governing oscillations are made to inform both the underlying theory of particle physics and our understanding of experimental data where properties of the neutrino are relevant.

The mixing angles  $\theta_{12}$ ,  $\theta_{13}$ , and  $\theta_{23}$  form the real part of the PMNS matrix and along with  $\delta_{\text{CP}}$  determine the flavour mixing pattern. Precision measurements of the mixing angles are important in determining the unitarity of the PMNS matrix. In § 2.2.3 the PMNS matrix is unitary by definition but this is under the implicit assumption of there being only three generations of leptons. A common feature of extensions to the SM is to include additional massive fermionic singlets, often sterile neutrinos. The inclusion of n sterile neutrinos that may mix with the SM neutrinos means the PMNS matrix is the  $3 \times 3$  submatrix of a larger  $(3 + n) \times (3 + n)$  unitary mixing matrix,

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & \dots & U_{e(3+n)} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} & \dots & U_{\mu(3+n)} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} & \dots & U_{\tau(3+n)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ U_{(3+n)1} & U_{(3+n)2} & U_{(3+n)3} & \dots & U_{(3+n)(3+n)} \end{pmatrix}.$$
(2.3.15)

Since a unitary matrix has orthonormal column and rows, it is clear the  $3 \times 3$  submatrix would not be unitary. As such, an important test of potential new physics is validating the unitarity of the PMNS matrix by testing the condition  $\sum_{m=1}^{3} U_{mp} U_{mq}^* = \delta_{pq}$ . This requires precision measurements of the mixing angles.

A mixing angle of particular importance is  $\theta_{23}$ . Determining the octant of  $\theta_{23}$ , mean-

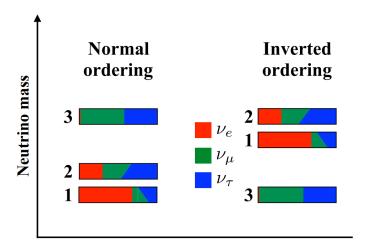


Figure 2.3: Illustration of the neutrino mass hierarchy. Coloured bands represent the proportion of each flavour eigenstate the mass eigenstates are composed of. For each band  $\delta_{CP}$  is varied from 0 (bottom) to  $\pi$  (top). Figure from [41].

ing  $\theta_{23} > 45^{\circ}$  or  $\theta_{23} < 45^{\circ}$  where 45° corresponds to maximal mixing, is necessary to understand the structure of the PMNS matrix. This result would be a strong model discriminator for many SM extensions that introduce new flavour symmetries, such as an  $\mu$ - $\tau$  symmetry [40]. Particular attention is given to the octant of  $\theta_{23}$  because experimentally it is a very difficult measurement owing to the presence of nearly degenerate solutions for the oscillation probability.

The determination of the sign of the mass splitting terms and so the order of mass eigenstates is of significant importance. Current experimental constraints on the mass splitting and the sign of  $\Delta m_{21}^2$  leave two possible scenarios for the mass ordering. Resolving the mass ordering as normal ordering (NO),  $m_1 < m_2 < m_3$ , or inverted ordering (IO),  $m_3 < m_1 < m_2$ , is known as the neutrino mass hierarchy problem. It is represented pictorially in Figure 2.3. One important consequence of the mass ordering is its effect on neutrinoless double  $\beta$ -decay ( $0\nu\beta\beta$ ). This is a double  $\beta$ -decay with no neutrinos in the final state. Such a process is only possible if a Majorana mass term exists so that an outgoing right-handed massive neutrino can couple as an incoming left-handed neutrino to another W boson. We see from Figure 2.4 that the effective Majorana mass  $m_{\beta\beta}$ , which determines the rate of  $0\nu\beta\beta$ , is highly dependent on the mass ordering. For this reason, resolving the mass ordering through oscillation measurements is an important part of the search for a Majorana neutrino.

Perhaps the most consequential of the oscillation parameters is the CP violating phase. The abundance of matter over antimatter in the observable universe is thought to be a consequence of some baryon/lepton asymmetry producing process in the very early universe known as baryogenesis/leptogenesis. In 1967, Sakharov proposed three conditions necessary for any theory of baryogenesis which also apply analogously to one of leptogenesis [43]:

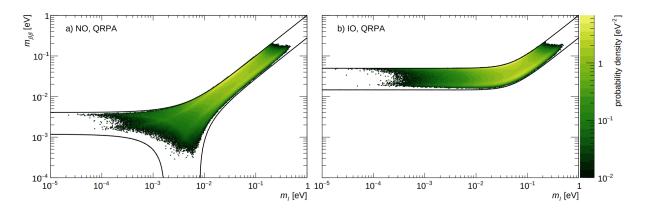


Figure 2.4: Effective Majorana neutrino mass  $m_{\beta\beta}$  as a function of lightest neutrino mass for normal ordering (left) and inverted ordering (right). The solid lines show allowed values from the best fit neutrino oscillation parameters and the colour scale the marginalised posterior distributions. Figure from [42].

- 1. **Baryon number violation**. To permit a process that can create baryons from anti-baryons and vice-versa.
- C and CP violation. To allow enhancement of baryon producing processes and suppression of anti-baryon producing processes.
- 3. Loss of thermal equilibrium. To prevent baryon number violating processes having the same rate as their inverse process.

The SM permits some CP violation through the quark sector but not enough to satisfy the second condition for models of baryogenesis. A leading model for leptogenesis generates a lepton asymmetry via CP violating decays of right-handed Majorana neutrinos which in turn generates the baryon asymmetry [44]. Such a model may account for the present day baryon asymmetry if the neutrino is Majorana and there is sufficient leptonic CP violation.

CP violation in neutrino oscillations implies that  $P_{\alpha \to \beta} \neq P_{\overline{\alpha} \to \overline{\beta}}$ . However, even with CP violation, invariance to CPT demands that  $P_{\alpha \to \alpha} = P_{\overline{\alpha} \to \overline{\alpha}}$ . Therefore, the CP asymmetry is given by  $A_{\text{CP}}^{\alpha\beta} = P_{\alpha \to \beta} - P_{\overline{\alpha} \to \overline{\beta}} \,\,\forall\,\,\alpha \neq \beta$ . This corresponds to experiments that can measure neutrinos in a new flavour to the production flavour, called appearance, in both neutrino and antineutrino modes. This is in contrast to experiments that measure only a deficit of neutrinos in the same flavour as the production flavour, called disappearance. When expressing  $A_{\text{CP}}^{\alpha\beta}$  in terms of the oscillation parameters, it is found that there is always a factor called the Jarlskog invariant J [45] given by

$$J = \sin \theta_{13} \cos^2 \theta_{13} \sin \theta_{12} \cos \theta_{12} \sin \theta_{23} \cos \theta_{23} \sin \delta_{CP}, \qquad (2.3.16)$$

which can be shown to be independent of the choice of basis for the mixing matrix we have through phase redefinitions of lepton fields. It is therefore the physically meaningful

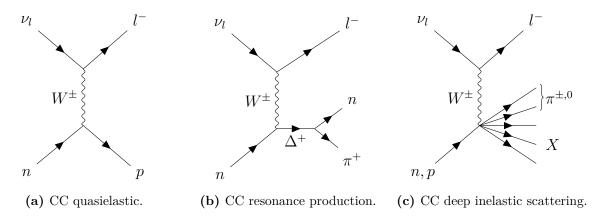
measure of leptonic CP violation. From J, we see that CP is conserved for  $\delta_{\rm CP} = 0, \pm \pi$  and maximally violated for  $\delta_{\rm CP} = \pm \pi/2$ . In addition, that by constraining  $\delta_{\rm CP}$  and the mixing angles in the standard PMNS parametrisation (2.2.28), the amount of leptonic CP violation can be measured.

### 2.4 Neutrino Interactions

As shown in § 2.3, the oscillation probability is always a function of the true neutrino energy. This means that in order to measure oscillation parameters the neutrino energy must be known for each event. Experimentally we cannot achieve monoenergetic neutrino fluxes. For accelerator neutrino experiments, this is due to pions and kaons being produced with their own energy spectra and then undergoing two-body decay in a beam that is not perfectly collimated, we will cover this in more detail in § 3.4. Because of this, it is necessary to reconstruct the neutrino energy for each event. This requires an interaction model, also called a cross section model, that relates a neutrino incident on a nucleus with a given energy to a set of final state particles that can be measured in a detector. By implementing these models in Monte Carlo (MC) simulations, they can be used to understand how reconstructed energy in the detector relates to the neutrino energy the oscillation probabilities depends on. Such MC implementations are known as neutrino event generators.

We consider the initial interaction of a neutrino on a nucleus as a scattering with a single nucleon suspended in the nuclear medium. To model this we must know the neutrino-nucleon cross section as a function of energy. For the energy ranges relevant to accelerator neutrino experiment, the total cross section will have contributions from three types of neutrino-nucleon interaction. These types of interaction are shown schematically in Figure 2.5.

Neutrinos with energies lower than the mass of their associated lepton may only interact with a nucleon via NC scattering. As the neutrino energy increases, CC interactions



**Figure 2.5:** Examples of the three main types of CC neutrino-nucleon interactions. Here  $l = e, \mu, \tau$  and X represents the hadronic system excluding pions.

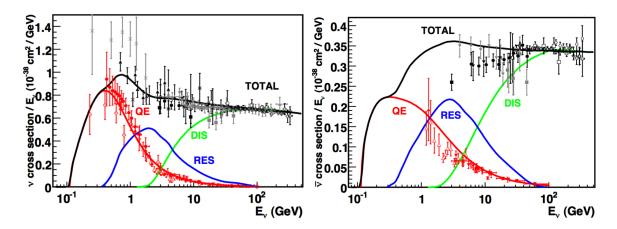
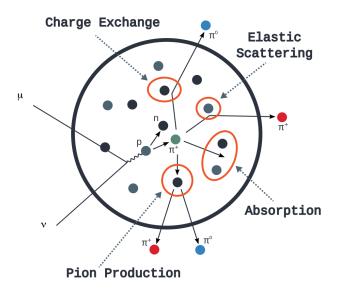


Figure 2.6: Predictions [46] and measurements [47] of total muon neutrino (left) and antineutrino (right) per nucleon CC cross section. Contributions from different interactions shown, QE is the CCQE interaction introduced in § 2.4. The error bars represent the uncertainties in the cross section measurements. Figure from [48].

accompany this NC interaction channel. The simplest CC interaction is quasielastic scattering (CCQE) with a nucleon as shown in Figure 2.5a. This occurs when neutrinos acquire sufficient energy to create the lepton's mass. At higher energies, an additional interaction channel becomes available to the neutrino known as resonance production (RES) as shown in Figure 2.5b. In this interaction the nucleon is excited into a baryonic resonance which quickly decays back to a nucleon and, most often, a single pion. Finally, at even higher energies, the neutrinos can scatter directly off quarks in a deep inelastic scattering (DIS) interaction as shown schematically in Figure 2.5c. A DIS interaction causes the break up of the nucleon containing the struck quark. This results in hadronisation and the formation of a shower of strongly interacting particles.

The contribution from each of these interaction types to the total neutrino cross section is shown in Figure 2.6. There is significant overlap between different interaction types and for the energy spectrum of many accelerator neutrino experiments, most notably DUNE, the total cross section will often be made up of multiple interaction types. For this reason, it is important the interaction model accurately predicts the relative contribution of each interaction type to the total cross section as a function of energy.

Modern accelerator neutrino experiments use detectors with a heavy nuclei medium to allow high interaction rates. For this reason, the interaction model must simulate neutrino-nucleus scattering rather than only neutrino-nucleon scattering. One part of this is to consider the initial state of the nucleon. Inside the nucleus, a nucleon is constantly changing its momentum as it moves in a nuclear potential, resulting in a smearing of the reconstructed energy that needs to be understood. This is typically described using a simple relativistic Fermi gas model [49, 50] with more advanced models being investigated such as the use of spectral functions [51]. In addition, the model of the initial state of the nucleus needs to also account for a Pauli blocking effect where interaction channels



**Figure 2.7:** Illustration of nucleon-nucleon and pion-nucleon final state interactions of the hadronic component of the initial interaction. Different interactions result in significantly different kinematics and multiplicities for the particles escaping the nucleus. Figure from [54].

that result in nucleons being in a state that is already occupied by another nucleon are forbidden. Lastly, it is important that correlations between nucleons are modelled properly to account for the interaction channel where the W in Figure 2.5a is absorbed by a pair of nucleons resulting in two outgoing nucleons and an enhanced cross section. These are known as two-particle-two-hole (2p2h) interactions [52]. These interactions are dominated by channels where the two nucleons interact via an exchange of a virtual pion known the meson exchange current.

After modelling the initial state of the nucleus to get an accurate description of the neutrino-nucleon interaction, final state interactions (FSI) must be considered. These are interactions the outgoing particles from the neutrino-nucleon interaction may undergo as they traverse the nuclear medium to be potentially emitted from the nucleus. Typically the lepton will escape the nucleus but the hadrons will interact as they propagate through the nucleus. They can change their momentum via scattering, undergo absorption, or collide with nucleons and generate additional particles. Some of the possible FSI effects are illustrated in Figure 2.7. One method of modelling FSI is the intranuclear cascade model [53].

The critical role of the interaction model is to provide a map from the true energy of the incident neutrino to the reconstructed energy in the detector. One method of energy reconstruction is to select CCQE interactions and consider the kinematics of the outgoing lepton only. This requires accurate knowledge of the smearing from the nucleon's initial motion, the 2p2h contributions, and FSI effects that can change pion multiplicity to produce QE-like events from non-QE interactions. Alternatively, there is the calorimetric energy reconstruction which looks to measure the energy of all the final state particles.

Final state particles meaning the particles that escape the nucleus following the neutrinonucleon interaction. This is the method that DUNE is designed to use. For this method to work well the interaction model must predict not just the fraction of neutrino energy absorbed by the nucleus one way or another, but an accurate prediction of the multiplicities and kinematics of final state particles. This is because the detector will have different efficiencies and resolutions for different particles and their energies. Most notably, it may not be possible to reconstruct the energy of neutral particles in the final state.

A key problem with the models employed by event generators to simulate the aforementioned nuclear effects is that there are significant uncertainties in the models themselves rather then just the parameters that govern them. This is expected to be a large source of systematic error and potentially bias in the precision measurement of oscillation parameters DUNE aims to make. Multiple experiments aim to study relevant interaction models including MiniBooNE [55], MicroBooNE [56], MINER $\nu$ A [57], and SBND [58]. By comparing neutrino-nucleus data to simulation at different energies they hope to improve modelling and properly parametrise uncertainties. Despite this, it is expected that interaction modelling will continue to be a dominant systematic uncertainty.

### 2.5 Neutrino Oscillation Measurements

Sensitivity to the full set of oscillation parameters requires access to multiple oscillation channels, a range of L/E, and exploitation of the matter effect. Accurate measurement of all parameters is therefore dependent on measurements from multiple experiments in distinct configurations. An overview of some of the important measurements from various experiments will be given in this section.

### 2.5.1 Solar and Long Baseline Reactor Neutrinos

The solar neutrino experiments, mentioned in §2.1, primarily measure the survival probability for electron neutrinos. Solar neutrinos are generated by thermonuclear reactions in the interior of the Sun with various reaction chains contributing to the total flux. The solar neutrinos have energy  $\sim 10 \, \text{MeV}$  and are, of course, at a baseline of approximately 1 AU. At this L/E, the electron neutrino survival probability is dominated by  $\theta_{12}$  and  $\Delta m_{21}^2$  with the sign of the latter being accessible through the contribution from matter effects with the Sun.

Long baseline reactor neutrino experiments are also sensitive to the same oscillation parameters through measuring the survival probability of reactor electron antineutrinos. The KamLAND experiment measured electron antineutrinos with an energy spectrum peaking at  $\sim 4\,\mathrm{MeV}$  from 56 nuclear reactors across Japan at a flux-weighted average baseline of  $\sim 180\,\mathrm{km}$  [59]. Importantly, KamLAND was able to measure the survival probability as a function of L/E as shown in Figure 2.8. This provided a precise mea-

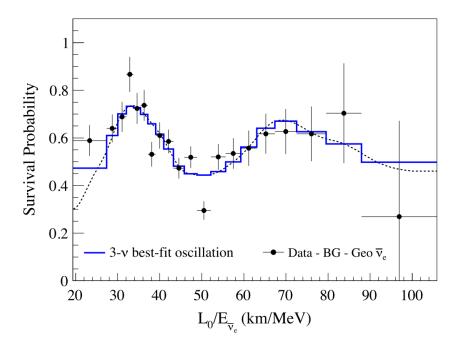


Figure 2.8: Ratio of observed signal  $\bar{\nu}_e$  with no oscillation expectation for the KamLAND experiment.  $L_0 = 180 \,\mathrm{km}$  is the flux-weighted average reactor baseline. Figure from [59]

surement of  $\theta_{12}$  and  $\Delta m_{21}^2$ .

Combining results from KamLAND with data from multiple solar neutrino experiments results in a strong constraint on  $\theta_{12}$  and  $\Delta m_{21}^2$  as shown in Figure 2.9. In the combined fit the sensitivities of the solar experiments and KamLAND complement each other, highlighting the importance of leveraging multiple experimental configurations. We also see from the fit that the sign of the mass splitting is resolved as  $\Delta m_{21}^2 > 0$  which means the mass eigenstate  $\nu_2$  must be heavier than  $\nu_1$ .

### 2.5.2 Short Baseline Reactor Neutrinos

KamLAND is one of many reactor neutrino experiments that aimed to measure a disappearance of electron antineutrinos from nuclear reactors. The other experiments are typically situated much closer to the reactors than KamLAND. At baselines of 0–2 km, the oscillations are most sensitive to the parameters  $\theta_{13}$  and  $\Delta m_{31}^2$ . The constraints on these parameters from the Daya Bay and RENO reactor neutrino experiments can be seen in Figure 2.10. We see that Daya Bay provides an accurate measurement of  $\theta_{13}$ . It should be noted that short baseline reactor experiments are not sensitive to the sign of the mass splitting through matter effects.

### 2.5.3 Atmospheric Neutrinos

Thus far, we have not mentioned the possibility of measuring atmospheric neutrinos. Atmospheric neutrinos are produced when cosmic rays strike the atmosphere to create a

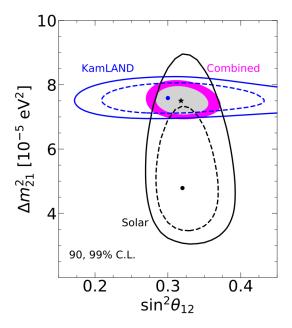
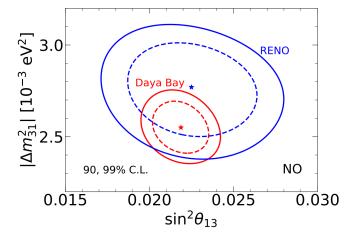


Figure 2.9: Allowed regions and best fit in  $\sin^2 \theta_{12} - \Delta m_{21}^2$  from solar neutrino experiments (black), KamLAND (blue), and a combined analysis of the two (shaded). The star is the best fit value of the combined analysis. The solar experiments considered are Homestake, GALLEX, GNO, SAGE, Borexino, Super-K, and SNO. Figure from [60].



**Figure 2.10:** Allowed regions and best fit in  $\sin^2 \theta_{13} - |\Delta m_{31}^2|$  with normal ordering assumed from Daya Bay (red) and RENO (blue) experiments. Figure from [60].

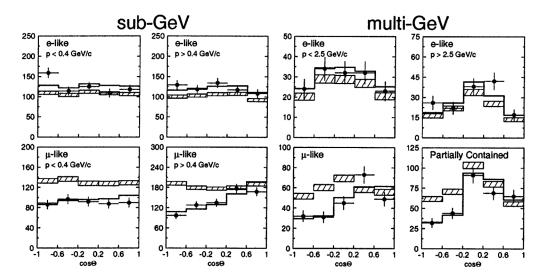


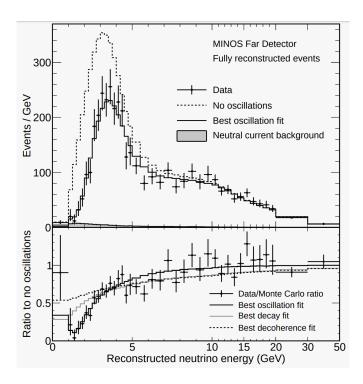
Figure 2.11: Super-K atmospheric neutrino data. Events are divided by e- or  $\mu$ -like and energy and are binned by zenith angle with  $\cos \Theta < 0$  corresponding to upward-going particles. The hatched region is the no oscillation prediction and the solid line is the best fit expectation for  $\nu_{\mu} \leftrightarrow \nu_{\tau}$  oscillations. Figure from [35].

hadronic shower containing pions. The charged pions then decay as  $\pi \to \mu \nu_{\mu}$  followed by decay of the muon as  $\mu \to e \nu_e \nu_{\mu}$  to produce a neutrino flux that is mostly  $\nu_{\mu}$  and  $\bar{\nu}_{\mu}$  and with energies ranging from a few MeV up to  $10^9$  GeV. Atmospheric neutrino oscillations are dominated by the parameters  $\theta_{23}$  and  $\Delta m_{32}^2$ , meaning the expected effect would be an absence of upward-going muon neutrinos, as these have travelled up through the Earth, and an excess of upward-going tau neutrinos.

In 1998 the Super-Kamiokande (Super-K) collaboration reported the results of an experiment that utilised a 50 kiloton water Cherenkov detector to reconstruct the flavour, energy, and direction of atmospheric neutrinos [35]. A clear oscillation signature was observed by binning observations into zenith angle as seen in Figure 2.11. The measurement of a disappearance in  $\mu$ -like interactions for upward-going neutrinos alongside a relatively constant rate of e-like interactions in all directions was strong evidence of  $\nu_{\mu} \rightarrow \nu_{\tau}$  oscillations. After almost 15 cumulative years of data-taking, Super-K's latest oscillation parameter constraints from atmospheric neutrinos only is, with normal hierarchy assumed,  $\sin^2\theta_{23} = 0.587^{+0.036}_{-0.069}$  and  $|\Delta m_{32}^2| = 2.50^{+0.12}_{-0.31} \times 10^{-3} \text{ eV}^2$  [61].

### 2.5.4 Accelerator Neutrinos

Accelerator neutrinos are produced by impacting a proton beam on a target material to produce mesons which can then decay into a beam of predominantly muon neutrinos and antineutrinos. The first accelerator neutrino experiment being the 1962 Brookhaven experiment mentioned in § 2.1. The prospect of using neutrino beams to study oscillations only became viable with the invention of the magnetic focusing horn [62] that allowed for the formation of a approximately collimated beam. Using different polarities for the



**Figure 2.12:** MINOS far detector muon neutrino disappearance measurement. The ratio shows fits for alternate explanations of the deficit: neutrino decay and decoherence. Figure from [69].

magnetic horn also allows for both neutrino and antineutrino beams with good purities. Long baseline neutrino experiments utilise such a focused neutrino beam and consist of a near detector (ND) to constrain systematic uncertainties by measuring the initial unoscillated beam flux and a far detector (FD) that measures the oscillated flux. They measure both the  $\nu_e$  appearance and disappearance channels from the initial  $\nu_\mu$  flux with both neutrino and antineutrino beams. This makes them sensitive to the oscillation parameters  $|\Delta m_{31}^2|$ ,  $|\Delta m_{32}^2|$ ,  $|\theta_{23}|$ ,  $|\theta_{13}|$ ,  $|\delta_{\rm CP}|$ , and potentially the mass ordering through matter effects.

The first long-baseline accelerator neutrino experiments were K2K [63] with a baseline of 250 km and MINOS [64] with a baseline of 735 km. Both experiments were able to make clear disappearance measurements as demonstrated for MINOS in Figure 2.12. This allowed for constraints on the values of  $\Delta m_{31}^2$  and  $\theta_{23}$  [65, 66]. Measurements of the appearance channel where also made which gave some sensitivity to  $\theta_{13}$  [67, 68].

The next generation of long-baseline detectors were T2K [70] and NOvA [71]. These experiments were designed to have an improved sensitivity to the appearance channel over their predecessors. With much greater statistics, both experiments were also able to make improved measurements of the disappearance channel resulting in better constraints on  $|\Delta m_{31}^2|$  and  $\theta_{23}$  as shown in Figure 2.13. The T2K experiment was able to confirm a non-zero  $\theta_{13}$  with conclusive statistical significance and produce a best fit of  $\sin^2 2\theta_{13} = 0.140^{+0.038}_{-0.032}$  assuming NO, fixed  $|\Delta m_{32}^2|$  and  $\theta_{23}$ , and CP conservation [72]. By making measurements of the appearance channel for both neutrinos and antineutrinos, NOvA

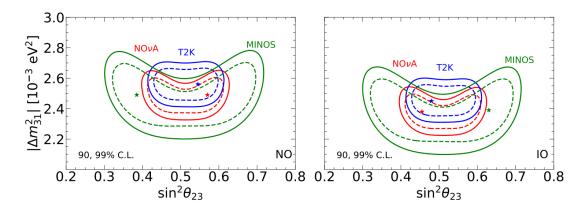


Figure 2.13: Allowed regions and best fit in  $\sin^2 \theta_{23} - |\Delta m_{31}^2|$  for NOvA (red), T2K (blue), and MINOS (green). Shown for normal hierarchy assumed (left) and inverted hierarchy assumed (right). Figure from [60].

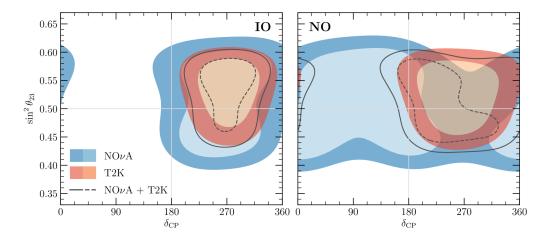


Figure 2.14: Allowed  $1\sigma$  and  $2\sigma$  regions and best fit in  $\delta_{\rm CP} - \sin^2 2\theta_{23}$  for NOvA, T2K, and their combination. Shown for normal hierarchy assumed (right) and inverted hierarchy assumed (left). Figure from [75]

and T2K made measurements of  $\delta_{\rm CP}$  to establish whether CP violation occurs in neutrino oscillations [73, 74], these are shown overlaid in Figure 2.14. The measurements exhibit tension for the NO case with the T2K data providing the best constraint and favouring maximal CP violation of  $\delta_{\rm CP} \sim 3\pi/2$ . However, neither experiment was able to exclude CP conserving values of  $\delta_{\rm CP}$  for either mass ordering with a credible interval greater than  $3\sigma$ , falling short of the  $5\sigma$  typically required for discovery.

Establishing conclusively if CP violation occurs in the leptonic sector and increasing the precision of oscillation parameter measurements is the purview of the upcoming long-baseline experiments DUNE [76] and Hyper Kamiokande [77], the former will be discussed at length in the next chapter.

**Table 2.2:** Best fit and  $3\sigma$  ranges for oscillation parameters from global fit of oscillation experiments. Shown for both normal and inverted ordering scenarios. Global fit data from NuFIT [75].

	Normal Ordering		Inverted Ordering	
	Best Fit Point $\pm 1\sigma$	$3\sigma$ range	Best Fit Point $\pm 1\sigma$	$3\sigma$ range
$\sin^2 \theta_{12}$	$0.304^{+0.013}_{-0.012}$	[0.269, 0.343]	$0.304^{+0.013}_{-0.012}$	[0.269, 0.343]
$\sin^2 \theta_{23}$	$0.570^{+0.018}_{-0.024}$	[0.407, 0.618]	$0.575^{+0.017}_{-0.021}$	[0.411, 0.621]
$\sin^2 \theta_{13}$	$0.02221^{+0.00068}_{-0.00062}$	[0.02034, 0.02430]	$0.02240^{+0.00062}_{-0.00062}$	[0.02053, 0.02436]
$\Delta m_{21}^2$	$7.42^{+0.21}_{-0.20} \times 10^{-5}  \text{eV}^2$	$[6.82, 8.04] \times 10^{-5} \mathrm{eV}^2$	$7.42^{+0.21}_{-0.20} \times 10^{-5} \mathrm{eV^2}$	$[6.82, 8.04] \times 10^{-5} \mathrm{eV}^2$
$\Delta m_{32}^{2}$	$2.514_{-0.027}^{+0.028} \times 10^{-3} \mathrm{eV}^2$	$[2.431, 2.598] \times 10^{-3} \mathrm{eV}^2$	$-2.497^{+0.028}_{-0.028} \times 10^{-3} \mathrm{eV}^2$	$[-2.583, -2.412] \times 10^{-3} \mathrm{eV}^2$
$\delta_{\mathrm{CP}}$	$195^{+51}_{-25}^{\circ}$	$[107, 403]^{\circ}$	$286^{+27^{\circ}}_{-32}$	[192, 360]°

#### 2.5.5 Current Oscillation Parameter Constraints

The best constraint on the parameters governing three flavour neutrino oscillations are obtained by combining data from the experiments discussed in this section along with many others not mentioned. This global fit of oscillation parameters is provided by NuFIT [75] and shown in Table 2.2.

Substantial improvements on these constraints are expected from the three main next-generation experiments: DUNE, Hyper-K, and the Jiangmen Underground Neutrino Observatory (JUNO) [78]. Hyper-K will probe CP-violation through high-statistics  $\nu_e/\bar{\nu}_e$  appearance measurements, while JUNO will determine the neutrino mass ordering via high-precision measurements of reactor antineutrino oscillations. These three experiments will provide powerful and complementary tests of the PMNS framework.

# The Deep Underground Neutrino Experiment

The Deep Underground Neutrino Experiment is a next-generation long-baseline neutrino oscillation experiment that primarily aims to make precision measurements of neutrino oscillation parameters with an emphasis on resolving the neutrino mass ordering and making a conclusive observation of CP violation in the lepton sector. This will be accomplished using large scale liquid argon time projection chambers (LArTPCs), the world's most intense neutrino beam, and an innovative near detector (ND) design. This chapter will start with an overview of the core components of DUNE. This will be followed by a discussion of some of DUNE's physics goals. LArTPC technology will subsequently be reviewed. The neutrino beam facility, the ND, and the far detector (FD) comprising the experiment will then be described.

# 3.1 Overview

As illustrated in Figure 3.1, DUNE consists of a neutrino beam generated by the Long-Baseline Neutrino Facility (LBNF) at the Fermilab, a ND  $\sim$ 600 m downstream, and a FD at a baseline of 1285 km and 1.5 km below the surface. In addition, there are multiple prototype detectors designed to test detector technologies, facilitate development of data acquisition and reconstructions algorithms, and study neutrino interactions on argon. The FD will host four detector modules each with a fiducial mass of at least 10 kt while the ND will consist of multiple unique detectors designed to constrain relevant systematic uncertainties and handle neutrino pile-up.

At the time of writing, the designs for the ND and FD complex are complete and excavation of the FD site is finished. Detector designs are in mature stages and construction for some components is underway. Construction of the full experiment is planned in a phased approach. Phase I will consist of two FD modules totalling 20 kt fiducial mass and is anticipated to start taking early physics data in the late 2020s from non-beam neutrino sources. As part of this phase, the beam and ND will come online by 2032. The full experiment design will be realised with Phase II which will include an additional two 10 kt FD modules, upgraded beam power, and an upgraded ND. It is expected that it will take up to 16 years of data taking following the completion of Phase I for DUNE to accomplish all of its core physics goals [80].

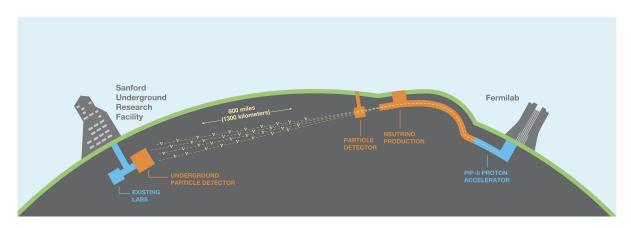


Figure 3.1: Illustration of the experimental configuration of DUNE. Figure from [79].

# 3.2 Physics Goals

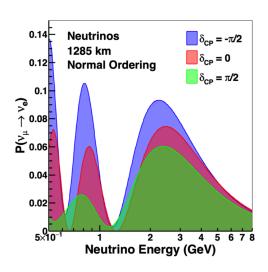
DUNE will be a long-baseline neutrino oscillation experiment, an observatory for neutrinos from astrophysical sources, and act as a nucleon-decay detector. This permits a rich physics programme. In this section, we will focus on the SM neutrino oscillation and supernova neutrino burst measurements. Not discussed is DUNE's ability to improve global nucleon decay sensitivities and the wide range of physics beyond the SM DUNE is sensitive to. A thorough review of these topics can be found in [41].

#### 3.2.1 Standard Model Neutrino Oscillations

Similarly to the accelerator neutrino experiments discussed in § 2.5.4, DUNE will measure both the electron (anti)neutrino appearance and muon (anti)neutrino disappearance channels. Good energy resolution, flavour identification, and high statistics will permit strong constraints on oscillation parameters that are either superior or complementary to existing measurements. The wide-band and high-energy neutrino beam will give DUNE unique sensitivity to both  $\delta_{CP}$  and the mass ordering through the appearance channel. The appearance probabilities plotted in Figure 3.2 highlight this sensitivity to  $\delta_{CP}$ .

DUNE's oscillation physics sensitivities from beam neutrinos are shown in Figure 3.3. They highlight some of the key expected results from DUNE:

- A conclusive determination of the neutrino mass ordering that will be possible for all true values of  $\delta_{CP}$ .
- A measurement of  $\delta_{CP}$  with a resolution of  $\sim 5-15^{\circ}$  that will give a conclusive determination of leptonic CP violation for a large range of true parameter space around maximal CP violation.
- A strong constraint on  $\theta_{23}$  with potential to resolve the octant depending on the true value.



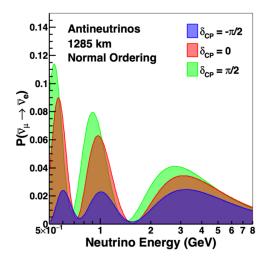


Figure 3.2: Appearance probability for neutrinos (left) and antineutrinos (right) for different possible values of  $\delta_{CP}$ . The baseline and energy range are chosen to be relevant for DUNE. Figure from [76].

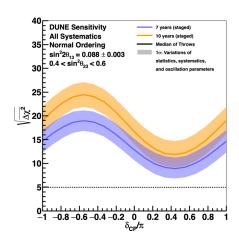
Aside from the beam neutrino measurements, the high fiducial mass of DUNE allows for oscillation physics programmes with other neutrino sources. Atmospheric neutrinos can be used to improve the sensitivity to CP violation [81, 82] and the mass ordering [83]. Solar neutrinos can provide new sensitivity to the parameters  $\Delta m_{21}^2$  and  $\theta_{12}$ . This would help investigate the slight tension between reactor and solar neutrino measurements shown in Figure 2.9. In addition, DUNE has a unique capability to constrain certain aspects of the solar neutrino flux [84].

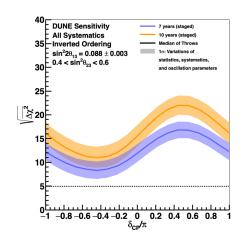
# 3.2.2 Supernova Neutrino Bursts

The sensitivity of DUNE to low energy neutrinos in the region of 5–30 MeV allows for the detection of the electron neutrino spectrum from a galactic core-collapse supernova. The dominant channel for detection by DUNE is the CC absorption of an electron neutrino on argon,

$$\nu_e + ^{40} \text{Ar} \to e^- + ^{40} \text{K}^*,$$
 (3.2.1)

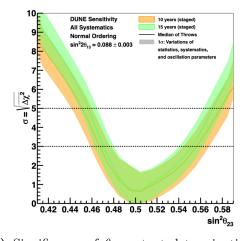
where  $^{40}$ K\* is an excited state. The interaction creates a short electron track that can be reconstructed alongside possible light and secondary particle signatures. Additionally, there are channels for the equivalent electron antineutrino capture and the elastic scattering of the neutrino with an electron. Core-collapse supernovae are expected to occur in the Milky Way once every few decades [85]. As of yet, there has been a single occurrence where a supernova neutrino burst was observed. A few dozen antineutrinos from SN1987A, an extragalactic supernova, were measured via inverse  $\beta$ -decay by neutrino observatories operational at the time [86, 87, 88]. During the several decades operational lifetime of DUNE there is a reasonable chances of another such event occurring. If it were to occur, DUNE would make a detailed measurement of the neutrino and antineutrino

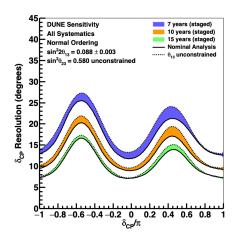




(a) Significance of mass ordering determination as a function of true  $\delta_{\rm CP}$  for true normal ordering.

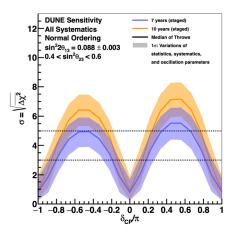
(b) Significance of mass ordering determination as a function of true  $\delta_{\rm CP}$  for true inverted ordering.





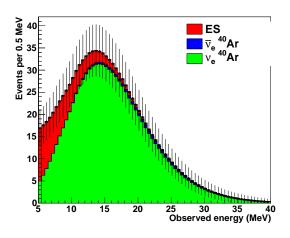
function of true  $\sin^2 \theta_{23}$ .

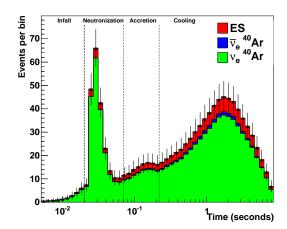
(c) Significance of  $\theta_{23}$  octant determination as a (d) Resolution of  $\delta_{CP}$  measurement as function of true  $\delta_{\rm CP}$ .



(e) Significance of CP violation determination, i.e.  $\delta_{\rm CP} \neq 0, \pi$ , as a function of true  $\delta_{\rm CP}$ .

Figure 3.3: Expected DUNE beam sensitivities to mass ordering (a,b), the octant of  $\theta_{23}$ (c), and CP violation (d,e) for equal parts neutrino and antineutrino beam running mode. Sensitivities are shown as a function of true  $\delta_{\rm CP}$  since its current value is unknown. Greater years of exposure correspond to increased statistics. The CP violation,  $\theta_{23}$  octant, and mass ordering sensitivities are given by a likelihood ratio expressed in terms of a test statistic for the two outcomes,  $\Delta \chi^2 = \chi_B^2 - \chi_A^2$ . Figures from [76].





- (a) Time integrated energy spectrum.
- (b) Time-dependent energy spectrum. Stages of the core-collapse relevant to the neutrino signal are shown.

**Figure 3.4:** Expected measured neutrino event rates for core-collapse supernova at 10 kpc for a 40 kt liquid argon detector. Events are broken down into electron elastic scattering (red), antineutrino CC absorption (blue), and neutrino CC absorption (green). Figure from [89].

spectrum as shown in Figure 3.4.

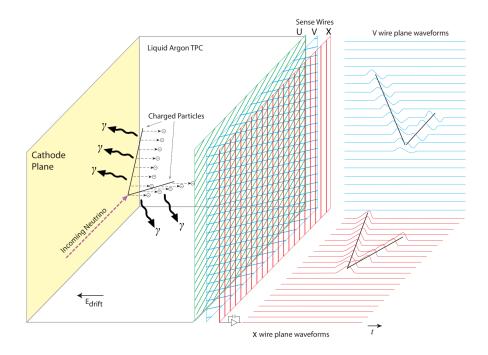
Measurement of the neutrino spectrum would provide a unique window to numerous astrophysical phenomena. In particular, it would be an important model discriminator for the dynamics of and mechanism for supernova explosions. The supernova neutrino burst is also prompt with respect to the electromagnetic signal. Observation of the burst can provide an early warning and some pointing information for astronomers. In addition, measurement of the neutrino burst provides sensitivity to the mass ordering through matter effects and to many models that extend the SM. A detailed overview of the potential physics programme can be found in [90].

# 3.3 Liquid Argon Time Projection Chamber Technology

The LArTPC is the primary detector technology for both the ND and FD of DUNE. It was introduced in 1977 by Rubbia [91], building on a proposal from Nygren [92], and first operated at scale as a neutrino detector for the ICARUS experiment in 2010 [93]. The LArTPC has since been employed at scale for the MicroBooNE experiment [94] and for ProtoDUNE-SP [95], a prototype of the DUNE FD.

# 3.3.1 Detection Principle

The operating principle of a typical LArTPC is shown in Figure 3.5. As the name suggests, the detection medium is purified LAr, necessitating a cryostat to reach  $< 87 \, \mathrm{K}$ 



**Figure 3.5:** Schematic of LArTPC operating principle. Shown with a horizontal electric field and wire plane charge readout. Photon detection system is not shown. Figure from [96].

temperatures. An incident neutrino may interact with an argon nuclei ejecting a set of final state particles from the nucleus. The charged particles will propagate through the LAr, ionising the argon atoms to produce free ionisation electrons. A uniform electric field applied using an anode and cathode plane will then drift the ionisation electrons through the LAr to the anode which is instrumented as a sensing plane. The design of the sensing plane can vary but it conventionally consists of several planes of sensing wires that the drifting electrons induce a current on as they approach. In addition to ionisations electrons, particles propagating in the LAr produce scintillation light which can be measured using photon detection systems around the detector.

The choice of LAr as a medium is driven by multiple factors. Argon is a noble element allowing ionisation electrons to drift freely through the LAr without being reabsorbed. LAr has a high electron mobility, reducing the electric field strength required to reach desired drift lengths. LAr has a low ionisation threshold. Furthermore, it has a high scintillation light yield and is transparent to its own scintillation light. The choice of using argon in its denser liquid state is driven by the requirement to maximise neutrino interaction rates.

In the conventional configuration, the sensing plane at the anode consists of three layers of wire planes. The first two planes are biased such that they are transparent to the electrons. They are called the *induction* planes. The last plane is biased such that it collects the drifting electrons. This plane is called the *collection* plane. The wires of the collection plane are strung vertically while the wires of the induction planes are angled at approximately  $\pm 45^{\circ}$  to the collection plane wires. Drifting electrons induce a

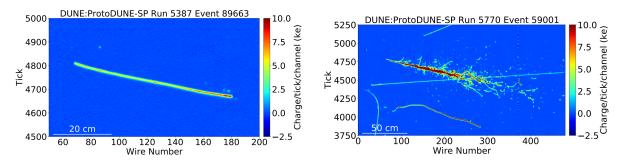
current on the wires which is digitised by an analogue-to-digital converter (ADC). The signal from each of the three wire planes corresponds to a two-dimensional ADC image in wire number and drift time. Obtaining the drift time requires measuring the time of the interaction in the detector. This is typically performed using scintillation light. Precision neutrino beam timing may also be used. The three two-dimensional wire plane images can then be combined into a single three-dimensional image of the event. With an appropriate wire pitch and drift length, a LArTPC can perform three-dimensional imaging of neutrino events with millimetre spatial resolution. Once calibrated, the ADC allows for a calorimetric energy estimation from the image. The scintillation light can also be used for complementary calorimetric information.

# 3.3.2 Particle Propagation in Liquid Argon

Neutrinos interacting in a LArTPC are reconstructed by observing the final state particles of the interaction with an argon nucleus. The path of the particles through the LAr and the energy they lose as they travel is measured using ionisation electrons and scintillation light. This allows the charged final state particles to be reconstructed accurately when they deposit all their energy in the LAr, i.e. they do not escape the detector. For neutrons in the final state this is much harder and they are typically treated as missing energy that needs to be corrected for in the neutrino energy estimation using a neutrino-nucleus interaction model. Neutrons are able to scatter with and be captured by argon nuclei which can produce scintillation light and sometimes eject a proton, but this happens on long timescales so the neutron will often escape the detector or the interaction cannot be associated with the original neutrino event.

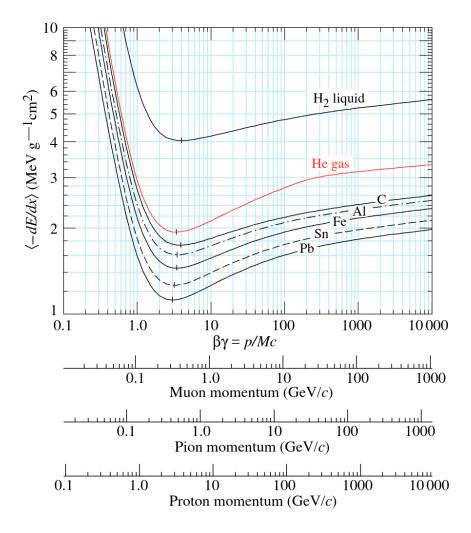
Charged particles propagating in LAr can be approximately split into two categories by their topology: track and shower. Tracks are typically produced by muons, protons, and charged pions. An example of a track most likely produced by a proton is shown in Figure 3.6a. The mean rate of energy lost in a track is well understood via the Bethe-Bloch equation [97] for moderately relativistic charged heavy particles passing through matter. Predictions for different particles are shown in Figure 3.7. It is clear that the sharp increase in energy lost at low momenta well describes the increase in deposited charge visible at the end of the proton track in Figure 3.6a. For a sufficiently long track, the distinct rate of energy loss for each particle can be used to identify the type of particle associated with a track. If a track particle does not interact inelastically and its type can be identified by other means, the Bethe-Bloch equation can also be used to estimate the initial energy by measuring the range of the track from start to stop. This is particularly useful for muon energy reconstruction.

An example of a shower can be seen in Figure 3.6b. Showers are generated by a cascading process of photons pair producing into  $e^+e^-$ , which then radiate photons via bremsstrahlung, which then pair produce into  $e^+e^-$ , and so on. This process continues

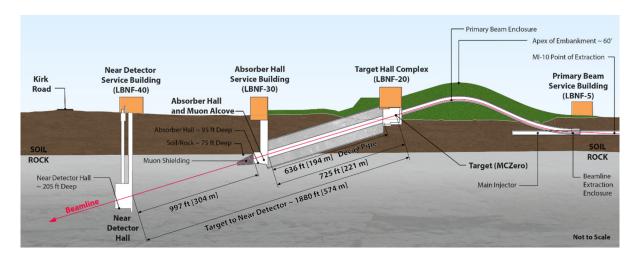


(a) Example of a track. 1 GeV proton candidate. (b) Example of a shower. 6 GeV electron candidate.

**Figure 3.6:** Event displays from ProtoDUNE-SP data showing examples of a track and shower topology. Particle beam is coming from the left. Figure from [98].



**Figure 3.7:** Mean rate of energy loss for charged particles propagating in different liquid media as a function of momentum calculated with the Bethe-Bloch equation. The rate for muons, protons, and pions is shown by adjusting the x-axis according to their masses. Figure from [99].



**Figure 3.8:** Diagram of the LBNF neutrino beamline. The beam starts on the right with extraction from the Main Injector. Figure from [101].

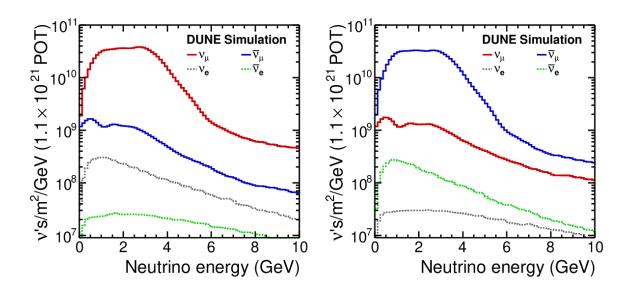
until enough energy has been lost to the LAr that energy loss of the electron via ionisation begins to dominate over bremsstrahlung. Showers are created by electrons, positrons and photons. A neutral pion will also produce showers through its decay to two photons. The inelastic nature of a shower means it is not possible to apply the Bethe-Bloch equation for energy estimation and instead an estimate is made by summing the ionisation energy visible in the shower.

# 3.4 The LBNF Neutrino Beam

The LBNF neutrino beam [100] provides the source of neutrinos for DUNE's long-baseline oscillation physics. The neutrinos are generated using a proton beam accelerated to between 60 GeV and 120 GeV by the Booster and Main Injector synchrotrons at Fermilab. Protons enter the Main Injector from the Booster in batches. Multiple batches are then extracted in spills of  $10 \,\mu s$  with an extraction cycle time of  $1.2 \, s$ .

After being extracted from the Main Injector, the proton spill is transported to the target hall. This is shown in Figure 3.8. In the target hall the protons collide with a cylindrical graphite target, producing a secondary beam of charged mesons, primarily pions. A single spill consists of  $7.5 \times 10^{13}$  protons-on-target (POT) and for energies of 120 GeV this corresponds to a beam power of 1.2 MW. With allowances for downtime, the 1.2 MW beam corresponds to  $1.1 \times 10^{21}$  POT per year. As part of DUNE's Phase II, a suite of upgrades to Fermilab's accelerator facilities known as the Proton Improvement Plan II [102] will increase the beam power to 2.4 MW.

The secondary beam of charged mesons is focused into a beam using three magnetic horns. The horns can be configured to focus positively charged particles while diverting negatively charged particles by running in forward horn current (FHC) mode and vice versa by running in reverse horn current (RHC) mode. Once focused, the mesons travel



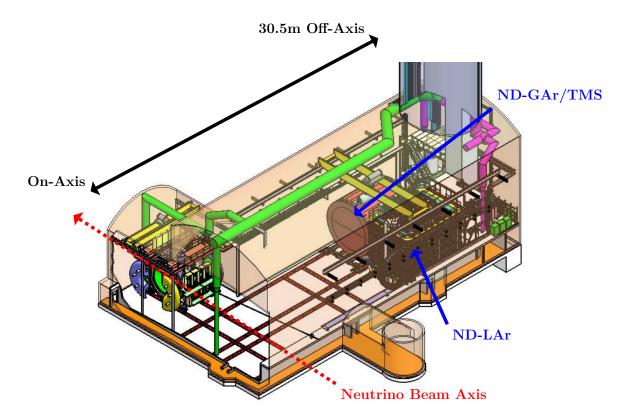
**Figure 3.9:** Composition of expected FHC (left) and RHC (right) flux spectra for the LBNF beamline time-integrated over 1 year normal running at 1.2 MW. Figure from [76].

through a 194 m decay pipe where they decay. The dominant decay modes are those of the pion and kaon to  $\mu^+\nu_{\mu}$  for FHC mode and  $\mu^-\bar{\nu}_{\mu}$  for RHC mode. At the end of decay pipe is a beam dump that absorbs any muons, protons, and undecayed mesons. The result is an approximately collimated neutrino beam dominated by  $\nu_{\mu}$  in FHC mode and  $\bar{\nu}_{\mu}$  in RHC mode. For this reason, FHC is synonymous with a neutrino beam and RHC with an antineutrino beam.

The predicted FHC and RHC neutrino flux from the LBNF beam is shown in Figure 3.9. The beam has a high muon neutrino purity but does contain some contamination. Electron neutrinos known as intrinsic  $\nu_e$  are present in the flux. These primarily come from muons decaying before the beam dump and uncommon charged kaon decays. There is also a wrong-sign component of the flux from decays of (positive)negative particles that make it past the focusing horn into the decay pipe in (RHC)FHC mode. The wrong-sign component of the flux makes up a larger proportion of the measured events in RHC mode than in FHC mode since the CC inclusive cross section for neutrinos is larger than that for antineutrinos.

# 3.5 The Near Detector

The ND utilises three detectors with distinct detector technologies to measure unoscillated neutrinos from the LBNF beam to constrain interaction model and flux systematic uncertainties. Figure 3.10 shows the layout of the ND hall. Starting upstream with respect to the neutrino beam, the first detector is ND-LAr. This is the primary LAr target of the ND, a necessity for constraining interaction model systematics relevant to the LAr detectors of the FD. Next is a downstream tracker that measures the momenta of muons



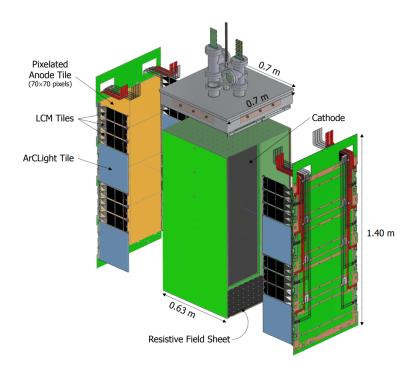
**Figure 3.10:** Diagram of ND hall with detectors in the far off-axis position. ND-GAr of Phase II (displayed) occupies the same position as TMS of Phase I (not displayed).

exiting ND-LAr. In Phase I this will be The Muon Spectrometer (TMS)<sup>1</sup>. In Phase II it is planned to upgrade TMS to ND-GAr which will, in addition to measuring exiting muon momentum, enable neutrino on argon measurements complementary to ND-LAr. Both ND-LAr and ND-GAr/TMS will be able to continuously move up to 30.5 m in the direction transverse to the neutrino beam. This is called moving off-axis. The capability of the ND to move off-axis is known as the Precision Reaction Independent Spectrum Measurement (PRISM or DUNE-PRISM). After ND-GAr/TMS, is the System for on-Axis Neutrino Detection (SAND) which will stay on-axis to measure the stability of the neutrino beam to reduce the impact of flux systematics.

# 3.5.1 ND-LAr

ND-LAr is a modular LArTPC with a fiducial mass of 67 t. The choice of a LArTPC is necessitated by the need to constrain cross section and detector uncertainties at the FD. The proximity of ND-LAr to the LBNF target hall means the detector will be subject to neutrino pile-up with  $\mathcal{O}(50)$  interactions in each beam spill. With the conventional LArTPC design discussed in § 3.3.1, it would not be possible to reliably disentangle single events from a spill lasting  $10\,\mu\mathrm{s}$  because of the comparatively slow drift time of  $\sim 1\,\mathrm{ms}$ . The ND overcomes this problem with a 3D pixel readout and a unique modular design

 $<sup>^{1}</sup>$ Following an ominous rebrand from the Temporary Muon Spectrometer



**Figure 3.11:** A diagram of the Module-0 demonstrator [106]. This is representative of the design of a single ND-LAr module which will be  $1 \text{ m} \times 1 \text{ m} \times 3 \text{ m}$ .

that segments the active volume into many small optically isolated drift regions.

The detector is divided into 35 optically isolated independent LArTPC modules in a  $5 \times 7$  grid. These comprise a detector with dimensions 5 m in the direction parallel to the neutrino beam, 7 m in the direction perpendicular, and 3 m vertically. A diagram of a single module prototype is shown in Figure 3.11. The cathode in the centre produces two 50 cm drift volumes for each module. The charge readout at each anode consists of LArPix pixel tiles [103]. These are grids of charge-sensitive self-triggering pixels that serve as anode surfaces to collect the drifted ionisation electrons. Each pixel is associated with a two-dimensional spatial coordinate and when combined with the drift time provides a natural three-dimensional representation of the event. The pixel pitch is planned to be  $< 4 \,\mathrm{mm}$ . The non-anode sides of the modules are instrumented with the light readout system which consists of both LCM [104] and ArCLight [105] tiles. Importantly, the light system will be able to combine timing and approximate spatial information to tag the time when clusters of charge are produced. When for each module this timing information is combined with the three-dimensional charge readout, ND-LAr will be able to make accurate associations of the charge within a pile-up to distinct neutrino events.

#### 3.5.2 TMS and ND-GAr

With a reasonable fiducial volume applied to the neutrino interaction vertex, the size of ND-LAr is sufficient to contain the hadronic system of all but the highest energy beam neutrino events. However, muons with momentum exceeding  $0.7 \, \text{GeV/c}$  have ranges greater than  $\sim 3 \, \text{m}$  and so will typically exit the detector. To accurately reconstruct

neutrino energy for the whole beam spectrum, the momentum of exiting muons must be measured. In Phase 1 this will accomplished with TMS. TMS is a magnetised calorimeter of alternating steel and plastic scintillator planes, much like the MINOS detector [107], that will be positioned downstream of ND-LAr. The muon momentum is measured using its range and the charge is determined via its curvature due to the magnetic field, allowing for rejection of the wrong sign background.

In Phase II, TMS will be replaced by the more capable detector ND-GAr. ND-GAr is a magnetised detector consisting of a high-pressure gaseous argon (GAr) TPC surrounded by a calorimeter. The fiducial mass of the GAr is nearly 1 ton. Electrons from charged particles ionising the argon atoms are drifted to a charge readout by an electric field parallel to the magnetic field. The curvature of the muon track in the magnetic field allows both the charge and momentum to be measured, ensuring ND-GAr fulfils its essential role as a muon spectrometer. In addition, ND-GAr will make measurements of neutrino interactions on argon independently of ND-LAr. Gaseous argon has a much lower energy threshold for charged particle tracking than LAr which will allow for better measurements of the final state particles near to the interaction vertex. This is particularly useful for constraining interaction model uncertainties that result in low energy pions and protons in the final state.

#### 3.5.3 PRISM

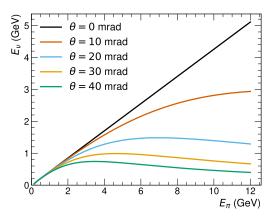
As mentioned previously, PRISM is the capability of ND-LAr and the muon spectrometer to move off-axis with respect to the neutrino beam. To understand why this is useful, we must consider the kinematics of the meson decays that produce the muon neutrino beam. The dominant decay is  $\pi \to \mu \nu_{\mu}$  which is shown in the lab frame in Figure 3.12. Applying energy and momentum conservation to this decay yields an equation for the energy of the outgoing neutrino,

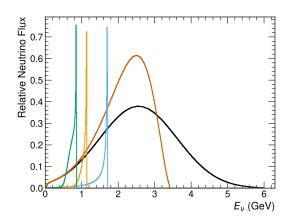
$$E_{\nu} = \frac{m_{\pi}^2 - m_{\mu}^2}{2(E_{\pi} - p_{\pi}\cos\theta)}.$$
 (3.5.1)

Assuming a perfectly collimated beam, going to off-axis positions is equivalent to sampling different decay angles  $\theta$ . Therefore, at different off-axis positions the kinematics of the pion decays that produce the neutrino beam can be significantly different as shown in Figure 3.13a. The effect this has on a neutrino flux is shown in Figure 3.13b. At greater angles  $\theta$ , corresponding to greater off-axis positions, the neutrino flux peaks at lower energies and its width reduces. This is a consequence of the maximum kinematically allowed energy reducing and thus the range of possible neutrino energies. This can be understood by considering the decay in the rest frame of the pion. Conservation laws in the rest frame yield for the neutrino energy,

$$E_{\nu} = \frac{m_{\pi}^2 - m_{\mu}^2}{2m_{\pi}} \approx 30 \text{ MeV}.$$
 (3.5.2)

Figure 3.12: Two-body decay of pion to a neutrino and muon in the lab frame.





- (a) Neutrino energy as a function of pion energy.
- (b) Expected flux from a Gaussian pion distribution. Off-axis angles have different normalisations.

Figure 3.13: Energy of the neutrino from two-body pion decay  $\pi \to \mu \nu_{\mu}$  and example of resulting neutrino flux. Shown for different neutrino-pion decay angles  $\theta$ . The decay angle of 40 mrad is equivalent to the ND in a  $\sim 23$  m off-axis position.

In the lab frame, the perpendicular component of the neutrino energy cannot exceed this maximum. Therefore, at higher decay angles the maximum neutrino energy is lower as it approaches 30 MeV at  $\theta = \pi/2$ .

The ability to access lower energy and narrower neutrino fluxes by putting detectors off-axis is well known and was exploited by both the T2K and NOvA experiments. While the ND and FD of DUNE are situated on-axis, PRISM will permit the exploitation of the off-axis effect at the ND. ND-LAr and the muon spectrometer will be movable by up to 30.5 m off-axis with a granularity of 10 cm and a sub centimetre precision. Such a movable detector was first proposed as nuPRISM for the J-PARC neutrino beamline [108]. In practice, PRISM gives the DUNE ND access to a continuous range of unique fluxes. It is planned that the ND will periodically move to different off-axis positions for data-taking to exploit this. We will discuss the impact of PRISM's off-axis capability on oscillation measurements in § 3.5.5.

## 3.5.4 SAND

The primary goal of the SAND detector is to provide a time-dependent constraint on the LBNF flux during the operation of DUNE by making on-axis beam neutrino measurements. Such a constraint is necessary since the neutrino flux depends on properties of the beamline hardware as well as on hadron production physics relevant to the interaction of protons with the target. For example, both the alignment of the focusing horns and the density of the graphite target can significantly change the flux at the ND. Over the lifetime of DUNE, many of these properties will change by unpredictable amounts. When combined with in situ beamline monitoring, SAND will allow for identification and modelling of beamline instabilities.

SAND will repurpose the magnet and electromagnetic calorimeter from the KLOE experiment [109]. The main inner target and tracking component of SAND will be a system of orthogonal straw tube trackers with a carbon target in addition to a thin LAr target.

## 3.5.5 Role in Oscillation Measurements

The measured quantities at the ND and FD are selected neutrino event rates at different reconstructed energies. Schematically, this event rate in reconstructed energy  $E_{\text{rec}}$  can be expressed at each detector as a convolution,

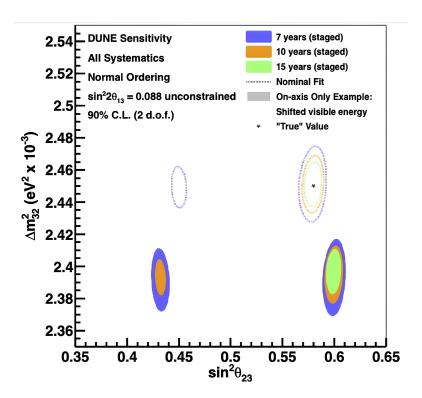
$$\mathcal{N}^{\text{ND}}(E_{\text{rec}}) = \int dE_{\nu} \, \Phi^{\text{ND}}(E_{\nu}) \times \sigma(E_{\nu}) \times D^{\text{ND}}(E_{\nu}, E_{\text{rec}}),$$

$$\mathcal{N}^{\text{FD}}(E_{\text{rec}}) = \int dE_{\nu} \, \Phi^{\text{FD}}(E_{\nu}) \times P_{osc}(E_{\nu}) \times \sigma(E_{\nu}) \times D^{\text{FD}}(E_{\nu}, E_{\text{rec}}),$$
(3.5.3)

where  $\Phi(E_{\nu})$  is the beam neutrino flux,  $P_{osc}(E_{\nu})$  is the oscillation probability for the relevant channel,  $\sigma(E_{\nu})$  is the neutrino-argon cross section, and  $D(E_{\nu}, E_{rec})$  are the detector effects composed of the efficiency and resolution. Extracting oscillation parameters requires a measurement of the oscillation probability. Due to different detector effects and the presence of oscillations at the FD, the ratio  $\mathcal{N}^{FD}(E_{rec}) / \mathcal{N}^{ND}(E_{rec})$  does not yield the oscillation probability. Instead, an accurate model of the flux, cross sections, and detector effects is required to fit the FD data for the oscillation probability.

The role of the ND is to make measurements of the unoscillated neutrino beam that inform the FD prediction for a given oscillation hypothesis. This is typically achieved by using ND data to fit a model for the flux and cross sections which is then used to make a prediction at the FD. This is especially important for the cross section model which, as discussed in §2.4, is poorly understood a priori. The high ND statistics permits strong model constraints for DUNE [110].

A challenge with using the ND in this way comes from there being no known complete neutrino-nucleus interaction model. A model with the correct form and many unknown parameters can be safely constrained by the ND. However, an incomplete model will dis-

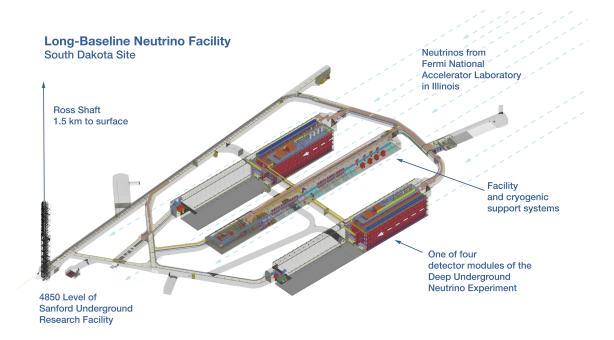


**Figure 3.14:** Allowed regions in  $\sin^2 \theta_{23} - \Delta m_{32}^2$  from fit using nominal MC (dashed) and using the mock dataset (shaded). Allowed regions drawn for different years of exposure. Figure from [110].

agree with the ND data even after its parameters have been constrained. This necessitates the addition of extra degrees of freedom to the model to force agreement with the data. When possible, these modifications will be theoretically well motivated but often will have to be empirical corrections. These corrections may result in good data-MC agreement at the ND, but are not appropriate when applied to the very different FD flux. This can incur an unknown bias in the oscillation parameters measured by the FD.

A study outlined in Chapter 4 of [110] clearly demonstrates this risk of a biased measurement. In this study, a mock dataset is generated by transferring 20% of proton kinetic energy to neutrons in the nominal simulation. This is a proxy for effects that are present in data but are not described by the interaction model used in simulation. The interaction model is then empirically reweighted to produce a good data-MC agreement for the on-axis ND measurement. Fitted oscillation parameters obtained using this reweighted interaction model for the FD data are shown in Figure 3.14. The result for the mock dataset is a clear bias in the measured oscillation parameters.

The mock dataset causes biased measurements because model corrections, especially those not well grounded theoretically, that result in good data-MC agreement for the on-axis ND measurement usually do not capture the correct relationship between true neutrino energy and reconstructed neutrino energy. PRISM can help remedy this through comparing ND data to MC at many off-axis positions with unique energy spectra. This will help identify any model corrections that would result in biased oscillation parameter

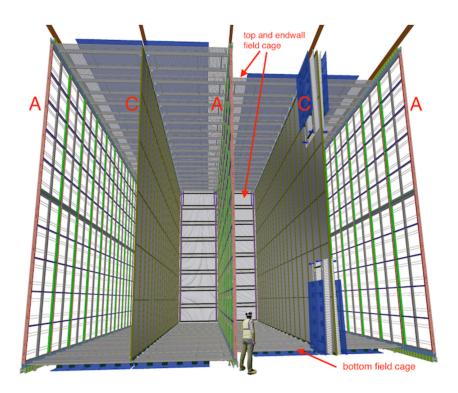


**Figure 3.15:** Diagram of the FD complex with the two Phase I detector modules shown. Figure from [79]

measurements when applied to an oscillated energy spectra. However, it will likely not be possible to produce a reasonable set of model corrections that simultaneously reproduce all relevant final state particles and their kinematics over all sampled off-axis energy spectra. In this case, PRISM can be utilised to perform an oscillation analysis that relies on linearly combining off-axis measurements to reproduce expected FD oscillated spectra. This way, the ND constraint is transferred directly to the FD along with any unknown cross section effects. This will be discussed in depth in Chapter 5.

# 3.6 The Far Detector

A diagram of the FD underground facility is shown in Figure 3.15. The facility will eventually house four detectors each contained in a cryostat with internal dimensions of 62.0 m deep along the beam direction, 15.1 m wide, and 14.0 m high. Each detector module will have a 10 kt fiducial mass. The first two LArTPC detector modules to be installed in Phase I will be a horizontal drift (HD) configuration module and a vertical drift (VD) configuration module. In Phase II a second optimised VD module will be installed along with a module of opportunity that is currently undergoing research and development. In this section, only the Phase I modules will be described with particular emphasis on the HD technology as it is most relevant to this thesis.



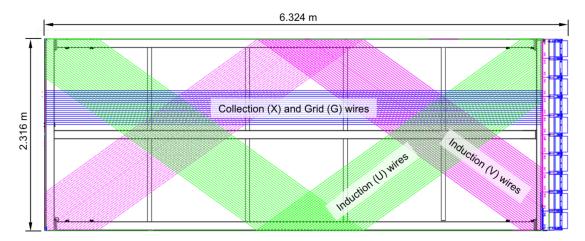
**Figure 3.16:** Design of the HD module drift regions. Anodes (A) and cathodes (C) are labelled. Figure from [111].

#### 3.6.1 Horizontal Drift

The HD detector is segmented into drift volumes by alternating anode cathode planes as shown in Figure 3.16. The maximum drift length will be  $3.5 \,\mathrm{m}$  and a  $500 \,\mathrm{V/cm}$  uniform electric field will be applied across each drift region. This corresponds to a maximum drift time of  $\sim 2 \,\mathrm{ms}$ . The detector readout is located at the anode planes which are composed of anode plane assembly (APA) units. One anode plane is made from a grid of APAs 2 high and 25 deep, totalling 150 APAs for the full detector.

A schematic of an APA is shown in Figure 3.17. The APA realises the conventional LArTPC readout discussed in §3.3.1. The wires of the collection plane, heretofore known as the Z plane, are strung vertically with a wire pitch of 4.8 mm. The two planes of induction wires U (outermost) and V (innermost) are wrapped around both sides of the APA such that they provide full coverage of the APA and cross each collection wire only once. This is achieved with a wire angle of  $\pm 35.7^{\circ}$  with respect to the collection plane wires. The induction wire pitch is 4.7 mm. The grid wires G are biased to act as a shielding plane against long range induction effects from the drifting electrons and are not read out.

The photon detection system is located in between the wire planes of the APA as shown in Figure 3.18. Silicon photomultipliers (SiPMs) are used as part of the X-ARAPUCA design [112, 113] as the photon detectors. The X-ARAPUCA is a photon trap coupled with several SiPMs. The design is economical and has a high detection efficiency.



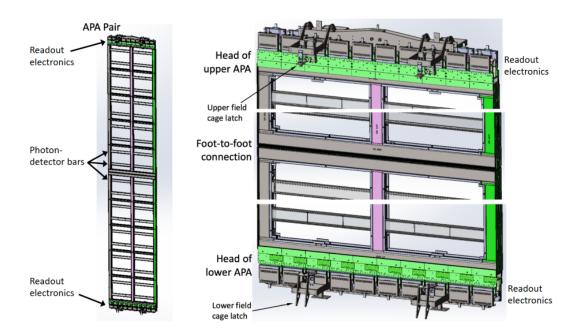
**Figure 3.17:** Schematic of a HD APA. The induction wires (U and V) are wrapped around both sides of the APA frame while the collection (X) and grid (G) wires are strung separately on both sides. Figure from [111].

The front-end readout electronics for the wires are located at the head of the APA frame as shown in Figure 3.18. These will operate submerged in LAr to take advantage of the improved gain and reduced electronics noise from the low temperatures. Noise is expected to be below the level of  $1000\,e^-$  per channel. For comparison, a minimum ionising particle corresponds to between  $20\,\mathrm{k}$  and  $30\,\mathrm{k}\,e^-$  per channel. A linear electronics response up to  $500\,\mathrm{k}\,e^-$  per channel is targeted to avoid saturation for the majority of beam neutrino events. To meet these design requirements, a 12-bit ADC is used with a 2 MHz sampling frequency. The digitised electronics response is shown in practice for ProtoDUNE-SP, a prototype HD detector, in Figure 3.19. The waveforms exhibit the expected bipolar and unipolar forms for induction and collection wires respectively with a good signal-to-noise ratio.

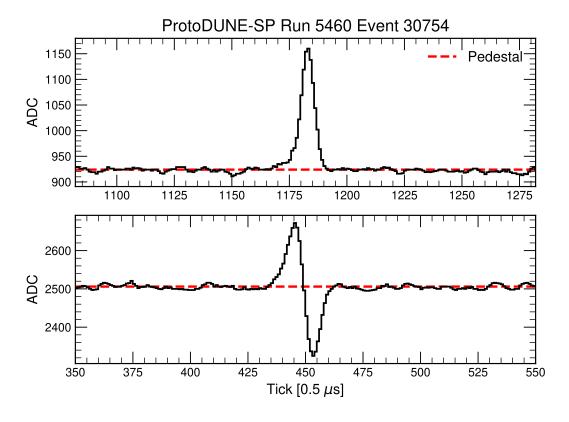
# 3.6.2 Vertical Drift

The VD detector module shares the same detection principles as the HD module but with some design differences. Notably, it utilises a vertical drift. This is accomplished with a single cathode plane that bisects the detector into a top and bottom drift region, each with a maximum drift lengths of 6.5 m. By applying a higher voltage than for the HD cathode, a uniform  $500\,\mathrm{V/cm}$  electric field is generated to drift electrons to the anodes at the top and bottom of the detector where the charge readout electronics are located.

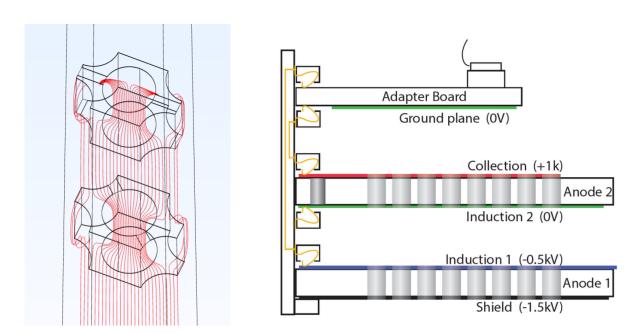
The anode planes are composed of a grid of  $3.4\,\mathrm{m} \times 3\,\mathrm{m}$  charge readout planes (CRPs). The CRPs are made of printed circuit board layers instrumented for charge readout. This is shown schematically in Figure 3.20. Appropriately biased perforations in the boards combined with layers of charge sensitive strips etched into the board produce a readout analogous to that of the HD module. The VD module will use the same X-ARAPUCA photon detection system as the HD module but installed on the vertical walls of the cryostat.



**Figure 3.18:** Schematic of a stacked APA pair as they would appear in the full HD anode plane. The photon detection bars are mounted across the width of the APA (left) and the readout electronics are located at the head of the APA (right). Figure from [111].



**Figure 3.19:** Digitised response of a single collection wire (top) and a single readout wire (bottom) for ProtoDUNE-SP data. Response is from drifted ionisation electrons most likely produced by cosmic muons.



**Figure 3.20:** Schematics of a VD CRP relevant to charge readout. Left: Field lines of the perforated printed circuit board design. Right: Printed circuit board layers of the CRP showing charge readout strips (coloured horizontal lines) and perforations (grey cylinders). Figures from [114].

# DUNE Simulation and Reconstruction

Predicting the neutrino event rate discussed in § 3.5.5 requires a full simulation chain encompassing the neutrino production at the LBNF beam through to the detector response at the near and far detector. A sequence of reconstruction algorithms is also required to reconstruct the energy and flavour of incident neutrinos. In this chapter, the simulation and reconstruction suite for the near and far detectors will be discussed. The discussion will cover the aspects of the simulation and reconstruction chain necessary to infer the properties of incident beam neutrinos from the charge they deposit in the detectors.

# 4.1 Event Generation in LAr

At both the ND and FD, the simulation chain starts by generating a neutrino interaction in LAr using a shared set of tools. This encompasses a simulation of the beam neutrino production, its interaction with an argon nucleus, the propagation and ionisation of the final state particles in the LAr, and the drift of the ionisation electrons to charge readout instrumentation.

The first step of the chain is the simulation of the LBNF beamline to generate neutrino fluxes for given horn polarities and POT. This is accomplished with the G4LBNF simulation tool [115] that conducts a detailed GEANT4 [116, 117, 118] simulation of the beamline from the primary proton beam through to the beam dump. This is the software used to generate the expected fluxes shown previously in Figure 3.9.

With the expected neutrino flux on hand, a simulation of the neutrino's interaction with the argon nuclei is performed using the GENIE generator [119]. This simulates the interaction of an incident neutrino with a nucleon followed by the nuclear effects discussed in §2.4. The output of the generator is a set of final state particles originating from an interaction vertex within the detector.

Next, the interactions of the final state particles as they propagate through the detector is simulated. This is performed with Geant4. The particles are propagated through the detector in discrete steps each with an associated energy loss. Any decays and secondary interactions that may occur during the propagation are also simulated with Geant4.

The set of energy depositions within the active volume of the detector must then be

processed into ionisation arriving at the charge readout. First, for each step through the active LAr, the number of ionisations is calculated using the ratio of the energy deposited to the average energy to produce an electron-ion pair for argon. This is subject to statistical fluctuations through a Fano factor [120] that empirically corrects for deviations from Poisson statistics due to correlations. To find the number of electrons that are free to drift the anode, a recombination factor is applied to account for electrons recombining with argon ions. This factor is computed using either Birks model [121] or the modified Box model [122]. After recombination, we are left with a number of ionisation electrons that drift towards a readout plane. The electrons are drifted in groups with the drift velocity and diffusion in the longitudinal and transverse directions simulated using measured properties of LAr at a given temperature and electric field strength. The reattachment of electrons to impurities in the LAr that may occur during the drift is simulated by applying a mean lifetime to the drifting electrons. The lifetime is measured using calibrations or with purity monitors in the LAr [123].

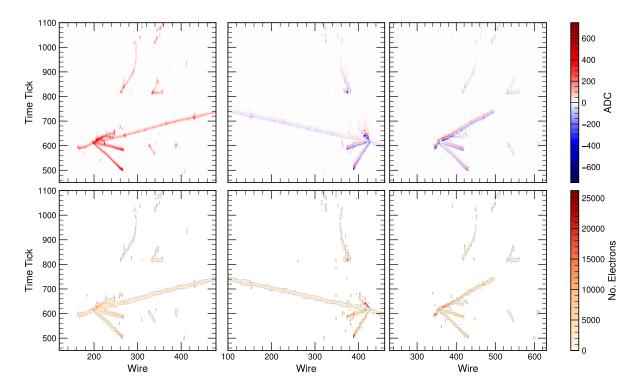
After the ionisation electrons arrive in the vicinity of the anode, the commonality between FD and ND is broken. The different charge readout at each detector necessitates different simulation and reconstruction chains.

# 4.2 The Far Detector

The FD simulation and reconstruction software is implemented in the art framework [124] and is built upon the LArSoft software suite [125]. art is an event-processing framework that ensures reproducibility and is interoperable with ROOT [126]. In this section, we will only give an overview of the software chain for the HD detector. The HD is simulated using a reduced geometry that is  $13.9 \,\mathrm{m}$  deep along the beam direction,  $7.3 \,\mathrm{m}$  wide, and  $12.0 \,\mathrm{m}$  high. This contains a single central anode plane of APAs in a  $2 \times 6$  grid and is known as the '1x2x6' geometry. The 1x2x6 geometry is used to reduce the computational cost of simulation at the expense of a small drop in the containment of beam neutrinos. Simulation of the full 10 kt geometry is currently under development and will be used for future analyses.

# 4.2.1 Detector Response Simulation

The Wire-Cell Toolkit [127] is used to simulate the response of the sense wires to drifted charge. The ionisation electrons are assumed to drift in a uniform field up to 10 cm from the wire planes where a more detailed simulation of the field in the proximity of the wire planes is performed. At this stage, the current induced on each wire from a group of electrons is calculated using a 2D field response function of a single wire to a single point-like electron. The field response is computed using the Garfield TPC drift simulation code [128] by simulating the drift paths of a single electron positioned at many discrete starting



**Figure 4.1:** Simulated detector response (top) and deconvolved charge (bottom) for an event in the Z (left), U (middle), and V (right) planes. Shown is a crop of a single APA. The simulated event is a 3.7 GeV  $\nu_{\mu}$  DIS interaction.

positions. The range of starting positions span 21 wires centred on the wire of interest and a step size of one tenth of a wire pitch is used. This generates a 2D field response for a wire that accounts for current induced by electrons incident on neighbouring wires as well as on itself.

Application of the field response function yields the current induced on each wire. Next, this is convolved with an electronics response function. This describes the amplification and shaping by the pre-amplifier, the effect of RC filters, and the digital sampling into a 12-bit ADC. Applying the electronics response function to the induced current gives the full simulated detector response to drifted ionisation electrons. An example of the detector response simulation is shown in the top row of Figure 4.1.

#### 4.2.2 Signal Processing and Hit Finding

Signal processing is the first stage of the reconstruction chain. The objective is to reconstruct the charge distribution of the incident electrons from the digitised waveform. This requires unfolding the convolution of the field response and electronics response functions to recover the number of ionisation electrons incident on each wire as a function of time. This process is known as *deconvolution*. By combining the field and electronics response into a single detector response function R(t, t'), the measured signal M(t') for a single

wire in the 1D case can be written as,

$$M(t') = \int_{-\infty}^{\infty} R(t, t') S(t) dt, \qquad (4.2.1)$$

where S(t) is the original charge signal. By assuming the detector response function to be time-invariant, R(t,t') = R(t'-t), we can take a Fourier transform of both sides to yield  $M(\omega) = R(\omega)S(\omega)$ , where  $\omega$  is an angular frequency. The desired signal is then simply

$$S(\omega) = \frac{M(\omega)}{R(\omega)}. (4.2.2)$$

Which can be then be returned to the original time domain via the inverse Fourier transform. In practice, the presence of electronics noise means the high frequency components of the noise spectrum will often be strongly amplified by (4.2.2). To address this, a software filter  $F(\omega)$  is introduced,

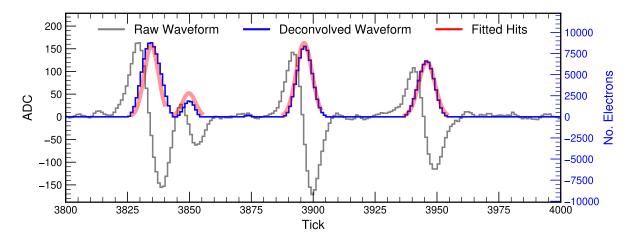
$$S(\omega) = \frac{M(\omega)}{R(\omega)} F(\omega). \tag{4.2.3}$$

To account for contributions to the measured signal of a given wire from electrons incident on neighbouring wires, (4.2.1) can be expanded to the 2D case to give a measured signal for the i-th wire,

$$M_i(t') = \int_{-\infty}^{\infty} (\dots + R_1(t'-t)S_{i-1}(t) + R_0(t'-t)S_i(t) + R_1(t'-t)S_{i+1}(t) + \dots) dt,$$
(4.2.4)

where  $S_i$  is the charge signal inside the boundaries of the *i*-th wire,  $R_0$  is the detector response function for charge incident on the wire of interest, and  $R_n$  is the detector response function for the charge incident on a wire n wires adjacent to the wire of interest. It is assumed that the detector response is translationally invariant, i.e. it is independent of the absolute position of the wire. The deconvolution then follows the same procedure as for the 1D case. An example of a deconvolved detector response is shown in the bottom row of Figure 4.1. More details on the signal processing procedure can be found in [129].

After deconvolution, the waveforms are parametrised into *hits* that are the input to downstream reconstruction algorithms. A hit finder module parametrises the deconvolved waveform by fitting it with a multiple Gaussian hypothesis. The resulting hits have a coordinate in wire number and time, a width in time, and a magnitude that can be calibrated to the energy loss of a particle in the LAr. An example of the fitted hits for a single induction wire can be seen in Figure 4.2. Due to induction wires being wrapped, an additional step is performed to determine which of the two drift volumes either side of an APA a hit is associated with. This disambiguation is performed by associating hits across all three planes.



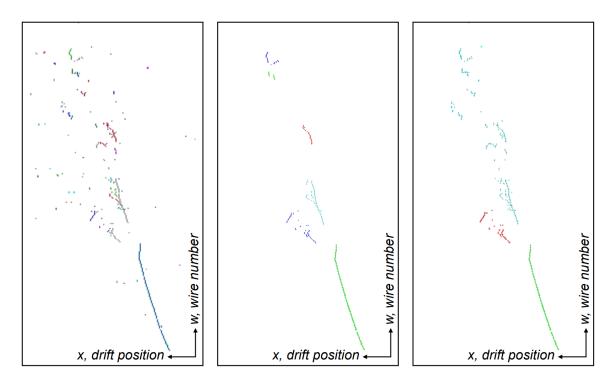
**Figure 4.2:** Example of hits on a single induction wire produced by multiple Gaussian fitting to the deconvolved waveform.

#### 4.2.3 Track and Shower Reconstruction

The high resolution and multiple 2D views of an event produced by a LArTPC allows for accurate reconstruction of the particle hierarchy of an event. This means the clustering of hits into track- and shower-like objects, the identification of primary and secondary interaction vertices, and the association of reconstructed objects with particles. The particle hierarchy summarises the interactions of the final state particles in the detector and allows the association of hits, and so energy depositions, with individual final state particles. The main event reconstruction software for the DUNE FD is the Pandora pattern-recognition software for LArTPCs [130]. Pandora employs a multi-algorithm approach and has been implemented with success on ProtoDUNE-SP [131] and MicroBooNE [132] data.

The multi-algorithm approach allows for flexibility in how Pandora reconstructs the particle hierarchy. The overall approach is as follows:

- 2D Clustering. Hits from each wire plane are considered separately and clustered based on their proximity and an assumption of a track-like topology. The algorithms at this stage are designed to prioritise purity over completeness, i.e. produce clusters that do not contain hits from more than one true particle.
- Track Reconstruction. Clusters from each of the three 2D wire plane are compared to identify triplets of clusters that belong to the same single particle. The combined independent clustering outcomes from the three planes are used to refine the clustering further. The result is a set of track-like cluster triplets, each representing a single reconstructed particle.
- Shower Reconstruction. The clusters hitherto assumed to be track-like are classified as track- or shower-like. Shower-like clusters then undergo merging in each 2D plane and are subsequently matched across the planes to produce reconstructed shower-like cluster triplets.



**Figure 4.3:** Illustration of the stages of Pandora reconstruction. The initial 2D clustering (left) is followed by track and shower reconstruction for individual particles (middle) which are then reconstructed into a particle hierarchy (right). The different colours represent distinct reconstructed clusters. Shown is a ProtoDUNE-SP beam particle data event. Figure from [131].

• Particle Hierarchy Creation. The 2D triplets are projected into 3D hits in the detector. Additional algorithms operate on the 3D hits to refine the clusters further. The reconstruction concludes by organising the resulting particles into a particle hierarchy, identifying the original final state particles and any subsequent particles generated via inelastic interactions.

Figure 4.3 illustrates the stages of this sequential particle hierarchy reconstruction.

With a reconstructed particle hierarchy, the 3D hits are parametrised into track and shower software objects that are associated with the original 2D hits. These are used as analysis inputs. The track object is a collection of trajectory points while the shower object is a cone. Both of these include calorimetric information and association maps to the original 2D hit objects.

# 4.2.4 Neutrino Signal Selection

A long-baseline oscillation analysis requires the ability to select muon neutrino or electron neutrino events at the FD. Since  $\nu_{\mu}/\nu_{e}$  interactions are discriminated using the lepton, this also requires rejecting NC interactions. The reconstructed particle hierarchy may be used to make this selection. However, the primary tool developed for making this selection is the convolutional visual network (CVN) [133]. The CVN is a convolutional neural network (CNN) trained on simulation to identify the flavour and interaction channel of

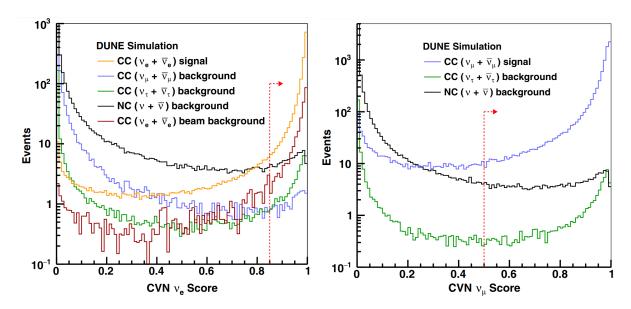


Figure 4.4: Distribution of CVN CC  $\nu_e$  (left) and CC  $\nu_\mu$  (right) score for a simulated FHC beam with 3.5 years exposure. Divided into contributions from the signal and relevant backgrounds. For the CC  $\nu_\mu$  score, backgrounds from CC  $\nu_e$  interactions are negligible and so not shown. The red arrows shows the selection cut that gets applied to the score. The distributions for the RHC beam are similar and use the same selection cuts. Figure from [133].

neutrino events from images of their interactions in the LArTPC. The input to the CVN is one  $500 \times 500$  pixel image of the hits from each of the three wire planes. The pixel maps are in wire number and time coordinates with the time coordinate being downsampled to match the spatial size of the wire pitch. The event will often need to be cropped to fit the  $500 \times 500$  pixel image. The centring of the crop is chosen using an approximation of the neutrino interaction vertex based on the known neutrino beam direction. The three views of the event are passed through independent encoder CNNs and the representations are then combined and pass through a single CNN head that predicts, among other outputs, a probability score for each neutrino flavour and for the interaction being NC.

Distributions of CVN scores are shown in Figure 4.4. The distributions show clear separation between the signal CC  $\nu_{\mu}/\nu_{e}$  and relevant backgrounds. By applying the selection cuts displayed on the figure, the CVN scores allows for selected neutrino samples with purities and efficiencies of  $\sim 90\%$ .

# 4.2.5 Neutrino Energy Estimation

The reconstruction of the incident neutrino's energy is performed by combining contributions from the leptonic and hadronic components of the interaction. The splitting of the interaction into these components uses the reconstructed particle hierarchy from Pandora.

For a CC muon neutrino interaction, the muon energy is measured in two ways depending on if the muon is fully contained in the detector or not. When the muon is contained, which is usually the case for beam neutrinos with a fiducial volume cut applied to the interaction vertex, the range of the track is used to reconstruct the energy. This follows from the principles discussed in § 3.3.2. When the muon track is not contained, the energy is estimated using the multiple Coulomb scattering (MCS) method first applied in the context of LArTPCs for the ICARUS detector [134]. MCS is the electromagnetic scattering of a charged particle with atomic nuclei as it traverses a medium. The MCS method is premised on the fact that this scattering is dependent on the momentum of the charged particle. The scatters cause perturbations in the muon track as it travels through the LAr. This can be measured by slicing the track into windows and fitting a line within each window to extract the scattering angle between subsequent windows. The scattering angles permit a fit for the muon's momentum. The reconstructed muon energies from both the range and MCS methods are made to match the simulated true muon energy using a linear MC correction.

For a CC electron neutrino interaction, the energy of the electron shower is measured by summing the charge of the collection plane hits associated with the shower. The ionisation charge is converted to energy deposited by the particle using correction factors for the electron lifetime and recombination effects discussed in §4.1. The resulting reconstructed shower energy has a linear MC correction applied to match the true electron energy in simulation.

For both electron and muon neutrino CC interactions, the energy of the hadronic system is reconstructed in the same way as for the electron using hits associated with the hadronic system. The reconstructed hadronic energy is the most sensitive to the neutrino-nucleus interaction model since the hadronic system includes contributions from neutral particles not visible to the charge readout. This missing energy has to be accounted for in the linear MC correction applied to the reconstructed hadronic energy. This folds into the reconstructed neutrino energy a strong dependence on the interaction model.

Figure 4.5 shows the energy resolutions for the discussed algorithms, with MC corrections applied, on simulated neutrino samples. The energy resolution is  $\sim 16\%$  for the muon neutrino sample and  $\sim 13\%$  for the electron neutrino sample. As expected, the resolution for the hadronic component of the interaction is significantly worse than as for the lepton.

# 4.3 The Near Detector

The heterogeneity of the ND and the novel pixel readout design of ND-LAr means the development of a complete simulation and reconstruction chain is somewhat more challenging than as for the FD. In this section, we will give an overview of the ND simulation and reconstruction available at the time of this thesis and relevant to the work within it. This comprises a ND facility geometry with ND-GAr, a full detector simulation of ND-LAr, and a parametrised reconstruction of neutrino events using ND-LAr in conjunction with ND-GAr. A full reconstruction chain for ND-LAr, the downstream tracker, and

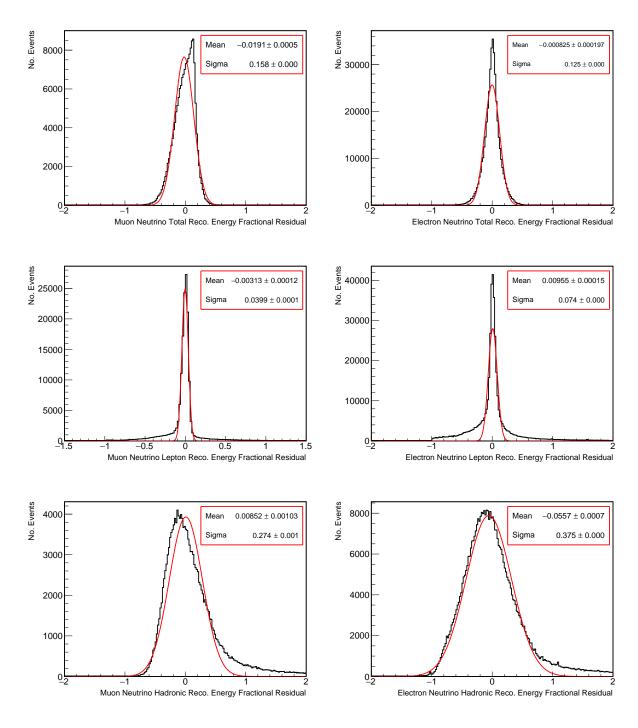


Figure 4.5: Energy resolutions for simulated FHC beam  $\nu_{\mu}$  (left) and  $\nu_{e}$  (right) with containment cuts applied. Shown for the reconstructed neutrino energy (top), reconstructed lepton energy (middle), and reconstructed hadronic energy (bottom). A Gaussian is fitted to each and the fit parameters displayed. For the  $\nu_{e}$  events, the neutrino energy resolution is better than the naive sum of the lepton and hadronic resolutions due to an anticorrelation between the two. This anticorrelation, which arises from hadronic energy being misassigned to the electron shower and vice versa, partially cancels in the sum taken when obtaining the neutrino energy.

matching between the two is currently under development using both a Pandora and an end-to-end deep-learning approach [135]. The ND software is not implemented in the art framework and does not make use of the LArSoft software suite as the FD does. Instead, ND software is currently framework-free with a miscellany of file formats.

## 4.3.1 Detector Response Simulation

For ND-LAr, the current induced on pixels at the anode and the subsequent electronics response is simulated with the larnd-sim library [136]. A pixillated charge readout requires a factor  $\mathcal{O}(10\text{--}100)$  more readout channels than a set of wire planes with the same pitch. This increased channel count motivates the parallelisation of detector simulation algorithms. For this reason, many of larnd-sim's algorithms are implemented on GPUs.

Calculating the current induced on each pixel follows a similar procedure to that of the FD described in § 4.2.1. The ionisation electrons are drifted uniformly until they come within 0.5 cm of the anode where a 3D field response function, the dimensions being time and two pixel coordinates, is used to find the induced current on each pixel. The field response is pre-calculated by simulating the drift paths of a single electron in a grid of starting points. The starting points have a 0.33 mm separation in both pixel directions and for each pixel the grid encompasses the 8 nearest neighbour pixels. The field response function is then used to calculate the induced current on the pixels from drifted ionisation charge. The induced current calculation for a pixel due to drifted charge from each GEANT4 step, a straight line of deposited energy converted into ionisation electrons, is simplified by discretising the step into many point-like charges and accounting for diffusion by making a number of random 3D perturbations to each discrete point.

Next, the electronics response of the self-triggering pixel pads is simulated. At each pixel the signal from collected electrons is amplified by a charge-sensitive amplifier. As the signal on the pad accumulates, the output voltage of the amplifier grows. If it reaches a set discrimination threshold, digitisation of the output voltage into an 8-bit ADC is triggered and the amplifier is reset. An ADC count representing the magnitude of the signal has then been recorded and the accumulated signal at the pixel discarded. Noise associated with the trigger reset, the discrimination threshold, and uncorrelated noise are included in the simulation. For each simulated event the resulting data *packets* consisting of an ADC count, a timestamp, and a channel ID are stored in an array in the HDF5 format [137].

An example detector response simulated using larnd-sim is shown in Figure 4.6. A simple reconstruction of the drift coordinate produces a voxelised image of the underlying event with a structure in the drift direction particular to the self-triggering pixel readout design.

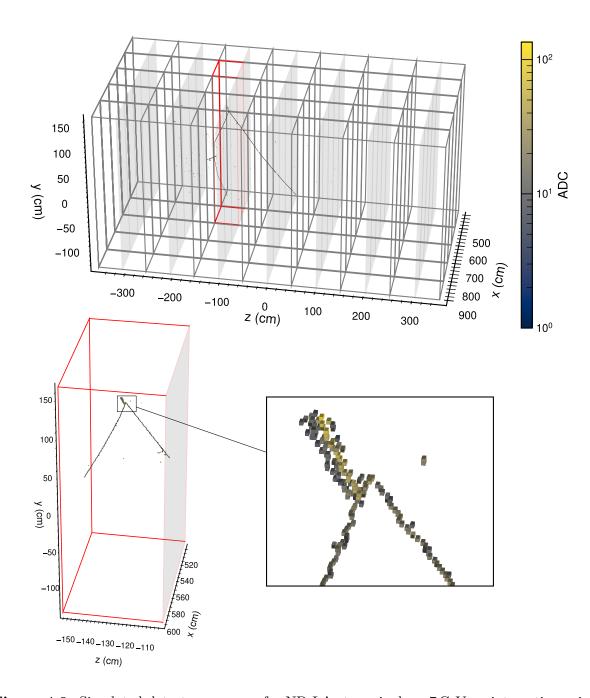


Figure 4.6: Simulated detector response for ND-LAr to a single  $\sim 7\,\mathrm{GeV}~\nu_{\mu}$  interaction using larnd-sim. The event within the full detector (top) with cathode planes (shaded fixed z), anode planes (unshaded fixed z), and module boundaries is displayed. The event within a single drift volume (bottom left) and a zoom on the interaction vertex (bottom right) is also shown. The drift coordinates (z) are obtained using the true interaction time. The drift coordinate is downsampled by a factor of 10 for visibility.

#### 4.3.2 Parametrised Reconstruction

A parametrised reconstruction of ND-LAr and ND-GAr is used as a substitute for a full reconstruction chain that is still under development. The ND parametrised reconstruction uses true event information from GENIE and GEANT4 to approximate the outputs of a realistic neutrino event reconstruction. This is applied to simulated neutrino events with interaction vertices in the fiducial volume of ND-LAr to facilitate realistic long-baseline oscillation analysis studies that require ND simulation.

The reconstruction of the muon depends on which volumes it propagates through during the Geant4 simulation of the ND facility. For CC muon neutrino events there are four possibilities:

- Tracker matched. The muon travels more than 50 cm through the active GAr volume of ND-GAr. In this case, the muon is matched perfectly to the neutrino event in ND-LAr, the sign is reconstructed perfectly, and the muon energy is reconstructed with a resolution of 2%.
- Electromagnetic calorimeter reconstructed. The muon travels more than 5 cm through the electromagnetic calorimeter that surrounds the TPC of ND-GAr and stops within it. In this case, the muon is matched perfectly to the neutrino event in ND-LAr, the sign is reconstructed perfectly, and the muon energy is reconstructed with a resolution of 10%.
- Contained in LAr. The muon stops in the active volume of ND-LAr. In this case, the charge reconstruction is only attempted for RHC mode with an efficiency of 75% based on Michel electron tagging. No charge reconstruction is implemented for the FHC beam since the antineutrino background to the neutrino signal is considered negligible. The muon energy is reconstructed with a 5% resolution.
- Exiting. The complement of the three previous possibilities. The muon cannot be properly reconstructed in this case and an energy estimate is made based on the range of the longest of the muon and pion tracks inside the active volume of ND-LAr.

For events not tracker matched or electromagnetic calorimeter reconstructed, there are probabilities that the event will be misreconstructed as an electron neutrino or a NC interaction. Misreconstruction as an electron neutrino may occur when there is a neutral pion in the final state and misreconstruction as a NC interaction when the muon track is shorter than 1 m.

The reconstruction of CC electron neutrino interactions depends only on the energy of the electron. An efficiency of 0% at  $300\,\mathrm{MeV}$  rising to 100% at  $700\,\mathrm{MeV}$  is used to decide if the electron is correctly reconstructed. If misreconstructed, the event is treated as a NC interaction. If correctly reconstructed, the electron energy is obtained with a resolution

that depends on the true energy as  $0.03 + 0.1/\sqrt{E_e}$ . These implementations of efficiency and energy resolution are designed to approximate the performance of FD simulation.

True NC interactions are usually correctly identified as NC. There is a possibility of being misidentified as a CC electron neutrino interaction when there is a neutral pion in the final state and as a CC muon neutrino interaction when there is a charged pion in the final state with  $a > 1 \,\mathrm{m}$  track.

The reconstruction of the hadronic energy is simply the sum of true deposited energy of all hadrons and their progeny within the active volume of ND-LAr. No energy thresholds or resolutions are applied. It is assumed that the hadronic energy deposits can be perfectly associated with their respective particle type and the reconstructed hadronic energy is split into contributions from different particle types. This approximation does not reflect the expected performance of full ND-LAr reconstruction which will have a worse hadronic energy resolution.

The parametrised reconstruction also reports the sum of the hadronic energy in a collar region of ND-LAr defined as the outer 30 cm of active volume. This is called the *hadronic veto*. It is used to veto events that are likely to have hadronic energy deposited outside of the detector. A veto energy of 30 MeV is typically used.

# PRISM Oscillation Analysis

5

By modelling the flux of neutrinos from the LBNF beam at different off-axis positions of the near detector, measurements from off-axis positions can be linearly combined to produce a near detector spectrum that very closely matches the oscillated far detector spectrum for any set of oscillation parameters. This PRISM capability enables an oscillation analysis where a prediction of the oscillated far detector spectrum is made via linear combinations and is fitted to far detector data to extract the oscillation parameters. The methodology of such an analysis will be discussed in this chapter.

The methodology for the oscillation analysis described in this chapter was developed by DUNE collaborators prior to the work described in this thesis. The development of the analysis is detailed in the thesis [138] which performs the first demonstration of a full long-baseline PRISM oscillation analysis.

## 5.1 Overview

A diagram outlining the steps to construct a prediction of the FD event rate through linear combination of off-axis ND event rates is shown in Figure 5.1. The ND data is used directly to make a prediction at the FD rather than being used to tune an interaction model. The oscillation analysis then proceeds by fitting the FD prediction, which is a function of the oscillation parameters, to the FD data. This approach to an oscillation analysis, as opposed to using the ND data to constrain model parameters, is known as extrapolation. The extrapolation approach is well-known and was taken for the NOvA oscillation analysis where simulation was used to convert measured ND  $\nu_{\mu}$  event rates to expected FD event rates [74]. The important distinction of the DUNE-PRISM extrapolation is that the linear combination of many off-axis measurements produces a FD prediction with an underlying neutrino energy spectrum that closely matches the underlying oscillated neutrino energy spectrum of the FD data. This ensures significant cancellation of many energy-dependent model uncertainties that would not be possible with access to only a single flux at the ND. PRISM can be used to generate a FD prediction for the four oscillation analysis channels of DUNE:  $\nu_{\mu}$  disappearance,  $\bar{\nu}_{\mu}$  disappearance,  $\bar{\nu}_{e}$  appearance, and  $\bar{\nu}_{e}$  appearance.

The PRISM oscillation analysis necessarily depends on the modelling of beam flux at different off-axis positions in order to calculate coefficients for the linear combination.

In addition, some amount of interaction model dependence may also enter the analysis through elements of the extrapolation procedure that rely on MC corrections. We will see how this occurs and how it can be minimised in the following sections.

## 5.2 Monte Carlo Data

The analysis is constructed from and applied to ND and FD simulation. The former will be referred to as MC and the latter, which is subject to statistical fluctuations, will be referred to as data. The MC is used to realise the steps outlined in Figure 5.1 which are then applied to ND data to produce a prediction which can be compared to FD data. Once the analysis has been validated, the effects of systematic uncertainties evaluated, and we have waited patiently for some number of years, it can be applied to real data from the ND and FD.

Unless stated otherwise, the ND and FD exposure in this chapter is assumed to be 3.5 years each for FHC and RHC mode with a beam power of 1.2 MW. Four FD modules are assumed for a FD exposure of 336 kt-MW-Yrs.

#### 5.2.1 Near Detector

A  $4 \,\mathrm{m} \times 2 \,\mathrm{m} \times 3 \,\mathrm{m}$  fiducial volume is defined within ND-LAr that must contain the true neutrino vertex for an event to be selected. The  $4 \,\mathrm{m}$  wide fiducial volume is then divided into 8 off-axis bins each 50 cm wide. As the ND physically moves to new off-axis positions, each off-axis bin of the fiducial volume will correspond to a new absolute off-axis position. In this way, the PRISM analysis is binned by off-axis position with a bin width of 50 cm over the full range of off-axis positions the ND may take. The true neutrino interaction vertex is used for this binning.

The relative exposure at each off-axis bin for a proposed schedule of off-axis detector positions over the course of a year is shown in Figure 5.2. This is the schedule used in this analysis. The ND MC does not follow this specific run plan and so each off-axis bin is weighted to match the relative exposure. The MC can then be scaled for any absolute exposure.

The current of the magnetic focusing horns is nominally set to 293 kA, but will run for one week per year at a special horn current of 280 kA with the detector in the on-axis position. The lowering of the current reduces the ability of the magnetic horns to focus high-energy mesons on-axis. This results in a dip in the neutrino flux in the region of 3–5 GeV which has been demonstrated to be beneficial to the PRISM linear combination. The 280 kA horn current data is included in the analysis in a single 4 m wide bin centred on-axis to suppress statistical fluctuations associated with the short exposure. It is treated as an additional off-axis bin. Further discussion of the ND off-axis schedule and of the special horn current run can be found in Chapter 4 of [110].

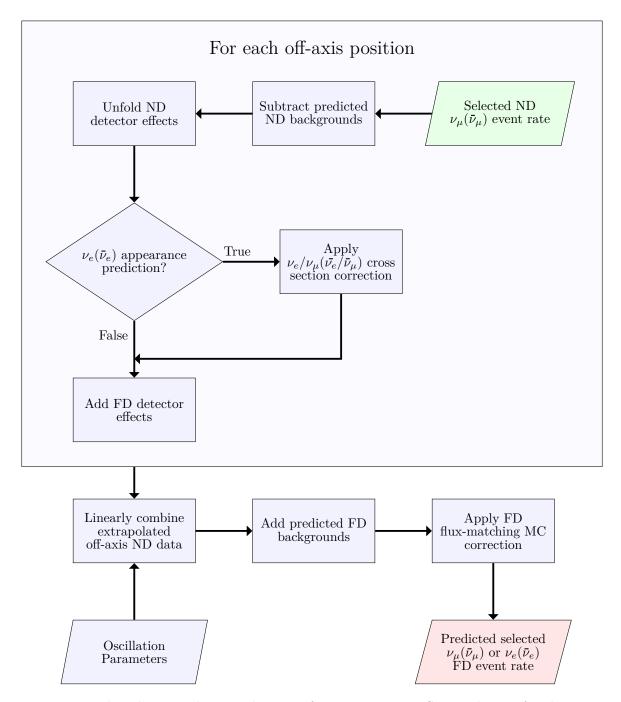
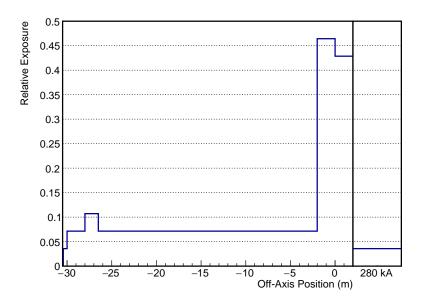


Figure 5.1: Flow diagram showing the steps for creating a PRISM prediction for the FD  $\nu_{\mu}$ ,  $\nu_{e}$ ,  $\bar{\nu}_{\mu}$ , or  $\bar{\nu}_{e}$  event rate. The steps start at the green input box and end at the red output box.



**Figure 5.2:** Relative exposure at each off-axis bin for the schedule of ND off-axis positions used in the analysis. Relative exposure also shown for the on-axis 280 kA horn current run. The same schedule is used for FHC and RHC beam.

As discussed in §4.3.2, full event reconstruction is not available for the ND. Using the reconstructed energies from the parametrised reconstruction, the ND analysis variable is defined as

$$E_{\text{rec}}^{\text{vis}} = E_{\text{rec}}^{\mu} + E_{\text{rec}}^{\pi^{\pm}} + E_{\text{rec}}^{\pi^{0}} + E_{\text{rec}}^{p} + E_{\text{rec}}^{\text{other}}.$$
 (5.2.1)

The reconstructed energy associated with neutrons is omitted as it is not expected to be accessible by a full event reconstruction. The analysis variable represents the reconstructed *visible* energy at the ND. In analogy to (5.2.1), a visible energy is used as the truth analysis variable. It is defined as

$$E_{\text{true}}^{\text{vis}} = E_{\text{true}}^{\mu} + T_{\text{true}}^{\pi^{\pm}} + T_{\text{true}}^{\pi^{0}} + T_{\text{true}}^{p} + T_{\text{true}}^{\text{other}}, \tag{5.2.2}$$

where  $T_{\text{true}}$  refers to true kinetic energy. Visible reconstructed and true energy variables are used to minimise the interaction model dependence of unfolding.

With reference to the parametrised reconstruction discussed in §4.3.2, CC  $\nu_{\mu}(\bar{\nu}_{\mu})$  events for FHC(RHC) mode are selected using the following criteria:

- True interaction vertex is within ND-LAr's fiducial volume.
- Muon is either tracker matched or contained in LAr.
- Event is reconstructed as a CC muon neutrino interaction.
- Muon charge is reconstructed as negative(positive).

• True hadronic energy in ND-LAr's hadronic veto region does not exceed 30 MeV.

Applying this selection produces the ND data spectra shown in Figures 5.3a and 5.3b. The selection can also be applied to the ND MC to yield the predicted background rates. The NC current background, shown in Figures 5.3c and 5.3d, is caused by tracks from charged pions being misidentified as muon tracks. Only NC interactions with sufficient energies to produce a charged pion track of sufficient length contribute to this background. The wrong-sign background, shown in Figures 5.3e and 5.3f, occurs when wrong-sign neutrinos have their charge incorrectly reconstructed. This may happen when the muon track is contained in the LAr rather than tracker matched. For clarity, one-dimensional selected ND data and predicted background event rates at two distinct off-axis detector positions are shown in Figure 5.4.

#### 5.2.2 Far Detector

MC samples for the FD are generated independently of any specific oscillation hypothesis. For the unoscillated FHC flux of beam muon neutrinos and intrinsic electron neutrinos three samples are generated: a non-swap sample where  $\nu_{\mu} \to \nu_{\mu}$  and  $\nu_{e} \to \nu_{e}$ , an electron-swap sample where  $\nu_{\mu} \to \nu_{e}$  and  $\nu_{e} \to \nu_{\tau}$ , and a tau-swap sample where  $\nu_{\mu} \to \nu_{\tau}$  and  $\nu_{e} \to \nu_{\mu}$ . The same three samples are generated for the RHC flux. These samples can then be combined to produce an oscillated FD flux for any oscillation hypothesis.

The FD MC is generated in the 1x2x6 geometry. The full reconstruction chain described in § 4.2 is used except with a 1D field response for the detector response and signal processing steps. The reconstructed neutrino energy,

$$E_{\rm rec}^{\nu} = E_{\rm rec}^{\rm lep} + E_{\rm rec}^{\rm had}, \tag{5.2.3}$$

is used as the analysis variable.

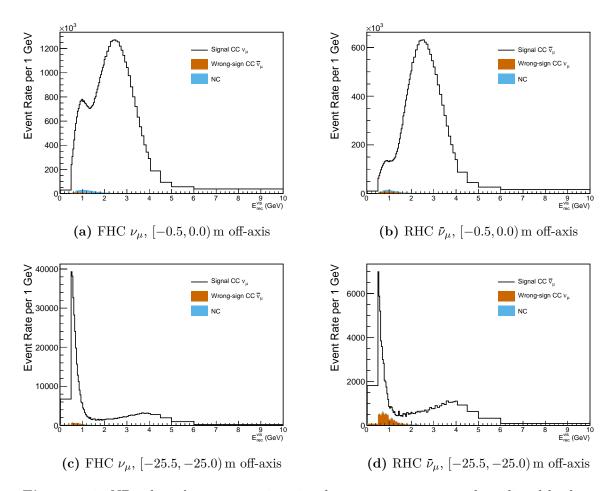
For a neutrino event to be selected in the FD, its true neutrino vertex must be within a fiducial volume defined as 50 cm from the active volume boundaries except for the downstream boundary where it is 150 cm. It is assumed the hadronic and leptonic components of the event can be properly reconstructed if the vertex is within this fiducial volume. The CVN score for the event must pass the cuts outlined in § 4.2.4 to be selected as either CC  $\nu_{\mu}(\bar{\nu}_{\mu})$  or CC  $\nu_{e}(\bar{\nu}_{e})$ .

The resulting selected FD spectra for a set oscillation hypothesis can be seen in Figure 5.5. The predicted background rates are also plotted. The RHC mode spectra have more significant contributions from backgrounds than the FHC mode spectra. This is due to the enhanced cross section for neutrinos over antineutrinos increasing the wrong-sign background while reducing the signal event rate for the RHC mode spectra. There is an intrinsic electron neutrino background which contributes significantly to the electron neutrino appearance spectra as it is impossible to distinguish from the signal. The

Figure 5.3: ND selected muon neutrino signal event rate spectra and predicted backgrounds.

(f) Predicted RHC  $\bar{\nu}_{\mu}$  wrong-sign background.

(e) Predicted FHC  $\nu_{\mu}$  wrong-sign background.



**Figure 5.4:** ND selected muon neutrino signal event rate spectra and predicted backgrounds for two example off-axis positions. For (a) and (b), the dip at 1–1.5 GeV is due to events where the muon exits ND-LAr and then does not have sufficient residual range to be tracker matched or electromagnetic calorimeter reconstructed by ND-GAr.

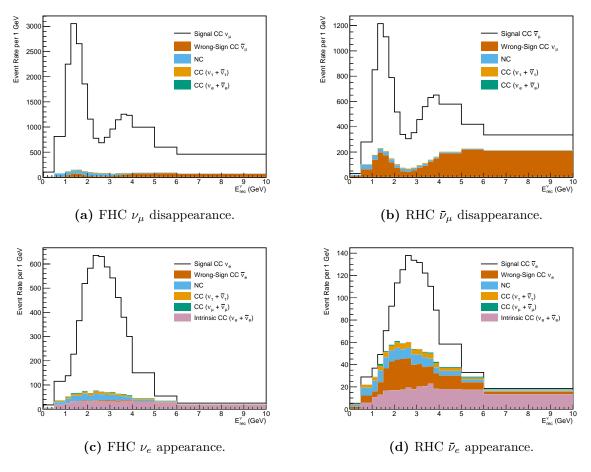


Figure 5.5: FD selected neutrino event rate spectra with contributions from predicted backgrounds. Shown for selected  $\nu_{\mu}$ ,  $\bar{\nu}_{\mu}$ ,  $\nu_{e}$ , and  $\bar{\nu}_{e}$  which comprise the four analysis channels. NuFIT 4.0 oscillation parameters [139] are assumed.

other background contributions are from NC interactions and non-signal CC neutrino interactions.

# 5.3 Extrapolation of Off-Axis Measurements

To compare a prediction composed of ND measurements with FD data, differences between the ND and FD must be corrected for. The prediction must be in the  $E_{\rm rec}^{\nu}$  variable (5.2.3) of the FD rather than the original  $E_{\rm rec}^{\rm vis}$  variable (5.2.1) of the ND. To make this correction, differences in backgrounds, efficiency, and resolution between the detectors must be accounted for as outlined in the loop of Figure 5.1.

#### 5.3.1 Near Detector Backgrounds

First, the predicted ND background event rate in each bin of  $E_{\text{rec}}^{\text{vis}}$  must be subtracted from the data. The background rates are shown in Figure 5.3 and the subtraction is applied to each off-axis bin. The resulting spectra are then treated as being pure  $\nu_{\mu}(\bar{\nu}_{\mu})$  event rates.

#### 5.3.2 Near Detector Effects

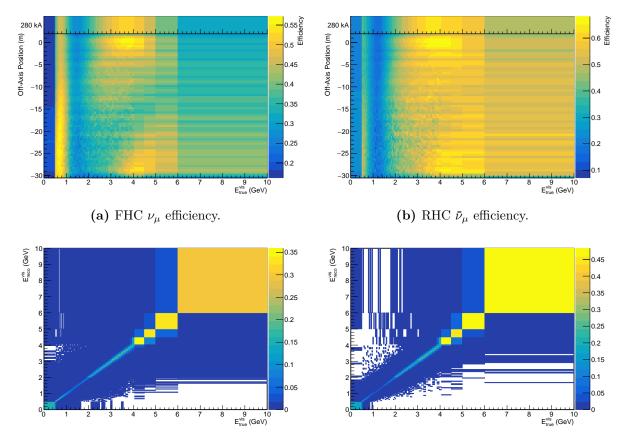
To remove detector effects is to map back to the true variable  $E_{\text{true}}^{\text{vis}}$  (5.2.2). This requires correcting for both selection efficiency and detector resolution. The selection efficiency for the i-th off-axis bin,  $\epsilon_i$ , is defined as the ratio of the total selected signal events with the total true signal events in the i-th off-axis bin. The selection efficiency must be calculated as function of  $E_{\mathrm{true}}^{\mathrm{vis}}$  since the reconstructed energy is not a good variable for events that are not selected. The selection efficiency at each off-axis position is shown in Figures 5.6a and 5.6b. The selection efficiencies have a number of notable features. The drop in efficiency at  $\sim 1 \text{ GeV}$  is where the muon will typically be energetic enough to exit ND-LAr but not to enter and be reconstructed by the downstream tracker. The rise and fall in efficiency that repeats every 2 m, resulting in horizontal bands in the figure, is due to each absolute off-axis position being composed of contributions from the 8 relative off-axis bins the fiducial volume is split into. Off-axis bins nearer the edge of the fiducial volume are less likely to pass the hadronic veto selection criterion. The more gradual change in the efficiency with off-axis position is due to using a true visible energy rather than a true neutrino energy variable. The topology of the interactions in a fixed  $E_{\text{true}}^{\text{vis}}$  bin is dependent on the underlying true neutrino energy spectrum which in turn depends on the off-axis position. Lastly, the RHC efficiency is generally higher than the FHC since antineutrino interactions tend to be more elastic so that the muon is more likely to reach the downstream tracker and there is less hadronic activity to contain.

The detector resolution can be described as a smearing matrix  $\mathbf{M}^{\text{ND}}$  that maps from true to reconstructed visible energy for the selected signal events,

$$E_{\text{rec},i}^{\text{vis}} = \sum_{j} M_{ij}^{\text{ND}} E_{\text{true},j}^{\text{vis}}.$$
 (5.3.1)

 $\mathbf{M}^{\mathrm{ND}}$  is constructed using ND MC under the assumption that the detector resolution is independent of off-axis position. First, an unnormalised smearing matrix is generated by binning the on-axis ND MC in  $E_{\mathrm{true}}^{\mathrm{vis}}$  and  $E_{\mathrm{rec}}^{\mathrm{vis}}$ . The selection efficiency is then included in this smearing matrix by enforcing that, for each off-axis position, the sum over the reconstructed energy bins for each true energy bin is equal to the efficiency associated with that true energy, i.e.  $\sum_i M_{ij}^{\mathrm{ND}} = \epsilon_j$ . We then have, for each off-axis position, an efficiency-normalised smearing matrix that describes the detector effects at the ND. The smearing matrices for the on-axis position are shown in Figures 5.6c and 5.6d.

To remove the detector effects, each off-axis spectrum is unfolded using the associated smearing matrix. The naive approach is to simply invert the matrix  $\mathbf{M}^{\text{ND}}$ . However, statistical fluctuations would produce large bin-to-bin fluctuations in the inverted matrix. This would in turn lead to very large variances in the solution for  $\vec{E}_{\text{rec}}^{\text{vis}}$ , resulting in large statistical errors and a strong dependence on systematic uncertainties that change the form of  $\mathbf{M}^{\text{ND}}$ . To remedy this, the unfolding is performed with regularisation to reduce



matrix.

(c) FHC  $\nu_{\mu}$  efficiency-normalised on-axis smearing (d) RHC  $\bar{\nu}_{\mu}$  efficiency-normalised on-axis smearing matrix.

Figure 5.6: MC components of ND detector effects correction procedure. ND selection efficiencies for each off-axis position and 280 kA horn current (top) and ND smearing matrices normalised to the on-axis position selection efficiency (bottom).

the variance at the cost of a small bias.

Tikhonov regularisation [140, 141] is applied by formulating the unfolding as the minimisation

$$\min_{E_{\text{true}}^{\text{vis}}} \left( \| \mathbf{M}^{\text{ND}} \vec{E}_{\text{true}}^{\text{vis}} - \vec{E}_{\text{rec}}^{\text{vis}} \|^2 + \| \mathbf{\Gamma} \vec{E}_{\text{true}}^{\text{vis}} \|^2 \right), \tag{5.3.2}$$

where  $\|\cdot\|$  denotes the Euclidean norm and  $\Gamma$  is a regularisation matrix.  $\Gamma$  is chosen to be the second-derivative finite difference operator,

$$\Gamma = \tau \begin{pmatrix}
1 & -2 & 1 & 0 & \dots & 0 & 0 & 0 \\
0 & 1 & -2 & 1 & \dots & 0 & 0 & 0 \\
0 & 0 & 1 & -2 & \dots & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & \dots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & 0 & \dots & 1 & -2 & 1 \\
0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & \dots & 0 & 0 & 0
\end{pmatrix},$$
(5.3.3)

where  $\tau$  is a parameter controlling the strength of the regularisation. The effect of the regularisation is to penalise curvature in the solution, producing a smooth unfolding. The parameter  $\tau$  is optimised to minimise the amount of bias introduced to the solution while still suppressing the variance. By differentiating with respect to  $E_{\text{true},i}^{\text{vis}}$ , it can be shown that (5.3.2) is solved by

$$E_{\text{true},i}^{\text{vis}} = \sum_{j} D_{ij} E_{\text{rec},j}^{\text{vis}}, \tag{5.3.4}$$

where  $D_{ij}$  are elements of the unfolding matrix **D** given by

$$\mathbf{D} = \left( \left( \mathbf{M}^{\text{ND}} \right)^{\text{T}} \mathbf{M}^{\text{ND}} + \mathbf{\Gamma}^{\text{T}} \mathbf{\Gamma} \right)^{-1} \left( \mathbf{M}^{\text{ND}} \right)^{\text{T}}.$$
 (5.3.5)

A comparison of unfolding with Tikhonov regularisation to a simple numerical matrix inverse can be seen in Figure 5.7. The fluctuations between neighbouring bins are strongly suppressed by the introduction of regularisation.

The unfolding matrices **D** are then used to correct for the ND detector effects at each off-axis position. Since the unfolding has the effect of migrating events across energy bins, a covariance matrix must be introduced into the analysis to account for correlations between bins. A diagonal covariance matrix for each off-axis position is defined and propagated through the unfolding.

#### 5.3.3 $\nu_e/\nu_\mu$ Cross Section Ratio

The PRISM analysis builds predictions from unoscillated beam  $\nu_{\mu}(\bar{\nu}_{\mu})$  data. For the  $\nu_{e}(\bar{\nu}_{e})$  appearance analysis this means a correction for the difference between the  $\nu_{\mu}(\bar{\nu}_{\mu})$  and

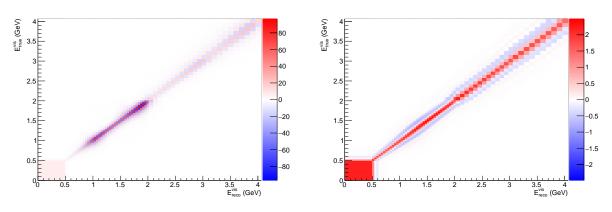


Figure 5.7: Comparison of unfolding matrices for matrix inverse (left) with Tikhonov regularised unfolding using  $\tau = 0.1$  (right). Both are for the smearing matrix in Figure 5.6c. A reduced energy range is used for visualisation.

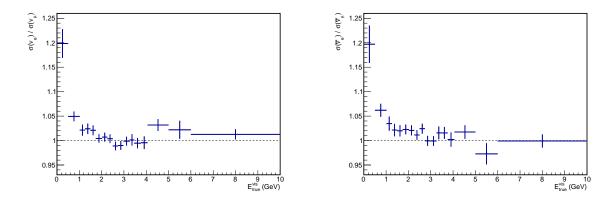


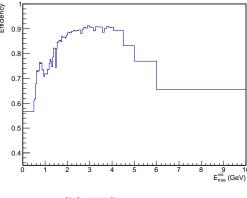
Figure 5.8:  $\nu_e/\nu_\mu$  (left) and  $\bar{\nu}_e/\bar{\nu}_\mu$  (right) cross section ratios. Error bars show statistical uncertainty from MC spectra used for the ratio. Coarser binning is used for visualisation.

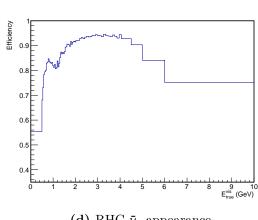
 $\nu_e(\bar{\nu}_e)$  cross sections must be applied. These cross section ratio corrections,  $\sigma(\nu_e)/\sigma(\nu_\mu)$  and  $\sigma(\bar{\nu}_e)/\sigma(\bar{\nu}_\mu)$ , are derived from MC as a function of  $E_{\rm true}^{\rm vis}$  so that they can be applied to each unfolded off-axis spectrum.

A muon neutrino and an electron neutrino event rate  $E_{\text{true}}^{\text{vis}}$  spectrum are generated with an identical flux. The cross section ratio is made by dividing the true signal  $\nu_{\mu}(\bar{\nu}_{\mu})$  event rate with the true signal  $\nu_{e}(\bar{\nu}_{e})$  event rate. The absence of detector effects means the fluxes cancel and the desired cross section ratio is obtained. The resulting cross section ratios can be seen in Figure 5.8. For an appearance analysis, each unfolded off-axis ND spectrum is weighted by this cross section ratio.

#### 5.3.4 Far Detector Effects

To complete the extrapolation to the FD reconstruction variable  $E_{\rm rec}^{\nu}$ , we need to introduce FD detector effects to the unfolded off-axis spectra. This follows a similar procedure to §5.3.2 except without the need to unfold. The selection efficiency at the FD is given as a function of  $E_{\rm true}^{\rm vis}$  by taking the ratio of selected signal events, i.e. events that pass



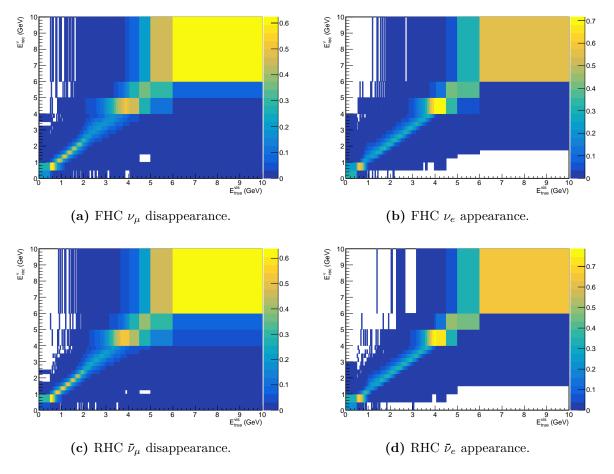


(c) RHC  $\bar{\nu}_{\mu}$  disappearance. (d) RHC  $\bar{\nu}_e$  appearance. **Figure 5.9:** FD MC selection efficiencies for the four analysis channels.

9 1 E<sub>true</sub> (GeV)

the neutrino flavour identification cut on the CVN's outputted score, with the total true signal events. This is done for each of the four analysis channels. The selection efficiencies can be seen in Figure 5.9.

A FD smearing matrix, M<sup>FD</sup>, is constructed from simulation. It maps from true visible energy to reconstructed neutrino energy for the selected signal events, analogous to (5.3.1). A separate smearing matrix is made for each of the four analysis channels and their associated signal. As with the ND smearing matrix, the selection efficiency is included in the  $\mathbf{M}^{\mathrm{FD}}$  matrices by enforcing that the sum over the reconstructed energy bins for each true energy bin is equal to the selection efficiency at that true energy. With this normalisation,  $\mathbf{M}^{\mathrm{FD}}$  can be used to smear each unfolded off-axis ND spectrum to the selected signal event rate  $E_{\rm rec}^{\nu}$  spectrum for the desired analysis channel. The covariance matrix introduced for the ND unfolding stage is further propagated through the application of the FD smearing matrix. The normalised FD smearing matrices for each of the four analysis channels can be seen in Figure 5.10. The FD smearing matrices are considerably less diagonal than their ND counterparts in Figures 5.6c and 5.6d. This is due to the FD MC using a full reconstruction chain rather than a parametrised reconstruction. In a future iteration of the analysis with access to ND MC that utilises a full reconstruction



**Figure 5.10:** FD MC smearing matrices for the four analysis channels normalised to the selection efficiency for the associated analysis channel.

chain, it is expected that the ND smearing matrices will be less diagonal.

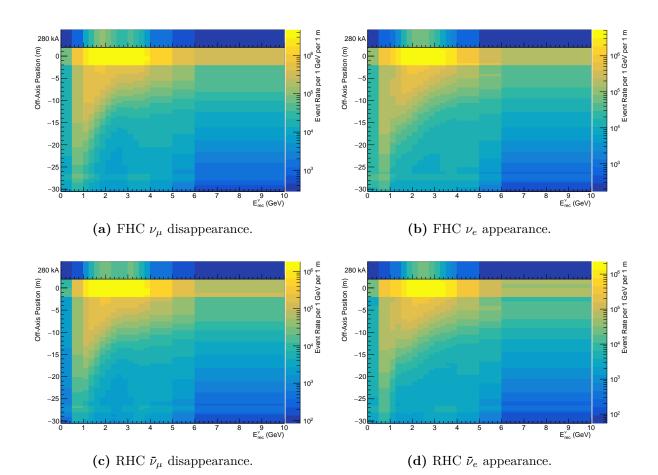
The histograms in Figure 5.11 show the unfolded ND off-axis spectra after the FD smearing matrix has been applied. They show the cumulative effects on the original spectra, shown in Figures 5.3a and 5.3b, of the procedures outlined in this section.

## 5.4 Generating a PRISM Prediction

With each off-axis ND spectrum extrapolated to the FD analysis variable they can now be linearly combined to approximate a single oscillated spectrum. Following some further corrections to this spectrum, we will then be left with our final FD prediction.

#### 5.4.1 Linear Combination

The extrapolated off-axis ND data event rates, as displayed in Figure 5.11, are represented as a matrix  $\mathbf{N}^{\text{data}}$  with elements  $N_{ij}^{\text{data}}$  corresponding to the *i*-th off-axis position and the *j*-th reconstructed neutrino energy bin. We then seek to linearly combine the rows of  $\mathbf{N}^{\text{data}}$  in such a way that they form an oscillated FD signal event rate spectrum for a given oscillation hypothesis,  $F_i^{\text{data}}$ . The coefficients of the linear combination are derived



**Figure 5.11:** ND event rate spectra after ND detector effects unfolding and FD detector effects smearing. The appearance spectra also have the relevant cross section ratio applied.

from MC in a way that is entirely independent of the neutrino-nucleus interaction model, requiring accurate modelling of the neutrino beam only. They are calculated using true signal neutrino event rates from the ND and FD MC as a function of true neutrino energy. Since the interaction model is identical at both detectors, its effect cancels out and the calculation is invariant to the choice of interaction model. It would be equivalent to use the true signal neutrino fluxes, the event rates are used instead for convenience.

Analogously to  $\mathbf{N}^{\text{data}}$ , MC is used to generate a matrix of ND off-axis true signal event rates,  $\mathbf{N}^{\text{MC}}$ . The target oscillated FD true signal event rate is also generated from MC and is denoted  $\vec{F}^{\text{MC}}$ . Calculation of the linear combination coefficients is then the task of solving  $\vec{F}^{\text{MC}} = \mathbf{N}^{\text{MC}}\vec{c}$  for  $\vec{c}$ . Since  $\mathbf{N}^{\text{MC}}$  are  $\nu_{\mu}(\bar{\nu}_{\mu})$  event rates, the appearance channel must generate  $\vec{F}^{\text{MC}}$  using  $\nu_{\mu}(\bar{\nu}_{\mu})$  event rates to ensure cancellation of the interaction model.  $\vec{F}^{\text{MC}}$  for the appearance channel is  $\nu_{\mu}(\bar{\nu}_{\mu})$  event rates in the oscillated  $\nu_{e}(\bar{\nu}_{e})$  shape. This is done only for the purpose of calculating  $\vec{c}$ .

For the same reasons as for the ND unfolding discussed in § 5.3.2,  $\vec{c}$  is solved for using Tikhonov regularisation. Analogously to (5.3.2), the coefficients are obtained through the minimisation

$$\min_{\vec{c}} \left( \| \mathbf{N}^{MC} \vec{c} - \vec{F}^{MC} \|^2 + \| \mathbf{\Gamma} \vec{c} \|^2 \right). \tag{5.4.1}$$

The regularisation matrix is chosen to be the first-derivative finite difference operator,

$$\Gamma = \tau \begin{pmatrix}
1 & -1 & 0 & 0 & \dots & 0 & 0 & 0 \\
0 & 1 & -1 & 0 & \dots & 0 & 0 & 0 \\
0 & 0 & 1 & -1 & \dots & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & \dots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & 0 & \dots & 1 & -1 & 0 \\
0 & 0 & 0 & 0 & \dots & 0 & 1 & -1 \\
0 & 0 & 0 & 0 & \dots & 0 & 0 & 0
\end{pmatrix}.$$
(5.4.2)

By penalising gradients in the solution, this choice of regularisation matrix encourages flat solutions where adjacent coefficients have similar magnitudes. The solution to the minimisation (5.4.1) is

$$\vec{c} = \left( \left( \mathbf{N}^{\text{MC}} \right)^{\text{T}} \mathbf{W} \mathbf{N}^{\text{MC}} + \mathbf{\Gamma}^{\text{T}} \mathbf{\Gamma} \right)^{-1} \left( \mathbf{N}^{\text{MC}} \right)^{\text{T}} \mathbf{W} \vec{F}^{\text{MC}}, \tag{5.4.3}$$

where  $\mathbf{W}$  is a diagonal matrix introduced to weight down the residuals associated with energies below 0.5 GeV and above 5 GeV so that the coefficients prioritise accuracy in the intermediate energy ranges. This is desirable because the high and lower energies are hard to reproduce with the ND spectra and they are less important for DUNE oscillation sensitivity.

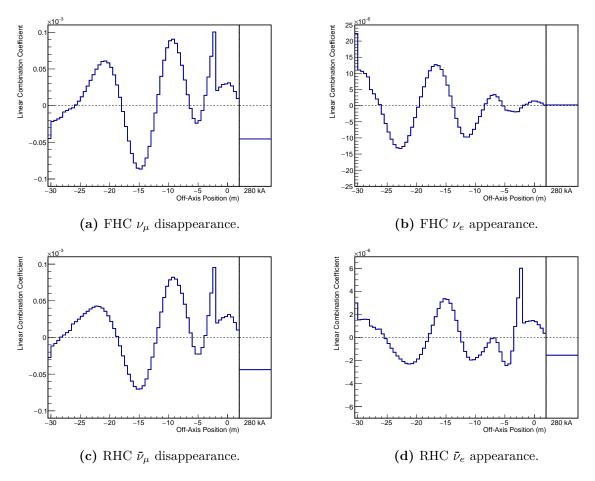


Figure 5.12: ND off-axis linear combination coefficients calculated using MC for each analysis channel. Each coefficient  $c_i$  is plotted at the *i*-th off-axis position corresponding to the row it multiplies in  $\mathbf{N}^{\mathrm{MC}}\vec{c}$ . NuFIT 4.0 oscillation parameters [139] are assumed.

The calculated linear combination coefficients for each of the four analysis channels is shown in Figure 5.12. The coefficients for the desired analysis channel can be then be applied to the unfolded ND event rates  $\mathbf{N}^{\text{data}}$  to yield a prediction of the oscillated FD signal event rate spectrum,  $\vec{F}^{\text{data}}$ , for the set of oscillation parameters used to make  $\vec{F}^{\text{MC}}$ . The covariance matrices from the extrapolated off-axis ND event rates are summed with a weighting of the associated linear combination coefficient squared to produce a single covariance matrix for the linearly combined FD prediction.

#### 5.4.2 Far Detector Backgrounds

The linear combination produces a prediction for the FD signal event rate only. We must add the predicted FD backgrounds, as shown in Figure 5.5, to the linearly combined spectrum. The result is a prediction of the total expected event rate at the FD.

#### 5.4.3 Flux-Matching Correction

Given sufficient ND statistics, the accuracy in which the linear combination can reproduce a desired oscillated flux is limited by two factors. One is that the off-axis positions only form an approximate basis for fluxes in the range of the beam neutrino energy, meaning that a perfect match in the linear combination is not expected. This is especially prevalent in the high energy tails of the oscillated spectra where the ND off-axis fluxes provide limited coverage. Another is the use of Tikhonov regularisation which introduces a small bias to the solution. The resulting inaccuracies in the prediction of the oscillated flux are corrected for using MC to derive a flux-matching correction.

The flux-matching correction is derived using the fractional difference between the linear combination and the target oscillated spectrum,  $\mathbf{N}^{\mathrm{MC}}\vec{c}$  and  $\vec{F}^{\mathrm{MC}}$  from § 5.4.1 respectively. Since this is a function of true neutrino energy, it is converted to an additive correction in reconstructed neutrino energy using the FD MC binned in true neutrino energy and reconstructed neutrino energy. This is then applied to the prediction of the oscillated spectrum from the linear combination alongside the FD backgrounds.

#### 5.4.4 Predictions

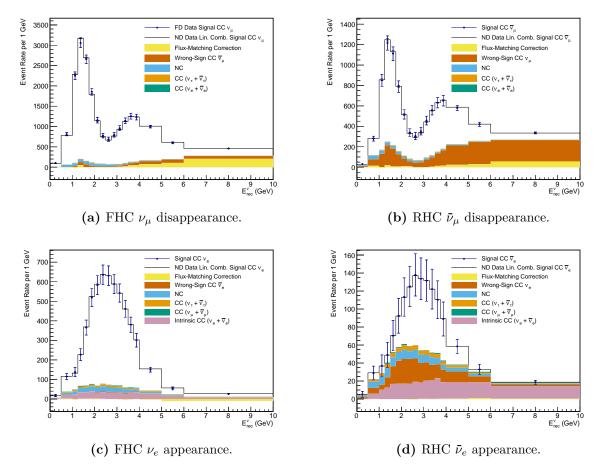
The final FD oscillated prediction is the composition of extrapolating the ND off-axis data to FD reconstructed neutrino energy, linear combination of the off-axis positions along with an associated flux-matching correction, and the addition of predicted FD backgrounds. The final predicted FD spectra are compared to the FD data for a set oscillation hypothesis in Figure 5.13. The PRISM prediction matches the data very well but is not a perfect match. This is a result of the regularisation in the unfolding procedure introducing a small bias. Although not currently implemented in the analysis, this may be corrected for using MC analogously to the flux-matching correction.

Changing the oscillation hypothesis used to create a PRISM prediction provides a new set of linear combination coefficients and FD background predictions. When applied to off-axis ND data, this produces a FD prediction for the new oscillation hypothesis. This allows the PRISM technique for generating a FD prediction to be used in an oscillation analysis.

### 5.5 Sensitivities

To fit the PRISM FD prediction,  $\vec{F}$ , to data,  $\vec{D}$ , the following chi-squared function is defined as a goodness of fit,

$$\chi^2 = (\vec{F} - \vec{D})^{\mathrm{T}} \mathbf{V}^{-1} (\vec{F} - \vec{D}), \tag{5.5.1}$$



**Figure 5.13:** Predicted FD oscillated spectra made with the PRISM method compared to FD data. The error bars show statistical error. The error on the prediction (black line) is the square root of the diagonal elements of the covariance matrix associated with the extrapolation. NuFIT 4.0 oscillation parameters [139] are assumed.

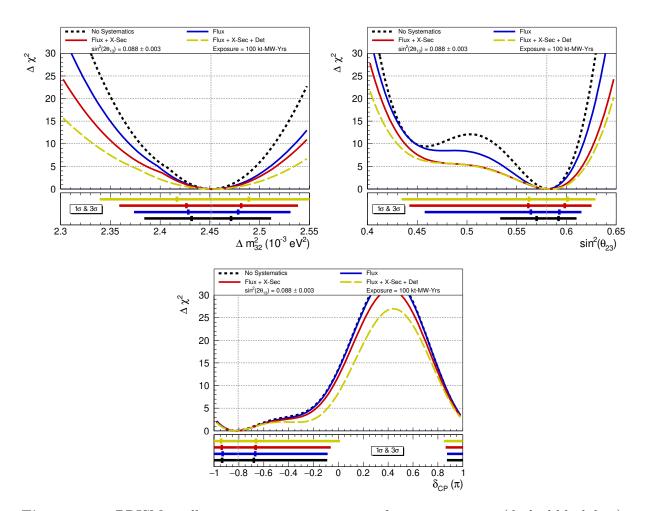
where  $\mathbf{V}$  is the combined covariance matrix of the FD data and the PRISM linear combination of the ND data. The FD prediction is a function of the oscillation parameters and a set of nuisance parameters controlling systematic uncertainties. A scan of the oscillation parameter(s) of interest is performed and at each point the  $\chi^2$  is minimised with respect to all other oscillation parameters and the nuisance parameters. Penalty terms are included in the  $\chi^2$  for the nuisance parameters and any oscillation parameters not being studied. These terms use a prior distribution to penalise the deviation of a parameter from its central value in units of the uncertainty. The sensitivities presented are from Asimov studies [142] where the ND and FD data are nominal MC with true oscillation parameters from NuFIT 4.0 [139] with normal ordering assumed. Since linearly combined ND MC is being fitted to FD MC, rather than FD MC to itself, the minimum of the  $\chi^2$  function is generally not zero. For this reason, the quantity  $\Delta\chi^2 = \chi^2 - \chi^2_{\min}$ , where  $\chi^2_{\min}$  is the smallest value of  $\chi^2$  at any of the fit points, is used to define confidence intervals. More details in the fitting procedure can be found in Chapter 6 of [138].

The systematic uncertainties are grouped into three categories: flux, cross section, and detector. The flux uncertainties arise from modelling of secondary hadrons produced when protons interact with the graphite target, and from uncertainties on the properties of LBNF components relevant to neutrino production, such as the magnetic horns and decay pipe. Cross section uncertainties are derived by varying parameters in the GENIE generator and applying weights to parametrise effects not modelled in GENIE. Detector uncertainties are implemented as variations in the energy scales and resolutions of both the ND and FD. More details on the implementation of systematic uncertainties can be found in Chapter 5 of [138].

#### 5.5.1 Four-Channel Oscillation Analysis

A four-channel oscillation analysis is performed using a joint fit of the  $\nu_{\mu}$ ,  $\bar{\nu}_{\mu}$ ,  $\nu_{e}$ , and  $\bar{\nu}_{e}$  predictions. The sensitivities with systematic uncertainties for  $\Delta m_{32}^2$ ,  $\sin^2 \theta_{23}$ , and  $\delta_{\rm CP}$  are shown in Figure 5.14. The true parameter values lie at the centre of the  $1\sigma$  intervals as expected. The constraints on the oscillation parameters are reasonably competitive with current measurements despite the sensitivities only being for a fraction the expected full exposure for DUNE.

The sensitivities show the impact of different types of systematic uncertainty. The flux uncertainty has a significant effect on the calculation of the linear combination coefficients which is the dominant cause of the reduction in sensitivity associated with flux uncertainties [138]. This is intrinsic to the PRISM extrapolation and can only be reduced with improved modelling of the neutrino beam flux. Cross section systematics also result in a significant sensitivity reduction. Since the linear combination is independent of cross section systematics, this is caused by the MC components of the analysis used to correct for different detector effects and backgrounds. It is possible to suppress this sensitivity

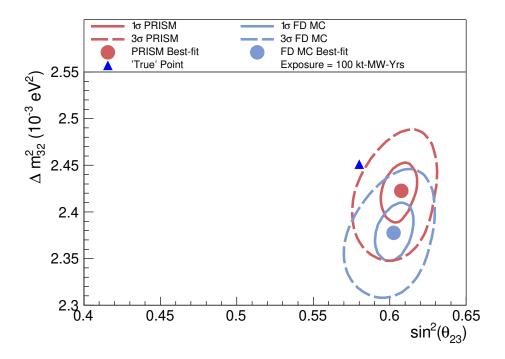


**Figure 5.14:** PRISM oscillation parameter sensitivities for no systematics (dashed black line), flux systematics (solid blue line), flux and cross section systematics (solid red line), and full systematics (dashed yellow line). The dashed vertical line shows the true value of the oscillation parameter of interest. Figure from [138].

reduction through the data-driven extrapolation techniques that will be discussed in § 5.6. The effect of detector systematics is due to the different detector designs of the ND and FD resulting in a set of uncorrelated detector systematics. Data-driven extrapolation techniques may be able to lessen their impact on the sensitivities but there will always be a considerable degradation that is irreducible.

#### 5.5.2 Missing Proton Energy Fake Data

The important advantage of using PRISM for an extrapolation analysis, as opposed to using the ND to constrain an interaction model, is that the validity of the resulting oscillation parameter constraints is much less dependent on mismodeling of the neutrino-nucleus interaction. The effect of mismodeling the interaction model on an on-axis constrained model oscillation analysis was discussed in § 3.5.5. Using a missing proton energy fake dataset where 20% of the proton energy is transferred to neutrons produced a significant bias in the measured oscillation parameters as shown in Figure 3.14. This fake data study



**Figure 5.15:** Allowed regions in  $\sin^2 \theta_{23} - \Delta m_{32}^2$  from fit to missing proton energy fake data. Shown for fit of PRISM prediction to FD fake data (red) and fit of FD MC to FD fake data (blue). Systematic uncertainties are not accounted for in either fits. Figure from [138].

is repeated for the PRISM extrapolation, the resulting sensitivity is shown in Figure 5.15. The PRISM fit is compared with a fit of FD MC to the FD fake data which is equivalent to the on-axis analysis used in Figure 3.14 without any ND constraint.

Although the PRISM fit is less biased than the fit to FD MC, the fake data still produces a significant bias in the measurement. This is due to MC components of the PRISM extrapolation being dependent on the interaction model. To study this, the ND selection efficiency and FD wrong-sign background predictions are generated using the fake data rather than the MC, i.e. the idealised case where these components of the analysis are entirely data-driven. This removes the effect of mismodeling of the proton energy from these components of the PRISM extrapolation. The result of this is shown in Figure 5.16 where the bias seen in Figure 5.15 is eliminated. This demonstrates that the PRISM analysis is capable of including unknown cross section effects in the ND constraint, therefore negating the risk of mismodeling the neutrino-nucleus interaction, if data-driven methods for the extrapolation can be developed. Such methods would reduce the dependence of the extrapolation on an interaction model.

## 5.6 Improvements to Extrapolation Method

The detector effects correction component of the PRISM extrapolation introduces a considerable amount of interaction model dependence into the analysis via the generation of smearing matrices and selection efficiencies using MC. This is antithetical to the aim

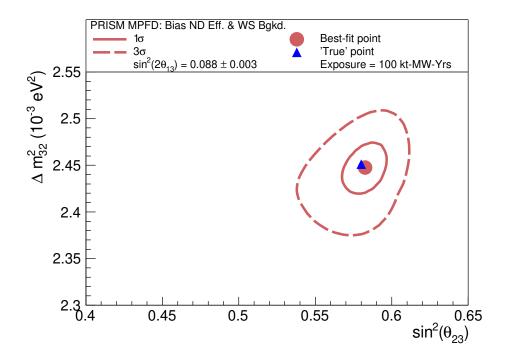


Figure 5.16: Allowed regions in  $\sin^2 \theta_{23} - \Delta m_{32}^2$  from fit of PRISM prediction to missing proton energy fake data. Full systematic uncertainties are accounted for. The ND efficiency and FD wrong-sign background corrections are derived from the fake data rather than nominal MC and their associated systematic uncertainties are not allowed to vary in the fit. Figure from [138].

of the PRISM oscillation analysis to be largely unaffected by any potential mismodeling of the neutrino-nucleus interaction. To address this, new methods for extrapolating the selection efficiencies and detector resolutions are under development and will be discussed in this section.

#### 5.6.1 Geometric Efficiency Correction

As discussed in § 5.5.2, the use of a MC correction for ND selection efficiency introduces significant dependence on the interaction model. At the ND the selection depends on the containment of the neutrino event, a hadronic veto must be satisfied and the muon must be tracker matched or contained. As shown in Figures 5.6a and 5.6b, the selection efficiencies are low and non-uniform at all energies and off-axis positions. This results in a large MC correction being applied that depends quite strongly on the interaction model since the containment at a given energy is dictated by the final state particles and their momenta. For this reason, a geometric efficiency correction is being developed to replace the MC correction.

The geometric efficiency correction approximates the probability for each individual selected ND data event to be selected again had it occurred somewhere else in the fiducial volume. The reciprocal of this probability is then used as an efficiency correction factor for that event. For each neutrino interaction in the fiducial volume, it is equally likely that

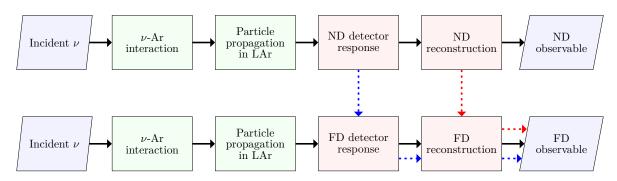
the same interaction would have occurred with the final state particles rotated around the beam axis or with the vertex elsewhere in the fiducial volume within the y-z plane, z being the detector coordinate most closely aligned with the beam axis, of the original vertex. The latter statement is based on the assumption that the change in flux in the y and z directions across the fiducial volume of ND-LAr is negligible<sup>1</sup>. Using this, the energy deposited in ND-LAr by each selected data event is subject to many random and uniform rotations around the beam axis and translations of the vertex to other points on the y-z plane. For each of these random throws, a hadronic and leptonic efficiency factor is calculated. The hadronic efficiency factor is 1 if the event passes the hadronic veto and 0 otherwise. The leptonic efficiency factor is the output of a neural network trained to predict the probability of a muon being tracker matched or contained based on its initial kinematics. By taking a product of these two efficiency factors at each throw and averaging the result over  $\mathcal{O}(1000)$  throws, an estimate of the probability that the initial selected data event would be selected again is returned.

The geometric efficiency correction allows for a efficiency correction factor for ND data events that is data-driven, i.e. it is a function of the final state particles present in the ND data rather than the MC. In its current stage of development, the true energy depositions in ND-LAr and the true neutrino vertex are used to make the random throws. In the future, these will be replaced with the 3D ND-LAr detector response and the reconstructed neutrino vertex. It is expected that when implemented in the analysis it will provide an accurate prediction for the selection efficiency. A small MC correction will likely be required to correct for imperfections in the method. The geometric efficiency correction will greatly reduce the interaction model dependence introduced into the extrapolation via the ND selection efficiency correction.

#### 5.6.2 Data-Driven Near-to-Far Translation

With the geometric efficiency correction implemented, a MC correction would still be required for the ND resolution, shown in their efficiency-normalised form in Figures 5.6c and 5.6d, and for the FD efficiency and resolution, shown in Figure 5.10. Although for the fake data study discussed in § 5.5.2 the MC based ND efficiency is the dominant source of bias, a considerable amount of interaction model dependence is introduced into the analysis via corrections for the detector resolutions and FD efficiency. Furthermore, the inclusion of full ND reconstruction will make the relation between the reconstructed variable and the true neutrino energy more complex and more dependent on the final state particles of the interaction than it is with the current parametrised reconstruction. This may increase the effect of cross section uncertainties on a MC correction. This motivates the development of a near-to-far translation that predicts relevant FD reconstructed vari-

<sup>&</sup>lt;sup>1</sup>For the off-axis positions of the ND, shifts in y and z result in up to 1.7 mrad and up to 0.3 mrad of off-axis effect respectively while shifts in x result in up to 7.0 mrad of off-axis effect.



**Figure 5.17:** Schematic of the processes that produce analysis variables from an incident neutrino at the ND (top) and FD (bottom). The ND and FD are LArTPCs with different instrumentation. This means some processes are identical (green boxes) while others are distinct (red boxes). The dashed arrows represent proposed near-to-far translations: the response translation (blue) and the reconstruction translation (red).

ables for each ND data event. When combined with the geometric efficiency correction, a near-to-far translation would also remove the necessity for an unfolding procedure in the analysis. The development of a near-to-far translation is the focus of this thesis and will be introduced in this section.

A comparison of the processes that generate reconstructed observables, energy estimations and selection criteria, from a neutrino incident at the ND or FD is shown in Figure 5.17. A MC correction for the detector effects that is generated by simulating this chain of processes back to the incident neutrino and applied to the binned ND observable of choice is highly dependent on the interaction model used in the simulation. However, as highlighted in the figure, the neutrino-nucleus interactions at the ND are from the same distribution as those at the FD. Therefore, the characteristics of ND data events can strongly inform the expected FD observables if properly utilised. A MC correction that is a function of a single reconstructed variable does not accomplish this. A datadriven correction that is a function of the ND data event in its entirety, and so of the underlying neutrino interaction, can. This is the purpose of a near-to-far translation. As indicated in Figure 5.17, a near-to-far translation may proceed in two ways: via a response translation or via a reconstruction translation. The response translation uses the ND detector response, i.e. what is measured in the detector before any processing, to predict the FD detector response to the same underlying neutrino interaction. The FD detector response can then be reconstructed to yield a prediction of the FD observable for the original ND data event. The reconstruction translation method is to directly predict the equivalent FD reconstruction from the reconstructed quantities of a ND data event. This is simpler than the former method but the full neutrino interaction may not be properly characterised by the ND reconstruction in all cases.

The near-to-far translation is implemented using machine learning (ML). ML is well suited to the task since the input data, especially in the case of using ND detector response, is high-dimensional and the mapping is non-linear. A paired dataset of ND and FD events

is used to train a model in a supervised manner. The paired dataset is generated using MC such that each pair shares the same underlying neutrino interaction. This allows a model to learn a mapping that takes the neutrino interaction from the input data and encompasses only the differences in detector response and, in the case of the reconstruction translation, the differences in reconstruction. The resulting mapping will be dependent on the detector simulation but have very little interaction model dependence.

The development of a paired dataset and its application in implementing a near-to-far translation will be discussed in the subsequent chapters. It will be useful to outline the scope of this work. Importantly, it is not expected that a perfect near-to-far translation is achievable. There will always be some residual that must be corrected for using MC. The goal is to make this MC correction small so that a significant fraction of the interaction model dependence may be removed from the analysis.

The development of a near-to-far translation in this thesis is focused on the disappearance analysis channel where muon neutrinos are measured at both the ND and FD. The appearance analysis channel where muon neutrinos at the ND are compared to electron neutrinos at the FD presents some additional difficulties. The outgoing lepton is of course different and the difference in lepton mass can alter the hadronic component of the interaction, especially at lower energies. However, they both share the same interaction channels and undergo the same nuclear effects. ND muon neutrino data events can still be used in data-driven predictions of electron neutrino interactions at the FD. A proposal for how this would be accomplished will be presented in §6.1.5.

# A Paired Near to Far Dataset

6

The implementation of a near-to-far translation using machine learning requires a training dataset to learn the mapping between ND and FD. The training data is made by constructing pairs of ND and FD response and reconstruction to the same underlying neutrino interaction and subsequent propagation in LAr. A collection of these near-far pairs form a paired dataset that can be used for training. The dataset encodes the differences between the detectors on the level of a single event. This permits learning a mapping between detectors that only relies on simulation of the detector readout since the underlying neutrino interaction at the ND and FD is identical for each pair. In this chapter, the procedure used to generate a paired dataset is outlined and the resulting dataset is discussed.

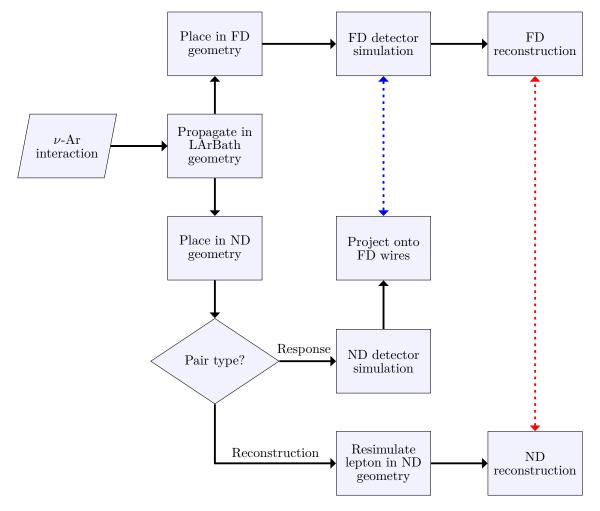
This chapter is focused on the FHC disappearance analysis channel, i.e. generating a sample of muon neutrinos at the ND and muon neutrinos at the FD. The steps required to generate paired data for the other analysis channels will be discussed. This chapter also makes the simplification of considering the FD in the reduced size 1x2x6 horizontal drift configuration only. This is in line with the DUNE's most recent long-baseline oscillation analysis [76] where the simplification is used since simulation of this FD configuration is the most mature.

## 6.1 Generating Near-Far Pairs

Generating paired ND and FD simulation requires the harmonisation of simulation from multiple formats and frameworks. To facilitate this, interfaces between a paired HDF5 [137] file structure and the ND and FD software are developed. This enables sharing of data across the simulation and reconstructions chains of the two detectors. Using this HDF5 file structure to link disparate software, paired events may be generated. An overview of the steps taken to generate a paired event is shown in Figure 6.1. Each step will be detailed in this section.

#### 6.1.1 Shared Neutrino Interaction

The first step is to obtain the underlying interaction within the detector that will be shared between ND and FD. A neutrino interaction with argon is simulated using GE-



**Figure 6.1:** Flow diagram showing the steps for creating a paired near-far event. ND and FD simulation is applied to the same simulated neutrino interaction and subsequent propagation in LAr. The steps yield a response pair (blue dashed line) and a reconstruction pair (red dashed line).

NIE. The true energy and interaction vertex of the incident neutrino is drawn from the FHC beam flux of the ND in the on-axis position. The flux at the on-axis position is a convenient choice to cover the full range of neutrino energies expected in on- and off-axis ND measurements. The beam flux is precomputed using the G4LBNF simulation of the beamline.

The final state particles of the interaction are propagated through LAr until they deposit enough energy to stop ionising. This is accomplished using edep-sim [143], a wrapper of Geant4, with the *LArBath* geometry. The LArBath geometry is an 800 m<sup>3</sup> volume of LAr centred on the coordinates of the on-axis position of ND-LAr. This is sufficient to contain the final state particles from interactions of incident beam neutrinos at all energies. The output of the simulation is a collection of *segments*, linear steps through the LAr each with an associated deposited energy.

The collection of segments are then placed in both the FD and ND geometries to be processed further with their respective detector simulations. The ND geometry used has ND-GAr as the muon spectrometer. Since the neutrino interaction is generated using the ND beam flux, the interaction point and orientation does not need to be changed when placing the segments in the ND.

When creating a reconstruction pair, an additional processing step is required at the ND. The primary lepton propagation is simulated in the ND geometry, rather than the LArBath geometry, using edep-sim. The original lepton segments are then replaced with this new simulation. As demonstrated in Figure 6.2, this is necessary to accurately reproduce the reconstruction of the lepton with ND-GAr. A process which is highly dependent on the path taken by the lepton through air, ND-GAr's cylindrical magnet, and its magnetised GAr drift chamber. By resimulating the muon in this manner, the near-to-far translation learnt with the paired dataset will no longer be a function of the physics underpinning the muon's propagation since this will come from two distinct simulations at the ND and FD. The propagation in LAr of the hadronic system and, importantly, the neutrino-nucleus interaction will remain the same at the two detectors.

When placing the same segments in the FD geometry an Earth's curvature correction is applied to the positions of the segments. This is to account for a different beam direction at the FD caused by the contentious effect of the Earth being spheroidal while the two detectors are locally level. The correction is a rotation of 11.57° upwards about an axis intersecting the interaction position and parallel to the detector's drift direction. The rotated segments are then placed in the FD geometry for a randomly and uniformly sampled interaction position in the fiducial volume (as defined in §5.2.2). An additional processing step of evenly splitting the FD segments into steps of length at most 0.4 mm is performed to ensure compatibility with the FD detector simulation which requires a smaller segment size than that required by the ND detector simulation. When generating a response pair, all FD segments with ND counterpart segments outside of the active

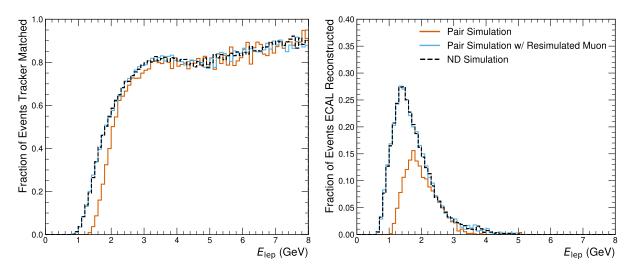


Figure 6.2: Fraction of CC  $\nu_{\mu}$  ND events selected via tracker matching (left) or reconstruction with the electromagnetic calorimeter (right) at different true lepton energies. Both selection criteria come from the parametrised reconstruction. Standard ND simulation is compared to the paired data procedure with and without resimulating the muon in the ND geometry.

volume of ND-LAr are removed.

The result of these steps is a neutrino event at both the ND and FD that shares the same neutrino-nucleus interaction, the same propagation of the hadronic particles in LAr, and, for the response pair, the same propagation of the muon in LAr. The events will undergo the detector simulation and reconstruction of their respective detector.

The use of the LArBath geometry assumes that both the FD and ND-LAr can be approximated as only LAr. Differences in propagation can arise from small amounts of other material, mostly fibreglass, in the instrumentation and high voltage systems of the detectors. The abundance of materials with a higher density than LAr in the inactive volumes between ND-LAr drift modules is expected to result in a suppression of the energy deposited in the active volumes on the order of a couple of percent when compared to LArBath. Although for now this is treated as negligible, it is significant enough to motivate the future development of a correction to be used in the paired dataset generation procedure.

#### 6.1.2 Detector Simulation and Reconstruction

In the ND, the parametrised reconstruction discussed in § 4.3.2 is applied directly to the segments in the detector. The reconstruction yields the criteria for a selection based on hadronic and leptonic containment and a reconstructed neutrino energy.

To generate response pairs, the detector simulation of the ND discussed in §4.3.1, larnd-sim, is applied to the segments. This simulates the ionisation and drift in the LAr and the response of the charge readout. The output is a collection of data packets along with a trigger packet that contains an initial timestamp for the interaction.

In the FD, the standard detector simulation and reconstruction algorithms discussed

in Chapter 4 are applied to the energy depositions. These include ionisation and electron drift in the LAr, the detector simulation and deconvolution with the 2D field response and the electronics response, and the full reconstruction chain. The reconstruction produces CVN scores for selection and a reconstructed neutrino energy.

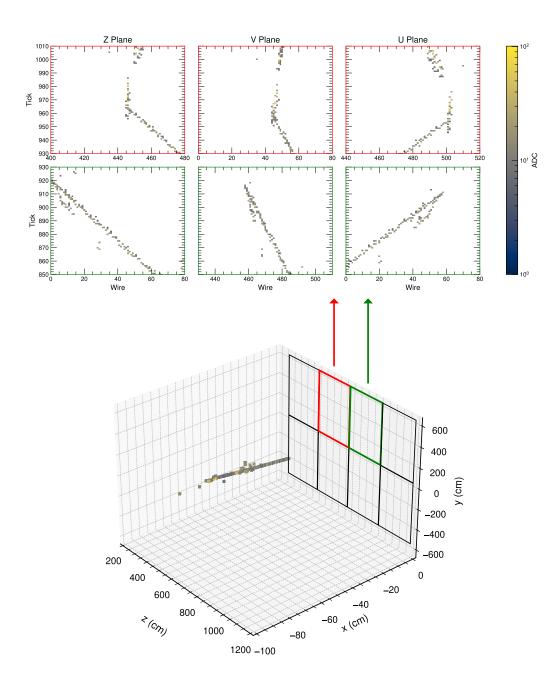
#### 6.1.3 Near-Far Pairing

A reconstruction pair is immediately obtained by associating the ND reconstructed variables with the FD reconstructed variables. The pair includes all reconstructed variables available from the ND parametrised reconstruction along with the high-level FD reconstruction.

To produce a response pair, some additional manipulations are required to put both detector responses in the same coordinate system. The aim is to transform the ND detector response into the wire and time tick coordinate system of the FD such that the two detector responses are spatially aligned. First, the drift coordinate of the ND packets is calculated using the initial interaction time, given by a trigger packet, to yield 3D positions in ND-LAr. Unlike the FD, the charge readout of ND-LAr does not include shielding against long-range induction effects from the drifting electrons, resulting in the calculated drift coordinate being shifted towards the anode due to the slightly earlier pixel triggers. This is corrected for with a small empirical shift away from the anode. The 3D packets are then moved into the FD geometry where they are aligned with their FD counterpart. The alignment is composed of a translation such that the ND vertex is moved to the FD vertex, an Earth's curvature correction, and the removal of any packets that lie outside of instrumented detector volumes after these transformations. The aligned packets then have the centres of their 3D positions projected onto the wires they would be drifted to and the time tick of their expected arrival. This produces three 2D views of the 3D packets in the wire and time tick coordinate system of the FD. This projection is illustrated in Figure 6.3. A response pair is then formed by associating each non-empty wire plane projection of the ND packets with the FD response at that wire plane. The transformations prior to the projection ensures that ND and FD response at each wire plane is aligned.

#### 6.1.4 Oscillation Analysis Compatible Pairs

An important benefit of the reconstruction translation is the simplicity of its implementation. It can be applied as a post-processing step on existing MC productions as it only requires ND reconstructed variables as input. This is in contrast to the response translation which would need to be integrated into production workflows since it is unfeasible to save detector responses for large ND MC samples. This motivates the generation of a paired dataset with the same detector simulation and reconstruction algorithms used in the PRISM oscillation analysis discussed in the previous chapter. The reconstruction



**Figure 6.3:** Example of a 3D ND-LAr detector response that has been moved to the FD (bottom) and projected onto wire planes (top). For this example, the 3D detector response spans two drift volumes and so is encoded by projections to the wires at the backward facing side of two APAs (drawn in red and green). The projection to each of the wire planes associated with the two APAs produces a 2D image in wire and time.  $80 \times 80$  crops of the full projections are shown for visualisation purposes.

translation learnt with such a dataset may be integrated into the current PRISM analysis rather than a future iteration. The MC sample used in the current PRISM analysis is known as the technical design report (TDR) sample in reference to being generated for DUNE's TDR [41]. The TDR sample dates back to 2018 and the simulation used to generate it has some characteristics that require amending the previously discussed steps.

Compared to the current simulation assumed in the previous sections, the TDR simulation employs a 1D field response for the detector simulation and subsequent deconvolution at the FD. Most relevant to the pair generation procedure, the TDR simulation predates a substantial refactor of the particle propagation simulation at the FD. A consequence of this is that there is no appropriate software object to load the segments from the shared neutrino interaction into. The nearest such object describes groups of ionisation electrons after they have been drifted to wires. To navigate this, the current simulation is used for the ionisation and drift of electrons up to the wire planes. The parameters governing these effects are configured to their TDR simulation values. From here, the TDR simulation chain can proceed with the remainder of the detector simulation and the reconstruction. The paired HDF5 file structure allows for this interoperability between distinct software versions.

#### 6.1.5 Steps to a Complete Four-Channel Sample

The implementation of a near-to-far translation for a complete four-channel analysis requires generating additional paired datasets to encompass  $\nu_e$  appearance and different detector configurations at the FD. To learn a response translation, the steps discussed previously are sufficient for both the FHC and RHC disappearance analysis channels. For a reconstruction translation, an additional RHC paired dataset would need to generated. This would require swapping the FHC beam flux used to generated the initial neutrino interaction with the RHC beam flux.

To learn a near-to-far translation for the appearance channel, the FD event of a nearfar pair must be the  $\nu_e$  interaction equivalent of the  $\nu_{\mu}$  interaction at the ND. The proposed way to construct this equivalent  $\nu_e$  interaction at the FD is to replace the muon with an electron simulated using the same kinematics. The hadronic system of the near-far pair would remain the same between the ND and FD.

Substituting the muon with an electron does not produce a perfectly equivalent interaction at the FD. Although the processes governing the neutrino-nucleon interaction and the subsequent FSI are the same for  $\nu_e$  and  $\nu_\mu$  interactions, the smaller mass of the electron means that on average  $\nu_e$  interactions transfer more energy to the nucleon or quark. The effect this has on some properties of the interaction is shown in Figure 6.4. The  $\nu_e$  is more likely to interact with very high inelasticity and slightly favours DIS over QE interactions. However, this does not appear to cause significant deviation in the interaction's final state particles, suggesting that, for a given incident neutrino energy, the

hadronic system of a  $\nu_{\mu}$  interaction describes the hadronic system of a  $\nu_{e}$  interaction well. This motivates a paired dataset where at the FD the muon is swapped with an electron. The mapping learnt from this paired dataset will permit a data-driven prediction of  $\nu_{e}$  interactions at the FD using the hadronic component of  $\nu_{\mu}$  interactions at the ND. A small MC correction would be required to account for differences in the interactions, some of which have been highlighted in Figure 6.4.

A complete near-to-far translation will require a paired dataset with the simulated FD in the VD configuration also. This requires swapping the HD configuration FD detector simulation and reconstruction for those of the VD configuration. For the response pair, the projection of ND packets to wires would be replaced with a projection to charge sensitive strips, also generating three paired 2D images for each drift volume.

# 6.2 Detector Response Pairs

An example of a response pair for the wire planes of one drift volume is shown in Figure 6.5. On average each ND-LAr event spans 4 FD drift volumes and produces 9 paired 2D wire plane responses<sup>1</sup>. The alignment procedure discussed in §6.1.3 ensures the ND response overlays the FD response at the appropriate wire and at a time just prior to the formation of the associated FD pulses. Visible in the figure at a tick of approximately 190 is a gap in the projected ND packets caused by the crossing of an anode plane in ND-LAr.

To confirm the ND response is consistently being projected to the wire plane with the correct FD response, the relationship between the summed ADC for both detector responses is examined in Figure 6.6. There is a clear linear relationship between the projected ND ADC and the FD ADC at the same wire plane, showing that the ND responses are being matched with the correct FD event and aligned properly. The smearing can be attributed to effects such as electronics noise and non-active regions in ND-LAr.

## 6.3 Reconstruction Pairs

The reconstruction pairs are validated more thoroughly to ensure the near-far pairing procedure accurately reproduces pure ND and FD simulation. The validation is performed using 597,320 samples from each of the paired dataset, the FD MC, and the on-axis ND MC. The near-far pairs are generated using the TDR simulation, as outlined in §6.1.4, while the FD and ND MC are sampled from the datasets used in the PRISM analysis. The paired dataset has been generated using the ND beam flux since the mapping learnt from it will be applied to ND data. As shown in Figure 6.7, this means the neutrino flux spectrum of the paired dataset matches that of the ND MC but not the FD MC. The FD

The wrapping of the induction wires around the APA means the average is not necessarily  $4 \times 3 = 12$  2D wire plane responses.

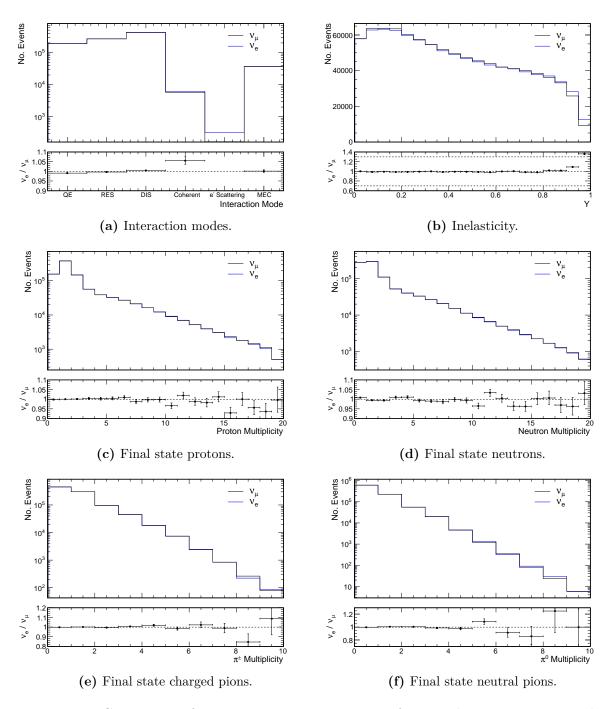
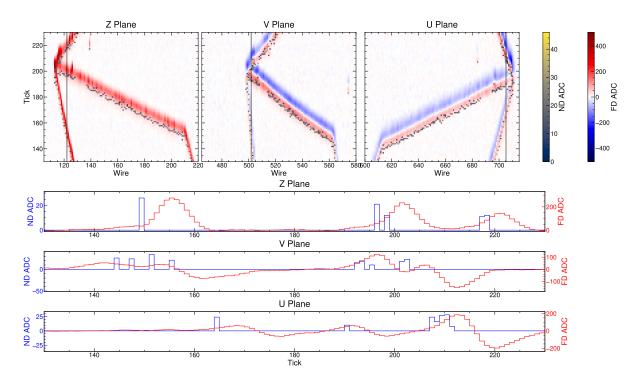
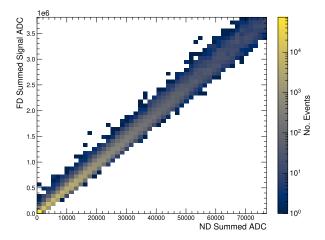


Figure 6.4: Comparison of neutrino interaction properties for  $\nu_{\mu}$  and  $\nu_{e}$  interactions simulated with GENIE. Both neutrino interaction samples have the same spectrum of incident neutrino energy. The error bars of the residuals represent statistical error.



**Figure 6.5:** Example of a response pair. Crops of the overlaid ND and FD detector responses for the wire planes of one APA are shown (top row). The detector responses at the wires marked with a grey band are also shown (bottom three rows).



**Figure 6.6:** Relationship between summed ND and FD ADC at the collection planes of response pairs. FD summed signal ADC is defined as the sum of all ADC over a threshold of 20.

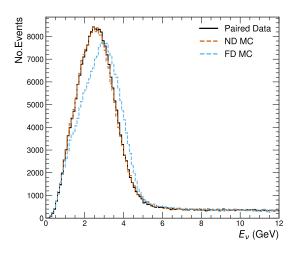


Figure 6.7: Beam neutrino flux spectrum of paired dataset, ND MC, and FD MC.

spectrum peaks at a higher energy than the ND spectrum due to the FD being at a much greater distance and so subtending a smaller angle around the on-axis position.

Since the paired data is produced using the same flux as the ND MC, the distributions of ND reconstructed variables should match each other. This is shown in Figure 6.8. In this figure, the selection for FHC CC  $\nu_{\mu}$  events, as outlined in § 5.2.1, is applied using the reconstruction from the paired data or ND MC. Therefore, each distribution depends on its associated reconstructed variable as well as the reconstructed variables that define the selection criteria. We see that the reconstruction from the paired data matches the pure ND MC closely except for a small difference in normalisation. This indicates the ND side of the near-far pair making procedure accurately reproduces the ND MC.

The discrepancy in normalisation is due to events being more likely to pass selection cuts in the ND MC than in the paired data. This arises from the discrepancy in the hadronic veto shown in Figure 6.9. Events generated using ND MC have a slightly lower hadronic veto value than events generated using the paired data procedure, and so are  $\sim 5\%$  more likely to meet the associated selection criterion. This is an expected consequence of the LArBath geometry. By not simulating the denser non-LAr volumes in ND-LAr, hadronic particles have a slightly increased range, making them more likely to enter the veto region at the detector edges. Since this effect is small, it is not expected to have a significant effect on the performance of any near-to-far translation trained using this paired data. However, as discussed at the end of § 6.1.1, correcting for this is a target for future developments.

Since the flux of the FD MC differs from that used for the paired data, the FD reconstruction of the paired data is validated as a function of the true neutrino energy. This is shown in Figure 6.10, where bins of true neutrino energy have been normalised to unity. The discontinuity at 8 GeV reconstructed lepton energy is caused by the MCS energy estimation method being used when the muon is identified as uncontained. The MCS reconstruction is limited to predicting muon energies of less than 10 GeV since

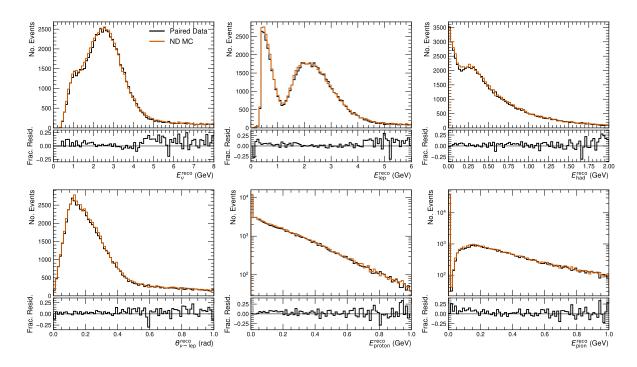


Figure 6.8: Distributions of reconstructed neutrino energy (top left), leptonic energy (top middle), hadronic energy (top right), neutrino-lepton opening angle (bottom left), proton energy (bottom middle), and pion energy (bottom right) for selected  $\nu_{\mu}$  events at the ND.

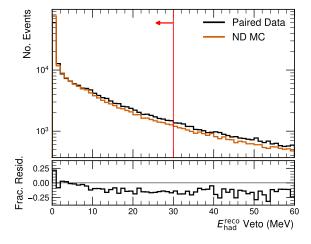
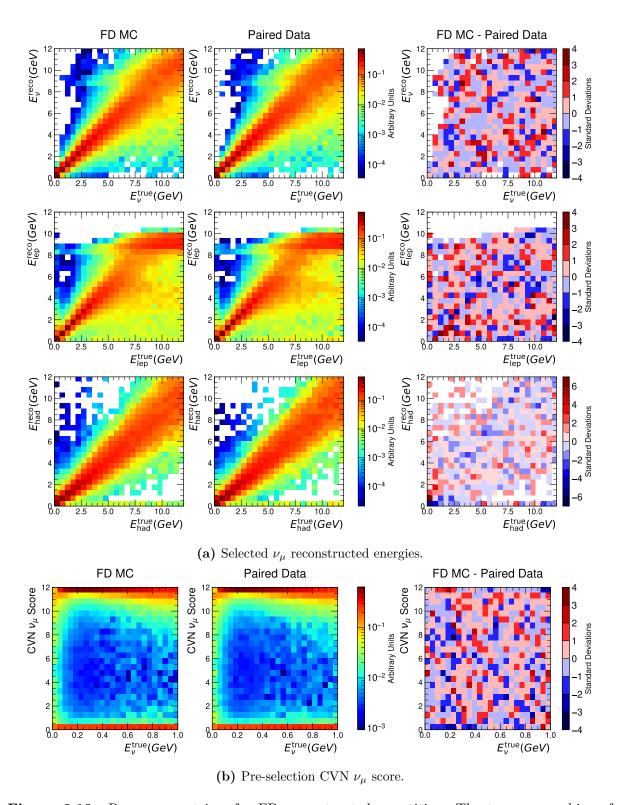


Figure 6.9: Distribution of hadronic veto for all fiducial volume ND events. The red line shows the selection cut typically used in the PRISM analysis,  $E_{\text{had}}^{\text{reco}}$  Veto  $\leq 30 \,\text{MeV}$ .

it is not considered reliable for energies higher than this. The population at  $\sim 5\,\text{GeV}$  reconstructed and  $\gtrsim 7.5\,\text{GeV}$  true lepton energy can be attributed to a range-based energy estimation being applied to muons that are incorrectly identified as contained.

For the neutrino energy, lepton energy, and CVN  $\nu_{\mu}$  score, the difference plots are consistent with statistical fluctuations. For the hadronic energy at true energy < 1 GeV, there is a migration to slightly lower reconstructed energies for the paired data compared to the FD MC. This is a small effect that has no discernible impact on the reconstructed neutrino energy. Its cause is not clear. Overall, the reconstructed energies and the CVN  $\nu_{\mu}$  score match well between the FD MC and the paired data which, along with the ND MC validation, indicates the near-far pair making procedure accurately reproduces MC at both detectors.



**Figure 6.10:** Response matrices for FD reconstructed quantities. The true energy bins of the FD MC and Paired Data histograms are normalised to unity. The difference histograms are in units of statistical uncertainty that has been propagated through the normalisation and subtraction.

Reconstruction Translation

7

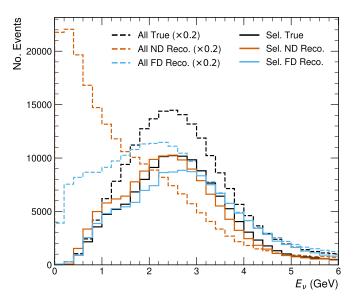
The dataset of near-far pairs discussed in the previous chapter enables the training of ML models to perform a near-to-far translation. The simplest translation is to predict, for each event, a vector of desired FD reconstruction variables given a vector of ND reconstruction variables as input. The reconstruction pairs can be used as training data for a ML model developed to perform this translation. The model would ultimately be applied to ND data. Such a ML approach provides a high-dimensional correction for detector effects for each event that would not be possible using the binned approach presented in § 5.3. By having many ND reconstruction variables as inputs to the model, the prediction will be a function of the underlying neutrino interaction in so far as the interaction is described by the ND reconstruction. This should yield a detector effects correction much less dependent on the interaction model used in the simulation.

The development and early implementation of the reconstruction translation will be discussed in this chapter. Only the translation for the FHC disappearance channel is considered.

## 7.1 Training Data

A dataset of near-far reconstruction pairs generated with the TDR simulation is used as a training dataset. The dataset consists of 1,198,619 pairs in total. The ND events that any model trained with this dataset will be applied to are those that pass selection cuts. In addition, the reconstruction variables are only good descriptors of the interaction when the event passes selection cuts for containment. For these reasons, the ND selection cuts used in the PRISM analysis, outlined in § 5.2.1, are applied to the dataset. No cuts on the FD variables are required since the pairing process places interaction vertices in the fiducial volume of the FD. The effect of the selection on the dataset is shown in Figure 7.1. There are 147,703 reconstruction pairs after the selection is applied. Of these, 70,000 are used for training and the remainder for test and validation.

The ND and FD reconstruction variables selected from the near-far pairs as the input and targets of the training are outlined in Table 7.1. The choice of ND reconstruction input variables is limited by the parametrised reconstruction which provides only a high level description of the neutrino interaction. When full ND reconstruction is available,



**Figure 7.1:** Effect of selection cuts applied to the training dataset on true and reconstructed neutrino energies.

lower level track and shower reconstruction variables that better characterise the event topology may be included. The distances to the fiducial volume boundary of the FD pair's interaction vertex are included in the input since containment in the 1x2x6 geometry is not guaranteed from the fiducial volume cut which can affect the reconstruction. The FD reconstruction outputs include the desired FD analysis variable fd\_numu\_nu\_E and FD selection criterion fd\_numu\_score. The FD reconstructed hadron and lepton energies are also included as outputs since doing so was found to improve performance on the main task of predicting fd\_numu\_nu\_E and fd\_numu\_score.

## 7.2 Autoregressive Transformer for Translation

The reconstruction translation is a regression task from a vector of ND variables  $\mathbf{x}_{\text{ND}}$  to a vector of FD variables  $\mathbf{x}_{\text{FD}}$ . The mapping  $\mathbf{x}_{\text{ND}} \to \mathbf{x}_{\text{FD}}$  is the correction for the different detector responses and subsequent reconstructions of the ND and FD. This encompasses differences in the non-linear transformation of the underlying neutrino interaction by detector components and differences in the finite resolution of the measured quantities of each detector. The same problems that arise in unfolding procedures [144] are present in the reconstruction translation. Namely, each  $\mathbf{x}_{\text{ND}}$  does not necessarily correspond to a unique  $\mathbf{x}_{\text{FD}}$ . To model the distribution of  $\mathbf{x}_{\text{FD}}$  for a given  $\mathbf{x}_{\text{ND}}$ , a conditional generative approach that models  $p(\mathbf{x}_{\text{FD}}, \mathbf{z} \mid \mathbf{x}_{\text{ND}})$ , where  $\mathbf{z}$  is a noise vector, is considered.

### 7.2.1 Architecture

Inspired by the recent wider successes of generative pretrained transformer (GPT) architectures [145] and their application to physics simulations at the Large Hadron Collider

	77 · 11 37	D
Type	Variable Name	Description
Input	Ev_reco	ND reconstructed neutrino energy
Input	Elep_reco	ND reconstructed lepton energy
Input	eRecoP	ND reconstructed total proton energy
Input	eRecoPipm	ND reconstructed total $\pi^{\pm}$ energy
Input	eRecoPiO	ND reconstructed total $\pi^0$ energy
Input	theta_reco	ND reconstructed $\nu$ -lepton angle
Input	muon_tracker	ND muon tracker-matched flag
Input	muon_contained	ND muon contained flag
Input	Ehad_veto	ND energy deposited in hadronic veto region
Input	fd_vertx_fv_dist	FD vertex shortest distance to fid. vol. $x$ boundary
Input	fd_verty_fv_dist	FD vertex shortest distance to fid. vol. $y$ boundary
Input	fd_vertz_fv_fdist	FD vertex distance to fid. vol. front $z$ boundary
Input	fd_vertz_fv_bdist	FD vertex distance to fid. vol. back $z$ boundary
Target	fd_numu_score	FD CVN CC $\nu_{\mu}$ score
Target	fd_numu_nu_E	FD reconstructed neutrino energy
Target	fd_numu_lep_E	FD reconstructed lepton energy
Target	fd_numu_had_E	FD reconstructed hadron energy

**Table 7.1:** Input and target variables for the near-to-far reconstruction translation.

[146], an autoregressive transformer is developed for the reconstruction translation. The autoregressive aspect refers to interpreting the ND and FD reconstruction vectors as sequences and predicting the first FD reconstruction variable in the sequence,  $x_{\text{FD},1}$ , given  $\mathbf{x}_{\text{ND}}$ , then the second FD reconstruction variable in the sequence,  $x_{\text{FD},2}$ , given  $\mathbf{x}_{\text{ND}}$  and  $x_{\text{FD},1}$ , and so on until the full sequence of FD reconstruction has been predicted. More formally, the joint probability is factorised as

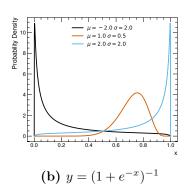
$$p(\mathbf{x}_{\mathrm{FD}}|\mathbf{x}_{\mathrm{ND}}) = \prod_{i=1}^{i \le N_{\mathrm{FD}}} p(x_{\mathrm{FD},i}|\mathbf{x}_{\mathrm{FD}}^{(i)}, \mathbf{x}_{\mathrm{ND}}), \tag{7.2.1}$$

where  $\mathbf{x}_{\mathrm{FD}}^{(i)}$  denotes the FD sequence up to and including the *i*-th element,  $(x_{\mathrm{FD},i},\ldots,x_{\mathrm{FD},1})$ , and  $N_{\mathrm{FD}}$  is the total number of FD reconstruction variables. The order of reconstruction variables presented in Table 7.1 is used to make the sequences. The two FD reconstruction variables relevant for the PRISM analysis are placed at the start of the sequence to ensure they are not impacted by bad predictions of the other FD variables.

The network learns each conditional probability as a Gaussian mixture

$$p(x_{\text{FD},i}|\mathbf{x}_{\text{FD}}^{(i-1)},\mathbf{x}_{\text{ND}}) = \sum_{i} w_j^{(i-1)} \mathcal{N}(x_{\text{FD},i};\mu_j^{(i-1)},\sigma_j^{(i-1)}),$$
(7.2.2)

where  $w_j$  is a probability for each Gaussian in the mixture and the superscript (i-1) denotes that the learnable parameter encodes the conditional dependence of  $x_{\text{FD},i}$  on  $\mathbf{x}_{\text{FD}}^{(i-1)}$  and  $\mathbf{x}_{\text{ND}}$ . By predicting the parameters governing the mixture with a sufficient



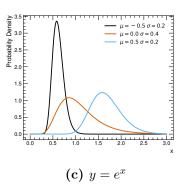


Figure 7.2: Transformed Gaussian functions used in the Gaussian mixture to model distributions of FD reconstruction variables. (a) is used to predict reconstructed energies, (b) is used to predict the CVN  $\nu_{\mu}$  score.

number of Gaussian distributions, the network should be able to accurately reproduce the conditional probabilities encoded by the training dataset.

To tailor the Gaussian mixture to the FD reconstruction variables, a change of variable can be applied to the Gaussian distribution. For the CVN  $\nu_{\mu}$  score, the Gaussian is transformed as  $y = (1 + \exp(-x))^{-1}$ , i.e. a sigmoid, to ensure the mixture is bounded between 0 and 1. For the reconstructed energies, the transformation to a log-normal distribution,  $y = \exp(x)$ , to ensure positive energies was experimented with. However, the asymmetry of the log-normal distribution was found to introduce bias into the predicted probability distributions. Non-transformed Gaussians are instead used for the reconstructed energies. Examples of the aforementioned distributions are shown in Figure 7.2.

The network trained to predict the parameters of the Gaussian mixture uses a transformer architecture <sup>1</sup> based on the GPT models. A diagram of the network architecture is shown in Figure 7.3. The backbone of the network is the transformer block which uses the attention mechanism to construct a representation of  $(\mathbf{x}_{FD}^{(i-1)}, \mathbf{x}_{ND})$  that captures the correlations between variables necessary to accurately generate the conditional probability via the parameters of the Gaussian mixture (7.2.2). The attention is implemented as a multi-head masked self-attention layer. This combines multiple *heads* of self-attention layers [147]. A single-head self-attention applied to a sequence  $x_i$  of dimension T starts by embedding the sequence into a d-dimensional embedding space using a learnable linear transformation. This embedding is denoted  $x_{i\alpha}$ . Three matrices are then computed using learnable transformations denoted  $W_{\alpha\beta}$ . These are

$$Q_{i\alpha} = W_{\alpha\beta}^Q x_{i\beta}, \quad K_{i\alpha} = W_{\alpha\beta}^K x_{i\beta}, \quad V_{i\alpha} = W_{\alpha\beta}^V x_{i\beta}, \tag{7.2.3}$$

known as the query, keys, and values respectively. These are then combined into the

<sup>&</sup>lt;sup>1</sup>Other non-transformer architectures may also be employed with minimal change to the overall reconstruction translation framework described in this chapter.

scaled dot product attention,

Softmax<sub>j</sub> 
$$\left(\frac{Q_{i\beta}K_{j\beta}}{\sqrt{d}}\right)V_{j\alpha}.$$
 (7.2.4)

The numerator,  $\mathbf{Q}\mathbf{K}^{\top}$ , encodes the similarity of all sequence elements with each other by taking the dot product in the embedding space of the query and keys. The result is a  $T \times T$  attention matrix. This is scaled by  $1/\sqrt{d}$  and the softmax function is used to normalise the attention values of each row to unity. The original representation  $x_{i\alpha}$  is then updated with the learnt attentions by multiplication with the values.

The self-attention is masked to enforce causality. This is only required for the FD variables where the autoregressive approach means a given FD variable should only depend on previous elements in the sequence, these being all ND variables and the FD variables that will have a prediction, and not future elements, these being FD variables that will not yet have a prediction. The causality mask is applied to the attention matrix  $\mathbf{Q}\mathbf{K}^{\mathsf{T}}$ . As an example, for T=5 and  $\dim(\mathbf{x}_{\mathrm{ND}})=3$ , the mask would be

To construct the multi-head masked self-attention layer, multiple masked self-attention heads are computed in parallel and their outputs concatenated and combined with a final learnable transformation. This is accomplished by splitting the embedding space such that, for n heads, the query, key, and values matrices of each head have dimension  $T \times d/n$  rather than  $T \times d$  as for the single-head.

The representation learnt by the transformer blocks is passed through a final linear layer to predict the parameters of the Gaussian mixture. The weight vector is transformed with a softmax function to convert the weights to probabilities of each Gaussian in the mixture. The vector of Gaussian widths is exponentiated to ensure they are non-negative. The probabilities, widths, and means are then combined to produce the conditional probability distribution.

The model is implemented in PyTorch [148] making use of its distributions library. Some of the important hand-tuned hyperparameters are listed in Table 7.2.

## 7.2.2 Training and Sampling

As indicated in Figure 7.3, the full sequence of ND and FD reconstruction variables with the final FD variable omitted are inputted to the model during training. A prediction for all FD conditional probability distributions is made in parallel. A loss function is

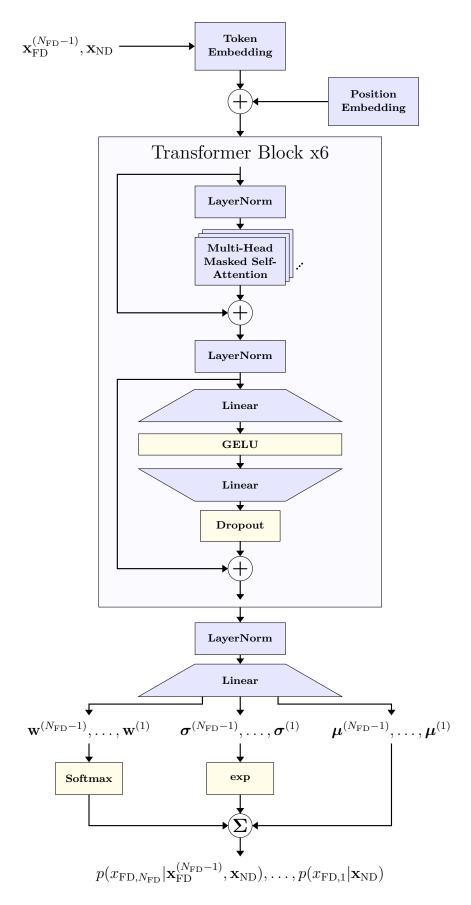
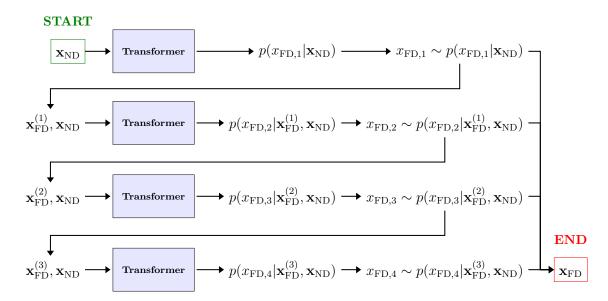


Figure 7.3: Diagram of the GPT transformer architecture used to predict conditional probability distributions. The transformer training inputs and outputs are shown. Layers with learnable parameters are shaded blue while fixed layers are shaded yellow.

Value
6
6
192
0.1
64

**Table 7.2:** Transformer model hyperparameters.



**Figure 7.4:** Schematic of the autoregressive sampling process of the transformer for four predicted FD variables.  $x \sim f(x)$  denotes a realisation of a random variable with the distribution f(x).

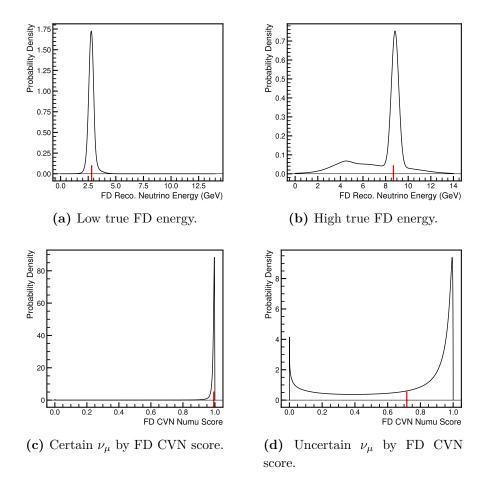
constructed by evaluating each distribution at the true value of the FD variable,

$$\mathcal{L} = \frac{1}{N_{\text{FD}}} \sum_{i=1}^{i \le N_{\text{FD}}} \left( -\log p(x_{\text{FD},i} | \mathbf{x}_{\text{FD}}^{(i-1)}, \mathbf{x}_{\text{ND}}) \right).$$
 (7.2.6)

Minimisation of this loss encourages the predicted conditional distributions to reproduce the densities of the FD variable in the training dataset.

The model is trained until the validation loss indicates convergence. This is typically reached before 8 epochs. The AdamW optimiser [149, 150] is used, gradient clipping is enforced, and training is performed with a batch size of 64.

Sampling from the trained model is done by predicting each FD conditional probability distribution and sampling from them in an autoregressive manner. This is shown schematically in Figure 7.4 for the prediction of the four FD variables used in the reconstruction translation. In the event of sampling a negative value for a FD reconstructed energy, the distribution is resampled until a non-negative value is returned.

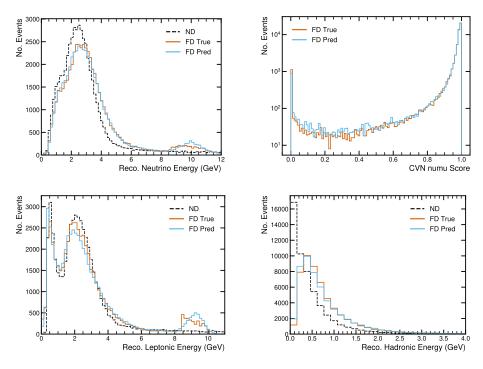


**Figure 7.5:** Example conditional probability distributions predicted by the model. The true value of the FD variable is marked in red.

#### 7.2.3 Results

The trained model is evaluated on a test dataset of 53,570 near-far pairs. The predicted conditional probability distributions for the FD reconstruction variables are illustrated in Figure 7.5 for some example events. These distributions highlight some qualitative features of the predictions. For low FD reconstructed neutrino energy, the predicted distributions are tight around the true value. While for higher energies, the distributions are wider around the true value and sometimes bimodal. This is likely due to the effects of hadronic and leptonic containment in the 1x2x6 FD that are present even with the fiducial volume cut applied. In Figure 7.5b, the peak around the true energy may be the case where the muon is properly contained while the shoulder at lower energies is the case where the muon exits the detector. The CVN  $\nu_{\mu}$  score distributions are typically sharply peaked near 1 since the application of the ND selection cuts means the training dataset CVN  $\nu_{\mu}$  scores are sharply peaked near 1. Despite this, Figure 7.5d demonstrates that the distribution can be moved towards lower scores for an event with lower true CVN  $\nu_{\mu}$  score.

Figure 7.6 shows the distributions of FD reconstruction variables predicted by the model. The predicted distributions generally match the true distributions of the test set

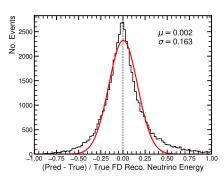


**Figure 7.6:** Comparison of predicted and true FD reconstruction distributions. The distribution of the equivalent ND reconstruction variable is shown where applicable.

well, especially for the two variables relevant to the PRISM analysis, the reconstructed neutrino energy and CVN  $\nu_{\mu}$  score. This is likely because these two variables are the first to be predicted in the autoregressive sampling and so are less susceptible to variations in the sampled value of previously predicted FD reconstruction variables. The predicted conditional probability distributions for the reconstructed neutrino and leptonic energies are able to account for the bump at 10 GeV caused by MCS reconstruction of uncontained muons reasonably well.

It is important that the model does not learn the unconditional distributions of the FD reconstruction variables rather than distributions properly conditioned on the ND reconstruction variables. The residuals of the model predictions for each event in the test set are shown in Figure 7.7. The model predictions are unbiased and, for the FD reconstructed neutrino energy, have a fairly poor resolution. This is expected due to there being no unique FD reconstruction vector for a given ND reconstruction vector. To further understand if the model has any unconditional dependence associated with the neutrino energy spectrum of the training dataset, the test dataset can be weighted using the true neutrino energy. The model predictions for the test dataset weighted to the FD oscillated spectrum for NuFIT 4.0 oscillation parameters [139] are shown in Figure 7.8. The predicted distributions continue to match the true distributions well with the weighting, suggesting that the predicted probability distributions of the FD reconstruction variables are conditional on the ND reconstruction.

The conditioning of the predicted FD reconstructed neutrino energy is examined in Figure 7.9 by marginalising over all but one of the ND variables. The distributions of



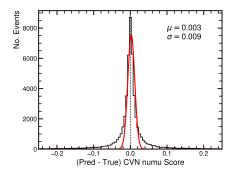
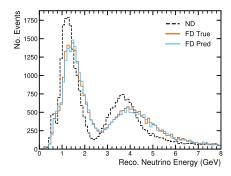
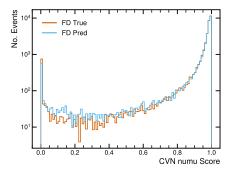


Figure 7.7: Residuals for FD reconstructed neutrino energy (left) and CVN  $\nu_{\mu}$  score (right) with fitted Gaussians.





**Figure 7.8:** Comparison of predicted and true FD reconstruction distributions for the test set weighted to a FD oscillated flux.

the model predictions match those of the test dataset closely. The sharp features of the distributions at high leptonic energies caused by an uncontained muon are present in the network prediction but slightly smeared, perhaps due to the limitation of modelling the distribution with a finite number of Gaussians. Accurately reproducing the dependence of the FD reconstructed neutrino energy on many ND reconstruction variables demonstrates the capability of the translation to make better use of the neutrino interaction in data than unfolding and smearing in a single reconstructed variable does.

If the use of composite FD analysis variables in the PRISM analysis is desired it would be beneficial to have accurate correlations between model predictions. This would mean analysis variables can be reliably constructed without requiring retraining. This is examined for FD reconstructed leptonic and hadronic energies in Figure 7.10. The correlations are reasonably accurate. It could likely be improved by explicitly enforcing a constraint on the sum of the energies during training.

# 7.3 Implementation in PRISM Analysis

Since the reconstruction translation only requires high-level reconstruction variables as input, it can be applied directly to the existing ND analysis files used in the PRISM analysis. By doing so, the efficacy of the reconstruction translation within the PRISM oscillation analysis is evaluated. The detector resolution correction via a smearing matrix

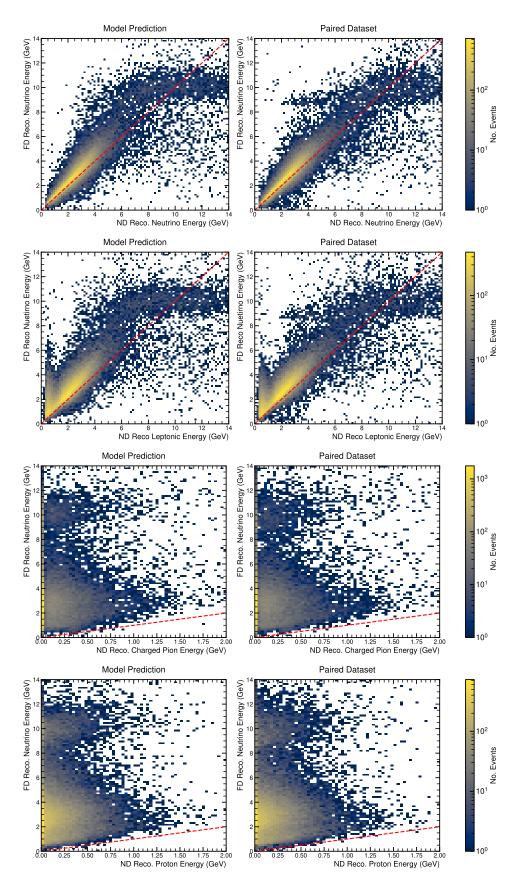
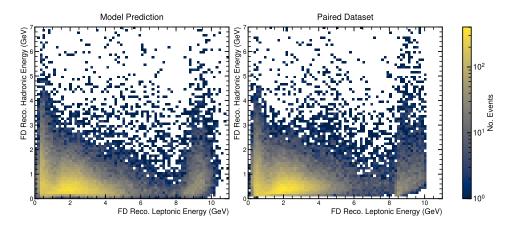


Figure 7.9: Comparison of 2D distributions of predicted and true FD reconstructed neutrino energy with some of the ND reconstruction variables the prediction is conditioned on: neutrino energy (top), leptonic energy (top middle), total charged pion energy (bottom middle), and total proton energy (bottom). The red dashed line is identity.



**Figure 7.10:** Comparison of 2D distribution of predicted and true FD reconstructed leptonic and hadronic energies.

described in §5.3.2 is replaced with the use of the predicted FD reconstructed neutrino energy. The impact of this on the FHC  $\nu_{\mu}$  disappearance analysis channel is assessed.

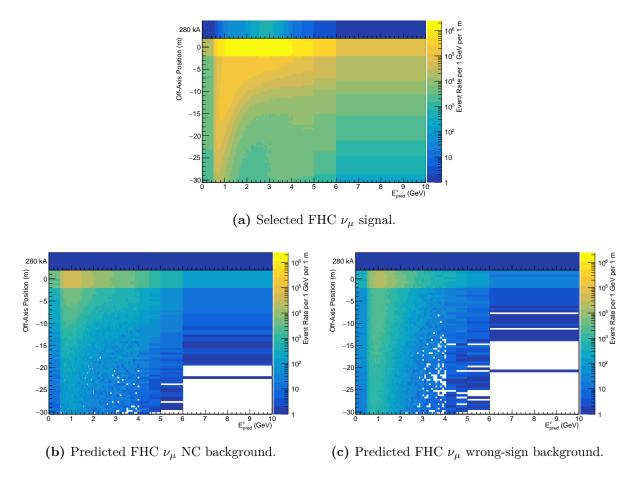
### 7.3.1 Monte Carlo Data

The MC used in the PRISM analysis discussed in Chapter 5 was made with the same TDR simulation used for training the reconstruction translation. The analysis files comprise approximately 150 million ND events stored in ROOT trees. The translation input, shown in Table 7.1, is constructed using the ND reconstruction of the event along with a randomly and uniformly sampled FD vertex within the fiducial volume. With this input, the model is used to make a prediction for the FD reconstructed neutrino energy and the CVN  $\nu_{\mu}$  score which is added to the analysis file via a friend tree. Model inference for one event takes approximately 20 ms with a single CPU core.

### 7.3.2 PRISM Prediction

Ideally, the reconstruction translation would be implemented alongside the geometric efficiency correction, discussed in §5.6.1, as an event-level correction for the detector effects. This would remove the requirement to unfold the ND spectra to a true energy. Each ND data event would have a ND selection efficiency correction factor, a predicted FD reconstructed neutrino energy, and a predicted FD CVN score that the FD selection cut can be applied to. However, the geometric efficiency correction is not at a stage where it may be implemented in the PRISM analysis and so the ND selection efficiency must be corrected for with MC alongside the event-level detector resolution correction from the translation. The FD selection efficiency is also corrected for using MC to focus only on the effect of a data-driven detector resolution correction.

First, the selected ND data spectra are taken in the predicted FD reconstructed neutrino energy variable,  $E^{\nu}_{\text{pred}}$ , provided by the translation. The ND MC is also taken in  $E^{\nu}_{\text{pred}}$  to produce the predicted background rates in this variable. These spectra are shown



**Figure 7.11:** ND selected muon neutrino signal event rate spectra and predicted backgrounds in the predicted FD neutrino energy variable.

in Figure 7.11. The predicted MC background rates are subtracted from the selected ND data to yield pure  $\nu_{\mu}$  event rates at each off-axis position.

To correct for selection efficiencies the ND spectra must be unfolded. A smearing matrix  $\mathbf{M}$  that maps from true visible energy to predicted FD reconstruction neutrino energy,

$$E_{\text{pred},i}^{\nu} = \sum_{j} M_{ij} E_{\text{true},j}^{\text{vis}}, \tag{7.3.1}$$

is constructed using ND MC. For each off-axis position, an efficiency normalised smearing matrix,  $\mathbf{M}^{\text{eff}}$ , is obtained by enforcing that the sum over the  $E^{\nu}_{\text{pred}}$  bins for each true energy bin is equal to the selection efficiency associated with that true energy. The same ND selection efficiencies shown in Figure 5.6a are used. The matrices  $\mathbf{M}$  and  $\mathbf{M}^{\text{eff}}$  for the on-axis position are shown in Figure 7.12.

The efficiency-normalised smearing matrices  $\mathbf{M}^{\text{eff}}$  are used to unfold each off-axis spectrum. Unfolding is performed with Tikhonov regularisation, as described in §5.3.2. The FD selection efficiency, shown in Figure 5.9a, is then applied to each unfolded ND spectrum. The smearing matrix  $\mathbf{M}$  is applied to the ND and FD efficiency corrected unfolded spectra to return to the original  $E_{\text{pred}}^{\nu}$  variable. This completes the detector effects corrected.

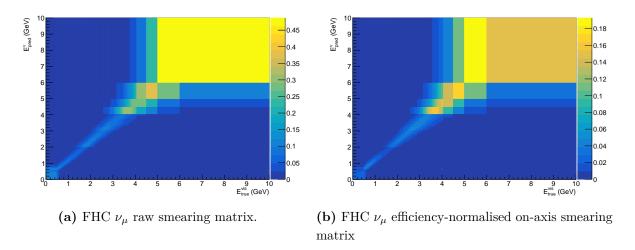


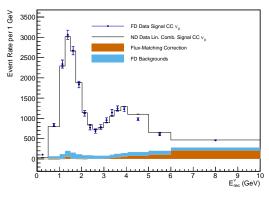
Figure 7.12: ND smearing matrices used to make selection efficiency correction.

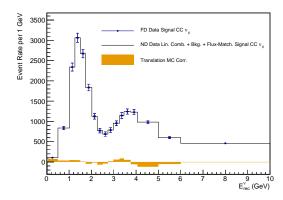
tion of the off-axis ND spectra. At this stage, the model predicted energy variable  $E^{\nu}_{\text{pred}}$  is considered equivalent to reconstructed neutrino energy at the FD, which is the analysis variable of the FD data, and thus is relabelled  $E^{\nu}_{\text{rec}}$ .

The formation of a PRISM prediction from the off-axis ND spectra that have been extrapolated to the FD analysis variables follows the methods of linear combination, FD background addition, and flux-matching correction as described in § 5.4. The resulting prediction is shown in Figure 7.13a. The Euclidean distance between the prediction and data histogram is 232.5. Relative to the prediction using the standard PRISM detector resolution, as shown in Figure 5.13a, which has a Euclidean distance of 161.0, the prediction matches the data reasonably closely. This demonstrates the good accuracy of the near-to-far translation. Since the match is slightly worse than as for the standard PRISM detector resolution correction, an additional MC component is added to the prediction to correct for the limited accuracy of the near-to-far translation. This small MC correction, shown in Figure 7.13b, is an additive correction that is simply the residual in each bin of  $E_{\rm rec}^{\nu}$  of the PRISM prediction for a given set of linear combination coefficients and the data. The result is a largely data-driven detector resolution correction with a small MC component that will be subject to systematic uncertainties.

### 7.3.3 Sensitivities

The PRISM prediction using the reconstruction translation is fitted to FD data (the FD MC for NuFIT 4.0 oscillation parameters) to yield oscillation parameter sensitivities. These sensitivities are compared to those produced using the standard PRISM detector effects correction described in §5.3. For a fair comparison, a final MC correction of the prediction-data residual is also included in the prediction using standard PRISM detector effects correction. To study the interaction model dependence, cross section systematic uncertainties are included in the fit. For both extrapolation methods, the ND selection





(a) Uncorrected prediction.

(b) Prediction with MC correction for reconstruction translation inaccuracy.

Figure 7.13: PRISM FD FHC  $\nu_{\mu}$  spectrum prediction using the reconstruction translation compared to FD data. NuFIT 4.0 oscillation parameters are used. The error bars show statistical error.

efficiency is frozen with respect to the cross section systematics to focus only on the difference in detector resolution correction. The smearing matrix  $\mathbf{M}$  (7.3.1) is also frozen as it will not be required following the integration of an event-level ND selection efficiency correction from the geometry efficiency correction. Therefore, cross section systematics enter the fit via the backgrounds, the flux-matching correction, the FD selection efficiency, and the final prediction-data residual MC correction. In addition, for the standard PRISM procedure they enter via the ND and FD smearing matrices shown in Figure 7.14.

Sensitivities for  $\Delta m_{32}^2$  and  $\sin^2\theta_{23}$  are shown in Figure 7.15. For both fits without systematics, the impact of statistical uncertainty is greater for the translation. This is because the smearing matrix used to make the ND selection efficiency correction, the mapping from  $E_{\rm true}^{\rm vis}$  to FD  $E_{\rm pred}^{\nu}$  shown in Figure 7.12, in the absence of the geometric efficiency correction is less diagonal that its counterpart matrix in the standard procedure, the mapping from  $E_{\rm true}^{\rm vis}$  to ND  $E_{\rm rec}^{\nu}$  shown in Figure 7.14. This smearing matrix contributes to the covariance matrix when it is used to unfold ND spectra.

Despite the reduced sensitivity in the statistics-only fit, the inclusion of cross section systematics in the  $\sin^2\theta_{23}$  fit causes the sensitivities to match between the two methods. This indicates a slight reduction in interaction model dependence from the reconstruction translation in the measurement of this parameter. For the  $\Delta m_{32}^2$  fit, the reconstruction translation results in a more substantial sensitivity improvement when cross section systematics are included. The  $3\sigma$  and  $5\sigma$  confidence intervals cover 4.1% and 5.3% less parameter space respectively.

The presented sensitivities demonstrate that the reconstruction translation is a viable improvement on the standard PRISM detector resolution correction. It results in an extrapolation that is less dependent on known cross section systematics and that should by extension be more robust to fake data studies of the like discussed in § 5.5.2. Any

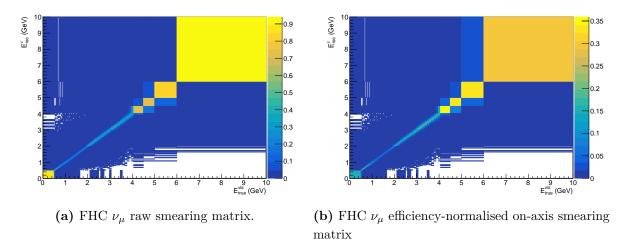


Figure 7.14: ND smearing matrices used to make selection efficiency and resolution correction for the standard extrapolation procedure.

improvements to the reconstruction translation model that reduces the size of the MC correction would further increase the advantage of this method. Possible improvements include a substantially larger training dataset, a training dataset with uniform neutrino energy spectrum, and utilising more descriptive input variables from full ND reconstruction.

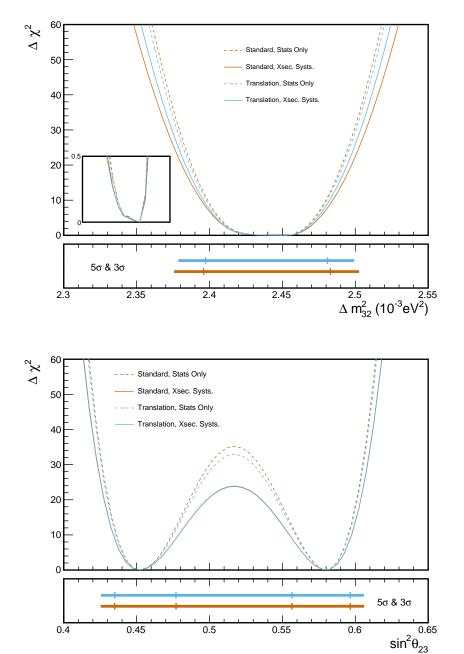


Figure 7.15: PRISM oscillation parameter sensitivities for  $\Delta m_{32}^2$  (top) and  $\sin^2\theta_{23}$  (bottom) with and without cross section systematic uncertainties. Shown for both the standard extrapolation procedure and the extrapolation using the reconstruction translation. The bottom pane shows the  $3\sigma$  (vertical markers) and  $5\sigma$  (lines) confidence intervals. The inset on the upper plot is a zoom-in on the best fit point to demonstrate that the reconstruction translation does not introduce any bias.

8

A superior but more computationally intensive method of the near-to-far translation is to use response pairs to train a ML model that predicts the FD detector response given a ND detector response. With this response translation, the FD prediction is a function of the full phase space of the neutrino interaction present in data. The topology of the interaction in the predicted FD response is set by the inputted ND response with the translation only correcting for the detector effects.

In this chapter, the method of a response translation is outlined and the development of ML models to implement it is discussed. The translation is validated through reconstruction of predicted FD responses and a study into the dependence of the method on cross section systematics is performed.

## 8.1 Proposed Method

Pile-up removed selected ND data

Muon tracker reconstruction

ND-LAr response

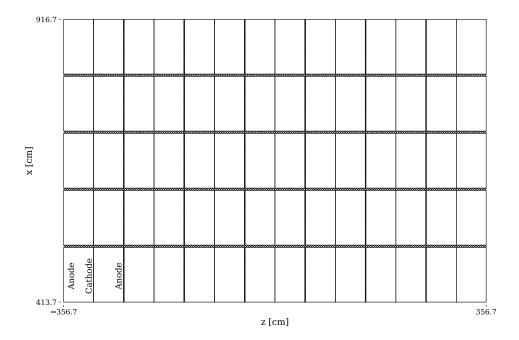
Correct for inactive regions

Predict FD response

FD analysis variables

**Figure 8.1:** Flow diagram of the high-level steps of the response translation that predicts FD analysis variables from a ND data event. The steps shaded in green require the development of new algorithms.

The diagram in Figure 8.1 outlines the steps of the proposed response translation. The translation starts with selected single neutrino interactions in the ND. Obtaining single interactions requires some level of reconstruction to remove pile-up. Once this has been done, the raw detector response, rather than a reconstructed object, for single neutrino interactions is returned. Two algorithms must be developed to translate each event to the



**Figure 8.2:** Top-down view of the ND-LAr drift modules. The hatched regions are the inactive volumes at the anode ( $\sim 1.3 \, \text{cm}$  wide) and between adjacent drift volumes ( $\sim 4.14 \, \text{cm}$  wide).

FD. The first is a correction for the inactive volumes of the modular ND-LAr. The second is a prediction of the FD electronics response equivalent to the ND electronics response in the data. Equivalent meaning the FD response as if the energy depositions present in the ND interaction had been deposited in the LAr of the FD.

The inactive volumes of ND-LAr are a result of the modular nature of the detector. Between each drift module there are thin volumes of LAr and other material in the anode where charge is not drifted to a readout plane. These are shown schematically in Figure 8.2. The FD also has inactive volumes at the anode and between adjacent APAs but these are spaced much further apart than at the ND. For this reason, a correction for the inactive ND volumes must be made for each event. It is proposed to do this by using computer vision algorithms to predict the electronics response in the inactive volumes as if the charge in these volumes had been deposited in LAr and drifted to a readout plane. This process of predicting the detector response in the inactive volumes by interpolation of the response in the neighbouring active volumes is referred to as infill. An infill with adequate accuracy will result in contiguous detector response across the inactive volumes so that there is little impact on the FD track and shower reconstruction of the translated detector response. In addition, it will give an approximate correction for the missing deposited energy based on the topology of the event. Performing the infill on the raw detector response prevents smearing from the ND reconstruction being folded into the translation.

After infilling the ND-LAr detector response, the equivalent FD response needs to be predicted. Currently, a muon reconstructed in ND-GAr/TMS is not considered. The FD

response is predicted only for the hadronic component of the event and the fraction of the muon track contained in ND-LAr. Steps to predict the FD response for the entire event by including the ND-GAr/TMS reconstruction are outlined at the end of this section. For the component of the event contained in ND-LAr, the 3D detector response first needs to be transformed into the coordinate system of the FD. This requires a reconstructed interaction vertex and is accomplished with the following steps:

- Rotate the ND event about a vertical axis intersecting the interaction vertex such that the beam direction is consistent with an on-axis measurement.
- Place in the FD geometry with a randomly and uniformly sampled interaction vertex within the fiducial volume.
- Apply the Earth's curvature correction rotation.
- Project the 3D detector response onto the FD wire planes.

These transformation were discussed in more detail in §6.1.3. With the ND event in the FD coordinate system, the prediction of the equivalent FD response is treated as an image-to-image translation problem and performed using ML. A model is trained using a dataset of response pairs, the creation of which was discussed in Chapter 6 where an example pair was shown in Figure 6.5. A separate model is used for each of the three types of wire plane: U, V, and Z.

The predicted FD response can then be loaded into the appropriate FD software object and passed through any FD reconstruction chain. From this reconstruction, any FD analysis variable can be obtained.

The response translation will be applied to selected ND data events. For events with a muon contained in ND-LAr, the electronics response translation alone is sufficient. For the many events selected with a muon reconstructed in the ND-GAr/TMS downstream tracker, a prescription to include in the translated FD response the component of the muon that is outside of ND-LAr is required. This may be done using the muon kinematics reconstructed by the downstream tracker and ND-LAr to simulate the remaining part of the muon in the FD. The new muon would be simulated at the reconstructed ND-LAr exit point of the original muon just prior to the projection of the 3D response to the FD wire planes. The FD detector simulation would be applied to the new energy depositions and the resulting response concatenated with the model predicted FD response. When a full reconstruction of ND-LAr together with the downstream tracker becomes available, implementing a complete response translation of events with a muon in the downstream tracker via this method will be an important next step.

## 8.2 ND-LAr Inactive Volume Infill

A ML model is trained to infill the ND event, predicting the ND-LAr detector response in the inactive volumes. The model is applied directly to the 3D ND-LAr detector response to minimise any directional ambiguities. The development of this model is discussed in this section.

While it may in principle be possible to incorporate the infill into the 2D near–far image translation task described in the next section, this approach was found to perform poorly in practice. For this reason, the dedicated 3D infill model described in this section was developed.

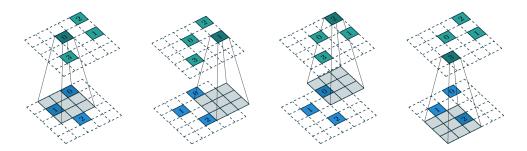
### 8.2.1 Sparse Tensor Networks

The key building block of the infill network will be the convolutional neural network (CNN) layer. CNNs typically employ kernels with two spatial dimensions and are applied to pixels. A CNN may also operate in 3D by applying kernels with three spatial dimensions to voxels. However, data from ND-LAr is encoded by 3D images with  $\mathcal{O}(10^9)$  voxels and so applying such 3D convolutions is prohibitively expensive both in terms of memory and processing time. For comparison, 3 FD APAs would encode a much larger volume with just  $\mathcal{O}(10^7)$  pixels from the wire readout. The challenge of ND-LAr's high spatial resolution is addressed by encoding the data as sparse tensors and utilising sparse tensor network libraries to perform convolutions. Since ND-LAr data has high spatial sparsity, i.e. many of the voxels are associated with zero ADC, a sparse convolution which ignores zero voxels will greatly reduce the number of computations. This is true for LArTPC data more generally and has been recognised as a powerful approach to applying ML in this domain [151].

Sparse tensor network libraries implement all of the standard neural network layers to operate on sparse tensors and produce sparse tensors as output. The PyTorch-based Minkowski Engine [152] sparse tensor network library is used in this work. In Minkowski Engine, a sparse tensor is represented in the coordinate list format,

$$\mathbf{C} = \begin{bmatrix} x_1^1 & x_1^2 & \dots & x_1^D \\ \vdots & \vdots & \ddots & \vdots \\ x_N^1 & x_N^2 & \dots & x_N^D \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} f_1^1 & f_1^2 & \dots & f_1^{N_f} \\ \vdots & \vdots & \ddots & \vdots \\ f_N^1 & f_N^2 & \dots & f_N^{N_f} \end{bmatrix},$$
(8.2.1)

where  $\mathbf{C} \in \mathbb{Z}^{N \times D}$  is a coordinate matrix,  $\mathbf{F} \in \mathbb{R}^{N \times N_f}$  is a feature matrix, D is the dimension of the space, N is the number of non-zero elements in the sparse tensor, and  $N_f$  is the number of channels. The coordinate list format is implemented as an unordered map with the D-dimensional integer coordinate as the key and the row index as the value. This allows for efficient random access to coordinates and their feature vectors along with



**Figure 8.3:** Illustration of the generalised convolution applied to a sparse tensor for four output coordinates (green) and three input coordinates (blue). Figure from [154].

access to neighbouring coordinates of a given coordinate as required for convolutions.

Minkowski Engine's convolutional layer generalises the conventional convolution on dense tensors to one that can be applied to sparse tensors [153]. A convolution consists of kernel weights,  $\mathbf{W} \in \mathbb{R}^{K^D \times N^{\text{out}} \times N^{\text{in}}}$ , that map an input feature matrix,  $\mathbf{F}^{\text{in}} \in \mathbb{R}^{N \times N^{\text{in}}}$ , to an output feature matrix,  $\mathbf{F}^{\text{out}} \in \mathbb{R}^{N \times N^{\text{out}}}$ , using a kernel of size K. Denoting the row of the feature matrix for the spatial coordinate  $\mathbf{u} \in \mathbb{Z}^D$  as  $\mathbf{f}^{\text{in/out}}_{\mathbf{u}} \in \mathbb{R}^{N^{\text{in/out}}}$  and the submatrix of the kernel weights for a given kernel coordinate  $\mathbf{i} \in \mathbb{Z}^D$  as  $\mathbf{W}_{\mathbf{i}} \in \mathbb{R}^{N^{\text{out}} \times N^{\text{in}}}$ , the generalised convolution can be written as

$$\mathbf{f}_{\mathbf{u}}^{\text{out}} = \sum_{\mathbf{i} \in \mathcal{K}(\mathbf{u}) \cap \mathcal{C}^{\text{in}}} \mathbf{W}_{\mathbf{i}} f_{\mathbf{u}+\mathbf{i}}^{\text{in}} \text{ for } \mathbf{u} \in \mathcal{C}^{\text{out}},$$
 (8.2.2)

where  $C^{\text{in}}$  is the set of input coordinates,  $C^{\text{out}}$  is the set of desired output coordinates, and  $\mathcal{K}(\mathbf{u}) \cap C^{\text{in}}$  is the intersection of the set of kernel coordinates for the kernel at the spatial coordinate  $\mathbf{u}$  with  $C^{\text{in}}$ . In this way, all operations that involve purely non-active coordinates are skipped. An illustration of the generalised convolution is shown in Figure 8.3 which, from left to right, has intersections  $\mathcal{K}(\mathbf{u}) \cap C^{\text{in}}$  given by  $\{(-1,0),(1,1)\}$ ,  $\{\}$ ,  $\{(-1,0)\}$ , and  $\{(1,0)\}$ . The high efficiency of the generalised convolution when applied to sparse input data makes the application of CNNs to ND-LAr data tractable. Further information on sparse tensor networks and their implementation in Minkowski Engine can be found in Chapter 4 of [154].

### 8.2.2 Data Preparation

Different configurations of particle generation and propagation are passed through the larnd-sim detector simulation to produce multiple training datasets. These are summarised in Table 8.1 where they are in approximate order of infill difficulty. It is hoped that by training on each of them consecutively the model will reach a better parameter space minimum. This technique is known as *curriculum learning* [155]. A dataset of 30,000 events is made for each configuration.

The configurations with the prefix dummy are made by choosing two spatial coordinates randomly and uniformly with the associated constraint stated in the description. Energy

Name	Particles	Energy	Description
dummy_fixzfixy	_		Straight line, z & y constant
dummy_fixz			Straight line, z constant
dummy	_		Straight line
single_mu	$1~\mu^-$	$0.2–3\mathrm{GeV}$	Single muon
multi_mu	$4~\mu^-$	$0.2–3\mathrm{GeV}$	$4 \mu^-$ from same vertex
multi_mu_pi	$2~\mu^-,~1~\pi^-$	$0.2–3\mathrm{GeV}$	$2 \mu^- \& 1 \pi^-$ from same vertex
nu	Beam flux	Beam flux	On-axis ND FHC beam neutrino flux

**Table 8.1:** Configurations of energy depositions in LAr used to make training datasets.

depositions are distributed in a straight line between the two coordinates with constant segment size and energy deposited at each segment. The configurations of  $\mu^-$  and  $\pi^-$  are made using the Geant General Particle Source with initial directions in a cone oriented along the neutrino beam direction. The energies are randomly and uniformly sampled from the specified energy ranges.

A coordinate system for the sparse tensor is made by discretising the 3D larnd-sim input into voxels. For the directions in the plane of the anode, the natural voxel size is the pixel pitch, 0.38 cm. In the drift direction, the  $0.1 \mu s$  temporal resolution of the electronics gives a spatial resolution for the readout of 0.016 cm. The voxel size in this direction is set to 0.16 cm to minimise the distance between subsequent self-triggered packets while ensuring that no two packets occupy the same voxel. Each drift module is therefore discretised into  $0.38 \times 0.38 \times 0.16$  cm voxels. At the inactive volumes, the discretisation uses voxel sizes as close to this as possible while maintaining an integer number of voxels. For the inactive volumes between adjacent drift modules, the gap is 11 voxels across with a voxel size of  $0.377 \times 0.38 \times 0.16$  cm. For the inactive volumes at the anode, the gap is 8 voxels across with a voxel size of  $0.38 \times 0.38 \times 0.1625$  cm. The complete discretisation of the active and inactive regions of ND-LAr consists of  $1324 \times 800 \times 4465$  voxels for  $\mathcal{O}(10^9)$ voxels in total. The voxel occupancy of the beam neutrino dataset for this discretisation is shown in Figure 8.4. This highlights the extreme sparsity of the ND-LAr response, demonstrating the suitability of a sparse tensor representation.

The feature matrix uses one channel,  $\mathbf{F} \in \mathbb{R}^{N \times 1}$ , which is the pedestal subtracted ADC of the ND-LAr response. Combined with the aforementioned voxelisation scheme, sparse tensors for ND-LAr events are formed.

To make infill training data from the simulated ND-LAr response, a self-supervised approach is taken by applying a randomly shifted mask of the inactive volumes to the 3D images. The masked image is then the input and the unmasked image is the target. An example is shown in Figure 8.5. With this training data, the model predicts the response in masked regions with the same structure as the inactive volumes of the detector. Although not utilised in this work, the mask may also be applied to ND-LAr response from data rather than simulation. Doing so in a future application of the infill network when ND-

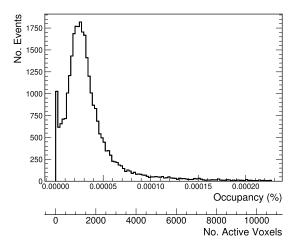


Figure 8.4: Occupancy of the beam neutrino dataset.

LAr data is available would remove any effect the data-MC discrepancy may have on the infill.

#### 8.2.3 Architecture

A model is constructed based on the U-Net architecture [156] using convolutional layers. A U-Net is an autoencoder network consisting of an encoder and decoder. The encoder uses strided convolutions to downsample the image to a tensor with small spatial size and a large feature dimension. This allows features of the image to be learnt at different spatial scales and gives the network a large receptive field. The decoder uses strided transposed convolutions to upsample the tensor back to the original spatial dimension of the input. To help in restoring the input image resolution, the U-Net uses skip connections. These are concatenations of the tensor in each layer of the encoder to the tensor in the decoder layer with the same spatial dimension. The U-Net is a clear choice for the infill task since the downsampling means the network has a large spatial context around the masked regions while the upsampling with skip connections allows restoration of the original image dimension with minimal loss of resolution.

The U-Net is made with sparse tensor network layers from Minkowski Engine. An illustration of the architecture is shown in Figure 8.6. The core convolutional block used for feature extraction consists of a convolutional layer, a normalisation layer, and a non-linearity. This block can be made to downsample or upsample by replacing the convolution with a strided convolution or a strided transposed convolution respectively. An encoder block consists of a downsampling block preceded and succeeded by a convolutional block. A decoder block consists of an upsampling block followed by a concatenation in the feature dimension from the skip connection in the feature dimension and is also preceded and succeeded by a convolutional block. The U-Net is composed of an initial convolutional block, 6 encoder blocks, 6 decoder blocks, and a final convolutional block. Relevant hyperparameters obtained by hand-tuning are listed in Table 8.2. The final hard

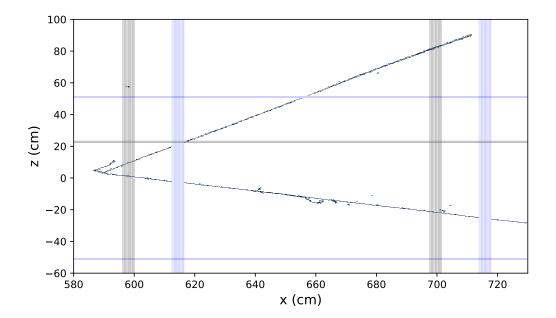


Figure 8.5: Example of masking used to make infill input-target training data. A crop of the x-z projection of the 3D data is shown. The grey shaded region is a mask produced by randomly shifting the inactive volumes displayed as blue shaded regions.

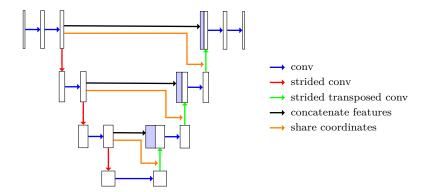


Figure 8.6: Schematic of U-Net architecture implemented with Minkowski Engine. The width and height of the blocks are indicative of the tensor's feature and spatial dimension respectively.

hyperbolic tangent non-linearity maps negative values to zero to facilitate the prediction of a masked voxel as empty.

Convolutions in the network that do not change the spatial dimension use the input coordinates as the output coordinates,  $C^{\text{out}} = C^{\text{in}}$  in (8.2.2). The strided convolutions of the encoder have new downsampled output coordinates,  $|\mathcal{C}^{\text{out}}| < |\mathcal{C}^{\text{in}}|$ , that are automatically generated by Minkowski Engine. The transposed convolutions of the decoder have output coordinates explicitly set to the coordinates of the corresponding sparse tensor at the encoder. This coordinate sharing is shown in Figure 8.6. As a consequence, the network cannot generate new coordinates and so any coordinates in ND-LAr's inactive regions that are relevant to the infill must be included in the input to the network. The method for doing this is discussed in the next section. Although it is possible for the transposed convolution to generate new coordinates [158], it was found to be challenging

Hyperparameter	Value
No. encoder/decoder layers	6
Max feature dimension	512
Normalisation	Instance
Non-linearity	Exp. linear unit [157]
Upsampling/downsampling conv. stride	2
Conv. Kernel size	3
Initial/final conv. kernel size	5
Final non-linearity	Hard tanh w/ min. val. 0

Table 8.2: Infill U-Net hyperparameters.

to limit the number of coordinates generated, and so the GPU memory requirement, by successive upsampling.

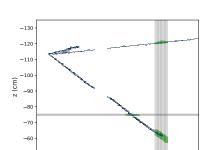
### 8.2.4 Training and Sampling

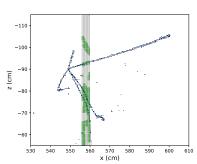
For each training epoch, a new randomly shifted mask of the inactive volumes is used for each event to make a unique input and target for the infill task. The random shift is constrained such that the mask is at least two mask widths from the detector's inactive volumes. This is to ensure the context necessary for the infill is not disrupted by missing response at the original inactive regions.

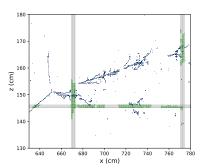
As noted in §8.2.3, the upsampling convolutions of the network cannot generate new coordinates. Any coordinates in the masked region of the output must be explicitly included as coordinates, with zero associated ADC in the feature matrix, in the input. Including the coordinates of the entire masked volume is untenable as a masked volume corresponding to a single gap between drift volumes or at an anode is composed of  $\mathcal{O}(10^6)$  voxels. Instead, candidate coordinates in the masked volumes are generated using a series of reflections to extend the detector response from the regions neighbouring the masked volumes into the masked volumes. This is done by finding reflection targets, unmasked coordinates within one mask width of the start or end of a masked region. A reflection target (x, y, z) undergoes a parity transformation through each unmasked coordinate (x', y', z') adjacent to the masked region being considered and within a maximum distance to the reflection target,

$$(x, y, z) \to (2x' - x, 2y' - y, 2z' - z).$$
 (8.2.3)

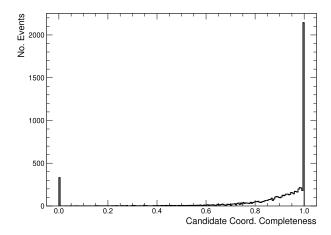
These transformed coordinates lie within the masked region and are the candidate infill coordinates. A small smearing is applied to them in each direction. The transformations extend tracks into the masked region in all possible directions given by combinations of coordinates adjacent to and in the vicinity of the masked region. Some examples of the resulting candidate coordinates are shown in Figure 8.7.







**Figure 8.7:** Examples of candidate coordinates (green shading) generated in the masked regions. The candidate coordinates are inputted to the infill network as voxels with zero ADC alongside the unmasked ND-LAr response. Shown are crops of the x-z projection of the 3D events.



**Figure 8.8:** Fraction of active coordinates in masked volumes where a candidate coordinate is generated via reflections of the unmasked coordinates.

As demonstrated in Figure 8.8, the majority of the masked detector response is included in the candidate coordinates. Detector response outside the candidate coordinates, and so not possible to reconstruct with the infill network, is by definition activity without a clear signature in the unmasked volumes. This method of generating candidate coordinates typically produces  $\mathcal{O}(10^4)$  coordinates, two to three orders of magnitude lower than including entire masked volumes in the input tensor, which can be included in the input to the infill network without exceeding memory limitations.

An important consideration when constructing a loss function for training is the sparsity of the detector response in the drift direction z. Successive data packets associated with a continuous charge deposition are separated by a few voxels in the z-direction due to the self-triggering. A composite loss function is designed to accommodate the difficulty in reliably predicting the correct specific infill voxel from the unmasked detector response. A convolution with kernel size  $1 \times 1 \times 3$  is applied to the target image to produced smeared

features that are denoted with a prime. This is used to define two loss terms,

$$\mathcal{L}_{\text{voxel}} = \frac{1}{|\mathcal{C}^{\text{mask},0}|} \sum_{\mathbf{i} \in \mathcal{C}^{\text{mask},0}} \left( f_{\mathbf{i}}^{\text{pred}} - f_{\mathbf{i}}^{\text{target}} \right)^2, \tag{8.2.4}$$

$$\mathcal{L}_{\text{sum}} = \frac{1}{(|\mathcal{C}^{\text{mask},1}|)^2} \left( \sum_{\mathbf{i} \in \mathcal{C}^{\text{mask},1}} f_{\mathbf{i}}^{\text{pred}} - \sum_{\mathbf{i} \in \mathcal{C}^{\text{mask},1}} f_{\mathbf{i}}^{\text{target}} \right)^2, \tag{8.2.5}$$

where  $f_{\mathbf{i}}$  is the ADC at coordinate  $\mathbf{i}$  and, with  $\mathcal{C}^{\text{mask}}$  being the set of all masked coordinates, the coordinate sets used in the summations are  $\mathcal{C}^{\text{mask},0} = {\mathbf{i} \in \mathcal{C}^{\text{mask}} | f_{\mathbf{i}}^{\prime \text{target}} = 0}$  and  $\mathcal{C}^{\text{mask},1} = {\mathbf{i} \in \mathcal{C}^{\text{mask}} | f_{\mathbf{i}}^{\prime \text{target}} \neq 0}$ . The loss function is then  $\mathcal{L} = \mathcal{L}_{\text{voxel}} + \alpha \mathcal{L}_{\text{sum}}$  where  $\alpha$  is a hyperparameter that is set to 2 after tuning. This loss is designed to smear out any voxel-to-voxel variations and prioritise the aggregate properties of the predicted tracks.

Models are trained on 20,000 events until the validation loss indicates convergence. The stochastic gradient descent optimiser with momentum is used and training is performed with a batch size of 12. Using the datasets in the order presented in Table 8.1, a model is trained with the dummy\_fixzfixy dataset and the learnt weights are used to initialise training with the succeeding dummy\_fixz dataset. These learnt weights are used to initialise training with the succeeding dummy dataset and so on until until training with the nu dataset yields the final model.

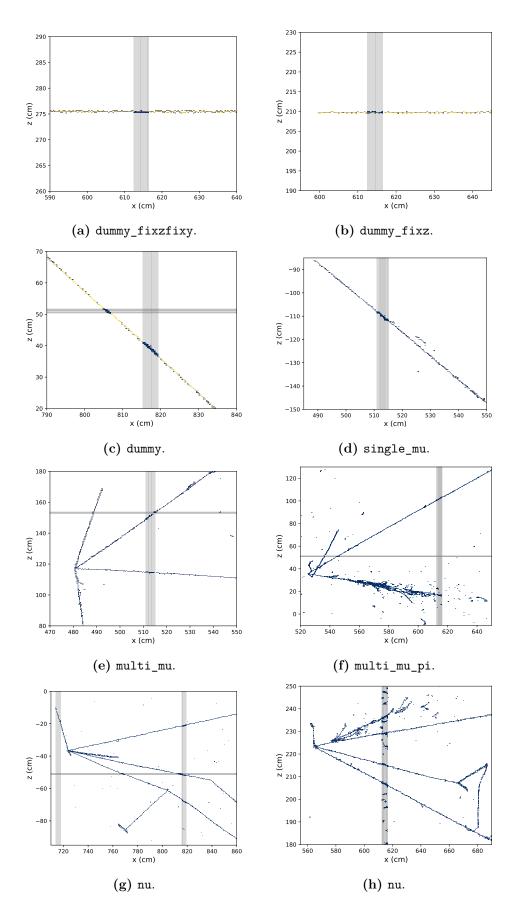
To use the trained model for infill of the inactive volumes, the model is sampled with no shift being applied to the inactive volume mask, i.e. the masked volume is the inactive volume. The predicted voxel response in the masked volume is mapped back to 3D detector coordinates and concatenated with the original active volume response. The result is an infilled 3D ND-LAr detector response.

#### 8.2.5 Results

Examples of the infill with the mask positioned at the inactive volumes for each stage of training with datasets of increasing complexity are shown in Figure 8.9. The infill generally has good accuracy on clear and isolated tracks. In regions with showers and overlapping activity, there are many false positive voxels in the infill.

The purity and completeness of the model at each successive training stage with a different dataset is shown in Table 8.3. These metrics consider the predicted voxel occupied if the ADC is non-zero and likewise for the target voxel. Completeness is the fraction of true occupied voxels the model predicts and purity is the fraction of predicted occupied voxels that are present in the truth. The final model has a reasonably high completeness and a low purity. As illustrated in Figure 8.9, the low purity is due to the infill not reproducing the sparsity of the detector response packets in the z-direction as well as many false positive predictions in regions of high activity, as in Figure 8.9h.

The summed ADC at the masked volumes for the predicted and target image is compared in Figure 8.10. The total predicted ADC in the masked volumes is several times



**Figure 8.9:** Examples of predicted inactive volume infill for each stage of training used to produce the final neutrino event infill model. Shown are crops of the x-z projection of the 3D events.

Model	Mean Purity	Mean Completeness
dummy_fixzfixy	55.26%	93.07%
dummy_fixz	43.76%	81.50%
dummy	16.38%	79.38%
single_mu	21.01%	71.87%
multi_mu	4.47%	82.15%
multi_mu_pi	6.35%	62.10%
nu	14.08%	74.76%

**Table 8.3:** Voxel occupancy purity and completeness for models at each stage of training.

larger than the target ADC. Although much better than the simple baseline infill, the infill is inaccurate in this regard.

The performance of the infill in isolation is fairly poor. However, it does still accomplish the goal of ensuring contiguous detector response across the inactive volumes. It is hoped that in the next stage of predicting the FD response from the infilled ND response, the impure infill will be sufficient.

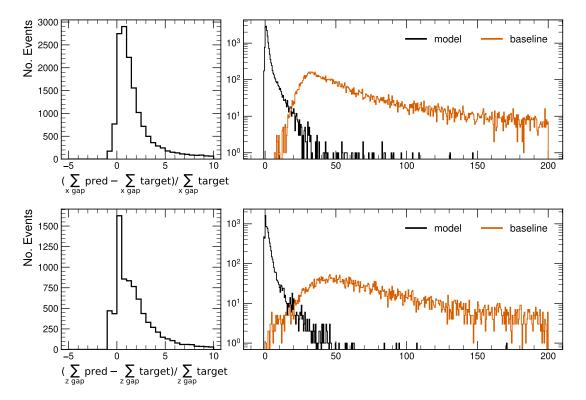
#### **Near-Far Image Translation** 8.3

Following the infill of the ND-LAr detector response, the response translation is completed by transforming the ND event into the FD coordinate system and predicting the equivalent FD response. The prediction of a FD response given a ND response is treated as an image-to-image translation between two domains of distinct detector technology. The development of a model to perform the style-transfer between the near and far is discussed in this section.

#### **Data Preparation** 8.3.1

As outlined in §8.1, the response-pairs discussed in Chapter 6 are used to create training data. Before projecting the ND-LAr response to the wire and time tick coordinates of the FD in the pair making procedure, the infill network is used to predict the detector response in the inactive volumes. This predicted detector response is projected alongside the original active volume response. With this amendment, the response-pairs are used to generate datasets of 50,000 near-far images corresponding to the response at a single APA for each of the U, V, and Z wire planes. The images are kept at their original resolution of  $480 \times 4492$  for the collection plane Z and  $800 \times 4492$  for the induction planes U and V. Unlike the TDR simulation used for the reconstruction translation dataset, the response pairs use a more recent version of FD software that utilises a 2D field response in the detector simulation.

Each pixel of the FD image has a single channel, the ADC. The ND image is prepared



**Figure 8.10:** Fractional residual of the summed ADC in separate masked volumes for predicted and target. Shown for masked volumes associated with adjacent drift volumes in the x-direction (top) and with the anode in the z-direction (bottom). The baseline is to set the value of all candidate masked coordinates to the mean value of the unmasked detector response.

with multiple channels:

### Ch. 1. ADC.

- Ch. 2. ND drift distance. This encodes drift electron attenuation and diffusion effects in the ND.
- Ch. 3. FD drift distance. This encodes drift electron attenuation and diffusion effects in the FD.
- Ch. 4. Shortest distance from the centroid of the ND packet to the wire it is projected to. This encodes some of the effect of downsampling to the lower resolution FD coordinate system.
- Ch. 5. Flag for if the ND packet is a prediction of the infill network.
- Ch. 6. Count of distinct ND packets that are projected onto the same pixel.

Since the ND has a slightly higher spatial resolution than the FD, multiple ND packets will sometimes be projected onto the same pixel. In these cases, the ADC channel is the sum of the constituent ND packets and channels 2 to 5 are the ADC weighted average of the constituent ND packets. An example image illustrating each channel is shown in Figure 8.11.

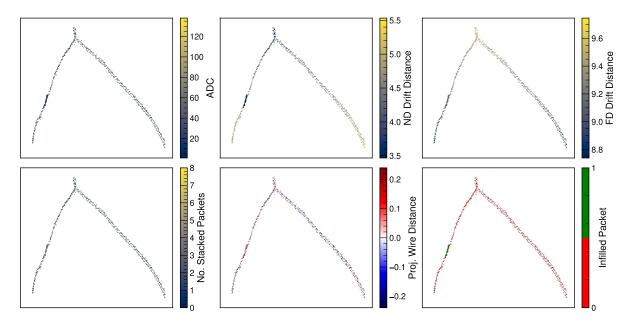


Figure 8.11: Each channel for a cropped example ND image from the Z plane dataset.

### 8.3.2 Architecture

A model for the response translation is developed that is loosely based on two important works in the field of paired image-to-image translation [159, 160]. The model architecture consists of a deep residual convolutional network (ResNet) [161] sandwiched between encoder and decoder convolutional layers. This is illustrated in Figure 8.13. The encoder and decoder blocks, which use the same strided and transposed convolutions discussed in §8.2.3 for downsampling and upsampling respectively, provide the network with a moderate receptive field while the ResNet blocks are the main feature extracting component. This architecture is chosen over the U-Net architecture used for the infill network since it is not desirable to have a large receptive field that grants each pixel the full spatial context of the interaction for predicting the FD response. Only the region local to the each pixel is required to encode the difference between electronics response between the two detectors.

Relevant hyperparameters for the model obtained by hand-tuning are listed in Table 8.4. The final non-linearity uses a hyperbolic tangent function that is clamped such that the predicted pedestal subtracted FD 12-bit ADC is within its allowed range as shown in Figure 8.12.

### 8.3.3 Training

To control the impact on the training loss of large regions of the FD images that are electronics noise only, a signal mask is applied to the predicted and target images as part of the loss function calculation. The signal mask is defined as the non-zero pixels of the ND image after a convolution with a kernel of size  $8 \times 60$  for the collection plane dataset and  $8 \times 100$  for the induction plane datasets. This is a conservative definition of the signal

Hyperparameter	Value
Padding	Reflection
Conv. kernel size	3
Dropout probability	0.5
Max feature dimension	512
Upsampling/downsampling conv. stride	2
Initial/final conv. kernel size	7
Final non-linearity	Tanh clamped to 12-bit ADC

**Table 8.4:** Response translation network hyperparameters.

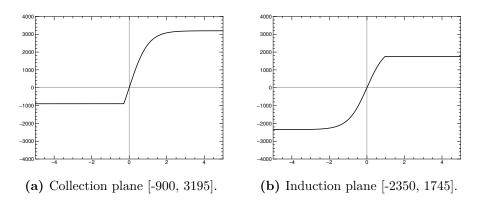


Figure 8.12: Tanh functions clamped for the pedestal subtracted FD ADC ranges of the collection plane and induction planes.

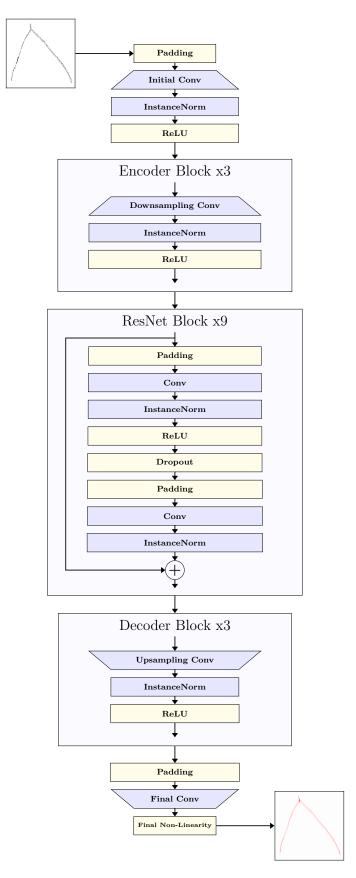
region for a single pixel. When such a mask is applied, all FD response induced by drift electrons is kept along with a significant, but much reduced, contribution from regions of electronics noise only. An example signal mask is shown in Figure 8.14.

Pixel-wise and wire-wise loss terms are defined as

$$\mathcal{L}_{\text{pix}} = \frac{1}{N_{\text{pix}}} \sum_{i,j} |y_{i,j} - x_{i,j}|, \qquad (8.3.1)$$

$$\mathcal{L}_{\text{wire}} = \begin{cases}
\frac{1}{N_{\text{wire}}} \sum_{i} \left| \sum_{j} y_{i,j} - \sum_{j} x_{i,j} \right| & \text{for Z dataset} \\
\frac{1}{2N_{\text{wire}}} \left( \sum_{i} \left| \sum_{j} y_{i,j} [x_{i,j} < 0] - \sum_{j} x_{i,j} [x_{i,j} < 0] \right| \\
+ \sum_{i} \left| \sum_{j} y_{i,j} [x_{i,j} \ge 0] - \sum_{j} x_{i,j} [x_{i,j} \ge 0] \right| \right) & \text{for U,V datasets} \end{cases} , (8.3.2)$$

where i and j are the non-zero wire and time tick indices of the signal mask,  $x_{i,j}$  and  $y_{i,j}$ are the target and predicted ADC at a pixel respectively,  $N_{pix}$  is the number of non-zero pixels in the signal mask, and  $N_{\text{wire}}$  is the number of wires in the signal mask with at least one non-zero pixel. The pixel-wise term is a mean absolute error (MAE) typical of computer vision tasks while the wire-wise loss encourages the ADC aggregated over a wire to be accurate. For the induction plane datasets, the wire-wise loss is split into



**Figure 8.13:** Diagram of the ResNet architecture used for the prediction of FD response from ND response. Layers with learnable parameters are shaded blue while fixed layers are shaded yellow. The width of a layer is indicative of the size of the feature dimension at that point in the network.



**Figure 8.14:** An example ND image (left), the associated signal mask (middle), and the associated FD image (right). Shown for a crop of near-far pair of the U plane dataset.

**Table 8.5:** MAE at the signal mask for each of the trained models over their corresponding test datasets. The MAE of the target FD images with themselves using distinct random seeds for simulating electronics noise is included for reference.

Model	MAE at Signal Mask	MAE due to Electronics Noise
U Plane	$5.32 \pm 0.71$	$4.91 \pm 0.01$
V Plane Z Plane	$5.22 \pm 0.67$ $6.26 \pm 2.25$	$4.91 \pm 0.01 4.25 \pm 0.01$

components of positive and negative target ADC to account for the bipolar nature of the signals. A loss function for training is constructed as  $\mathcal{L} = \mathcal{L}_{pix} + \alpha \mathcal{L}_{wire}$  where  $\alpha$  is a hyperparameter that is set to 0.05 after tuning.

Models for each dataset are trained on 30,000 images until the validation loss indicates convergence, taking approximately 50 epochs. The Adam optimizer [150] with weight decay is used and training is performed with a batch size of 2, a limitation of the high image resolutions.

#### 8.3.4 Results

The predicted FD response for a ND event of each wire plane dataset is shown in Figure 8.15. Qualitatively, the networks are successful in performing the style-transfer to the FD response. The predicted waveforms have the correct shape and are clearly conditioned on the ND response. The single-wire responses in Figure 8.15b demonstrates that the networks can further refine the infill by ignoring spurious infilled ND response.

Using test datasets of 20,000 paired images, the MAE at the non-zero pixels of the signal mask over each of the trained models is shown in Table 8.5. Comparing them to the MAE expected from different realisations of the electronics noise, these results demonstrate an accurate prediction of the FD response at the individual pixel level.

The key metrics for the similarity of the predicted and target response come from applying the reconstruction. The reconstructed quantities relevant to oscillation analyses must match closely between the predicted and target FD responses for a ND event. This is examined in the next section.

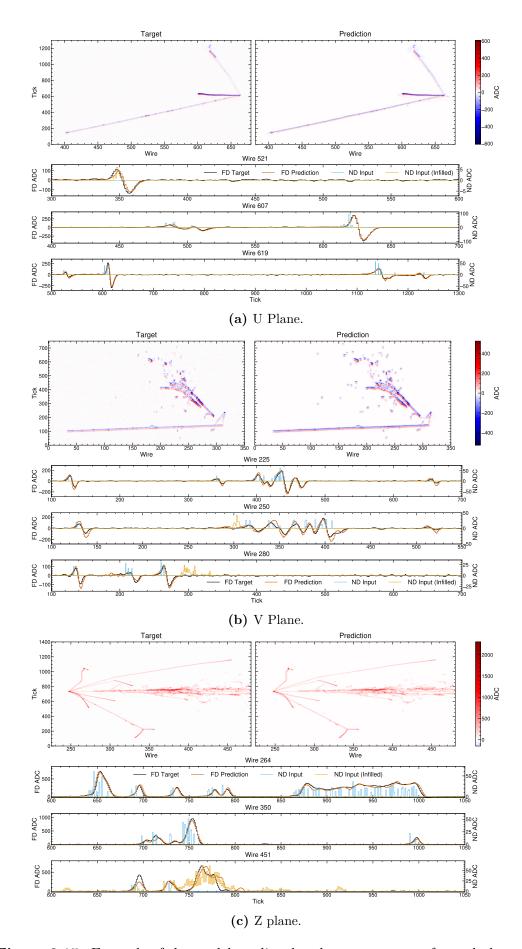


Figure 8.15: Example of the model predicted and target response for each dataset.

# 8.4 Validation with Reconstruction

With trained infill and response translation models, the response translation outlined in Figure 8.1 can be implemented to generate predictions of the FD reconstruction. In this section, these predictions are compared to target FD reconstruction.

## 8.4.1 Generating Samples

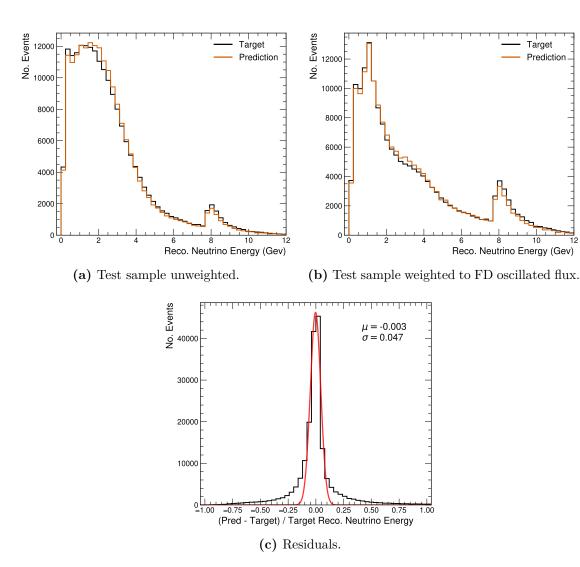
Following the steps outlined in § 6.1, a sample of on-axis FHC ND beam neutrino interactions is generated up to the stage of ND-LAr detector simulation. The infill network is applied to the 3D detector response and the output is transformed and projected to the FD wire and tick coordinate system as outlined in § 8.1. The relevant response translation model is applied to each readout plane of each APA of the projection to produce a prediction of the full FD response to the neutrino interaction contained in ND-LAr. These predicted images are converted to the software object that encodes the electronics response at each wire of the FD. Signal processing and the subsequent reconstruction chain is then applied. The resulting reconstructed quantities are the predicted FD reconstruction.

To generate the target FD reconstruction, the deposited energy in the LArBath geometry for each ND neutrino interaction is placed in the FD using the same 3D transformations applied to the associated ND-LAr detector response. Since muons reconstructed in ND-GAr/TMS are not currently included in the response translation, a mask is applied to remove energy depositions at the FD that are outside of ND-LAr in the associated ND event. The detector simulation, signal processing, and reconstruction chain is then applied to the energy depositions in the FD to produce the target FD reconstruction. Masking out the energy deposits not contained in ND-LAr means that a perfect response translation will yield predicted FD reconstruction that is equal to the target FD reconstruction up to small variations due to electronics noise.

A sample of  $\sim 200,000$  prediction-target FD reconstruction pairs is generated. Due to the mask applied in the generation of the FD target prediction, these may be compared without requiring the application of any ND containment cuts.

### 8.4.2 Results

The predicted reconstructed neutrino energy is compared to the target in Figure 8.16. The response translation produces an accurate prediction of the reconstructed neutrino energy with little bias. A comparison of the residuals (Figure 8.16c) with the equivalent result from the reconstruction translation (Figure 7.7) demonstrates the advantage of the much larger phase space of the ND event used in the response translation approach. There is very little ambiguity in the FD response prediction which facilitates accurate event-level predictions. An accurate event-level prediction ensures a close match between the predicted and target reconstructed energy distributions with minimal dependence on the

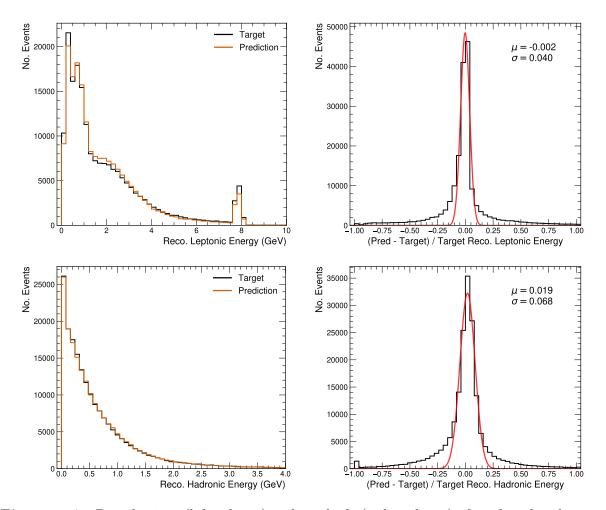


**Figure 8.16:** Distributions and residuals of predicted and target FD reconstructed neutrino energy. The distributions in (b) are weighted to a FD oscillated spectrum for NuFIT 4.0 oscillation parameters [139]. The oscillated distributions look unusual due to the removal of energy depositions outside ND-LAr discussed in § 8.4.1.

underlying true neutrino energy distribution of the training dataset. This is demonstrated in Figure 8.16b.

The hadronic and leptonic components of the reconstructed neutrino energy are shown in Figure 8.17. The reconstructed hadronic energy demonstrates that the calorimetric information of the predicted FD response is accurate. To obtain the leptonic energy, hits are clustered in each plane and matched between them to form reconstructed tracks. The muon track is identified and its energy estimated by range or MCS. The results for the reconstructed leptonic energy show that the predicted FD response is of a sufficient quality and is consistent across the three planes so that this process may proceed and produce an accurate reconstructed energy for the muon.

Figure 8.18 compares the predicted CVN  $\nu_{\mu}$  score with the target. There is a clear shift towards lower scores in the prediction. Although many events have accurate predicted



**Figure 8.17:** Distributions (left column) and residuals (right column) of predicted and target FD reconstructed leptonic (top row) and hadronic (bottom row) energy.

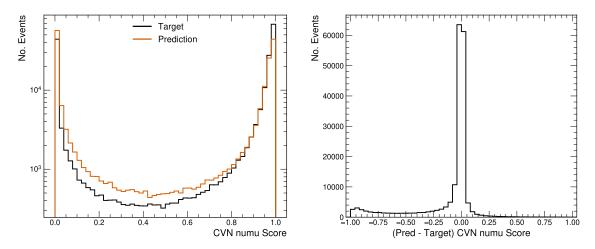


Figure 8.18: Distributions (left) and residuals (right) of predicted and target FD CVN  $\nu_{\mu}$  score.

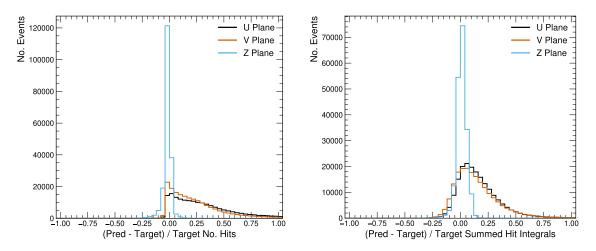
scores, this would result in a significant number of events being incorrectly predicted to not be classified as muon neutrino interactions at the FD. Despite the predicted FD response being similar enough to the target to accurately reconstruct particle tracks, there is a discrepancy that results in a deviation of the CVN score.

The CVN is a CNN that operates on images of the hits from each plane. Properties of the hits in each plane are examined in Figure 8.19. In the Z plane, the predicted unipolar response produces accurate hits. Since only the Z plane is used for calorimetry, this is consistent with the reconstructed energies in Figures 8.16 and 8.17. For the U and V planes, the predicted bipolar responses produce significantly more hits than in the case of the target response. The origin of this is not clear from the test results presented in § 8.4.2. The hits are formed by fitting the waveforms to detector response after deconvolution, as discussed in § 4.2.2, where a subtle discrepancy in the shape of induction plane signal may have a significant effect. Such a discrepancy in the predicted bipolar induction response is likely the cause of the additional hits which in turn activate learnt features of the CVN that reduce the  $\nu_{\mu}$  score.

Overall, the product of the infill and response translation networks is a predicted FD response that yields accurate reconstructed energies for each event. The predicted response at the Z plane closely resembles the FD simulation. The induction plane responses, while superficially resembling FD simulation, are not sufficiently accurate under the signal processing which leads to discrepancies in the CVN outputs.

# 8.5 Cross Section Systematics

Incorporating the response translation into the PRISM oscillation analysis is not currently possible due to the lack of a prescription for a muon reconstructed in ND-GAr/TMS and the requirement of ND-LAr detector response. Instead, a toy study into the effect of interaction model uncertainties on a detector effects extrapolation performed with the response



**Figure 8.19:** Residuals of predicted and target total number of hits (left) and sum of all hit integrals (right) in each plane.

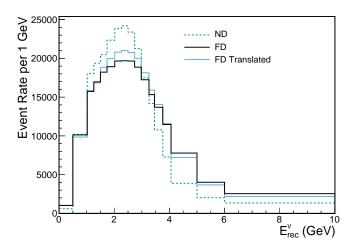
translation, as opposed to the current unfolding and smearing procedure is conducted. The study, discussed in this section, considers a single extrapolated ND spectrum under cross section systematic shifts to quantify the impact of a response translation on the overall aim of reducing the interaction model dependence of a detector effects correction.

### 8.5.1 Monte Carlo Data

To study the effect of the response translation in isolation, ND and FD MC that only differ by detector resolution are generated. The geometry used for FD simulation is restricted to a bounding box with ND-LAr dimensions such that the current response translation may predict the entire FD response.

A sample of on-axis FHC ND beam neutrino interactions is generated to obtain both ND-LAr detector response and parametrised reconstruction. The detector response is translated to the FD and reconstructed to provide a prediction of the reconstructed neutrino energy at the FD. This is saved alongside the ND event in the analysis file.

To generate a FD sample with the same flux, each GENIE neutrino interaction vertex, defining the final state particles and their kinematics, from the ND sample is placed in the fiducial volume of the FD. A bounding box with ND-LAr dimensions is placed relative to each FD interaction vertex at the same distances as the associated ND interaction vertex is to the edges of ND-LAr. An Earth's curvature correction rotation is applied to each FD neutrino interaction vertex and its bounding box. The neutrino interactions in the FD are then propagated through the LAr and any energy depositions outside of the bounding box are removed. The detector simulation and reconstruction are applied to the remaining energy depositions. This procedure yields the FD reconstruction of an incident neutrino that may have been "seen" by only ND-LAr. The resulting reconstructed neutrino energy is added to the analysis file in the same entry as the associated ND event. Each analysis file entry therefore encodes: the true neutrino interaction, the ND



**Figure 8.20:** Selected reconstructed neutrino energy event rate ND, true FD, and translated FD spectra.

parametrised reconstruction, the *translated* FD reconstructed neutrino energy obtained from the response translation, and the *true* FD reconstructed neutrino energy obtained from resimulating the interaction in the FD.

The reconstructed and true neutrino energies,  $E_{\rm rec}^{\nu}$  and  $E_{\rm true}^{\nu}$ , are used as analysis variables. The CC  $\nu_{\mu}$  ND selection criteria, described in §5.2.1, is applied to yield contained and selected ND events and their associated true and translated FD  $E_{\rm rec}^{\nu}$ . The MC sample generated has  $\sim 80,000$  selected events. Their reconstructed spectra are shown in Figure 8.20.

### 8.5.2 Extrapolating to the Far Detector

Extrapolating the ND spectrum to the FD in this contrived setup requires only a correction for the detector effects. The response translation provides this by binning the events of the ND spectrum in translated  $E_{\rm rec}^{\nu}$ . This is compared to the standard detector effects correction described in § 5.3. For the standard detector effects correction, smearing matrices for both detectors are constructed. These are shown in Figure 8.21. Since no selection efficiency correction needs to be performed, the sum over their reconstructed energy bins is normalised to unity. The ND smearing matrix is unfolded using Tikhonov regularisation, this is shown in Figure 8.22. Extrapolation to the FD is the consecutive application of the unfolded ND smearing matrix and the FD smearing matrix to the ND  $E_{\rm rec}^{\nu}$  spectrum. The extrapolation via both the response translation and the smearing is shown in Figure 8.23.

The extrapolation via the translation requires a MC correction to match the true FD spectrum well. An additive MC correction is included as shown in Figure 8.24. For a fair comparison, a similar, but much smaller, MC correction is added to the standard extrapolation procedure as well.

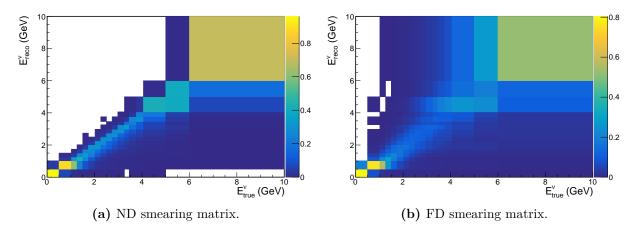


Figure 8.21: ND and FD smearing matrices.

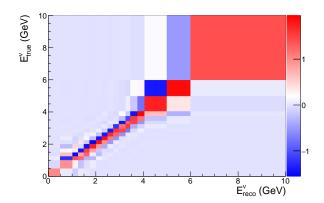
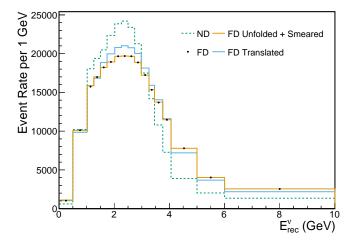


Figure 8.22: ND smearing matrix unfolded with Tikhonov regularisation ( $\tau = 0.1$ ).



**Figure 8.23:** ND spectra extrapolated to the FD using the response translation and the standard smearing procedure.

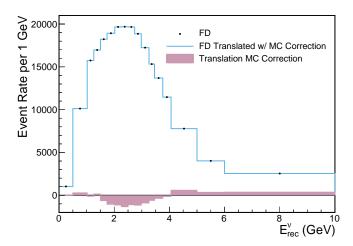


Figure 8.24: Response translation extrapolation with a MC correction.

### 8.5.3 Effect of Cross Section Systematics

To compare the response translation with the method of unfolding and smearing, the effect of cross section systematics on the extrapolated spectra is considered. Should a detector resolution correction be entirely independent of the interaction model, the extrapolated spectrum will match the FD spectrum under any shift in the interaction model as both spectra are composed of the same neutrino-argon interactions. Any discrepancy between the two that is introduced by a shift in the interaction model is attributed to the detector resolution correction.

Systematic uncertainties in the interaction model are expressed as event weights for various cross section dials. Each dial represents a systematic uncertainty and provides a set of  $\pm 1\sigma$ ,  $\pm 2\sigma$ , and  $\pm 3\sigma$  weights for each event based on properties of its neutrino-nucleus interaction. The effect of a shift in a given dial is the application of these weights to the events that comprise a spectrum. The cross section dials provided by the GENIE generator and the DUNE collaboration are implemented in the CAFAna framework [162]. There are a total of 54 dials that are broadly categorised as: quasielastic (QELike), resonance production (RES), deep inelastic scattering (DIS), final state interactions (FSI), and neutral current (NC).

The effect of each cross section dial on the extrapolation is evaluated separately by applying the reweighting to shift both the ND and FD  $E^{\nu}_{\rm rec}$  spectrum. For the translation, the translated FD  $E^{\nu}_{\rm rec}$  spectrum is composed of a prediction for each ND event and so undergoes the same event reweighting. The nominal additive MC correction, as it appears in Figure 8.24, is then applied to yield the shifted extrapolated spectrum associated with the cross section dial. To obtain the shifted extrapolated spectrum for the standard extrapolation procedure, the nominal unfolding and smearing matrix, as they appear in Figure 8.22 and 8.21b respectively, are applied to the reweighted ND  $E^{\nu}_{\rm rec}$  spectrum followed by the addition of the small MC correction. Since the extrapolated spectra

are made to match the FD spectrum perfectly in the nominal case via the additive MC correction, any fractional difference between the two after reweighting is induced by the cross section dial and is a measure of the interaction model dependence introduced by the detector resolution correction.

The fractional difference in each  $E_{\text{rec}}^{\nu}$  bin induced by all cross section dials added in quadrature is shown in Figure 8.25. The extrapolation using the response translation deviates under cross section systematics much less than the standard extrapolation procedure. There is a significant reduction in the impact from DIS dials, a moderate reduction from RES and NC dials, and negligible change in the impact from FSI and QE-Like dials. This demonstrates that, even with a moderate MC correction, the response translation yields a significant reduction in the interaction model dependence of the detector resolution correction. This reduction is a consequence of the data-driven aspect of the response translation. The predicted FD response is unique for each ND event and encompasses the full neutrino interaction, meaning that the predicted FD responses that comprise the extrapolated spectrum will share the same interaction physics present in ND measurements.

Interaction model dependence enters the response translation through error in the FD response prediction which necessitates the inclusion of a correction derived purely from MC. In Figure 8.25, this results in significant contributions from FSI and QELike dials for the response translation. Further improvements to the accuracy of the predicted FD response would reduce the magnitude of the MC correction and suppress these contributions.

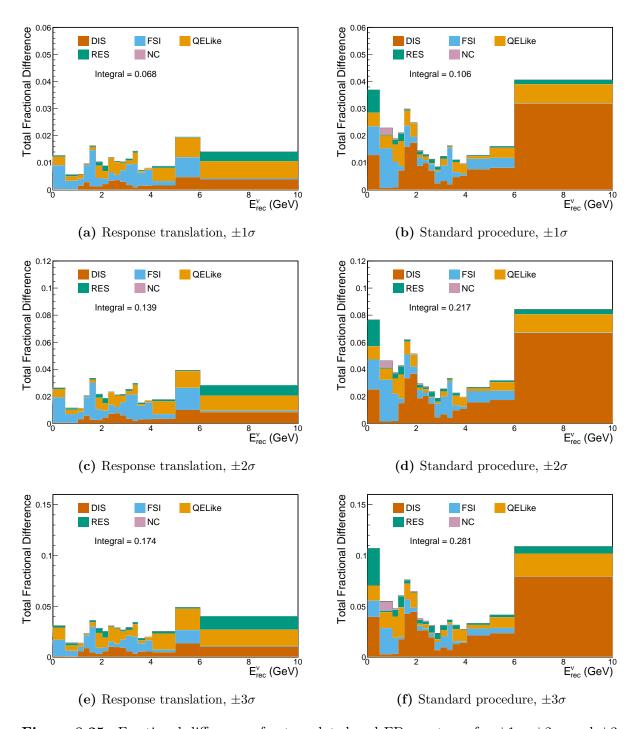


Figure 8.25: Fractional difference of extrapolated and FD spectrum for  $\pm 1\sigma$ ,  $\pm 2\sigma$ , and  $\pm 3\sigma$  shifts of all 54 cross section dials added in quadrature. Shown for the detector resolution correction using the response translation and using the standard procedure of unfolding and smearing. The total fractional difference is composed of the largest absolute fractional difference of each plus-or-minus cross section dial shift.

Conclusions and Outlook

9

This thesis presents two novel detector resolution corrections for the PRISM oscillation analysis based on translating each ND event to the FD. The method for constructing a PRISM FD prediction from off-axis ND measurements is discussed and the near-to-far translation is motivated. A paired dataset of neutrino interactions at the ND and FD is generated via careful manipulation of simulation to facilitate the training of ML models for the translation. Models for the translation of the reconstructed quantities and of the detector response are developed. Both are shown to be more robust to changes in the interaction model than the standard method of detector resolution correction.

A framework for paired dataset generation is developed that simulates single neutrino interactions up to the last point of similarity between the two detectors and subsequently applies the separate detector simulation and reconstruction chains. This framework is used to generate samples suitable for training translation models for the FHC disappearance analysis channel. These samples are shown to closely match standard ND and FD simulation.

Development of the paired dataset is ongoing with current efforts towards using the framework to produce a RHC disappearance sample and extending the framework to produce appearance samples by resimulating the ND muon as an electron in the FD, a method that is shown to be a valid approach in principle. An assumption of the paired dataset generation is that charged particle propagation in both detectors can be approximated using a monolithic LAr volume. Any additional energy lost to materials denser than LAr in inactive volumes enclosed by the active volume is treated as negligible. There is ongoing work to understand the validity of this assumption and formulate an event-level correction to the ND energy depositions where this approximation is most tenuous.

A transformer is trained on the paired dataset to predict FD analysis variables given the ND reconstruction of each event. This model predicts a conditional probability distribution as a Gaussian mixture which it subsequently samples from. The predicted distributions of FD reconstruction matches the truth closely and the conditioning on ND variables is reproduced accurately. The reconstruction translation is implemented in the PRISM oscillation analysis as the detector resolution correction used to generate the predicted FD spectrum. The resulting oscillation parameter sensitivities for a fit with cross

section systematic uncertainties demonstrate a reduction in interaction model dependence for the translation compared to the standard correction and yield a marginal sensitivity improvement.

The clearest path to improving the performance of the reconstruction translation is to retrain with a substantially larger training dataset. To facilitate this, work is currently underway to apply ND selection cuts during the paired dataset generation to avoid running FD simulation for events that will not be used in the training. The translation may also be improved by training with a dataset of neutrino interactions with a uniform true energy spectrum. This should reduce any dependence on the underlying true neutrino energy distribution to yield improved performance for the set of oscillation hypotheses and systematic shifts it is applied to. In addition, as DUNE's ND reconstruction chain matures, a study into the appropriate set of input ND variables from full reconstruction, as opposed to the parametrised reconstruction currently available, will be required.

A near-to-far translation of the detector electronics responses is developed utilising computer vision. A sparse tensor network is trained on masked ND-LAr simulation to infill detector response in the inactive volumes. The infilled ND-LAr response is then projected into the FD readout coordinate system and translated to the equivalent FD response by a model trained using the paired dataset. The style transfer to the FD is successful, producing realistic FD waveforms that are well-conditioned on the ND input and have inactive regions appropriately infilled. When the translated response is passed to reconstruction, the reconstructed neutrino energy is accurate while the CVN score is not. This is understood to be because the predicted induction plane responses lack the structure expected by the deconvolution performed as part of the signal processing. The prediction of the reconstructed neutrino energy from the response translation is used to extrapolate a ND spectrum to the FD. The behaviour of this spectrum under cross section systematic shifts is compared to a spectrum extrapolated with the standard PRISM detector resolution correction. The study shows a significant reduction in interaction model dependence when using the response translation.

Before attempting to integrate the response translation in the PRISM analysis, the predictions of the responses should be improved with emphasis on the induction planes. This may be accomplished by employing recent advances in image-to-image tasks based on diffusion [163, 164]. In addition, it may be beneficial to predict the response after signal processing where the waveforms are unipolar in all readout planes.

In summary, this thesis describes the design and implementation of a technique that reduces the interaction model dependence of the PRISM oscillation analysis and in turn contributes to addressing the challenge of systematic uncertainties in DUNE's precision oscillation measurement programme. The reconstruction translation is shown to be performant when integrated into the analysis and is undergoing active development with the aim of enhancing the robustness and sensitivity of the full four channel analysis. A model

predicting the FD electronics response given a ND interaction is developed to demonstrate the potential of the more ambitious response translation.

# References

- [1] N. Bohr. "On the Constitution of Atoms and Molecules". In: *Phil. Mag. Ser.* 6 26 (1913), pp. 1–24. DOI: 10.1080/14786441308634955.
- [2] E. Rutherford. "Bakerian Lecture: Nuclear constitution of atoms". In: *Proc. R. Soc. Lond. A* 97.686 (1920), pp. 374–400. DOI: 10.1098/rspa.1920.0040.
- [3] C. D. Ellis and W. A. Wooster. "The average energy of disintegration of radium E". In: Proc. Roy. Soc. Lond. A 117.776 (1927), pp. 109–123. DOI: 10.1098/rspa. 1927.0168.
- [4] W. Pauli. Pauli letter collection: letter to Lise Meitner. https://cds.cern.ch/record/83282. Typed copy. 1930.
- [5] W. Bothe and H. Becker. "Künstliche Erregung von Kern- $\gamma$ -Strahlen". In: Zeitschrift fur Physik 66.5-6 (May 1930), pp. 289–306. DOI: 10.1007/BF01390908.
- [6] J. Chadwick. "Possible Existence of a Neutron". In: Nature 129 (1932), p. 312. DOI: 10.1038/129312a0.
- [7] E. Fermi. "An attempt of a theory of beta radiation. 1." In: Z. Phys. 88 (1934),
   pp. 161–177. DOI: 10.1007/BF01351864.
- [8] C. S. Wu et al. "Experimental Test of Parity Conservation in  $\beta$  Decay". In: *Phys. Rev.* 105 (1957), pp. 1413–1414. DOI: 10.1103/PhysRev.105.1413.
- [9] C. L. Cowan et al. "Detection of the free neutrino: A Confirmation". In: Science 124 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103.
- [10] M. Goldhaber, L. Grodzins, and A. W. Sunyar. "Helicity of Neutrinos". In: Phys. Rev. 109 (3 Feb. 1958), pp. 1015–1017. DOI: 10.1103/PhysRev.109.1015.
- [11] R. P. Feynman and M. Gell-Mann. "Theory of the Fermi Interaction". In: Phys. Rev. 109 (1 Jan. 1958), pp. 193–198. DOI: 10.1103/PhysRev.109.193.
- [12] E. C. G. Sudarshan and R. E. Marshak. "Chirality Invariance and the Universal Fermi Interaction". In: *Phys. Rev.* 109 (5 Mar. 1958), pp. 1860–1862. DOI: 10. 1103/PhysRev.109.1860.2.
- [13] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Phys. Rev. Lett.* 9 (1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36.
- [14] K. Kodama et al. "Observation of tau neutrino interactions". In: *Phys. Lett. B* 504 (2001), pp. 218–224. DOI: 10.1016/S0370-2693(01)00307-0. arXiv: hep-ex/0012035.

- [15] S. L. Glashow. "Partial Symmetries of Weak Interactions". In: *Nucl. Phys.* 22 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2.
- [16] Abdus Salam and John Clive Ward. "Electromagnetic and weak interactions". In: *Phys. Lett.* 13 (1964), pp. 168–171. DOI: 10.1016/0031-9163(64)90711-5.
- [17] J. Goldstone. "Field Theories with Superconductor Solutions". In: *Nuovo Cim.* 19 (1961), pp. 154–164. DOI: 10.1007/BF02812722.
- [18] Jeffrey Goldstone, Abdus Salam, and Steven Weinberg. "Broken Symmetries". In: *Phys. Rev.* 127 (1962), pp. 965–970. DOI: 10.1103/PhysRev.127.965.
- [19] Peter W. Higgs. "Broken Symmetries and the Masses of Gauge Bosons". In: *Phys. Rev. Lett.* 13 (1964). Ed. by J. C. Taylor, pp. 508–509. DOI: 10.1103/PhysRevLett. 13.508.
- [20] F. Englert and R. Brout. "Broken Symmetry and the Mass of Gauge Vector Mesons". In: *Phys. Rev. Lett.* 13 (1964). Ed. by J. C. Taylor, pp. 321–323. DOI: 10.1103/PhysRevLett.13.321.
- [21] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. "Global Conservation Laws and Massless Particles". In: *Phys. Rev. Lett.* 13 (1964). Ed. by J. C. Taylor, pp. 585– 587. DOI: 10.1103/PhysRevLett.13.585.
- [22] Abdus Salam. "Weak and Electromagnetic Interactions". In: Conf. Proc. C 680519 (1968), pp. 367–377. DOI: 10.1142/9789812795915\_0034.
- [23] Steven Weinberg. "A Model of Leptons". In: Phys. Rev. Lett. 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.
- [24] Gerard 't Hooft. "Renormalizable Lagrangians for Massive Yang-Mills Fields". In:  $Nucl.\ Phys.\ B$  35 (1971). Ed. by J. C. Taylor, pp. 167–188. DOI: 10.1016/0550–3213(71)90139–8.
- [25] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. "Search for neutrinos from the sun". In: *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205.
- [26] John N. Bahcall. "Solar neutrinos. I: Theoretical". In: Phys. Rev. Lett. 12 (1964), pp. 300-302. DOI: 10.1103/PhysRevLett.12.300.
- [27] P. Anselmann et al. "GALLEX results from the first 30 solar neutrino runs". In: *Phys. Lett. B* 327 (1994), pp. 377–385. DOI: 10.1016/0370-2693(94)90744-7.
- [28] Y. Fukuda et al. "Solar neutrino data covering solar cycle 22". In: *Phys. Rev. Lett.* 77 (1996), pp. 1683–1686. DOI: 10.1103/PhysRevLett.77.1683.
- [29] A. I. Abazov et al. "Search for neutrinos from sun using the reaction Ga-71 (electron-neutrino e-) Ge-71". In: *Phys. Rev. Lett.* 67 (1991), pp. 3332–3335. DOI: 10.1103/PhysRevLett.67.3332.

- [30] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. "Remarks on the unified model of elementary particles". In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: 10. 1143/PTP.28.870.
- [31] B. Pontecorvo. "Inverse beta processes and nonconservation of lepton charge". In: Zh. Eksp. Teor. Fiz. 34 (1957), p. 247.
- [32] V. N. Gribov and B. Pontecorvo. "Neutrino astronomy and lepton charge". In: *Phys. Lett. B* 28 (1969), p. 493. DOI: 10.1016/0370-2693(69)90525-5.
- [33] Q. R. Ahmad et al. "Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 89 (2002), p. 011301. DOI: 10.1103/PhysRevLett.89.011301. arXiv: nucl-ex/0204008.
- [34] K. Eguchi et al. "First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance". In: Phys. Rev. Lett. 90 (2 Jan. 2003), p. 021802. DOI: 10.1103/ PhysRevLett.90.021802.
- [35] Y. Fukuda et al. "Evidence for oscillation of atmospheric neutrinos". In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.
- [36] L. Wolfenstein. "Neutrino oscillations in matter". In: Phys. Rev. D 17 (9 May 1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369.
- [37] S. P. Mikheev and A. Yu. Smirnov. "Resonant amplification of neutrino oscillations in matter and solar neutrino spectroscopy". In: *Nuovo Cim. C* 9 (1986), pp. 17–26. DOI: 10.1007/BF02508049.
- [38] Carlo Giunti and Chung W. Kim. "Neutrino Oscillations In Matter". In: Fundamentals of Neutrino Physics and Astrophysics. Oxford University Press, Mar. 2007. DOI: 10.1093/acprof:oso/9780198508717.003.0009.
- [39] Vernon D. Barger et al. "Matter Effects on Three-Neutrino Oscillations". In: *Phys. Rev. D* 22 (1980), p. 2718. DOI: 10.1103/PhysRevD.22.2718.
- [40] P. F. Harrison and W. G. Scott. "mu tau reflection symmetry in lepton mixing and neutrino oscillations". In: *Phys. Lett. B* 547 (2002), pp. 219–228. DOI: 10. 1016/S0370-2693(02)02772-7. arXiv: hep-ph/0210197.
- [41] Babak Abi et al. "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics". In: (Feb. 2020). arXiv: 2002. 03005 [hep-ex].
- [42] Matteo Agostini, Giovanni Benato, and Jason Detwiler. "Discovery probability of next-generation neutrinoless double-  $\beta$  decay experiments". In: *Phys. Rev. D* 96.5 (2017), p. 053001. DOI: 10.1103/PhysRevD.96.053001. arXiv: 1705.02996 [hep-ex].

- [43] A. D. Sakharov. "Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe". In: *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967), pp. 32–35. DOI: 10.1070/PU1991v034n05ABEH002497.
- [44] M. Fukugita and T. Yanagida. "Baryogenesis Without Grand Unification". In: Phys. Lett. B 174 (1986), pp. 45–47. DOI: 10.1016/0370-2693(86)91126-3.
- [45] C. Jarlskog. "Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconservation". In: *Phys. Rev. Lett.* 55 (1985), p. 1039. DOI: 10.1103/PhysRevLett.55.1039.
- [46] D. Casper. "The Nuance neutrino physics simulation, and the future". In: Nucl. Phys. B Proc. Suppl. 112 (2002). Ed. by J. G. Morfin, M. Sakuda, and Y. Suzuki, pp. 161–170. DOI: 10.1016/S0920-5632(02)01756-5. arXiv: hep-ph/0208030.
- [47] J. A. Formaggio and G. P. Zeller. "From eV to EeV: Neutrino Cross Sections Across Energy Scales". In: *Rev. Mod. Phys.* 84 (2012), pp. 1307–1341. DOI: 10. 1103/RevModPhys.84.1307. arXiv: 1305.7513 [hep-ex].
- [48] Fundamental Physics at the Intensity Frontier. May 2012. DOI: 10.2172/1042577. arXiv: 1205.2671 [hep-ex].
- [49] R. A. Smith and E. J. Moniz. "NEUTRINO REACTIONS ON NUCLEAR TARGETS". In: Nucl. Phys. B 43 (1972). [Erratum: Nucl.Phys.B 101, 547 (1975)],
   p. 605. DOI: 10.1016/0550-3213(75)90612-4.
- [50] A. Bodek and J. L. Ritchie. "Further Studies of Fermi Motion Effects in Lepton Scattering from Nuclear Targets". In: Phys. Rev. D 24 (1981), p. 1400. DOI: 10. 1103/PhysRevD.24.1400.
- [51] O. Benhar et al. "Spectral function of finite nuclei and scattering of GeV electrons". In: Nucl. Phys. A 579 (1994), pp. 493–517. DOI: 10.1016/0375-9474(94)90920-2.
- [52] M. Martini et al. "A Unified approach for nucleon knock-out, coherent and incoherent pion production in neutrino interactions with nuclei". In: *Phys. Rev. C* 80 (2009), p. 065501. DOI: 10.1103/PhysRevC.80.065501. arXiv: 0910.2622 [nucl-th].
- [53] G. D. Harp. "Extension of the isobar model for intranuclear cascades to 1 GeV". In: *Phys. Rev. C* 10 (1974), pp. 2387–2396. DOI: 10.1103/PhysRevC.10.2387.
- [54] L. Alvarez-Ruso et al. "NuSTEC White Paper: Status and challenges of neutrino-nucleus scattering". In: *Prog. Part. Nucl. Phys.* 100 (2018), pp. 1–68. DOI: 10.1016/j.ppnp.2018.01.006. arXiv: 1706.03621 [hep-ph].
- [55] A. A. Aguilar-Arevalo et al. "First Measurement of the Muon Neutrino Charged Current Quasielastic Double Differential Cross Section". In: *Phys. Rev. D* 81 (2010), p. 092005. DOI: 10.1103/PhysRevD.81.092005. arXiv: 1002.2680 [hep-ex].

- [56] P. Abratenko et al. "First Measurement of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon at  $E_{\nu} \sim 0.8$  GeV with the Micro-BooNE Detector". In: *Phys. Rev. Lett.* 123.13 (2019), p. 131801. DOI: 10.1103/PhysRevLett.123.131801. arXiv: 1905.09694 [hep-ex].
- [57] G. A. Fiorentini et al. "Measurement of Muon Neutrino Quasielastic Scattering on a Hydrocarbon Target at  $E_{\nu} \sim 3.5$  GeV". In: *Phys. Rev. Lett.* 111 (2013), p. 022502. DOI: 10.1103/PhysRevLett.111.022502. arXiv: 1305.2243 [hep-ex].
- [58] Dominic Brailsford. "Physics Program of the Short-Baseline Near Detector". In: J. Phys. Conf. Ser. 888.1 (2017), p. 012186. DOI: 10.1088/1742-6596/888/1/012186.
- [59] S. Abe et al. "Precision Measurement of Neutrino Oscillation Parameters with KamLAND". In: *Phys. Rev. Lett.* 100 (2008), p. 221803. DOI: 10.1103/PhysRevLett. 100.221803. arXiv: 0801.4589 [hep-ex].
- [60] P. F. de Salas et al. "2020 global reassessment of the neutrino oscillation picture". In: JHEP 02 (2021), p. 071. DOI: 10.1007/JHEP02(2021)071. arXiv: 2006.11237 [hep-ph].
- [61] K. Abe et al. "Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV". In: *Phys. Rev. D* 97.7 (2018), p. 072001. DOI: 10. 1103/PhysRevD.97.072001. arXiv: 1710.09126 [hep-ex].
- [62] S. van der Meer. "A Directive Device for Charged Particles and Its use in an Enhanced Neutrino Beam". In: (Feb. 1961). DOI: 10.5170/CERN-1961-007.
- [63] Yuichi Oyama. "K2K (KEK to Kamioka) neutrino oscillation experiment at KEK PS". In: YITP Workshop on Flavor Physics. Jan. 1998. arXiv: hep-ex/9803014.
- [64] "Neutrino Oscillation Physics at Fermilab: The NuMI-MINOS Project". In: (May 1998).
- [65] P. Adamson et al. "Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam". In: *Phys. Rev. Lett.* 101 (2008), p. 131802. DOI: 10.1103/PhysRevLett.101.131802. arXiv: 0806.2237 [hep-ex].
- [66] M. H. Ahn et al. "Measurement of Neutrino Oscillation by the K2K Experiment". In: Phys. Rev. D 74 (2006), p. 072003. DOI: 10.1103/PhysRevD.74.072003. arXiv: hep-ex/0606032.
- [67] P. Adamson et al. "Electron Neutrino and Antineutrino Appearance in the Full MINOS Data Sample". In: *Phys. Rev. Lett.* 110.17 (2013), p. 171801. DOI: 10.1103/PhysRevLett.110.171801. arXiv: 1301.4581 [hep-ex].
- [68] M. H. Ahn et al. "Search for electron neutrino appearance in a 250 km long baseline experiment". In: *Phys. Rev. Lett.* 93 (2004), p. 051801. DOI: 10.1103/PhysRevLett.93.051801. arXiv: hep-ex/0402017.

- [69] P. Adamson et al. "Measurement of the Neutrino Mass Splitting and Flavor Mixing by MINOS". In: *Phys. Rev. Lett.* 106 (2011), p. 181801. DOI: 10.1103/PhysRevLett.106.181801. arXiv: 1103.0340 [hep-ex].
- [70] K. Abe et al. "The T2K Experiment". In: Nucl. Instrum. Meth. A 659 (2011), pp. 106–135. DOI: 10.1016/j.nima.2011.06.067. arXiv: 1106.1238.
- [71] D. S. Ayres et al. "NOvA: Proposal to Build a 30 Kiloton Off-Axis Detector to Study  $\nu_{\mu} \rightarrow \nu_{e}$  Oscillations in the NuMI Beamline". In: (Mar. 2004). arXiv: hep-ex/0503053.
- [72] K. Abe et al. "Observation of Electron Neutrino Appearance in a Muon Neutrino Beam". In: *Phys. Rev. Lett.* 112 (2014), p. 061802. DOI: 10.1103/PhysRevLett. 112.061802. arXiv: 1311.4750 [hep-ex].
- [73] K. Abe et al. "Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations". In: *Nature* 580.7803 (2020). [Erratum: Nature 583, E16 (2020)], pp. 339–344. DOI: 10.1038/s41586-020-2177-0. arXiv: 1910.03887 [hep-ex].
- [74] M. A. Acero et al. "Improved measurement of neutrino oscillation parameters by the NOvA experiment". In: *Phys. Rev. D* 106.3 (2022), p. 032004. DOI: 10.1103/PhysRevD.106.032004. arXiv: 2108.08219 [hep-ex].
- [75] Ivan Esteban et al. "The fate of hints: updated global analysis of three-flavor neutrino oscillations". In: *JHEP* 09 (2020), p. 178. DOI: 10.1007/JHEP09(2020) 178. arXiv: 2007.14792 [hep-ph].
- [76] B. Abi et al. "Long-baseline neutrino oscillation physics potential of the DUNE experiment". In: Eur. Phys. J. C 80.10 (2020), p. 978. DOI: 10.1140/epjc/s10052-020-08456-z. arXiv: 2006.16043 [hep-ex].
- [77] K. Abe et al. "Physics potential of a long-baseline neutrino oscillation experiment using a J-PARC neutrino beam and Hyper-Kamiokande". In: *PTEP* 2015 (2015), p. 053C02. DOI: 10.1093/ptep/ptv061. arXiv: 1502.05199 [hep-ex].
- [78] Angel Abusleme et al. "JUNO physics and detector". In: *Prog. Part. Nucl. Phys.* 123 (2022), p. 103927. DOI: 10.1016/j.ppnp.2021.103927. arXiv: 2104.02565 [hep-ex].
- [79] Fermilab Creative Services. *DUNE at LBNF*. URL: https://wms.fnal.gov/gallery/view?id=39 (visited on 08/2024).
- [80] A. Abed Abud et al. "Snowmass Neutrino Frontier: DUNE Physics Summary". In: (Mar. 2022). arXiv: 2203.06100 [hep-ex].
- [81] Kevin James Kelly et al. "Sub-GeV Atmospheric Neutrinos and CP-Violation in DUNE". In: *Phys. Rev. Lett.* 123.8 (2019), p. 081801. DOI: 10.1103/PhysRevLett. 123.081801. arXiv: 1904.02751 [hep-ph].

- [82] Animesh Chatterjee and Albert De Roeck. "Neutrino oscillations with atmospheric neutrinos at large liquid argon TPCs". In: *Phys. Lett. B* 855 (2024), p. 138838. DOI: 10.1016/j.physletb.2024.138838. arXiv: 2402.16441 [hep-ph].
- [83] Christoph A. Ternes et al. "Neutrino mass ordering at DUNE: An extra  $\nu$  bonus". In: *Phys. Rev. D* 100.9 (2019), p. 093004. DOI: 10.1103/PhysRevD.100.093004. arXiv: 1905.03589 [hep-ph].
- [84] Francesco Capozzi et al. "DUNE as the Next-Generation Solar Neutrino Experiment". In: *Phys. Rev. Lett.* 123.13 (2019), p. 131803. DOI: 10.1103/PhysRevLett. 123.131803. arXiv: 1808.08232 [hep-ph].
- [85] Karolina Rozwadowska, Francesco Vissani, and Enrico Cappellaro. "On the rate of core collapse supernovae in the milky way". In: *New Astron.* 83 (2021), p. 101498. DOI: 10.1016/j.newast.2020.101498. arXiv: 2009.03438 [astro-ph.HE].
- [86] E. N. Alekseev et al. "Possible Detection of a Neutrino Signal on 23 February 1987 at the Baksan Underground Scintillation Telescope of the Institute of Nuclear Research". In: JETP Lett. 45 (1987). Ed. by J. Tran Thanh Van, pp. 589–592.
- [87] R. M. Bionta et al. "Observation of a Neutrino Burst in Coincidence with Supernova SN 1987a in the Large Magellanic Cloud". In: Phys. Rev. Lett. 58 (1987), p. 1494. DOI: 10.1103/PhysRevLett.58.1494.
- [88] K. Hirata et al. "Observation of a Neutrino Burst from the Supernova SN 1987a". In: Phys. Rev. Lett. 58 (1987). Ed. by K. C. Wali, pp. 1490–1493. DOI: 10.1103/ PhysRevLett.58.1490.
- [89] Vitaly A. Kudryavtsev. "Underground physics with DUNE". In: J. Phys. Conf. Ser. 718.6 (2016), p. 062032. DOI: 10.1088/1742-6596/718/6/062032. arXiv: 1601.03496 [physics.ins-det].
- [90] B. Abi et al. "Supernova neutrino burst detection with the Deep Underground Neutrino Experiment". In: Eur. Phys. J. C 81.5 (2021), p. 423. DOI: 10.1140/epjc/s10052-021-09166-w. arXiv: 2008.06647 [hep-ex].
- [91] C. Rubbia. "The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors". In: (May 1977).
- [92] D. R. Nygren. "The Time Projection Chamber: A New 4 pi Detector for Charged Particles". In: *eConf* C740805 (1974). Ed. by J Kadyk et al., p. 58.
- [93] C. Rubbia et al. "Underground operation of the ICARUS T600 LAr-TPC: first results". In: JINST 6 (2011), P07011. DOI: 10.1088/1748-0221/6/07/P07011. arXiv: 1106.0975 [hep-ex].
- [94] R. Acciarri et al. "Design and Construction of the MicroBooNE Detector". In: JINST 12.02 (2017), P02017. DOI: 10.1088/1748-0221/12/02/P02017. arXiv: 1612.05824 [physics.ins-det].

- [95] A. Abed Abud et al. "Design, construction and operation of the ProtoDUNE-SP Liquid Argon TPC". In: JINST 17.01 (2022), P01005. DOI: 10.1088/1748-0221/17/01/P01005. arXiv: 2108.01902 [physics.ins-det].
- [96] Babak Abi et al. "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV: Far Detector Single-phase Technology". In: JINST 15.08 (2020), T08010. DOI: 10.1088/1748-0221/15/08/T08010. arXiv: 2002.03010 [physics.ins-det].
- [97] H. Bethe. "Theory of the Passage of Fast Corpuscular Rays Through Matter". In: Annalen Phys. 5 (1930), pp. 325–400. DOI: 10.1002/andp.19303970303.
- [98] B. Abi et al. "First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform". In: JINST 15.12 (2020), P12004. DOI: 10.1088/1748-0221/15/12/P12004. arXiv: 2007.06722 [physics.ins-det].
- [99] R. L. Workman et al. "Review of Particle Physics". In: *PTEP* 2022 (2022). DOI: 10.1093/ptep/ptac097.
- [100] James Strait et al. "Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 3: Long-Baseline Neutrino Facility for DUNE June 24, 2015". In: (Jan. 2016). arXiv: 1601. 05823 [physics.ins-det].
- [101] V. Papadimitriou et al. "Design of the LBNF Beamline". In: 6th International Particle Accelerator Conference. 2015, THPF120. DOI: 10.18429/JACoW-IPAC2015-THPF120.
- [102] M. Ball et al. "The PIP-II Conceptual Design Report". In: (Mar. 2017). Ed. by Valeri Lebedev. DOI: 10.2172/1346823.
- [103] D. A. Dwyer et al. "LArPix: Demonstration of low-power 3D pixelated charge readout for liquid argon time projection chambers". In: *JINST* 13.10 (2018), P10007. DOI: 10.1088/1748-0221/13/10/P10007. arXiv: 1808.02969 [physics.insdet].
- [104] N. Anfimov et al. "Development of the Light Collection Module for the Liquid Argon Time Projection Chamber (LArTPC)". In: *JINST* 15.07 (2020), p. C07022. DOI: 10.1088/1748-0221/15/07/C07022.
- [105] M. Auger et al. "ArCLight—A Compact Dielectric Large-Area Photon Detector". In: *Instruments* 2.1 (2018), p. 3. DOI: 10.3390/instruments2010003. arXiv: 1711. 11409 [physics.ins-det].
- [106] Adam Abed Abud et al. "Performance of a modular ton-scale pixel-readout liquid argon time projection chamber". In: (Mar. 2024). arXiv: 2403.03212 [physics.ins-det].

- [107] I. Ambats et al. "The MINOS Detectors Technical Design Report". In: (Oct. 1998).

  DOI: 10.2172/1861363.
- [108] S. Bhadra et al. "Letter of Intent to Construct a nuPRISM Detector in the J-PARC Neutrino Beamline". In: (Dec. 2014). arXiv: 1412.3086 [physics.ins-det].
- [109] A. Aloisio et al. "The KLOE detector: technical proposal". In: (Jan. 1993).
- [110] V. Hewes et al. "Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report". In: *Instruments* 5.4 (2021), p. 31. DOI: 10.3390/instruments5040031. arXiv: 2103.13910 [physics.ins-det].
- [111] Babak Abi et al. "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV: Far Detector Single-phase Technology". In: JINST 15.08 (2020), T08010. DOI: 10.1088/1748-0221/15/08/T08010. arXiv: 2002.03010 [physics.ins-det].
- [112] A. A. Machado and E. Segreto. "ARAPUCA a new device for liquid argon scintillation light detection". In: *JINST* 11.02 (2016), p. C02004. DOI: 10.1088/1748-0221/11/02/C02004.
- [113] A. A. Machado et al. "The X-ARAPUCA: An improvement of the ARAPUCA device". In: Apr. 2018. DOI: 10.1088/1748-0221/13/04/C04026. arXiv: 1804. 01407 [physics.ins-det].
- [114] Adam Abed Abud et al. "The DUNE Far Detector Vertical Drift Technology. Technical Design Report". In: JINST 19.08 (2024), T08004. DOI: 10.1088/1748-0221/19/08/T08004. arXiv: 2312.03130 [hep-ex].
- [115] "Long-Baseline Neutrino Facility (LBNF)/DUNE Conceptual Design Report Annex 3A Opt: Optimized Neutrino Beamline". In: (Sept. 2017). URL: https://docs.dunescience.org/cgi-bin/sso/ShowDocument?docid=4559.
- [116] J. Allison et al. "Recent developments in Geant4". In: *Nucl. Instrum. Meth. A* 835 (2016), pp. 186–225. DOI: 10.1016/j.nima.2016.06.125.
- [117] John Allison et al. "Geant4 developments and applications". In: *IEEE Trans. Nucl. Sci.* 53 (2006), p. 270. DOI: 10.1109/TNS.2006.869826.
- [118] S. Agostinelli et al. "GEANT4–a simulation toolkit". In: Nucl. Instrum. Meth. A 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.
- [119] C. Andreopoulos et al. "The GENIE Neutrino Monte Carlo Generator". In: Nucl. Instrum. Meth. A614 (2010), pp. 87–104. DOI: 10.1016/j.nima.2009.12.009. arXiv: 0905.2517 [hep-ph].
- [120] T. Doke et al. "Estimation of Fano factors in liquid argon, krypton, xenon and xenon-doped liquid argon". In: *Nucl. Instrum. Meth.* 134 (1976), pp. 353–357. DOI: 10.1016/0029-554X(76)90292-5.

- [121] J. B. Birks. "Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations". In: *Proc. Phys. Soc. A* 64 (1951), pp. 874–877. DOI: 10.1088/0370-1298/64/10/303.
- [122] J. Thomas and D. A. Imel. "Recombination of electron-ion pairs in liquid argon and liquid xenon". In: Phys. Rev. A 36 (1987), pp. 614–616. DOI: 10.1103/PhysRevA. 36.614.
- [123] Yiwen Xiao. "Purity monitoring for ProtoDUNE-SP". In: J. Phys. Conf. Ser. 2156 (2021), p. 012211. DOI: 10.1088/1742-6596/2156/1/012211.
- [124] C. Green et al. "The Art Framework". In: J. Phys. Conf. Ser. 396 (2012). Ed. by Michael Ernst et al., p. 022020. DOI: 10.1088/1742-6596/396/2/022020.
- [125] E. L. Snider and G. Petrillo. "LArSoft: Toolkit for Simulation, Reconstruction and Analysis of Liquid Argon TPC Neutrino Detectors". In: J. Phys. Conf. Ser. 898.4 (2017). Ed. by Richard Mount and Craig Tull, p. 042057. DOI: 10.1088/1742-6596/898/4/042057.
- [126] I. Antcheva et al. "ROOT: A C++ framework for petabyte data storage, statistical analysis and visualization". In: *Comput. Phys. Commun.* 180 (2009), pp. 2499–2512. DOI: 10.1016/j.cpc.2009.08.005. arXiv: 1508.07749 [physics.data-an].
- [127] Wire-Cell. Wire-Cell Toolkit. URL: https://wirecell.bnl.gov/.
- [128] R. Veenhof. "GARFIELD, recent developments". In: Nucl. Instrum. Meth. A 419 (1998). Ed. by M. Krammer et al., pp. 726–730. DOI: 10.1016/S0168-9002(98) 00851-1.
- [129] C. Adams et al. "Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation". In: *JINST* 13.07 (2018), P07006. DOI: 10.1088/1748-0221/13/07/P07006. arXiv: 1802.08709 [physics.ins-det].
- [130] J. S. Marshall and M. A. Thomson. "The Pandora Software Development Kit for Pattern Recognition". In: Eur. Phys. J. C 75.9 (2015), p. 439. DOI: 10.1140/epjc/s10052-015-3659-3. arXiv: 1506.05348 [physics.data-an].
- [131] Adam Abed Abud et al. "Reconstruction of interactions in the ProtoDUNE-SP detector with Pandora". In: Eur. Phys. J. C 83.7 (2023), p. 618. DOI: 10.1140/epjc/s10052-023-11733-2. arXiv: 2206.14521 [hep-ex].
- [132] 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: Eur. Phys. J. C 78.1 (2018), p. 82. DOI: 10.1140/epjc/s10052-017-5481-6. arXiv: 1708.03135 [hep-ex].

- [133] B. Abi et al. "Neutrino interaction classification with a convolutional neural network in the DUNE far detector". In: *Phys. Rev. D* 102.9 (2020), p. 092003. DOI: 10.1103/PhysRevD.102.092003. arXiv: 2006.15052 [physics.ins-det].
- [134] A. Ankowski et al. "Measurement of through-going particle momentum by means of multiple scattering with the ICARUS T600 TPC". In: Eur. Phys. J. C 48 (2006), pp. 667–676. DOI: 10.1140/epjc/s10052-006-0051-3. arXiv: hep-ex/0606006.
- [135] Francois Drielsma et al. "Scalable, End-to-End, Deep-Learning-Based Data Reconstruction Chain for Particle Imaging Detectors". In: 34th Conference on Neural Information Processing Systems. Feb. 2021. arXiv: 2102.01033 [hep-ex].
- [136] Adam Abed Abud et al. "Highly-parallelized simulation of a pixelated LArTPC on a GPU". In: JINST 18.04 (2023), P04034. DOI: 10.1088/1748-0221/18/04/P04034. arXiv: 2212.09807 [physics.comp-ph].
- [137] The HDF Group. HDF5. URL: https://www.hdfgroup.org/solutions/hdf5/.
- [138] Ciaran Hasnip. "DUNE-PRISM A New Method to Measure Neutrino Oscillations". PhD thesis. Oxford U., 2023.
- [139] Ivan Esteban et al. "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of  $\theta_{23}$ ,  $\delta_{CP}$ , and the mass ordering". In: *JHEP* 01 (2019), p. 106. DOI: 10.1007/JHEP01(2019)106. arXiv: 1811.05487 [hep-ph].
- [140] A.N. Tikhonov. "On the solution of improperly posed problems and the method of regularization". In: *Sov. Math.* 5 (1963), p. 1035.
- [141] David L. Phillips. "A Technique for the Numerical Solution of Certain Integral Equations of the First Kind". In: J. Assoc. Comput. Machinery 9.1 (1962), pp. 84– 97. DOI: 10.1145/321105.321114.
- [142] Glen Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: Eur. Phys. J. C 71 (2011). [Erratum: Eur.Phys.J.C 73, 2501 (2013)], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an].
- [143] C. McGrew. edep-sim: An Energy Deposition Simulation. https://github.com/ClarkMcGrew/edep-sim.
- [144] Volker Blobel. "An Unfolding method for high-energy physics experiments". In: Conference on Advanced Statistical Techniques in Particle Physics. Aug. 2002, pp. 258–267. arXiv: hep-ex/0208022.
- [145] Alec Radford et al. "Language Models are Unsupervised Multitask Learners". In: 2019.
- [146] Anja Butter et al. "Jet Diffusion versus JetGPT Modern Networks for the LHC". In: (May 2023). arXiv: 2305.10475 [hep-ph].

- [147] Ashish Vaswani et al. "Attention is all you need". In: *Proceedings of the 31st International Conference on Neural Information Processing Systems.* NIPS'17. Long Beach, California, USA: Curran Associates Inc., 2017, pp. 6000–6010.
- [148] PyTorch: An Imperative Style, High-Performance Deep Learning Library. 2019. arXiv: 1912.01703 [cs.LG]. URL: https://arxiv.org/abs/1912.01703.
- [149] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. 2019. arXiv: 1711.05101 [cs.LG]. URL: https://arxiv.org/abs/1711.05101.
- [150] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. 2017. arXiv: 1412.6980 [cs.LG]. URL: https://arxiv.org/abs/1412.6980.
- [151] Laura Dominé and Kazuhiro Terao. "Scalable deep convolutional neural networks for sparse, locally dense liquid argon time projection chamber data". In: *Phys. Rev. D* 102 (1 July 2020), p. 012005. DOI: 10.1103/PhysRevD.102.012005. URL: https://link.aps.org/doi/10.1103/PhysRevD.102.012005.
- [152] Christopher Choy, JunYoung Gwak, and Silvio Savarese. "4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2019, pp. 3075–3084.
- [153] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. "3D Semantic Segmentation with Submanifold Sparse Convolutional Networks". In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2018, pp. 9224–9232. DOI: 10.1109/CVPR.2018.00961.
- [154] Christopher Choy. "High-dimensional convolutional neural networks for 3D perception". PhD thesis. Stanford University, 2020.
- [155] Yoshua Bengio et al. "Curriculum learning". In: *Proceedings of the 26th Annual International Conference on Machine Learning*. ICML '09. Montreal, Quebec, Canada: Association for Computing Machinery, 2009, pp. 41–48. URL: https://doi.org/10.1145/1553374.1553380.
- [156] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. "U-Net: Convolutional Networks for Biomedical Image Segmentation". In: *CoRR* abs/1505.04597 (2015). arXiv: 1505.04597. URL: http://arxiv.org/abs/1505.04597.
- [157] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. "Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)". In: *arXiv* preprint (Nov. 2015). arXiv: 1511.07289 [cs.LG].
- [158] JunYoung Gwak, Christopher Choy, and Silvio Savarese. "Generative sparse detection networks for 3d single-shot object detection". In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16. Springer. 2020, pp. 297–313. arXiv: 2006.12356 [cs.CV].

- [159] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. "Perceptual losses for real-time style transfer and super-resolution". In: *European Conference on Computer Vision*. 2016.
- [160] Phillip Isola et al. "Image-to-Image Translation with Conditional Adversarial Networks". In: CVPR (2017).
- [161] Kaiming He et al. "Deep Residual Learning for Image Recognition". In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2016, pp. 770–778. DOI: 10.1109/CVPR.2016.90.
- [162] C. Backhouse. "The CAFAna framework for neutrino analysis". In: *Snowmass 2021*. Mar. 2022. arXiv: 2203.13768 [hep-ex].
- [163] Chitwan Saharia et al. Palette: Image-to-Image Diffusion Models. 2022. arXiv: 2111.05826 [cs.CV]. URL: https://arxiv.org/abs/2111.05826.
- [164] Gaurav Parmar et al. One-Step Image Translation with Text-to-Image Models. 2024. arXiv: 2403.12036 [cs.CV]. URL: https://arxiv.org/abs/2403.12036.