

Searches for Bound Neutron-Antineutron Oscillation in Liquid Argon Time Projection Chambers

A thesis submitted to the University of Manchester
for the degree of Doctor of Philosophy
in the Faculty of Science and Engineering

· 2017 ·

Jeremy E. T. Hewes

School of Physics and Astronomy

Contents

1	Introduction	15
2	Theory	17
2.1	Phenomenology	17
2.1.1	Neutron-antineutron oscillation in field-free vacuum . .	20
2.1.2	Neutron-antineutron oscillation in magnetic fields . . .	21
2.1.3	Neutron-antineutron oscillation in the nucleus	22
2.2	Bound to free lifetime conversion	24
2.3	Baryon number violation	28
3	Measurement History	31
3.1	Free neutron searches	31
3.2	Bound neutron searches	33
3.3	Antineutron annihilation branching ratios	36
3.3.1	Measurements of $\bar{p}p$ annihilation	37
3.3.2	Measurements of $\bar{p}n$ annihilation	38
3.3.3	Crystal Barrel data and new branching ratios	40
4	Liquid Argon Time Projection Chambers	43

4.1	LArTPC detector principles	43
4.1.1	Time projection chamber	45
4.1.2	Light detection	47
4.1.3	Cryogenics and argon purity	48
4.2	Particle signatures	48
4.2.1	Charged leptons and mesons	50
4.2.2	Nucleons	51
4.2.3	Electromagnetic showers	52
4.2.4	Reconstructing π^0 mesons	52
4.3	LArTPC detectors	53
4.3.1	History	53
4.3.2	The MicroBooNE detector	54
4.3.3	The DUNE far detector	57
5	GENIE Event Generator	61
5.1	Simulation methods	61
5.1.1	The initial state	63
5.1.2	Simulating the oscillating neutron	63
5.1.3	Simulating the annihilating nucleon	64
5.1.4	Simulating the remnant nucleus	65
5.1.5	Simulating annihilation products	65
5.1.6	Final state interactions	66
5.2	Simulation results	67
5.2.1	Super-Kamiokande comparison	73
5.3	Discussion	75

5.3.1	Branching ratio corrections	75
5.3.2	Validating the phase space approximation	75
5.4	Conclusions	78
6	Convolutional Neural Networks	81
6.1	CNN operating principles	81
6.2	Network layer types	83
6.2.1	Convolution layer	84
6.2.2	ReLU layer	85
6.2.3	Pooling layer	86
6.2.4	Fully connected layer	86
6.3	Network training	87
6.3.1	Solver	87
6.3.2	Classification	87
6.3.3	Loss	88
6.3.4	Gradient descent	88
6.3.5	Learning rate	90
6.4	Tools	90
6.4.1	Caffe	90
6.4.2	VGG16 network architecture	91
6.4.3	LArCV	91
7	$n - \bar{n}$ Oscillation Search in DUNE	93
7.1	Atmospheric neutrino backgrounds	93
7.1.1	Neutrino interactions	95
7.2	Monte Carlo production	100

7.2.1	Event generation	101
7.2.2	Physics simulation	101
7.2.3	Detector simulation	102
7.2.4	Wire reconstruction	103
7.3	Image preparation	103
7.4	Network training	111
7.5	Network performance	112
7.5.1	Network tests	112
7.5.2	MC truth feature identification	119
7.6	Discussion	124
7.6.1	APA stitching	124
7.6.2	Effect of containment	128
7.7	Sensitivity calculation	130
7.7.1	Three-plane vs one-plane training	134
8	$n - \bar{n}$ Oscillation Search in MicroBooNE	139
8.1	Cosmogenic background simulation	140
8.2	Event selection	140
8.3	Monte Carlo production	142
8.4	Network training	143
8.5	Network performance	143
8.6	MicroBooNE sensitivity to $n - \bar{n}$ oscillation	150
8.7	Data-MC agreement	152
9	Conclusions	161
9.1	Event generator development	161

9.2	Convolutional neural networks	162
9.3	Future challenges	163
A	Antiproton scattering data	165
A.1	Antiproton annihilation in hydrogen	165
A.1.1	Cresti <i>et al.</i>	165
A.1.2	Baltay <i>et al.</i>	165
A.1.3	Pavlopoulos <i>et al.</i>	167
A.1.4	Backenstoss <i>et al.</i>	167
A.2	Antiproton annihilation in deuterium	169
B	Neutrino oscillations	171
	References	172

Word count: 23,530

Abstract

The next-generation Deep Underground Neutrino Experiment’s liquid argon detector represents an opportunity to probe previously unexplored parameter space for beyond-Standard Model processes. One such process is baryon number violating neutron-antineutron oscillation, the observation of which would have profound implications on our understanding of the origin of the matter-antimatter asymmetry in the universe, and provide strong hints as to the nature of neutrino mass.

A GENIE $n - \bar{n}$ oscillation event generator was developed and officially released, taking into account various nuclear effects and final state interactions. Previous searches for the process are summarised, as are the sources of antiproton scattering used to derive $n - \bar{n}$ branching ratios.

The viability of machine learning image processing techniques to identify simulated signal $n - \bar{n}$ events and reject potential atmospheric neutrino backgrounds in DUNE is explored. Images are produced using simulated $n - \bar{n}$ and atmospheric neutrino events in DUNE, and a convolutional neural network is trained to distinguish the two. The network’s ability to accept signal and reject background corresponds to a free $n - \bar{n}$ lifetime sensitivity of 1.6×10^9 s at 90% confidence level, a factor of 5 improvement on the current limit from Super-Kamiokande.

These machine learning techniques are applied to data from the on-surface MicroBooNE detector, and the network is found to be highly sensitive to differences between data and Monte Carlo (MC) simulations. Recommendations are made for further studies into the use of such techniques, and potential avenues for overcoming challenges in data-MC disagreement are presented.

Declaration

I declare that no portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or institute of learning.

Chapter 3 discusses $n - \bar{n}$ branching ratio calculations and antiproton annihilation data. My work attempts to reconcile existing branching ratios with antiproton annihilation data, but I am not responsible for producing any of the branching ratio tables discussed. The branching ratios shown in Table 3.2 are taken from a recent Super-Kamiokande search [1], and updated ratios shown in Table 3.5 are reproduced from a recent PhD thesis by J. Gustafson [2].

Chapter 5 describes the development of a simulation package for $n - \bar{n}$ oscillation. This event generator was developed within the existing GENIE framework, but all development of the standalone $n - \bar{n}$ package was my own work, in terms of both physics simulations and its integration into the GENIE environment. During development of this $n - \bar{n}$ event generator existing GENIE code was used as a reference, particularly the proton decay application due to its similar nature as a standalone application.

Chapters 7 and 8 discuss the use of convolutional neural networks (CNNs) for event selection in DUNE and MicroBooNE, respectively. Simulation and reconstruction are carried out using LArSoft [3], which comprises contributions from many collaborators across several experiments. The open-source Caffe framework was used [4], including modifications made by MicroBooNE collaborators at MIT and Columbia University [5]. CNN image production code was my own work, but makes use of image production and interface utilities developed by these collaborators.

GPU computing resources at Nevis Laboratories were used for CNN-based studies. All CNN training and testing, and all studies into CNN performance and feature identification, are my own work. All code for DUNE and MicroBooNE lifetime sensitivity calculation was developed by me, and uses the open-source Cubature tool [6] for multi-dimensional integration.

Copyright Statement

The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the “Copyright”) and he has given the University of Manchester certain rights to use such Copyright, including for administrative purposes.

Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made **only** in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of such copies made.

The ownership of certain Copyright, patents, designs, trademarks and other intellectual property (the “Intellectual Property”) and any reproductions of copyright works in the thesis, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.

Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see <http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=24420>), in any relevant Thesis restriction declarations deposited in the University Library, the University Library’s regulations (see <http://www.library.manchester.ac.uk/about/regulations>) and in

the University's policy on Presentation of Theses.

Acknowledgments

I would like to thank my PhD advisor, Dr. Georgia Karagiorgi, without whose guidance and support my work on $n - \bar{n}$ oscillation would not have been possible. Thanks also to my advisors Prof. Stefan Söldner-Rembold and Dr. Justin Evans at the University of Manchester for guiding my early work on LBNE, and for their continued support over the past four years. I greatly appreciate the help Andy Furmanski provided during my time at Fermilab, as well as all of my colleagues at the University of Manchester.

I spent two years of my time as a graduate student working at Fermilab, and it was an overwhelmingly positive experience. My work on $n - \bar{n}$ oscillation is indebted to the conveners of the MicroBooNE APE group — Jonathan Asaadi, Serhan Tufanli and especially Yun-Tse Tsai — whose many insights have been invaluable. Thanks to the MicroBooNE spokespeople, Sam Zeller and Bonnie Fleming, for cultivating an environment of positivity and inclusiveness within the collaboration, for which I am particularly grateful.

Thanks to the GENIE developers, especially Costas Andreopoulos, for their assistance with incorporating my $n - \bar{n}$ code into an official GENIE release. My work on $n - \bar{n}$ oscillation in DUNE was carried out in the nucleon decay and atmospheric working group, and the conveners Jen Raaf and Michel Sorel provided much feedback which was essential to my analysis. My colleagues at Nevis Labs, particularly Kazu Terao and Vic Genty, provided both expertise and tools regarding convolutional neural networks without which my work would not have been possible. Thanks also to all of my other friends and collaborators at Fermilab, of whom there are too many to list here.

Love to my parents Catherine and Simon, and my sisters Rowena and

Olivia, for their unwavering support, and for providing me with the opportunities that led me here. For all our adventures across the US, thank you to Kate, Andrew, Ben and Catherine — sorry the AC in my car never worked. Thanks to Rach and Abi for providing me with a home base in Manchester, and for their boundless generosity and hospitality. Thanks to Doggi and Ginny for their emotional support.

Thanks to Amy, Elena, Raquel and all my other friends in Chicago for all the time we spent together. For all the important weekly distractions from work, and snacks, thanks to Dave, Jasmine and Lois. Thanks to the Manchester arts community, which is unparalleled in both talent and inclusiveness — I will be sorry to leave it. Love to all the friends and family I don't have the space to name here.

Thank you to Austin, Ali, Jack, Keith, Art, Dre, Andi and Janine for the work you do and the warmth you do it with, which proved a source of great comfort in distant and unfamiliar places.

Finally, and most importantly, holla at my dog Kitty, for everything.

Blank page

1 Introduction

The conservation of baryon number \mathcal{B} in the Standard Model of particle physics is an accidental symmetry. It is not required by any fundamental laws, and its removal does not break the Standard Model in any considerable manner. Furthermore, the removal of this symmetry provides fresh avenues for the appearance of new physics, and new opportunities to search for such new physics experimentally.

There is much precedent for the oscillation of neutral particles. The oscillation of flavoured neutral mesons such as the K^0 into their own antiparticles is a matter of record, as is the flavour oscillation of neutrinos. Permitting non-conservation of \mathcal{B} opens up a new type of neutral particle oscillation: $n - \bar{n}$ oscillation, the oscillation of a neutron into an antineutron [7].

In the Standard Model of particle physics, both lepton number \mathcal{L} and baryon number \mathcal{B} are independently conserved. Many beyond Standard Model (BSM) theories — for instance, those that predict the decay of the proton — allow for \mathcal{L} and \mathcal{B} to be violated independently, but preserve their combination $\mathcal{B} - \mathcal{L}$ as a symmetry, allowing the proton to decay into a leptonic final state [8]. However, motivations exist to move beyond even this model, to models that discard such conservations altogether and allow \mathcal{B} and \mathcal{L} conservation to be violated without any such concessions.

The rise in use of liquid argon time projection chambers (LArTPCs) presents new opportunities for experimental searches for $n - \bar{n}$ oscillation. The future Deep Underground Neutrino Experiment (DUNE) [9] offers the promise of a high-resolution detector, shielded from backgrounds due to its extreme depth, capable of sensitivity beyond current limits due to its large mass.

This thesis describes work carried out to quantify the prospects for pushing the boundaries of parameter space in terms of $n - \bar{n}$ oscillation lifetime. It provides a brief summary of the theory and phenomenology of $n - \bar{n}$ oscillation, and summarises previous experimental measurements — both past searches for $n - \bar{n}$ oscillation, and also a review of the antiproton scattering data from which $n - \bar{n}$ oscillation branching ratios are derived. It then summarises the development of a GENIE event generator to simulate the $n - \bar{n}$ oscillation process, accounting for various nuclear effects and final state interactions.

Finally, the implementation of a novel event selection technique for $n - \bar{n}$ oscillation searches is presented, in the form of a convolutional neural network (CNN). Following a brief summary of CNN operating principles, a study to classify LArTPC events using CNNs is described. The network is used to distinguish simulations of $n - \bar{n}$ signal events and atmospheric neutrino backgrounds; the performance of the network is discussed, and its resolving power is described in terms of a 90% confidence level free $n - \bar{n}$ lifetime sensitivity in DUNE. This is the first $n - \bar{n}$ lifetime sensitivity calculation carried out for DUNE. Finally, a preliminary search for $n - \bar{n}$ oscillation using data from the on-surface MicroBooNE detector is presented, and challenges in overcoming data-MC disagreements for CNNs as an analysis tool are discussed.

2 Theory

The signature of $n - \bar{n}$ oscillation itself is independent of the gauge theories which allow the process to occur. This is in contrast to, for instance, proton decay, in which different models predict different final states, leading to a wide variety of model-dependent topologies; no such variation between models exists for $n - \bar{n}$ oscillation. Consequently, the phenomenology of the process and the theories that predict it can be discussed independently.

This chapter summarises the phenomenology behind this process in several contexts: in a field-free vacuum, a magnetic field, and bound in a nucleus. For the latter, it also summarises the calculation of a nuclear suppression factor, and describes some of the beyond the Standard Model theories that predict this process.

2.1 Phenomenology

At the basic level, we can construct an effective Hamiltonian H_{eff} to describe neutron-antineutron oscillation [7, 10]. The off-diagonal transition matrix elements are real and described as

$$\langle \bar{n} | H_{\text{eff}} | n \rangle = \langle n | H_{\text{eff}} | \bar{n} \rangle \equiv \delta m , \quad (2.1)$$

where $|n\rangle$ and $|\bar{n}\rangle$ are eigenstates for the neutron and antineutron, respectively, and δm is a small but non-zero quantity that allows transition between the two.

The diagonal matrix elements have a complex component and are written

$$\langle n | H_{\text{eff}} | n \rangle = M_{11}, \quad \langle \bar{n} | H_{\text{eff}} | \bar{n} \rangle = M_{22} , \quad (2.2)$$

where M_{11} and M_{22} are a general form for the diagonal matrix elements.

The imaginary component of the diagonal matrix elements is $\text{Im}(M_{jj}) = -i\lambda/2$ for $j = 1, 2$, where $\lambda = \tau_n^{-1} = 1/880 \text{ s}^{-1}$, the inverse of the mean free neutron lifetime τ_n . The full effective Hamiltonian in the (n, \bar{n}) basis is

$$\mathcal{M} = \begin{pmatrix} M_{11} & \delta m \\ \delta m & M_{22} \end{pmatrix} . \quad (2.3)$$

At this point, we also define the difference between the two mass states as

$$\Delta M \equiv M_{11} - M_{22} . \quad (2.4)$$

To obtain the energy eigenvalues, we must diagonalise the Hamiltonian matrix in Equation 2.3 [11]. First, we separate this matrix into two components,

$$\mathcal{M} = \begin{pmatrix} \frac{1}{2}(M_{11} + M_{22}) & 0 \\ 0 & \frac{1}{2}(M_{11} + M_{22}) \end{pmatrix} + \begin{pmatrix} \frac{\Delta M}{2} & \delta m \\ \delta m & -\frac{\Delta M}{2} \end{pmatrix} . \quad (2.5)$$

This is more succinctly expressed as

$$\mathcal{M} = \frac{1}{2}(M_{11} + M_{22})I + \frac{1}{2}\Delta M K , \quad (2.6)$$

where I is a 2×2 identity operator and K is an operator defined as

$$K = \begin{pmatrix} 1 & \frac{2\delta m}{\Delta M} \\ \frac{2\delta m}{\Delta M} & -1 \end{pmatrix} . \quad (2.7)$$

The K operator has its own eigenvalues $\kappa_{1,2}$, defined as

$$K |n_{1,2}\rangle = \kappa_{1,2} |n_{1,2}\rangle . \quad (2.8)$$

The energy eigenvalues can then be described in terms of the κ eigenval-

ues,

$$E_{1,2} = \frac{1}{2}(M_{11} + M_{22}) + \frac{1}{2}\Delta M \kappa_{1,2} . \quad (2.9)$$

We then introduce an angle θ , defined as

$$\tan(2\theta) = \frac{2\delta m}{\Delta M} . \quad (2.10)$$

By substituting Equation 2.10 into Equation 2.7, we can now express K as

$$K = \begin{pmatrix} 1 & \tan(2\theta) \\ \tan(2\theta) & -1 \end{pmatrix} . \quad (2.11)$$

We now find the eigenvalues of K by writing out the characteristic equation

$$\text{Det}[K - \kappa I] = \kappa^2 - 1 - \tan^2(2\theta) = 0 . \quad (2.12)$$

Solving this equation yields the eigenvalues of K ,

$$\kappa_{1,2} = \pm \sqrt{1 + \tan^2(2\theta)} = \pm \sqrt{1 + \left(\frac{2\delta m}{\Delta M}\right)^2} = \pm \frac{\sqrt{(\Delta M)^2 + 4(\delta m)^2}}{\Delta M} . \quad (2.13)$$

We can then substitute Equation 2.13 into Equation 2.9 to provide the energy eigenvalues

$$E_{1,2} = \frac{1}{2} \left[M_{11} + M_{22} \pm \sqrt{(\Delta M)^2 + 4(\delta m)^2} \right] , \quad (2.14)$$

with the energy difference between the two states

$$\Delta E = E_1 - E_2 = \sqrt{(\Delta M)^2 + 4(\delta m)^2} . \quad (2.15)$$

The oscillation probability of a purely neutron state as a function of time is written as

$$\begin{aligned}
P(n(t) = \bar{n}) &= |\langle \bar{n} | n(t) \rangle|^2 = \sin^2(2\theta) \sin^2[(\Delta E)t/2] e^{-\lambda t} \\
&= \left[\frac{(\delta m)^2}{(\Delta M/2)^2 + (\delta m)^2} \right] \sin^2 \left[\sqrt{(\Delta M/2)^2 + (\delta m)^2} t \right] e^{-\lambda t} .
\end{aligned} \tag{2.16}$$

We can consider ΔM to be a term that generalises any differences in interactions between neutrons and antineutrons — as will be shown in the following sections, this difference can arise in vastly different ways depending on the circumstances.

There is a difference in scale between the off-diagonal mixing term δm and the mass difference ΔM . The existing upper bound on δm is on the order of 10^{-29} MeV [7], approximately 32 orders of magnitude smaller than the neutron mass. Therefore for any experimental environment, δm can be assumed to be orders of magnitude smaller than ΔM , and thus $\Delta E \approx \Delta M$. This allows Equation 2.16 to be simplified to

$$P(n(t) = \bar{n}) = \left(\frac{2\delta m}{\Delta E} \right)^2 \sin^2 \left(\frac{\Delta E \cdot t}{2} \right) e^{-\lambda t} . \tag{2.17}$$

2.1.1 Neutron-antineutron oscillation in field-free vacuum

In a vacuum, free of any ambient magnetic fields or potentials, the energy expectation values are

$$M_{11} = M_{22} = m_n - \frac{i\lambda}{2} , \tag{2.18}$$

and so the mass difference $\Delta M = 0$.

The values M_{11} and M_{22} from Equation 2.18 are substituted into the general form of the matrix from Equation 2.3 to produce the free matrix

$$\mathcal{M}_f = \begin{pmatrix} m_n - i\lambda/2 & \delta m \\ \delta m & m_n - i\lambda/2 \end{pmatrix} . \tag{2.19}$$

As before, in the general case, diagonalising this matrix provides the mass eigenstates

$$|n_{\pm}\rangle = \frac{1}{\sqrt{2}}(|n\rangle \pm |\bar{n}\rangle) . \quad (2.20)$$

The energy eigenvalues in this situation are

$$E_{1,2} = (m_n \pm \delta m) - \frac{i\lambda}{2} . \quad (2.21)$$

In this case there is maximal mixing, with $\theta = \pi/4$, as defined in Equation 2.10. Since the quasi-free case corresponds to the case where $\Delta E \cdot t \ll 1$, the small-angle approximation can be used to simplify the oscillation probability in Equation 2.17 to

$$P(n(t) = \bar{n}) \approx \left(\frac{t}{\tau_{n\bar{n}}} \right)^2 e^{-\lambda t} , \quad (2.22)$$

where $\tau_{n\bar{n}}$ is the free $n - \bar{n}$ oscillation lifetime,

$$\tau_{n\bar{n}} = \frac{1}{\delta m} . \quad (2.23)$$

Section 3.2 in the following chapter describes experimental searches for $n - \bar{n}$ oscillation using neutron beams in such quasi-free conditions.

2.1.2 Neutron-antineutron oscillation in magnetic fields

In environments where neutrons experience a constant magnetic field \vec{B} , the opposite signs of the neutron and antineutron magnetic dipole moments $\mu_{n,\bar{n}}$ introduce a new term into the effective Hamiltonian matrix, compared to the free Hamiltonian in Equation 2.19. The resulting matrix is

$$\mathcal{M}_B = \begin{pmatrix} m_n - \vec{\mu}_n \cdot \vec{B} - i\lambda/2 & \delta m \\ \delta m & m_n + \vec{\mu}_n \cdot \vec{B} - i\lambda/2 \end{pmatrix} . \quad (2.24)$$

The difference between the two diagonal mass terms, previously defined in Equation 2.4, is then

$$|\Delta M| = 2\vec{\mu}_n \cdot \vec{B} . \quad (2.25)$$

The energy eigenvalues can be described by substituting M_{11} , M_{22} and ΔM into the general form in Equation 2.14, to give

$$E_{1,2} = m_n \pm \sqrt{(\vec{\mu}_n \cdot \vec{B})^2 + (\delta m)^2} - \frac{i\lambda}{2} . \quad (2.26)$$

For free neutrons, the presence of a magnetic field suppresses the rate at which $n - \bar{n}$ oscillation occurs. If the magnetic field is sufficiently small, and free neutron time of flight is much smaller than the decay lifetime of the neutron, the $n - \bar{n}$ oscillation probability is independent of magnetic field. Experimental searches using neutron beams must use magnetic shielding to counteract any ambient magnetic fields, and prevent any associated suppression of $n - \bar{n}$ oscillation.

2.1.3 Neutron-antineutron oscillation in the nucleus

Experimental searches for bound $n - \bar{n}$ oscillation are the focus of the work discussed in this thesis. For neutrons bound inside the nucleus, the mechanism of $n - \bar{n}$ oscillation is entirely different to the field-free and magnetic field cases.

The effective masses of the neutron and antineutron are modified by the potentials they experience — the neutron potential is real, $V_n = V_{nR}$, but the antineutron potential has an imaginary component, associated with its ability to annihilate with other nucleons, and is written $V_{\bar{n}} = V_{\bar{n}R} - iV_{\bar{n}I}$.

We can use these potentials to define effective masses for the bound neutron and antineutron,

$$m_{n,\text{eff}} = m_n + V_n, \quad m_{\bar{n},\text{eff}} = m_n + V_{\bar{n}} . \quad (2.27)$$

These effective masses can then be used to construct an effective Hamiltonian matrix for $n - \bar{n}$ oscillations in matter,

$$\mathcal{M}_m = \begin{pmatrix} m_{n,\text{eff}} & \delta m \\ \delta m & m_{\bar{n},\text{eff}} \end{pmatrix}. \quad (2.28)$$

The potentials V_{nR} , $V_{\bar{n}R}$ and $V_{\bar{n}I}$ are all $O(100)$ MeV, and so $n - \bar{n}$ mixing is heavily suppressed [7]. The mixing angle θ is defined as

$$\tan(2\theta) = \frac{2\delta m}{|m_{n,\text{eff}} - m_{\bar{n},\text{eff}}|} = \frac{2\delta m}{\sqrt{(V_{nR} - V_{\bar{n}R})^2 + V_{\bar{n}I}^2}} \ll 1. \quad (2.29)$$

Diagonalising \mathcal{M}_m provides the energy eigenvalues

$$E_{1,2} = \frac{1}{2} \left[m_{n,\text{eff}} + m_{\bar{n},\text{eff}} \pm \sqrt{(m_{n,\text{eff}} - m_{\bar{n},\text{eff}})^2 + 4(\delta m)^2} \right]. \quad (2.30)$$

Focusing on the first mass eigenstate, and substituting in the effective mass terms from Equation 2.27, we obtain

$$E_1 \approx m_n + V_n - i \frac{(\delta m)^2 V_{\bar{n}I}}{(V_{nR} - V_{\bar{n}R})^2 + V_{\bar{n}I}^2}. \quad (2.31)$$

The imaginary component of this expression is responsible for generating the nuclear instability due to nucleon-antinucleon annihilation. The rate of this particular process, in which the neutron spontaneously oscillates and subsequently annihilates to produce pions, is

$$\Gamma_m = \frac{1}{\tau_m} = \frac{2(\delta m)^2 |V_{\bar{n}I}|}{(V_{nR} - V_{\bar{n}R})^2 + V_{\bar{n}I}^2}. \quad (2.32)$$

Due to large differences in neutron and antineutron energy states and long lifetimes, we can make the opposite assumption to the quasi-free case in Equation 2.22 — that $\Delta E \cdot t \gg 1$. We assume that the \sin^2 term in Equation 2.17 averages to 1/2, and the oscillation probability is reduced to

$$P(n(t) = \bar{n}) \approx \frac{1}{2} \left(\frac{2\delta m}{\Delta E} \right)^2 e^{-\lambda t}. \quad (2.33)$$

2.2 Bound to free lifetime conversion

Friedman and Gal perform in-depth calculations of the suppression of $n - \bar{n}$ oscillation [12]. They parameterise the oscillation of the free neutron by splitting the degenerate mass states into two eigenstates $m \pm \delta m$. In the nuclear medium, this splitting is entirely suppressed, but is replaced by two distinct widths — one for the nuclear disappearance lifetime of an antineutron, and one for the lifetime of a bound neutron, described as

$$\Gamma_{\text{m}} \approx \left(\frac{4\delta m}{\Gamma_{\bar{n}}} \right) \delta m , \quad (2.34)$$

where Γ_{m} is the nuclear disappearance width per nucleon and $\Gamma_{\bar{n}} \approx 320$ MeV is the annihilation width of \bar{n} inside the nucleus [12]. Taking into account the different nuclear potentials V_n and $V_{\bar{n}}$, they use a closure approximation [13] to refine Equation 2.34 to

$$\Gamma_{\text{m}}^{\text{closure}} \approx \Gamma_{\bar{n}} \frac{(\delta m)^2}{\langle V_{\bar{n}\text{I}} \rangle^2 + \langle V_{\bar{n}\text{R}} - V_{n\text{R}} \rangle} . \quad (2.35)$$

Here $\langle V_{\bar{n}\text{I}} \rangle = \Gamma_{\bar{n}}/2$ is an average of the imaginary part of the \bar{n} nuclear potential.

The relationship between the free $n - \bar{n}$ oscillation lifetime $\tau_{n\bar{n}}$ and the $n - \bar{n}$ oscillation lifetime of a neutron inside the nucleus T_{m} is

$$\tau_{n\bar{n}} \approx 2 \sqrt{\frac{\hbar T_{\text{m}}}{\Gamma_{\bar{n}}}} . \quad (2.36)$$

Using the single particle shell model, we can write the single particle wavefunctions for the neutron $u_{\nu\ell j}(r)$ and antineutron $w_{\nu\ell j}(r)$ in the nucleus,

$$\begin{aligned} -\frac{\hbar^2}{2\mu} u_{\nu\ell j}''(r) + \frac{\hbar\ell(\ell+1)}{2\mu r^2} u_{\nu\ell j}(r) - V_{n\text{R}}(r) u_{\nu\ell j}(r) \\ - E_{\nu\ell j} u_{\nu\ell j}(r) + \delta m w_{\nu\ell j}(r) = 0 , \end{aligned} \quad (2.37)$$

$$\begin{aligned}
& -\frac{\hbar^2}{2\mu} w''_{\nu\ell j}(r) + \frac{\hbar\ell(\ell+1)}{2\mu r^2} w_{\nu\ell j}(r) - [V_{\bar{n}R}(r) - iV_{\bar{n}I}(r)] w_{\nu\ell j}(r) \\
& - E_{\nu\ell j} w_{\nu\ell j}(r) + \delta m u_{\nu\ell j}(r) = 0 , \quad (2.38)
\end{aligned}$$

where the $\nu\ell j$ subscripts refer to the quantum numbers for a given single-particle state, $-V_{nR}(r)$ and $-(V_{\bar{n}R}(r) - iV_{\bar{n}I}(r))$ are the nuclear potentials experienced by neutrons and antineutrons, respectively, and μ is the effective mass of a single nucleon. The wavefunctions $u_{\nu\ell j}(r)$ and $w_{\nu\ell j}(r)$ are constant inside the nucleus and decay with r outside.

The energy values in Equations 2.37 and 2.38 are defined as

$$E_{\nu\ell j} = B_{\nu\ell j} + \frac{i\Gamma_{\nu\ell j}}{2} , \quad (2.39)$$

where the real component $B_{\nu\ell j}$ is the binding energy for a given $\nu\ell j$ single-particle state, and the imaginary component $\Gamma_{\nu\ell j}$ is the disappearance width.

Multiplying Equation 2.37 by $u_{\nu\ell j}^*(r)$, and its complex conjugate by $u_{\nu\ell j}(r)$, and then subtracting them, we arrive at the expression

$$i\Gamma_{\nu\ell j}|u_{\nu\ell j}(r)|^2 + \delta m(w_{\nu\ell j}(r)u_{\nu\ell j}^*(r) - w_{\nu\ell j}^*(r)u_{\nu\ell j}(r)) = 0 , \quad (2.40)$$

where Equation 2.39 has been substituted for $E_{\nu\ell j}$, and $w_{\nu\ell j}^*(r)$ is the complex conjugate of the antineutron single particle wavefunction $w_{\nu\ell j}(r)$.

By using the identity that for any complex vector v ,

$$\text{Im}(v) = \frac{1}{2i}(v - v^*) , \quad (2.41)$$

we can remove radial dependence by integrating Equation 2.40 over r from $0 \rightarrow \infty$, at which point it becomes

$$\Gamma_{\nu\ell j} = -\frac{2\delta m \int \text{Im}(w_{\nu\ell j}(r) u_{\nu\ell j}^*(r)) dr}{\int |u_{\nu\ell j}(r)|^2 dr} . \quad (2.42)$$

Using the condition that at $t = 0$ the initial state is purely neutrons, with

no antineutron component, we can assume that $|w/u| = \mathcal{O}(\delta m/B)$, where B represents the binding energy (ie. $\Gamma_{\bar{n}}$). Using Equation 2.42, we then conclude that $\Gamma_{\nu\ell j}$ is of order $(\delta m)^2/\Gamma_{\bar{n}}$, and neglect the $\Gamma_{\bar{n}}$ term in Equation 2.39 which is of order $\delta m/B^2$. The δm and $E_{\nu\ell j}$ terms in Equations 2.37 and 2.38 reduced to a term representing the neutron single-particle binding energy, $B_{\nu\ell j}$. The solutions $u_{\nu\ell j}$ are real functions, and Equation 2.42 describes the relationship between two negligible terms which we neglect. We then use these assumptions to rewrite the radial wave equations as

$$-\frac{\hbar^2}{2\mu} u_{\nu\ell j}''(r) + \frac{\hbar^2 \ell(\ell+1)}{2\mu r^2} u_{\nu\ell j}(r) - V_{nR}(r) u_{\nu\ell j}(r) + B_{\nu\ell j}^{(n)} u_{\nu\ell j}(r) = 0. \quad (2.43)$$

The antineutron radial wave equation is now written with a reduced wavefunction $v_{\nu\ell j}(r) = w_{\nu\ell j}/\delta m$,

$$-\frac{\hbar^2}{2\mu} v_{\nu\ell j}''(r) + \frac{\hbar^2 \ell(\ell+1)}{2\mu r^2} v_{\nu\ell j}(r) - [V_{nR}(r) - iV_{\bar{n}I}(r)] v_{\nu\ell j}(r) + B_{\nu\ell j}^{(n)} v_{\nu\ell j}(r) + v_{\nu\ell j}(r) = 0. \quad (2.44)$$

By applying the same method to Equations 2.43 and 2.44 as was used to obtain Equation 2.42, we obtain

$$-2(\delta m)^2 \int \text{Im}(v_{\nu\ell j}(r) u_{\nu\ell j}^*(r)) dr = 2(\delta m)^2 \int V_{\bar{n}I}(r) |v_{\nu\ell j}(r)|^2 dr. \quad (2.45)$$

We can then substitute the right hand side of Equation 2.45 into Equation 2.42, to obtain

$$\begin{aligned}\Gamma_{\nu\ell j} &= \frac{2 (\delta m)^2 \int V_{\bar{n}1}(r) |v_{\nu\ell j}(r)|^2 dr}{\int u_{\nu\ell j}^2(r) dr} \\ &= -\frac{2 (\delta m)^2 \int u_{\nu\ell j}(r) \text{Im}(v_{\nu\ell j}(r)) dr}{\int u_{\nu\ell j}^2(r) dr} .\end{aligned}\quad (2.46)$$

The averaged disappearance width per neutron is

$$\Gamma_m = \frac{1}{N} \sum n_{\nu\ell j} \Gamma_{\nu\ell j} , \quad (2.47)$$

where $n_{\nu\ell j}$ is the number of neutrons in each $\nu\ell j$ single particle state, and $N = \sum n_{\nu\ell j}$ is the total number of neutrons in the nucleus. The bound $n - \bar{n}$ oscillation width Γ_m is then related to the free $n - \bar{n}$ oscillation lifetime $\tau_{n-\bar{n}}$ via a conversion factor R , defined as

$$R = \frac{\hbar}{\Gamma_m \tau_{n\bar{n}}^2} , \quad (2.48)$$

where R has units of inverse time.

The nuclear disappearance lifetime T_m is

$$T_m = R \tau_{n\bar{n}}^2 . \quad (2.49)$$

To relate the bound $n - \bar{n}$ oscillation lifetime to the equivalent lifetime of a free neutron, theoretical conversion parameters are calculated by finding numerical solutions to the above equations. Since the nuclear potential experienced by the neutron and antineutron varies for different nuclei, this conversion factor must be calculated independently for each nuclear isotope.

These calculations of nuclear conversion factor are reproduced from [12]; their calculation of this conversion factor was carried out for ^{16}O and ^{56}Fe , but not directly for ^{40}Ar . However, in private communication the authors of this paper estimate that the ^{40}Ar could be assumed to be the same as the ^{56}Fe value — $R = 0.666 \times 10^{23} \text{ s}^{-1}$ — with an additional 10% uncertainty on top of the existing 30% uncertainty.

2.3 Baryon number violation

In the Standard Model, quarks have a baryon number \mathcal{B} of $1/3$, and baryons have $\mathcal{B} = 1$; their antiparticles correspondingly have $\mathcal{B} = -1/3$ and $\mathcal{B} = -1$. The transition of a neutron into an antineutron involves the simultaneous transition of six quarks into antiquarks. Baryon number is $\mathcal{B} = 1$ in the initial state and $\mathcal{B} = -1$ in the final state, a change of $|\Delta\mathcal{B}| = 2$.

In the Standard Model Lagrangian, baryon number is conserved, but it is a so-called “accidental symmetry” — it is conserved, but this conservation is not associated with any continuous global symmetry in the Standard Model. It is therefore possible to construct Grand Unified Theories (GUTs) which permit the non-conservation of baryon number.

The principal motivation for searching for baryon number violation is the search for the source of the universe’s matter-antimatter asymmetry, via a process called baryogenesis. This process was first proposed by Andrei Sakharov in 1967 [14], who also suggests the three conditions necessary for baryogenesis to occur:

- Baryon number violation
- C and CP violation
- Interactions outside of thermal equilibrium

Baryon number violation is therefore a vital component for explaining the matter-antimatter asymmetry of the universe through baryogenesis. The possibility of $n - \bar{n}$ oscillations was first suggested not long afterwards, in a paper by V. A. Kuzmin [15]. Experimental observation of $n - \bar{n}$ oscillation would constitute definitive proof of \mathcal{B} non-conservation, and have serious implications for our understanding of the origin of the universe.

The left-right symmetric extension to the Standard Model has the gauge group $SU(3)_c \otimes SU(2)_L \otimes SU(2)_R \otimes U(1)_{\mathcal{B}-\mathcal{L}}$, introducing $\mathcal{B} - \mathcal{L}$ as a gauge symmetry of the $U(1)_{\mathcal{B}-\mathcal{L}}$ group [7]. Such models can give rise to proton decay, which permits non-zero $\Delta\mathcal{B}$ and $\Delta\mathcal{L}$ but leaves their combination $\Delta(\mathcal{B} - \mathcal{L})$ unchanged.

The electric charge can be expressed as

$$Q = I_{3L} + I_{3R} + \frac{\mathcal{B} - \mathcal{L}}{2} , \quad (2.50)$$

where I_{3L} and I_{3R} are the left and right handed isospins [16]. At distance scales below the electroweak scale, $I_{3L} = 0$ and so

$$\Delta I_{3R} = -\Delta \frac{\mathcal{B} - \mathcal{L}}{2} , \quad (2.51)$$

suggesting the possibility of spontaneously broken $\mathcal{B} - \mathcal{L}$ by units of 2.




Figure 2.1: A diagram for $n - \bar{n}$ oscillation in a G_{422} left-right symmetric partial unification model (from [7]).

Furthermore, the $\mathcal{B} - \mathcal{L}$ symmetry and colour symmetry can be combined to form the group $G_{422} = \text{SU}(4) \otimes \text{SU}(2)_L \otimes \text{SU}(2)_R$, as in the partial unification model proposed by Mohapatra and Marshak [16]. This model unifies the quarks and leptons into the fields $\psi : (2, 1, 4)$ and $\psi^c : (2, 1, \bar{4})$; the Higgs fields are $\phi_1 : (2, 2, 1)$ and $\phi_{15} : (2, 2, 15)$ for fermions, and additional fields $\Delta^c : (1, 3, 10)$ and $\bar{\Delta}^c : (1, 3, \bar{10})$ for the spontaneous breaking of $\mathcal{B} - \mathcal{L}$ symmetry. A diagram for $n - \bar{n}$ oscillation in such a model is shown in Figure 2.1.

The spontaneous breaking of $\mathcal{B} - \mathcal{L}$ can give rise to Majorana neutrinos and the $\Delta\mathcal{L} = 2$ neutrinoless double β decay process, and also allows $\Delta\mathcal{B} = 2$

processes, giving rise to $n - \bar{n}$ oscillations. Due to this relationship between the violation of \mathcal{B} and \mathcal{L} , observation of $n - \bar{n}$ oscillation would also provide strong hints as to the nature of neutrino mass [7].

3 Measurement History

As discussed in Chapter 2, $n - \bar{n}$ oscillation manifests differently under different circumstances: the lifetime of free neutrons oscillating in a magnetic field is vastly different from bound neutrons oscillating inside a nucleus. This enables searches using different experimental setups — namely, searches in free neutron beams and in large-mass detectors.

Due to the scale of theory-derived conversion factors between free and bound lifetime limits, both experimental approaches offer competitive sensitivity despite vastly differing techniques. The following sections describe all $n - \bar{n}$ lifetime experimental searches to date (summarised in Table 3.1), and all limits discussed are lower limits at 90% confidence level. The current best limit on the lifetime of free $n - \bar{n}$ oscillation is 2.7×10^8 s, set by the Super-Kamiokande experiment in 2015 [1], converted from a bound lifetime using a factor derived from theory.

3.1 Free neutron searches

One approach to measuring $n - \bar{n}$ oscillations is observing a beam of quasi-free neutrons focused on a target and searching for nucleon-antinucleon annihilations, as such topologies would be indicative of a neutron oscillating into an antineutron during flight. This experimental technique was first proposed in 1980 [13], and carried out several times over the following two decades.

The first experiments to search for neutron-antineutron oscillation in free neutron beams were based at the 57 MW Reactor at the Institut Laue-Langevin (ILL), which provided an intense source of slow neutrons. Two ILL searches were published with beams of quasi-free neutrons: the first was

published in 1985, setting a lifetime limit of 10^6 s [17], and the second in 1994, with an improved limit of 8.6×10^7 s [18].

Both experiments use largely identical experimental approaches: neutrons are cooled to a temperature of 25 K, and sent through a drift vessel where they can propagate for $t \approx 0.1$ s. This vessel is separated out into an oscillation region, an annihilation region and a beam dump. Due to the significant effect of even infinitesimal magnetic fields on oscillation probability (as discussed in greater detail in Section 2.1.2), the oscillation region is protected by a μ -metal shield to cancel out the Earth's magnetic field and any other residual ambient magnetic fields. The annihilation region contains a carbon foil target and is surrounded by a detector consisting of three components: a vertex detector, a scintillation counter hodoscope and a calorimeter. A diagram of the experimental setup is provided in Figure 3.1.




Figure 3.1: Diagram of experimental setup for the neutron-antineutron oscillation search conducted with a beam of free neutrons at the Institut Laue-Langevin [18].

In addition to the ILL experiment, an additional search with free neutron beams was carried out at the University of Pavia in 1990, using the 250 kW Triga Mark II reactor [19]. This search used a similar experimental setup to the ILL search, and set a lifetime limit of 4.9×10^5 s, observing zero signal and zero background events.

ILL's search is the most recent attempt to constrain the lifetime of $n - \bar{n}$ oscillations using quasi-free neutron beams and still provides the most stringent direct limit on free $n - \bar{n}$ oscillation over twenty years later. However,

there are prospects for this limit to be surpassed in the foreseeable future with the upcoming European Spallation Source (ESS), a high-intensity pulsed spallation neutron source set to come online in 2019 [20].

An experiment has been proposed, which would take advantage of this neutron source to carry out another search for $n - \bar{n}$ oscillation, making use of newer technologies such as ellipsoidal reflectors that enhance the neutron beam intensity at the target to improve sensitivity. The experimental proposal claims a factor of ~ 1000 improvement in sensitivity over the ILL limit assuming a run time of 3-5 years at ESS [20], corresponding to a free lifetime sensitivity at the $\sim 10^{10}$ s level.

3.2 Bound neutron searches

As discussed in Section 2.1.3, bound neutrons can also spontaneously oscillate into antineutrons, but the lifetime is heavily suppressed in comparison to the free lifetime. However, since very large detectors can contain many orders of magnitude more neutrons than a quasi-free neutron beam, this suppressed lifetime can be offset with a far greater number density, offering lifetime limits competitive with quasi-free $n - \bar{n}$ oscillation experiments.

Several experiments have searched for bound $n - \bar{n}$ oscillations using large detectors and set lifetime limits. These bound lifetime limits can be expressed in terms of an equivalent free lifetime, using conversion factors derived from theory as described in Section 2.2. Different conversion factors are applied depending on the nucleus, as $n - \bar{n}$ rate suppression is dependent on the nuclear potential.

The first search for bound $n - \bar{n}$ oscillation was carried out at the Homestake mine in 1983 [21], shortly before the first free search at ILL. A 300-ton water Cherenkov detector was used to search for the process using neutrons bound in ^{16}O , setting a free-equivalent limit of 2×10^7 s.

This limit was surpassed by the Irvine-Michigan-Brookhaven collaboration in 1984 [22], which again used a water Cherenkov detector, with a much larger fiducial mass of 3300 tons. A limit of 1.1×10^8 s was set by defining an ‘isotropy angle’ that separated $n - \bar{n}$ signal from atmospheric neutrino

background according to directionality of Cherenkov light.

A third water Cherenkov measurement was made by the Kamiokande collaboration in 1986 [23]. This measurement was carried out by identifying events with Cherenkov rings from a central vertex, and making cuts on total net momentum and invariant mass. A free-equivalent lifetime limit of 1.2×10^8 s was set.

In 1990, the first non-Cherenkov bound $n - \bar{n}$ oscillation search was carried out in the Frèjus fine-grained tracker detector [24]. The Frèjus detector was a 900-ton iron calorimeter, at a depth of 1.8 km underground. Instead of Oxygen, this experiment used ^{56}Fe as its target nucleus, and collected data for a total of 1600 days. It measured an event rate consistent with its atmospheric neutrino background prediction, and set a lifetime limit of 1.2×10^8 s.

Over a decade later, the Soudan 2 collaboration set a new limit on the $n - \bar{n}$ oscillation lifetime [25], again with an iron tracking calorimeter. Using a 770 ton fiducial mass detector at a depth of 700 m, with an exposure of 5.56 kt·yrs, a lifetime limit of 1.3×10^8 s was set.

After another decade-long break, in 2015 the Super-Kamiokande collaboration performed a search for $n - \bar{n}$ oscillation with a water Cherenkov detector [1]. With a fiducial mass of 22.5 kt and an exposure of 1489 live-days, a limit of 2.7×10^8 s was set — a factor of two improvement over the previous best limit. To date, this is the strictest limit set on the neutron oscillation lifetime.

In 2017, the SNO collaboration published a search for $n - \bar{n}$ oscillation in their heavy water Cherenkov detector, considering D and ^{16}O nuclei [26]. Their analysis yielded fewer events than expected — 23 events, a 1.6σ deficiency compared to the expected background of 30.5 atmospheric neutrino events — and set a limit of 1.23×10^8 s on the free $n - \bar{n}$ oscillation lifetime.

Table 3.1: Chronological list of all searches for $n - \bar{n}$ oscillation, and the limits they set on the free lifetime of the process. For searches in bound nuclei, the target nucleus and corresponding rate suppression factor (as defined in Equation 2.49) are provided.

Experiment		Year	Type	Nucleus	$R [10^{23} s^{-1}]$	τ_{free} limit [s]
Homestake	[21]	1983	Bound	^{16}O	1.0	2×10^7
IMB	[22]	1984	Bound	^{16}O	1.0	1.1×10^8
ILL	[17]	1985	Free	—	—	1.0×10^6
Kamiokande	[23]	1986	Bound	^{16}O	1.0	1.2×10^8
Triga Mk. II	[19]	1989	Free	—	—	4.9×10^5
Frèjus	[24]	1990	Bound	^{56}Fe	1.4	1.2×10^8
ILL	[18]	1994	Free	—	—	8.6×10^7
Soudan 2	[25]	2002	Bound	^{56}Fe	1.4	1.3×10^8
Super-Kamiokande	[1]	2015	Bound	^{16}O	1.0	2.7×10^8
SNO	[26]	2017	Bound	$^2\text{H}/^{16}\text{O}$	0.25/1.0	1.2×10^8

Table 3.2: Neutron-antineutron oscillation final state branching ratios, as used in Super-Kamiokande’s 2015 search for $n - \bar{n}$ oscillation [1]. Branching ratios are provided independently for the $\bar{n}p$ and $\bar{n}n$ annihilation processes, so each column independently sums to 100%.

$\bar{n}p$		$\bar{n}n$	
Channel	Branching ratio	Channel	Branching ratio
$\pi^+\pi^0$	1%	$\pi^+\pi^-$	2%
$\pi^+2\pi^0$	8%	$2\pi^0$	1.5%
$\pi^+3\pi^0$	10%	$\pi^+\pi^-\pi^0$	6.5%
$2\pi^+\pi^-\pi^0$	22%	$\pi^+\pi^-2\pi^0$	11%
$2\pi^+\pi^-2\pi^0$	36%	$\pi^+\pi^-3\pi^0$	28%
$2\pi^+\pi^-2\omega$	16%	$2\pi^+2\pi^-$	7%
$3\pi^+2\pi^-\pi^0$	7%	$2\pi^+2\pi^-\pi^0$	24%
		$\pi^+\pi^-\omega$	10%
		$2\pi^+2\pi^-2\pi^0$	10%

3.3 Antineutron annihilation branching ratios

Unlike proton decay, in which individual branching ratios for the decay of the proton represent distinct theoretical models, the branching ratios for $n - \bar{n}$ oscillation are due to a single standard model process, the annihilation of an antineutron with another nucleon. While the specific partial unification model through which the process is produced will have an effect on the rate, the topology of the event itself is model independent. Correspondingly, the branching ratios for $n - \bar{n}$ oscillation represent the process that occurs *after* the oscillation itself — the annihilation of an antineutron with another nucleon, a Standard Model process that can be measured experimentally.

Since $n - \bar{n}$ oscillation is a process beyond the Standard Model that has never been observed, its topology cannot be directly studied using real data. However, since the topology is the annihilation of a nucleon and an antinucleon at rest, measurements of such processes can be used to model the $n - \bar{n}$ oscillation topology.

Since antineutron scattering data is scarce and antiproton scattering data is more common, we assume that $\bar{n}n$ annihilation final states will be identical to those of $\bar{p}p$ annihilation and that $\bar{n}p$ annihilation final states are identical to those of $\bar{p}n$ after accounting for charge conjugation. Using this assumption, at-rest antiproton annihilation data can be used to extrapolate the branching ratios for $n - \bar{n}$ oscillation.

The branching ratios shown in Table 3.2 originate from the $n - \bar{n}$ oscillation search published by the Super-Kamiokande experiment [1], derived from antiproton annihilation data. We find it instructive to attempt to reproduce these ratios from the original data, and to provide a reproduction of the original data in an easily accessible format.

By following the references for these branching ratios, attempts have been made to understand the source of these numbers. The full branching ratio tables from all data sources referenced in the following sections have been reproduced in Appendix A, along with a short description of each data set

considered.

Chapter 5 describes the use of these branching ratios to develop an $n - \bar{n}$ event generator. However, this section details a subsequent attempt to contextualise these branching ratios by comparing them to the data they are derived from, and suggests avenues through which they can be refined.

3.3.1 Measurements of $\bar{p}p$ annihilation




Figure 3.2: Branching ratios from antiproton annihilation data on hydrogen, as shown in Tables A.2, A.3, A.5 and A.7. Branching ratios from different experimental measurements are shown for each final state, as is the average of these ratios. The Super-Kamiokande branching ratios are also shown for comparison.

In Figure 3.2, the Super-Kamiokande branching ratios are directly compared to the four data sets they are derived from. The average is taken of these four data sets, and a comparison between the Super-Kamiokande branching ratios and this average is shown in Table 3.3. These ratios do not agree perfectly — for instance, the branching ratio for the $\pi^+\pi^-$ channel is

consistently larger by a factor of ~ 6 in the Super-Kamiokande branching ratios than in data. A similar effect can be observed for the $\pi^+\pi^-\omega$ channel, which is ≈ 2.5 larger than data. The exact methods used to produce these Super-Kamiokande ratios are unknown, and may consider effects not taken into account in this comparison.

Table 3.3: Comparison between $\bar{n}p$ branching ratios used by Super-Kamiokande and the average of branching ratios measured with antiproton annihilation on hydrogen. $n - \bar{n}$ branching ratios consider $\bar{n}p$ and $\bar{n}n$ annihilation independently, so these branching ratios sum to 100%.

Channel	Branching ratio [%]		Ratio [SK/data]
	SK	Data average	
$n\pi^0$	1.5	2.1	0.70
$\pi^+\pi^-$	2	0.3	5.8
$\pi^+\pi^-\pi^0$	6.5	8.7	0.74
$\pi^+\pi^-n\pi^0$	39	33.2	1.2
$2\pi^+2\pi^-$	7	9.0	0.78
$2\pi^+2\pi^-\pi^0$	24	25.8	0.93
$2\pi^+2\pi^-n\pi^0$	10	18.2	0.55
$\pi^+\pi^-\omega$	10	4.2	2.4

The Super-Kamiokande branching ratios omit many high-multiplicity final states. The full reproduced branching ratio tables from data, as shown in Tables A.1, A.4 and A.6, contain many final states with 7 or more pions, including $3\pi^+3\pi^-$ channels and channels containing up to $4\pi^0$. The branching ratios for these final states are small, but not negligible. It is possible that these were omitted from the Super-Kamiokande analysis due to their high multiplicity — the $2\pi^+2\pi^-3\pi^0$ channel, for instance, would produce 10 distinct Cherenkov rings, perhaps too many to efficiently reconstruct — but these events could be more effectively reconstructed in a liquid argon time projection chamber, and would be worthwhile to include.

3.3.2 Measurements of $\bar{p}n$ annihilation

In Figure 3.3 and Table 3.4, the Super-Kamiokande branching ratios are directly compared to the data set from which they are derived. Unlike the




Figure 3.3: Branching ratios from antiproton annihilation data on neutrons in deuterium, as shown in Table 3.4. The Super-Kamiokande branching ratios are also shown for comparison.

Table 3.4: A comparison between $\bar{n}p$ branching ratios used by Super-Kamiokande and Bettini *et al.* antiproton annihilation data collected at the CERN Proton Synchrotron [27]. Only channels for which both tables have ratios available are considered; the full tables for each experiment are provided in Tables 3.2 and A.8 respectively. $n - \bar{n}$ branching ratios consider $\bar{n}p$ and $\bar{n}n$ annihilation independently, so these branching ratios sum to 100%. In order to reflect the inversion of charge between $\bar{n}p$ and $\bar{p}n$ events, the final states have been inverted with respect to Table A.8.

Channel	Branching ratio [%]		Ratio [SK/data]
	SK	Bettini <i>et al.</i>	
$\pi^+ n \pi^0$ ($n > 0$)	19	16.4	1.2
$\pi^+ \pi^0$	1	0.7	1.4
$2\pi^+ \pi^- n \pi^0$ ($n \geq 0$)	58	59.7	0.97
$2\pi^+ \pi^- \pi^0$	22	21.8	1.0
$3\pi^+ 2\pi^- n \pi^0$	7	23.4	0.30

$\bar{p}p$ data, there is good agreement in branching ratios for all channels that can be directly compared. However, the $2\pi^+\pi^-2\omega$ channel is not accounted for in annihilation data — as discussed in Section 5.3.1, this final state is often forbidden during simulation due to lack of available energy. The branching ratio for the $3\pi^+2\pi^-n\pi^0$ final state is much larger in data, since the Super-Kamiokande branching ratios omit final states with $n > 1$.

3.3.3 Crystal Barrel data and new branching ratios

The branching ratios used by Super-Kamiokande and discussed above were derived from antiproton annihilation data at the Brookhaven AGS and the CERN Proton Synchrotron. However, more recent antiproton annihilation data is available, from which more accurate branching ratios and a larger number of exclusive channels can be produced.

A recent PhD thesis [2] has produced an updated list of $n - \bar{n}$ oscillation final states. It takes data from the Crystal Barrel spectrometer and ASTERIX at LEAR, which have measured many antiproton annihilation channels in greater detail [28–30].

These improved final states from [2] are reproduced in Table 3.5. These new branching fractions are in good agreement with previous sets of antiproton annihilation data summarised above, especially for channels such as $\pi^+\pi^-$.

Table 3.5: Updated $n - \bar{n}$ oscillation final state branching ratios, as reproduced from [2] and originally derived from Crystal Barrel and ASTERIX data [28–30]. $n - \bar{n}$ branching ratios consider $\bar{n}p$ and $\bar{n}n$ annihilation independently, so each column independently sums to 100%. Uncertainties on branching fractions are not provided, but in original data sources they are typically in the 10-20% [29, 31].

$\bar{n}p$		$\bar{n}n$	
Channel	Branching ratio	Channel	Branching ratio
$2\pi^0$	0.06%	$\pi^+\pi^0$	0.1%
$3\pi^0$	0.8%	$\pi^+2\pi^0$	0.7%
$4\pi^0$	0.3%	$\pi^+3\pi^0$	14.8%
$5\pi^0$	1.0%	$\pi^+4\pi^0$	1.4%
$6\pi^0$	0.01%	$2\pi^+\pi^-$	2.0%
$7\pi^0$	0.1%	$2\pi^+\pi^-\pi^0$	17.0%
$\pi^+\pi^-$	0.3%	$2\pi^+\pi^-2\pi^0$	10.8%
$\pi^+\pi^-\pi^0$	1.6%	$2\pi^+\pi^-3\pi^0$	30.1%
$\pi^+\pi^-2\pi^0$	13.0%	$3\pi^+2\pi^-$	5.5%
$\pi^+\pi^-3\pi^0$	11.2%	$3\pi^+2\pi^-\pi^0$	2.3%
$\pi^+\pi^-4\pi^0$	3.3%		
$\pi^+\pi^-5\pi^0$	1.4%		
$2\pi^+2\pi^-$	6.0%		
$2\pi^+2\pi^-\pi^0$	13.5%		
$2\pi^+2\pi^-2\pi^0$	16.6%		
$2\pi^+2\pi^-3\pi^0$	0.6%		
$3\pi^+3\pi^-$	2.2%		
$3\pi^+3\pi^-\pi^0$	2.0%		

Blank page

4 Liquid Argon Time Projection Chambers

Liquid argon time projection chambers are a relatively modern detector technology. The first neutrino experiment to utilise the technology was the Icarus experiment [32], measuring a beam of neutrinos sent from CERN to the Gran Sasso laboratory in Italy. Following the successful use of this technology to measure neutrino interactions, the US neutrino community began to adopt this technology, first with the Argoneut TPC [33] (later used as part of the LArIAT test beam experiment [34]), and later with the larger MicroBooNE experiment [35]. In the coming few years, MicroBooNE will be joined by the SBND TPC and the refurbished Icarus detector to form an array of short-baseline neutrino detectors [36] — and in the longer term, by the Deep Underground Neutrino Experiment [9], which will be constructed over the coming decade.

This chapter discusses the history of liquid argon detectors, summarises the concepts and high-level workings of the LArTPC detector technology, discusses particle interactions and signatures, and finally describes the currently operating MicroBooNE detector and the current nominal design for the future DUNE detector.

4.1 LArTPC detector principles

LArTPCs are 3D stereoscopic high-resolution imaging detectors that measure particle interactions by collecting ionisation electrons across a large volume of liquid argon. When a charged particle propagates through liquid argon,

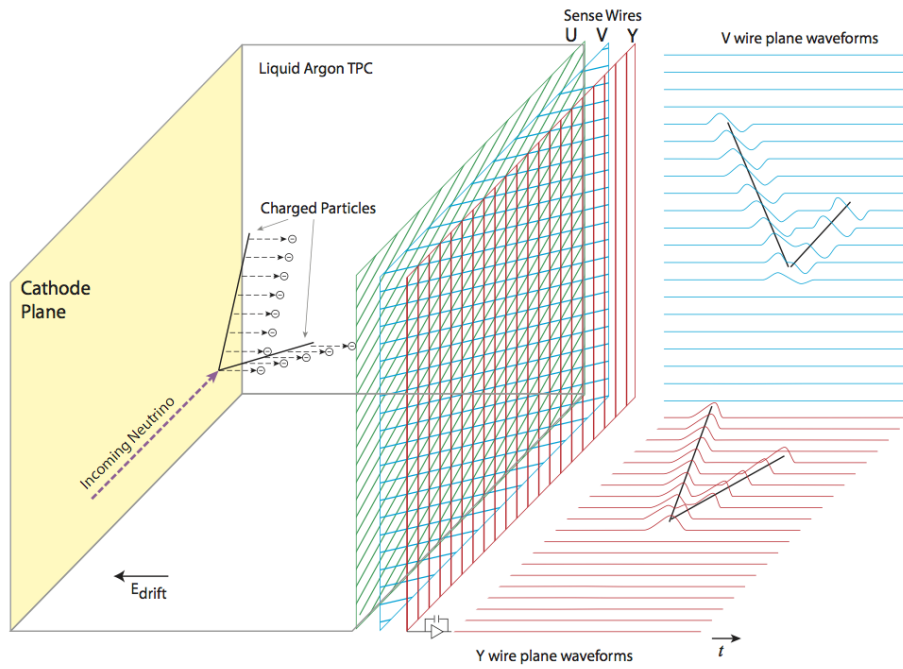


Figure 4.1: A diagram demonstrating the basic operating principles of a time projection chamber (from [35]). Charged particles produce ionisation electrons as they propagate through the liquid argon; an induced electric field then causes these electrons to drift to several planes of readout wires, where they are collected as electronic signals.

it leaves a train of ionisation electrons and positive ions as it travels and also produces scintillation light. An electric field applied across the medium causes ionisation electrons to drift. A series of wires can be used to collect these electrons and measure their charge, as shown in Figure 4.1. Purifying the argon reduces the rate at which electrons are captured by electronegative impurities and therefore increases the proportion of electrons collected. Additionally, photomultiplier tubes can be utilised to collect scintillation light.

These are the basic operating principles of a liquid argon time projection chamber. A combination of electron and light collection systems, supported by cryogenic and purification systems, allows for the precision detection of charged particles and reconstruction of their tracks.

4.1.1 Time projection chamber

The TPC itself is cuboid in shape; one side of the cryostat contains an anode plane assembly (APA), which comprises two or three planes of readout wires, each oriented at a different angle to produce a different two-dimensional image. Wire spacing is typically on the order of ~ 3 mm, providing mm-level spatial resolution. The first and last wire planes are positively and negatively biased, respectively, at $\mathcal{O}(100)$ V, to ensure ionisation electrons produce a signal on all wire planes before being collected.

The opposite side of the TPC comprises a flat metal cathode plane or grid, held at a large negative voltage in order to induce the drift field that causes ionisation electrons to drift towards the anode wires. The remaining four sides of the TPC comprise a series of metal field shaping tubes that act as a field cage. Each loop is connected to the others via variable resistors and held at a different electrical potential, maintaining a uniform drift field between the anode and cathode.

A typical value for the drift electric field strength is ≈ 500 V/cm. Given the size of LArTPCs, this generally equates to a total negative potential of $\mathcal{O}(10 - 100)$ kV on the cathode. The electric field strength must be chosen to balance spatial resolution against energy resolution — the weaker the drift field, the slower the drift of ionisation electrons, improving the spatial

resolution in the dimension perpendicular to the anode plane. However, ionisation electrons also recombine with impurities in the liquid argon with a characteristic lifetime, and so a longer drift time results in worse energy resolution.




Figure 4.2: Example of the signal induced on all three wire planes by ionisation electrons (from [37]). The x axis on the left figure shows distance parallel to the collection plane, while the y axis shows distance from collection plane in cm. The lines in the left plot represent electric field lines. The right plot shows an example of the waveforms produced on each wire plane: a bipolar but largely positive signal on the first induction plane, as electrons approach; a more evenly bipolar signal on the second induction plane, as the electrons pass by; and a unipolar negative signal on the collection plane, as the electrons are collected.

A bias voltage is applied across the three wire planes, causing electrons to drift towards the furthest plane. As they pass intermediate planes, they induce a current — first negative, as they approach, and then positive, as they pass by — before being collected by the final, positively biased plane

of wires. A TPC therefore generally comprises one or two ‘induction planes’ that produce bipolar waveforms, and a final ‘collection plane’ that produces largely unipolar waveforms (as shown in Figure 4.2).

Each wire plane provides a different 2D representation of particle interactions in the plane parallel to the APA. The location of deposited charge along the spatial dimension perpendicular to the anode plane can be inferred from electron drift time, if the t_0 of the interaction is known.

Typically, the readout wires are connected to cold electronics: signal processing electronics located within the cryostat in order to minimise electronic noise. These electronics consist of ASICs connected to cold motherboards that issue configuration commands to ASICs, as well as transfer the output. These motherboards are connected to cold cables on the interior of the cryostat, which are in turn connected to metal flanges that provide an interface to the exterior of the cryostat, where warm cables carry data for downstream DAQ processing.

4.1.2 Light detection

Particles traversing liquid argon produce light through scintillation and Cherenkov radiation. Scintillation light is emitted isotropically by the radiative decay of argon excimers, and liquid argon is a bright scintillator, yielding around 24,000 photons per MeV at 500 V/cm drift field [35], with a wavelength of 128 nm [38]. It is transparent to the scintillation light it produces. The scintillation light has two components, with decay times of 6 ns and 1.6 μ s, though the timescales of both components are far shorter than the typical electron drift time.

Since scintillation light is prompt in comparison to electron drift times, photon detection is used as a means of time-tagging events. Light detection provides a t_0 for particle interactions, which are used in conjunction with drift time to determine the final spatial dimension of particle interactions.

4.1.3 Cryogenics and argon purity

Since argon is gaseous at ambient temperatures, it must be continuously cryogenically cooled to remain in its liquid state. Liquid argon detectors make use of cryogenic systems to maintain an operating temperature of around 87 K, and must also minimise temperature gradient in order to achieve a uniform electron drift velocity [35].

The survival probabilities of ionisation electrons and scintillation photons (and therefore the detector’s ability to resolve charge and scintillation light) are dependent on the purity of liquid argon. For ionisation electrons, the primary sources of electron capture are electronegative contaminants such as O_2 , neutral molecules that attract ionisation electrons to form negative ions [39]. Nitrogen contamination must be minimised to prevent quenching and attenuation of scintillation photons [35]. Exact purity requirements depend on the specific detector design; for instance, a longer drift volume requires a higher electronegative purity to achieve the same energy resolution.

4.2 Particle signatures

The raw readout of a LArTPC is the charge on a wire, expressed in analogue-to-digital conversion (ADC) counts over a range of discrete time ticks (typically 2 MHz digitisation frequency [35], leading to $0.5 \mu s$ precision in time). Reconstruction techniques are used to process these raw waveforms, converting them into analysable information.

Several discrete stages of data processing are required in order to move from raw waveforms to fully reconstructed particles. The first of these is signal processing, which is performed on the raw waveform.

Signal processing is a reconstruction stage that accounts for the effects the detector itself has on waveforms — corrections for electronics, drift field, electron diffusion and recombination. Typically, signal deconvolution is applied to correct for these effects, resulting in a Gaussian wire signal — although prolonged activity on a wire ($\mathcal{O}(\text{ms})$) is indicative of a track parallel with the wire, in which case the Gaussian approximation is inaccurate and a fast

Fourier transform method is used instead [40].

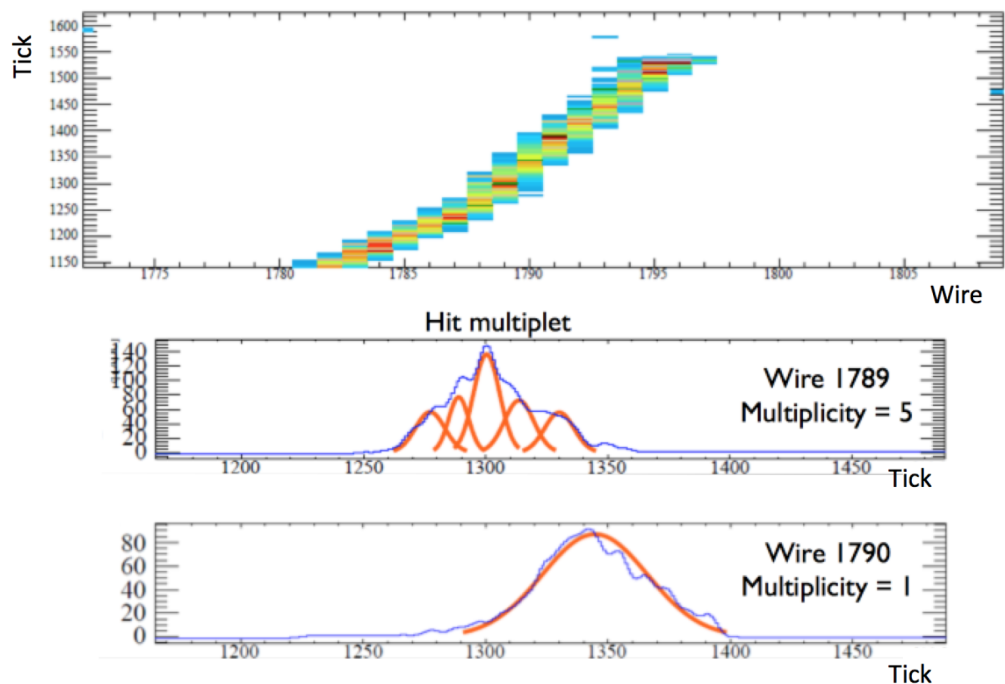


Figure 4.3: A track in a TPC (top) is detected as activity on wires, which are characterised as hits. Hits are described using Gaussian distributions — a simple signal (bottom) can be described using a single hit, while more complex activity on a wire (centre) must be described using a series of hits (from [40]). This figure shows a simulated track in the MicroBooNE TPC, using two wires in the first induction plane.

After signal processing, hit reconstruction is applied. All activity on the wire is described in terms of Gaussian distributions, as demonstrated in Figure 4.3. Once the processed waveforms have been converted into hits, object reconstruction can begin.

Pattern recognition is performed to cluster hits together and associate them with objects. Tools such as Pandora [41] cluster hits independently in each wire plane. These clusters are then matched across wire planes, and refined using information from all planes, producing clusters of 2D hits grouped over all three wire planes.

The wirecell method [42] takes a different approach — using medical

imaging techniques, the waveforms in each plane are combined into 3D hits inside the TPC, and these 3D hits are subsequently clustered into objects.

Once pattern recognition techniques have been applied, and hits have been clustered into track-like and shower-like objects, high-level reconstruction is carried out. Track and shower objects are created, with the geometry of a line and a cone respectively, and physics quantities such as specific energy loss dE/dx are computed, and particle identification performed.

4.2.1 Charged leptons and mesons

Charged leptons and mesons — most commonly μ^\pm leptons and π^\pm mesons — are identified by their long straight tracks inside LArTPCs. A muon will typically decay via $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$, while charged pions typically decay to muons. Typical energy thresholds in a LArTPC are ~ 10 MeV [43].

The rate at which ionisation occurs is described by the Bethe Bloch equation,

$$-\left\langle \frac{dE}{dx} \right\rangle = \frac{4\pi}{m_e c^2} \cdot \frac{n z^2}{\beta^2} \cdot \left(\frac{e^2}{4\pi\epsilon_0} \right) \cdot \left[\ln \left(\frac{2m_e c^2 \beta^2}{I \cdot (1 - \beta^2)} \right) - \beta^2 \right], \quad (4.1)$$

where c is the speed of light, z the particle's charge, ϵ_0 the vacuum permittivity, $\beta = v/c$, and n and I the medium's electron density and mean excitation potential respectively [39]. The density correction term that describes the charged particle's polarisation of the medium is omitted, as its effects are negligible in the particle energy regime typically observed in LArTPC neutrino experiments.

Following an energy calibration to convert each hit's ADC count into an equivalent charge, the rate of energy deposition along a charged particle's trajectory can be measured, and a track's particle type determined. For instance, the PIDA technique [44] averages the quantity

$$A_i = dE/dx_{calo,i} R_i^{0.42}, \quad (4.2)$$

where dE/dx_{calo} and R are the energy deposition and residual range for

each point along a track's trajectory, and the total score A is the average of A_i over each trajectory point i . As shown in Figure 4.4, μ s and π s are most difficult to resolve, while K and p particles have a more distinct signature.




Figure 4.4: The PIDA score for Monte Carlo charged particles in a LArTPC, as defined in Equation 4.2 (from [44]).

4.2.2 Nucleons

Protons are heavily ionising charged particles, and so proton tracks left in a TPC are shorter and deposit more energy than their hadronic and leptonic counterparts. The energy threshold for efficiently reconstructing a proton track is a kinetic energy of ≈ 20 MeV [45]. Since neutrons are neutral particles, they cannot be directly detected in the TPC.

4.2.3 Electromagnetic showers

Although electromagnetic showers are cascades of both photons and e^\pm particles, it is possible to identify the type of the initiating particle by studying the dE/dx at the start of the shower.

For photons, the first step in showering is to convert into an electron-positron pair. During the formation period for these particles, before they are able to initiate a cascade, they both behave as minimum ionising particles (MIPs). The energy loss at the start of the shower will be consistent with two MIPs. By contrast, for EM showers initiated by an electron or positron the energy loss at the start of the shower is consistent with a single MIP. Electrons and photons can be distinguished by measuring dE/dx for the first few cm of the shower, and determining whether it resembles one or two MIPs.

Automatically reconstructing EM showers in LArTPCs is not a trivial problem. Unlike tracks, which are easy to identify from a pattern recognition standpoint, showers are more complex. Although approximately conical in shape, defining this cone can be challenging, and the decision of whether to include hits at the edge of the cone can have a large effect on shower direction.

Many techniques are in development to solve this problem — for instance, one approach applies Gaussian blurring to hits in order to emphasise the broadly conical shape of a shower [46], while another applies a principle components analysis to clustered hits in order to achieve a similar goal. These tools, and shower reconstruction as a whole, are still under development [47].

4.2.4 Reconstructing π^0 mesons

At present, the greatest challenge for identifying $n - \bar{n}$ events using traditional reconstruction methods is the identification of $\pi^0 \rightarrow \gamma\gamma$ decays. To correctly identify and reconstruct the π^0 , we must first identify a pair of electromagnetic showers in the event, use their invariant mass to confirm that the two showers did originate from a π^0 and calculate the pion energy, using

$$E_{\pi^0} = \sqrt{2 E_1 E_2 (1 - \cos \theta_{12})} , \quad (4.3)$$

where E_1 and E_2 are the energies of the two showers, and θ_{12} is the angle between their directions.

The efficiency with which π^0 decays are reconstructed in a LArTPC strongly depends on the efficiency with which single showers are identified, and their calorimetry, direction and spatial resolution determined. The angular resolution of showers is of particular importance: mis-identifying the angle of an electromagnetic shower by even a few degrees has a large impact on the ability to correctly tag π^0 particles.

A study to reconstruct π^0 mesons from electromagnetic showers in the ICARUS T600 detector [48] found the π^0 mass resolution to be 27.3% (improved to 16.1% by limiting the event sample to cleaner events). However, this analysis was performed using only 212 events, and crucially, it was not carried out using automated reconstruction — for each event analysed, the hits were clustered into separate showers manually by analysers.

Since this approach is statistics-limiting, the ICARUS π^0 mass resolution is not representative of a generalised LArTPC's ability to reconstruct π^0 mesons. Studies in MicroBooNE using MC reco (as defined in Section 7.2.2) indicate π^0 mesons can be reconstructed with a purity (proportion of reconstructed π^0 s that correspond to a true π^0) of 82% and an efficiency (fraction of true π^0 s that were correctly reconstructed) of 61.2% in the context of a neutral current single- π^0 cross-section measurement [49].

4.3 LArTPC detectors

4.3.1 History

Argon-based detectors were first proposed by Carlo Rubbia in 1964 [50], and also by Willis & Radeka later in the same year [51], as a method of total ionisation particle detection. They serve as a modern equivalent to bubble chamber technology invented in the 1950s; the electronic instrumentation of a LArTPC provides the opportunity for automated reconstruction and event selection, greatly streamlining the process of data collection and analysis.

The first LArTPC to be constructed was the 3 ton ICARUS prototype,

commissioned in 1991 [52]. Following its success, a much larger version, the 600 ton ICARUS T600 was commissioned in 2001 [53]. The first US-based LArTPC was ArgoNeuT, first commissioned in 2008 [54].

4.3.2 The MicroBooNE detector




Figure 4.5: A schematic diagram demonstrating the cylindrical MicroBooNE cryostat and the rectangular field cage within. The TPC is approximately 10 metres long, and approximately 2.5 metres tall and wide (from [35]).

MicroBooNE is an on-surface 87 ton active volume LArTPC based at Fermilab in the United States. Its primary physics goals are to investigate the MiniBooNE low energy excess [55] and perform neutrino-argon cross-section measurements in the ~ 1 GeV range, and it also serves as R&D for the new generation of LArTPC experiments.




Figure 4.6: A cross-sectional diagram showing the internal configuration of the MicroBooNE detector. The x dimension is defined parallel to the electron drift direction (from [35]).

4.3.2.1 Time projection chamber

The MicroBooNE TPC consists of a single cathode plane and a single APA, with physical dimensions $10.4 \times 2.7 \times 2.3 \text{ m}^3$; a schematic diagram is shown in Figures 4.5 and 4.6. The cathode consists of 9 stainless steel sheets, supported by a frame of cylindrical and square tubes on the outer edge. The APA contains three wire planes: two induction planes of 2400 wires each, at an angle of $\pm 60^\circ$ to the vertical, and a collection plane of 3456 vertical wires, with a total of 8256 wires. Wires in the diagonal U and V planes are shorter in length at either end of the TPC, due to the rectangular shape of the APA.

The operational drift voltage in MicroBooNE is 275 V/cm, with a nominal overall drift high voltage of -70 kV on the cathode. The U, V and Y wire planes are biased to -200 V, 0 V and +400 V respectively.

4.3.2.2 Light collection




Figure 4.7: A schematic of the light collection system in MicroBooNE, showing both circular PMTs and light guide paddle PMTs (from [35]).

Light in MicroBooNE is collected using 32 photomultiplier tubes, each coated with a wavelength-shifting plate that shifts 128 nm vacuum ultraviolet scintillation light into the visible spectrum with a wavelength of $425 \pm 20 \text{ nm}$. In addition, four light-guide PMTs were installed in MicroBooNE in order to perform R&D for future LArTPCs. The locations of these PMTs within the cryostat are shown in Figure 4.7.

4.3.2.3 Purity monitoring

The purity of argon in MicroBooNE is measured using double-gridded ion chambers [56]. Electrons are produced at a cathode via the photoelectric effect using a UV laser, and drift to an anode where they are collected. The purity of the liquid argon is measured by taking the ratio of the charge liberated at the cathode to the charge collected at the anode,

$$Q_A/Q_C = e^{-t/\tau} , \quad (4.4)$$

where τ is the electron lifetime in liquid argon.

4.3.3 The DUNE far detector




Figure 4.8: The full DUNE far detector with four 10kt underground modules, in their expected configuration at SURF in Lead, SD (from [57]).

In contrast with MicroBooNE, the DUNE far detector is still in its design phase, and has not yet been constructed. It will consist of four 10 kt modules, at a depth of approximately 1.5 km underground at the Sanford Underground Research Facility in Lead, SD [9].

Two reference designs are currently being considered: a single phase design that shares high-level design concepts with existing LArTPCs like Mi-




Figure 4.9: The planned layout for alternating anode and cathode planes in a single DUNE module (from [57]). The anode wire planes on each side measures a single drift region, while the central anode wire plane measures signals from drift regions on both sides.

croBooNE, and a dual phase design that proposes several modifications to the current LArTPC design standard [57]. The current nominal DUNE design assumes the DUNE 40 kt fiducial mass will consist of two 10 kt single phase modules and two 10 kt dual phase modules. The work described in this thesis assumes a single-phase geometry for the DUNE far detector, and the following sections describe the current nominal single phase geometry, which is used to produce simulations for the studies in Chapter 7.

4.3.3.1 Time projection chamber

The external dimensions of a DUNE single-phase 10 kt module will be 12 m high, 14.5 m wide and 58 m long; the layout of the four modules is shown in Figure 4.8. Due to limitations in the size of individual drift volumes and APAs, each 10kt module will in turn consist of many TPC modules. Although these modules are often referred to as 10 kt modules, due to the 10 kt fiducial mass used for neutrino analyses, the active mass of each module is ≈ 12.1 kt [57]. Nucleon decay searches may be able to make use of this larger mass to improve sensitivity.

Each drift volume is approximately 3.5 m in distance between anode and cathode; each module will alternate between anode and cathode as demonstrated in Figure 4.9, with a total of three APAs and four readout volumes across the detector's width. Two modules will be stacked vertically, and the entire structure will be repeated 25 times along the beam direction, resulting in 200 total drift volumes and 150 wire plane readout cells.

The DUNE far detector will have three wire planes; the two induction planes will be angled $\pm 35.7^\circ$ from the vertical, with a final collection plane oriented vertically. Unlike MicroBooNE, in which the diagonal wire planes have short wires at each edge, DUNE's diagonal wires will wrap around at the edge of the APA, and so each induction wire will be instrumented on both sides of the collection plane. This allows the central APA in the detector to measure both drift volumes on either side using effectively fewer channels. For interactions that cross the anode plane this introduces additional ambiguities, as electrons will drift and be collected on the anode wires from both directions.

4.3.3.2 Light collection

The design of the DUNE photon detection system has not yet been finalised. According to the current reference design [57], the photon detection system will be mounted as modules onto the APA frames, with a light-guide and 12 silicon photomultipliers. Each APA will have ten such modules attached at regular intervals.

5 GENIE Event Generator

In order to determine DUNE’s ability to measure $n - \bar{n}$ oscillation, it is vital to have access to simulations of the process. Due to the lack of any readily available simulations, an event generator was developed, as part of this thesis work, as a standalone application within the GENIE simulation framework.

This chapter discusses the development of this event generator to simulate the experimental signature of nucleus-bound $n - \bar{n}$ oscillation. It summarises techniques used to simulate the process and then presents some truth-level topological studies, demonstrating the effects that simulation of various nuclear effects have on the final state. It also describes some approximations made, and suggests potential extensions.

5.1 Simulation methods

The development of an $n - \bar{n}$ event generator poses a specific set of challenges; since it is a beyond Standard Model process that has never been observed, there is no data that can be used as a reference. Instead, we must use available data that provides the closest approximation, and use physics simulations to adapt this data as accurately as possible to a signal simulation.

The full $n - \bar{n}$ interaction can be summarised as a series of stages:

- A neutron bound inside the nucleus, with Fermi momentum and binding energy, spontaneously oscillates into an antineutron.
- The antineutron annihilates with another nucleon, which also has Fermi momentum and binding energy.

- The products of this decay are produced inside the nucleus, according to some prescribed branching ratios.
- These decay products propagate through the nucleus — reinteracting with nucleons as they travel, or decaying — until they escape the nucleus.

To accurately simulate this entire process, several different aspects of nuclear modelling must be simulated. Since simulating these processes from first principles would be time-consuming and needlessly complex, the event generator framework GENIE is used, which provides the necessary physics utilities.

GENIE (Generates Events for Neutrino Interaction Experiments) [58]) is primarily a neutrino interaction event generator. As such, its primary function is to simulate the interaction between a neutrino and a nucleon — accounting for neutrino flux, neutrino-nucleon interaction cross-sections, final state branching ratios, intranuclear kinematics, and the intranuclear transport of interaction by-products. Several of these utilities can be adapted to be used to simulate $n - \bar{n}$ oscillation, namely the nucleon kinematics and intranuclear transport functions. This overlap extends beyond just $n - \bar{n}$ oscillations: a standalone GENIE application for nucleon decay simulation existed prior to the development of this $n - \bar{n}$ event generator.

The $n - \bar{n}$ oscillation generator was developed within a private repository of GENIE v2.9.0, and in the later stages of development became a GENIE incubator project. Upon completion, the generator was officially released in GENIE v2.12.0 [59].

GENIE simulations use an object called the Event Record (`genie::EventRecord`) to track the various stages of the interaction. This event record is a list of particles, with information on each particle's type, energy, momentum and position, and at which stage of the interaction it existed. The following few sections step through the simulation process in sequence, beginning with the initial state nucleus prior to $n - \bar{n}$ oscillation and ending with final state particles leaving the nucleus.

5.1.1 The initial state

The event generator is designed to take the initial state element and isotope as a user input. This is specified using the Particle Data Group nuclear code [60], expressed in the form $\pm 10LZZZAAAI$, where L is the total number of strange quarks, ZZZ the total number of protons, AAA the total baryon number, and I the isomer level (where $I = 0$ corresponds to the ground state). For instance, for an unexcited ^{40}Ar nucleus, the nuclear PDG code is 1000180400.

A restriction is placed on which nuclei can be provided as an initial state, since it must be included in GENIE's `PDGLibrary`. Since the event generator makes use of GENIE's nuclear utilities, the user is limited to those nuclei for which GENIE records this information. Once a valid nucleus is specified, this nucleus is added to the GENIE event record as the initial state (`kIStInitialState`), completely at rest.

5.1.2 Simulating the oscillating neutron

Once the initial state nucleus has been selected, the next step is to simulate individual nucleons inside this nucleus. GENIE does not independently simulate all nucleons inside the nucleus, and so an approximation must be made. Calling GENIE's function to provide a nucleon's Fermi momentum and binding energy does not account for any nucleons previously simulated. Since this event generator explicitly simulates only two nucleons — the oscillating neutron and the nucleon it annihilates with — this lack of correlation is not an issue for medium and large sized nuclei. However, for very small nuclei this approximation is insufficient, as it can introduce inconsistencies in momentum conservation between individual nucleons and the nucleus as a whole.

The oscillating neutron's position inside the nucleus is selected at random based on the density profile of nucleons within the nucleus, using GENIE's `utils::nuclear::Density(r,A)` function. For nuclei with a baryon number of 20 or greater, GENIE models the nuclear density with a Woods-Saxon distribution [61],

$$\rho(r) = \frac{\rho_0}{1 + e^{(r-R_0)/a}} , \quad (5.1)$$

where r is depth inside the nucleus; $R_0 = r_0 A^{\frac{1}{3}}$ is the nuclear radius, with $r_0 = 1.4$ fm in GENIE; ρ_0 is usually nuclear density at $r = 0$, but in GENIE is replaced with a normalisation constant to express nuclear density as a probability distribution; and a is a distance describing the “surface thickness” of the nucleus, set to $a = 0.54$ fm for ^{40}Ar in GENIE.

Similarly, nucleon Fermi momentum and binding energy are provided by the `genie::NuclearModelI` class, an interface to whichever nuclear model the user has enabled in GENIE’s configurable user physics options. By default, the Bodek-Ritchie Fermi gas model [62] is enabled in GENIE v2.12.0 — although other models are available, such as the local Fermi gas model and the effective spectral function model [63]. The Bodek-Ritchie model is used for all work discussed in this thesis.

The neutron that undergoes oscillation into an antineutron is added to the GENIE event record as part of the stable decayed state (`kIStDecayedState`). By this stage it has already oscillated into an antineutron, but due to GENIE bookkeeping it must be added to the event record as a neutron. This has no effect on later stages of simulation.

5.1.3 Simulating the annihilating nucleon

Once the oscillating neutron has been selected, the generator then considers the nucleon with which the antineutron annihilates. The annihilation nucleon is selected to be either a proton or a neutron, based on the proton-to-neutron ratio of the initial state nucleus omitting the neutron that has oscillated. For example, in ^{40}Ar there is an 18/39 chance of annihilating with a proton and a 21/39 chance of annihilating with a neutron.

The position of this second nucleon is assigned to be identical to the oscillating neutron’s position. This is a simplifying approximation — a more rigorous approach would be to simulate all nucleons in the nucleus according to the nuclear density particle and select the closest candidate, then select the annihilation vertex as a point equidistant between the two. However, since

this would yield an annihilation vertex distribution similar to the oscillating neutron's position, the position of the second nucleon is approximated to be identical to that of the oscillating neutron.

As with the oscillating neutron, the annihilation nucleon's removal energy and Fermi momentum are simulated using `genie::NuclearModelI`, and the particle is added to the event record as part of the stable decayed state (`kIStDecayedState`).

5.1.4 Simulating the remnant nucleus

With the initial nucleus and two interacting nucleons simulated, the next step is to simulate the remnant nucleus. This remnant nucleus is assigned the nuclear PDG code of the initial nucleus with the two interacting nucleons subtracted, and its momentum and energy are chosen to conserve energy and momentum with respect to the two annihilating nucleons. The remnant nucleus is added to the event record as part of the stable final state (`kIStStableFinalState`).

5.1.5 Simulating annihilation products

With the initial state nucleus, annihilating nucleons and remnant nucleus simulated, the next step is to simulate annihilation products. The annihilation vertex is approximated as the position of both annihilating nucleons.

A Monte Carlo method is used to select a final state based on the branching ratios shown in Table 3.2, which originate from a recent Super-Kamiokande $n - \bar{n}$ search [1]. A particle object is then created for each of the annihilation products listed in the chosen final state, with the corresponding PDG code assigned.

The net momentum of the two-nucleon system is calculated, and then the two annihilating nucleons are Lorentz boosted into this system's rest frame using ROOT's `TLorentzVector::Boost` function. The total available energy of the two particles in this frame is calculated, and then the ROOT `TGenPhaseSpace` class is used to distribute this energy to the annihilation products. This class performs a phase space decay, assigning each anni-

hilation product with energy and momentum based on the total available centre-of-mass energy. Finally, these decay particles are Lorentz boosted back into the lab frame, and added to the event record as hadrons inside the nucleus (`kIStHadronInTheNucleus`).

5.1.6 Final state interactions

Before hadrons produced during the nucleon-antineutron annihilation can be detected, they must escape the nucleus in which $n - \bar{n}$ oscillation occurred. For heavy nuclei such as ^{40}Ar , hadrons can travel up to 8 fm before escaping the nucleus [61], and so the probability of reinteracting is high. Final state interactions (FSI) have a large impact on the final state, and so simulating this step is vital.

The transport of final state hadrons out of the nucleus is handled by a subpackage of GENIE called INTRANUKE [58]. The branching fractions and reinteraction probabilities in this model are tuned to bubble chamber data on hydrogen and deuterium targets. It uses the free cross section to estimate the likelihood of reinteraction, defining the mean free path as

$$\lambda(E, r) = \frac{1}{\sigma_{hN, tot} \times \rho(r)} , \quad (5.2)$$

where $\sigma_{hN, tot}$ is the cross section and $\rho(r)$ the nuclear density; in the model used for this work, cross-sections are tuned to data primarily from bubble chamber data on hydrogen and deuterium targets. It also accounts for the formation time of hadrons, which manifests as a ‘free step’ at the start of a hadron’s lifetime, in which it will not interact.

The default model, and the model used for this work, is hA INTRANUKE [58]. Hadrons are propagated through the nucleus, until they either reinteract or escape the nucleus. If they reinteract, the module uses lookup tables from data to move immediately to simulated particles that exit the nucleus, unlike the full hN cascade model that simulates the interaction products and continues to propagate them through the nucleus. As a consequence, the hA module is less computationally intensive, and is less sensitive to uncertainties in the mean free path of pions in the nucleus.

The INTRANUKE module selects hadrons generated in the nucleus, by identifying all particles in the event record with status `kIStHadronInTheNucleus`. It then performs hadron transport on all of these particles, simulates the particles that exit the nucleus, and adds these particles to the event record with the status (`kIStStableFinalState`).

5.2 Simulation results

Figure 5.1 shows an example event display of an $n - \bar{n}$ event simulated by the GENIE event generator. This event was produced in GENIE and then run through subsequent MicroBooNE simulation stages to produce this image.

In order to determine the impact of nuclear effects on $n - \bar{n}$ final states, several Monte Carlo samples were produced in varying configurations. The $n - \bar{n}$ event generator makes use of three nuclear utility functions provided by GENIE:

- Binding energy
- Fermi momentum of the annihilating nucleons
- FSI that occur as hadrons exit the nucleus

Figures 5.2, 5.3, 5.4 and 5.5 demonstrate the effects of nuclear simulations in GENIE. In each of these figures, a distribution produced using default GENIE nuclear effects is compared to the same distribution with components of the nuclear simulation disabled, as follows:

- FSI disabled, binding energy and Fermi momentum enabled
- FSI and binding energy disabled, Fermi momentum enabled
- FSI and Fermi momentum disabled, binding energy enabled
- FSI, Fermi momentum and binding energy disabled
- (FSI enabled, binding energy and Fermi momentum disabled)







Figure 5.1: An example event display of a simulated $n - \bar{n}$ event in the MicroBooNE detector. The event was simulated using the GENIE event generator, and then propagated through the standard MicroBooNE simulation chain. The event is $n\bar{n} \rightarrow \pi^+\pi^-3\pi^0$; the π^+ is absorbed during final state interactions, and the event contains a short proton track ejected during FSI. Six electromagnetic showers emerge from the central vertex, as well as a long track produced by a π^- and a short track produced by a proton. The colour scale represents charge deposited on wires, ranging from blue for low charge deposition to red for high charge deposition.




(a) Binding energy and Fermi momentum enabled, FSI disabled.




(b) Binding energy enabled, Fermi momentum and FSI disabled.



(c) Fermi momentum enabled, binding energy and FSI disabled.




(d) Binding energy, Fermi momentum and FSI disabled.




(e) Binding energy and Fermi momentum disabled, FSI enabled.


Figure 5.2: Event pion multiplicity for various configurations of nuclear simulation. The nominal distribution, with all nuclear effects enabled, is provided for reference in red in all subfigures. Since the binding energy and Fermi momentum of initial state nucleons do not affect multiplicity of pions in the final state, all distributions except the nominal are identical except for statistical fluctuations.




(a) Binding energy and Fermi momentum enabled, FSI disabled.




(b) Binding energy enabled, Fermi momentum and FSI disabled.



(c) Fermi momentum enabled, binding energy and FSI disabled.




(d) Binding energy, Fermi momentum and FSI disabled.




(e) Binding energy and Fermi momentum disabled, FSI enabled.


Figure 5.3: Momentum of individual final state pions for various configurations of nuclear simulation. The nominal distribution, with all nuclear effects enabled, is provided for reference in red in all subfigures.




(a) Binding energy and Fermi momentum enabled, FSI disabled.




(b) Binding energy enabled, Fermi momentum and FSI disabled.



(c) Fermi momentum enabled, binding energy and FSI disabled.




(d) Binding energy, Fermi momentum and FSI disabled.




(e) Binding energy and Fermi momentum disabled, FSI enabled.


Figure 5.4: Event invariant mass of particles exiting the nucleus for various configurations of nuclear simulation. The nominal distribution, with all nuclear effects enabled, is provided for reference in red in all subfigures.




(a) Binding energy and Fermi momentum enabled, FSI disabled.




(b) Binding energy enabled, Fermi momentum and FSI disabled.



(c) Fermi momentum enabled, binding energy and FSI disabled.



(d) Binding energy, Fermi momentum and FSI disabled.



(e) Binding energy and Fermi momentum disabled, FSI enabled.

Figure 5.5: Absolute magnitude of momentum vector sum all particles exiting the nucleus for various configurations of nuclear simulation. The nominal distribution, with all nuclear effects enabled, is provided for reference in red in all subfigures.

Pion multiplicity in the final state (Figure 5.2) is heavily influenced by FSI, and is not affected at all by Fermi momentum or binding energy. The individual pion momentum distribution (Figure 5.3) is smeared towards lower momenta by FSI, and shifted to the left by binding energy. A small peak around 900 MeV in momentum is visible when Fermi momentum is disabled, corresponding to the two-body final states in which each pion inherits half of the invariant mass from the annihilation.

The invariant mass is calculated as

$$M = \sqrt{\sum E_i^2 - |\sum \vec{p}_i|^2}, \quad (5.3)$$

where the i index refers to pions in the final state, each of which has energy E_i and momentum \vec{p}_i .

FSI smear the invariant mass (Figure 5.4) down from ~ 2 GeV to lower energies. A small peak is visible around the mass of a single pion, corresponding to events where all but one pions are absorbed during FSI. As with pion momentum, binding energy introduces a characteristic shift to lower invariant mass, and Fermi momentum smears the distribution by ~ 100 MeV.

The total event net momentum is defined as the absolute magnitude of a vector sum of all final state pions. The total event net momentum distribution (Figure 5.5) is smeared to higher momenta by both FSI and Fermi momentum; the impact of Fermi momentum is more apparent in this distribution than in individual pion momentum or invariant mass. Net momentum is not affected by the simulation of binding energy.

5.2.1 Super-Kamiokande comparison

The $n - \bar{n}$ analysis published by the Super-Kamiokande collaboration [1] describes the tools used to simulate the process. For the signal simulation, they use an event generator designed by the IMB collaboration to search for $n - \bar{n}$ oscillation in Oxygen. The pion multiplicities and momenta produced by the GENIE event generator were compared to the Super-Kamiokande simulation, as shown in Table 5.1, as were the branching ratios for various FSI processes, as shown in Table 5.2.

GENIE predicts a much lower final state pion multiplicity than Super-Kamiokande’s event generator in ^{16}O (Table 5.1), and as a result the average pion momentum is larger. This is due to GENIE’s simulation of final state interactions predicting a higher probability of reinteraction. As demonstrated in Table 5.2, GENIE predicts a 34% probability of a final state particle escaping the nucleus without reinteracting, compared to 49% in Super-Kamiokande.

Table 5.1: Comparison of pion multiplicities and momenta between between Super-Kamiokande simulations in ^{16}O and GENIE $n - \bar{n}$ simulations in both ^{16}O and ^{40}Ar . The multiplicities and momenta quoted for Super-Kamiokande come from [1]. A 100,000 event ^{16}O sample was generated using the GENIE event generator, to allow direct comparison with Super-Kamiokande values. Also provided are the same values for a 1,000,000 event GENIE sample generated in ^{40}Ar .

	Super-K.	GENIE (^{16}O)	GENIE (^{40}Ar)
π multiplicity	3.5	2.37	2.94
π^\pm multiplicity	2.2	1.57	1.96
π^\pm mean mom. [MeV]	310	372	344
π^\pm RMS mom. [MeV]	190	190	190

Table 5.2: Comparison of branching fractions for various FSI processes experienced by intermediate state pions between Super-Kamiokande and GENIE $n - \bar{n}$ event generators. The branching ratios quoted for Super-Kamiokande come from [1], while the corresponding values for GENIE come from the same ^{16}O and ^{40}Ar samples described in Table 5.1.

	Super-K.	GENIE (^{16}O)	GENIE (^{40}Ar)
No FSI	49%	34.0%	15.6%
Absorption	24%	18.8%	24.0%
Nucleon interaction	3%	4.2%	5.3%
Scattering	24%	43.1%	55.1%

5.3 Discussion

5.3.1 Branching ratio corrections

The annihilation branching ratios *effectively* used within the GENIE event generator are not identical to the branching ratios published by previous $n - \bar{n}$ searches, shown in Table 3.2. These ratios are not explicitly modified, but it is sometimes necessary to reselect a final state during simulation.

Reselecting the final state is necessary due to event kinematics in one final state specifically — $\bar{n}p \rightarrow 2\pi^+\pi^-2\omega$. The branching ratios, which are derived from the annihilation of ‘at-rest’ antiprotons on deuterium targets, claim this final state occurs in 16% of $\bar{n}p$ annihilations. However, the total mass energy of this final state is 1984 MeV, 106 MeV greater than the mass energy in the initial state. For the ^{40}Ar nucleus, it is highly unlikely that the centre-of-mass energy available due to Fermi momentum of the two nucleons is large enough to account for both the binding energy and this mass difference.

When there is not enough available energy to perform the annihilation into this final state, the generator will select a different final state, using the same branching ratios. As a consequence, this final state is very heavily suppressed — the branching ratio falls from 16% to 0.003%, while the branching ratios for all other final states are scaled up proportionally. The effective branching ratios, including dynamic modifications due to event kinematics, are presented in Table 5.3.

This discrepancy occurs because ‘at-rest’ antiproton annihilation data is measured using low-energy antiprotons, which are not truly at rest. As discussed in Section 3.3.1, the data sets from which branching ratios are derived do not explicitly state their definition of ‘at-rest’, though in the equivalent data sets for $\bar{p}n$ annihilation data, ‘at-rest’ is defined by antiproton momenta as high as 250 MeV, sufficient to produce the final state in question.

5.3.2 Validating the phase space approximation

As discussed in Section 5.1.5, the process of dividing out energy produced in the annihilation between final state particles is performed via a phase

Table 5.3: Neutron-antineutron oscillation final state branching ratios generated by the GENIE event generator, including the dynamic scaling-down of final states suppressed by event kinematics.

$\bar{n} + p$		$\bar{n} + n$	
$\pi^+\pi^0$	1.2%	$\pi^+\pi^-$	2.0%
$\pi^+2\pi^0$	9.5%	$2\pi^0$	1.5%
$\pi^+3\pi^0$	11.9%	$\pi^+\pi^-\pi^0$	6.5%
$2\pi^+\pi^-\pi^0$	26.2%	$\pi^+\pi^-2\pi^0$	11.0%
$2\pi^+\pi^-2\pi^0$	42.8%	$\pi^+\pi^-3\pi^0$	28.0%
$2\pi^+\pi^-2\omega$	0.003%	$2\pi^+2\pi^-$	7.1%
$3\pi^+2\pi^-\pi^0$	8.4%	$2\pi^+2\pi^-\pi^0$	24.0%
		$\pi^+\pi^-\omega$	10.0%
		$2\pi^+2\pi^-2\pi^0$	10.0%

space approximation. This method does not model the underlying physics of each final state, such as correlations in energy and direction of final state products caused by intermediate states. Due to the high probability that final state particles will reinteract before exiting the nucleus, resulting in a change in energy and direction, absorption, and/or the production of new particles, most small modifications to final state particle kinematics caused by detailed modelling of intermediate states immediately after annihilation will be lost by the time these particles exit the nucleus. Based on this assumption, the phase space decay is taken to be an adequate method for distributing energy and momentum among final state particles.

A study was performed to test the validity of this assumption, using the final state $\bar{n}n \rightarrow \pi^+\pi^-\pi^0$. Events were generated in GENIE without final state interactions enabled, in order to compare directly with antiproton annihilation data on deuterium. Although the GENIE generator performs a phase space decay immediately from the two nucleons to the three-pion final state, a review of antiproton annihilation data (from which annihilation branching ratios are derived) instead claims that 55% of annihilations occur via the intermediate state $\bar{n}n \rightarrow \pi\rho \rightarrow 3\pi$. The ‘dipion mass’ (invariant mass of each two-pion combination) shows a peak at 770 MeV, the mass of the ρ meson, which is not modelled at all by the phase space approximation in GENIE. A comparison is shown in Figure 5.6.

Modifications were made to the GENIE event generator to make the dipion mass distribution more closely resemble the data. In accordance with antiproton scattering data [31], a phase space decay from the initial to final state is performed 45% of the time. In the remainder of events, a phase space decay is performed into the intermediate $\pi\rho$ final state, followed by a second phase space decay when the ρ decays into two pions. Gaussian smearing is applied to the modified GENIE distribution to account for detector uncertainties, and is compared directly with data in Figure 5.6.




Figure 5.6: Dipion mass distributions for the $\bar{n}n \rightarrow \pi^+\pi^-\pi^0$ final state. The green distribution in both subfigures comes from antiproton-deutrium annihilation data [31]. The blue distribution (a) is GENIE with a simple phase space approximation, while the red distribution (b) is a modified version of GENIE that accounts for the intermediate state with a ρ meson, with additional Gaussian smearing applied to pion momentum to account for detector uncertainties in data.

Event samples from both the original and modified GENIE generators are generated, and the dipion mass and pion momentum distributions of final state particles are compared both before and after final state interactions. The distributions before FSI are displayed in Figure 5.7, and the corresponding distributions after FSI are displayed in Figure 5.8.

As shown in Figures 5.7 and 5.8, the considerable difference in dipion mass distributions before FSI is drastically reduced with the addition of FSI. The mass peak around 770 MeV, which appears as a large δ function in the

dipion mass distribution in the modified generator before FSI, is reduced to a small excess after FSI. The pion momentum distributions are very similar both before and after FSI, since the intermediate state has a greater effect on correlated pairs of pions than on any individual pion. Based on this study, the phase space approximation without any modelling of intermediate states is taken to be sufficiently accurate for ^{40}Ar , given that any fine structure modelled in intermediate states is lost in FSI regardless.




Figure 5.7: The dipion mass distribution (a) and pion momentum distribution (b) of both the original (single phase space decay) and modified (phase space decay with intermediate state) GENIE event generators, before final state particles have undergone final state interactions.

5.4 Conclusions

The GENIE neutron-antineutron oscillation module provides a simple, robust method of simulating bound neutron-antineutron oscillation. The simulation of this process is vital for any searches for neutron-antineutron oscillation inside a relatively heavy nucleus, and this module is the first open-source, publicly available $n - \bar{n}$ event generator. Due to the flexibility of the generator, it facilitates not only the search for $n - \bar{n}$ in DUNE, but in any mid-to-heavy nucleus. Detailed modelling of intermediate states is shown to be unnecessary, at least for heavier nuclei, as all resulting features are washed out by final state interactions.




Figure 5.8: The dipion mass distribution (a) and pion momentum distribution (b) of both the original (single phase space decay) and modified (phase space decay with intermediate state) GENIE event generators, after the final state particles have propagated out of the nucleus and undergone final state interactions.

Blank page

6 Convolutional Neural Networks

Convolutional neural networks (CNNs) are an image processing tool that perform successive layers of convolutions on images in order to identify features and associate these features with labels. The LeNet CNN architecture was devised in 1998 as a tool for identifying handwritten characters [64], but the potential applications are far more widespread. Recent advances in graphical processing greatly improve the speed at which CNNs train and classify images, leading to their widespread adoption as a classification technique.

Since LArTPCs naturally produce high-resolution images as their output, CNNs are a natural fit as an event classification and analysis tool. Recent work by the MicroBooNE collaboration has demonstrated that CNNs are effective both at single-particle identification and at neutrino event identification [5], and provides analysis tools for performing CNN-based studies in LArTPCs.

This chapter describes the operating principles of CNNs, and summarises specific tools used. This serves as prelude to the following chapters, which describe the use of CNNs within the context of $n - \bar{n}$ oscillation searches in LArTPCs.

6.1 CNN operating principles

Convolutional neural networks are a form of artificial neural network (ANN), computational systems designed to mimic the connectivity of the human brain. These systems are comprised of many individual, highly-connected




Figure 6.1: A diagram of the basic architecture of a Feed-forward Neural Network (from [65]). Neurons in the input layer take input parameters and pass them to neurons in the hidden layer. These hidden neurons produce linear combinations of the input values by applying learned weights to each parameter, and passes these values to an output neuron which uses them to make a decision.

nodes (“neurons”) working together to process some input to produce a desired output [65].

The simplest form of ANN is the “feed-forward neural network” (FNN), a very basic example of which is shown in Figure 6.1. Inputs to the network are used to initialise a layer of input neurons, connected in turn to a layer of hidden neurons. Each of these hidden neurons applies learned weights to inputs from previous layers to produce an output value. Depending on the depth of the network, this hidden layer may be connected to one or more subsequent hidden layers. The final hidden layer is in turn connected to an output layer containing one or more neurons, which use the output of previous connected layers to make a decision.

Such networks are trained by varying these weights inside the hidden layers. The network will begin with random values for these parameters, and then assess the corresponding output and iterate, optimising towards the most desirable outcome. In supervised learning, the network associates each input with a given label, and assesses its own performance by its ability to infer this label from the output.

CNNs are a subcategory of ANNs designed for performing image pattern recognition. Since any moderately-sized image is comprised of an array of pixels with high dimensionality, a more traditional ANN approach quickly becomes impractical due to the number of neurons required. CNN architecture is specifically designed to deal with this issue.

6.2 Network layer types

The basic structure of a CNN is as follows (all layer types are described in more detail below):

- **Input:** The input layer consists of an array containing pixelwise image information.
- **Convolution:** Convolution layers convolve the input image in order to identify features.

- **Rectified linear unit (ReLU):** ReLU layers apply an elementwise activation function to the previous layer, and are typically used immediately after convolution layers.
- **Pooling:** Pooling layers perform downsampling on an activation map in order to manage memory overhead.
- **Fully connected:** Fully connected layers use activation maps from previous layers to learn non-linear combinations from previous layers.
- **Output:** Output layers compute a loss parameter or assign scores to an image.

In a typical network architecture, one or several convolution layers are performed in succession, followed by a ReLU layer and a pooling layer. This pattern of convolutions followed by pooling is repeated several times, followed by one or more fully connected layers, and then output layers. A brief summary of the function of each of these layers is provided below.

6.2.1 Convolution layer




Figure 6.2: The calculation performed in a CNN convolution layer. A single point on the output feature map is calculated by taking a small local region of the image, multiplying each element by a filter designed to pick out a specific image feature, and then taking a sum over all elements.

A convolution layer convolves a local region of the input image with a randomly generated **filter**. This filter is a rectangular multiplication matrix, each element a weight that is varied and optimised during network training. This filter is swept over small subsamples of the input image, multiplying each local region by the filter to produce a single output value. The **filter size** determines the size of this multiplication matrix, while the **stride** determines how many pixels the filter steps over between applications. Each convolution of the filter with a local image region produces a single output value, and the repeated application of this filter leads to the production of an output image, with each pixel value corresponding to an application of the convolution filter. This output image is known as an **activation map**; an example is shown in Figure 6.2.

Each convolution layer applies many different filters to produce an array of activation maps, each associated with a different feature of the input image. For subsequent convolution layers, each activation map is produced by applying the filter simultaneously to all activation maps from the previous layer.

Pixel values in activation maps for a given layer are calculated as

$$y_i = \theta_i x_i + b_i , \quad (6.1)$$

where i is used as an overall index for the layer, taking into account each image label k , the image depth d , the number of filters l , and the number of parameters associated with each filter n . x_i is the pixel value for a local region of the input image, and θ_i and b_i are, respectively, the weight and bias parameters which define the filters applied in a given convolution layer [66].

6.2.2 ReLU layer

Rectified Linear Unit (ReLU) layers modify an activation map by performing $\max(0, a)$ on each pixel a . Negative weights and biases in convolution filters can lead to negative values in activation maps, which can interfere destructively with positive activations in subsequent layers. For this reason, all negative values are instead set to zero, achieving so-called **sparse repre-**

sentations. This also assists in preventing the vanishing of weight gradients when training with a gradient descent method, described in Section 6.3.4 [67].

6.2.3 Pooling layer




Figure 6.3: The downsampling performed in a CNN pooling layer. In this example, the “max pool” approach is taken for a 2×2 filter size, in which the largest pixel value of each 2×2 square is taken as a single pixel value, downsampling the image by a factor 4.

In CNN terms, pooling refers to downsampling of an image array. When pooling is performed, some local region of the image is reduced into a single pixel value. All studies discussed in this thesis utilise a 2×2 “max pool” approach is taken, in which each pixel in the downsampled image is the maximum pixel value of each 2×2 square in the original image.

In the network architecture used for these studies, the number of activation maps increases at deeper network layers. In order to constrain the amount of memory required to perform these convolutions, pooling must be used to reduce the size of each individual activation map.

6.2.4 Fully connected layer

After all convolution layers have been performed, a fully connected layer is applied. This layer combines information from all activations in the previous layer by applying learned matrix multiplications and biases to them, in the same manner as a more traditional feed-forward neural network. It possesses a separate set of weights and biases for each input image label, allowing it

to independently learn features associated with different labels. Its number of outputs is equal to the number of image labels provided.

6.3 Network training

6.3.1 Solver

In Caffe (discussed further in Section 6.4.1), network learning is handled by a **solver**. The solver is responsible for optimising the network by improving a **loss** parameter (defined below in Section 6.3.3). This is achieved by assessing the network's current performance, and constructing gradients with respect to previous network iterations in order to choose network weights that minimise the loss.

The solver is an additional layer surrounding the network itself — it selects network weights, queries the network for the resulting loss and gradient, and then iteratively uses these to improve the network. It also produces periodic snapshots of the network, both as a **solverstate**, which is used to resume training, and as a **caffemodel**, which is used to construct a network for image classification.

6.3.2 Classification

Events are classified by the network using a **SoftmaxLayer** that expresses the network's response to each input label as a probability. This probability for each classification label k is written

$$p_k = \frac{e^{a_k}}{\sum_{k'}^K e^{a_{k'}}}, \quad (6.2)$$

where a_k are the outputs for each image label from the fully connected layer of the network, K is the total number of image classification labels, and k' is an additional index to indicate image label.

6.3.3 Loss

During learning, the network improves by minimising a parameter called the loss function, which abstracts the number of mistakes the network makes during image classification. This parameter is calculated using the Caffe layer `SoftmaxWithLossLayer`, and is derived from the softmax probability defined in Equation 6.2 as

$$J = -\frac{1}{K} \sum_k^K \log(p_k) , \quad (6.3)$$

where again k is an index indicating image label and K is the total number of image labels.

6.3.4 Gradient descent

The optimisation of network performance is achieved using a technique called **gradient descent**. As discussed in the previous section, the loss function, computed from network weights, is a scalar value abstracting mistakes made by the network, and must be minimised in order to optimise network performance. In order for the network to effectively learn, it must be able to use successive iterations to construct a loss gradient, and use this to infer a direction in which to improve. Gradient descent is performed by the network solver during training [68].

With each successive iteration of the network, denoted by an integer index t , the network parameters θ_i — *ie.* the weights and biases which comprise the filters in each convolution layer — are updated according to the expression

$$\theta_{t+1,i} = \theta_{t,i} - \eta \cdot \nabla_{\theta} J(\theta_{t,i}) , \quad (6.4)$$

where η is the configurable network learning rate (described in more detail in Section 6.3.5), and $\nabla_{\theta} J(\theta_{t,i})$ is the gradient of the loss with respect to the parameters of the network, more succinctly defined as

$$g_{t,i} = \nabla_{\theta_t} J(\theta_{t,i}) . \quad (6.5)$$

For each successive iteration, each network weight is modified according to its previous value and the gradient,

$$\theta_{t+1,i} = \theta_t - \eta \cdot g_{t,i} . \quad (6.6)$$

Various techniques for gradient descent exist — one such method is called **Adagrad**, in which the learning rate is allowed to vary for each weight,

$$\theta_{t+1,i} = \theta_{t,i} - \frac{\eta}{\sqrt{G_{t,ii} + \epsilon}} g_{t,i} . \quad (6.7)$$

Here $G_{t,ii}$ is a diagonal matrix, where each element is the sum of the squares of all gradients preceding the current iteration. This serves to reduce the learning rate for weights with a steep gradient, while allowing weights with a more shallow gradient to vary more freely. The ϵ term is a small smoothing term, included to avoid division by zero. From this point forward, θ will be considered to be a vector, and so i subscripts will be dropped.

The gradient descent technique **RMSProp** improves on **Adagrad** by replacing the gradient sum with an updating average, reducing memory overhead. The updating average is computed as

$$E[g^2]_{t,i} = \delta E[g^2]_{t-1,i} + (1 - \delta) g_{t,i}^2 , \quad (6.8)$$

where δ is a configurable parameter, the **rms_decay**, typically set to a value of 0.99. This updating average $E[g^2]_{t,i}$ is calculated using the average from the previous iteration $E[g^2]_{t-1,i}$ and the square of the gradients from the current iteration $g_{t,i}^2$. For ease, we now define the corresponding change in weight,

$$\Delta\theta_{t,i} = -\frac{\eta}{\sqrt{E[g^2]_{t,i} + \epsilon}} g_{t,i} , \quad (6.9)$$

allowing the network weights for each iteration to be expressed simply as

$$\theta_{t+1,i} = \theta_{t,i} + \Delta\theta_{t,i} . \quad (6.10)$$

6.3.5 Learning rate

The learning rate is a network parameter that describes how much the network weights are allowed to vary between successive iterations. Since training a network involves converging on a preferred region of parameter space, the learning rate is reduced as the number of iterations increases, taking smaller steps in parameter space as we approach the optimal network model.

In Caffe, several different types of **learning rate policy** are defined, which describe whether the learning rate is constant or changing over time. The policy used in these studies is **inv**, which defines the learning rate as

$$\alpha = \alpha_{\text{base}}(1 + \gamma n_{\text{iter}})^{-k}, \quad (6.11)$$

where α_{base} is the configurable base learning rate, n_{iter} is the number of training network iterations, and γ and k are additional configurable parameters. The learning rate is equal to the base learning rate initially, and gradually decrease as the network trains.

6.4 Tools

This section discusses the specific tools used in the studies described in this thesis: the CNN framework and network architecture, as well as the LArCV framework for storing and interfacing data.

6.4.1 Caffe

The Caffe CNN framework was used for these studies [4]. Caffe is written in C++ and includes Python bindings. Various layer types are implemented as C++ classes, allowing the architecture of a network to be defined using a text-based configuration file written in Google’s Protocol Buffer format [69]. Each layer of the network is configured by specifying the layer type, the layers connected above and below, and any configurable parameters associated with the chosen layer type. A modified version of Caffe was used, allowing for image inputs to be provided in LArCV format.

6.4.2 VGG16 network architecture

For this study, the VGG16 network architecture was used. This network is very deep, with 14 weighted layers — the fully network architecture is shown in Table 6.1. This network architecture was selected based on previous MicroBooNE studies to find effective LArTPC network architectures. The number and dimensionality of fully connected layers has been drastically reduced, to reflect the difference in purpose — the VGG16 networks are designed for a broad range of classes, while for these studies only two classes of physics event are considered.

6.4.3 LArCV

LArCV is a tool developed by MicroBooNE collaborators at Nevis Labs and MIT, designed for storing and interfacing images from LArTPCs with CNNs [70]. It uses the ROOT data format to save events in an eventwise manner, allowing various data types to be stored for each event.

For the purposes of these studies, only two data products are used: the `ROI` and `Image2D` classes. In this context, the `ROI` class is used only to label each event as containing either a signal or background image. The three images associated with the three LArTPC wire planes are stored as `Image2D` objects, and LArCV offers utilities such as image downsampling and zero-padding which are useful for preparing CNN inputs.

Table 6.1: VGG16 network architecture used for the studies described in this thesis, containing only a single fully connected layer. Each convolution layer uses a size of 3×3 for all its filters, and the number of filters increases with depth.

Network layer	Kernel size	No. filters
Input (RGB image)		
Convolution	3×3	64
Convolution	3×3	64
Max pool	2×2	
Convolution	3×3	128
Convolution	3×3	128
Max pool	2×2	
Convolution	3×3	256
Convolution	3×3	256
Convolution	3×3	256
Max pool	2×2	
Convolution	3×3	512
Convolution	3×3	512
Convolution	3×3	512
Max pool	2×2	
Convolution	3×3	512
Convolution	3×3	512
Convolution	3×3	512
Max pool	2×2	
Fully connected (2)		
Softmax		

7 $n - \bar{n}$ Oscillation Search in DUNE

A convolutional neural network (CNN) approach was utilised to determine its effectiveness as a tool for selecting rare events in LArTPCs. A CNN was trained on high-statistics simulation samples of labelled signal ($n - \bar{n}$) and background (atmospheric neutrino) events. Once trained, the CNN was then tested with independent signal and background simulation samples, classifying them with a score from 0 (background-like) to 1 (signal-like). Since a search for $n - \bar{n}$ oscillation is a counting experiment, we classify events using a trained CNN and select any which pass a certain score threshold.

This chapter describes a study performed in Monte Carlo for the DUNE far detector (assuming single phase, described in more detail in Chapter 4), to determine the effectiveness of a CNN at distinguishing signal $n - \bar{n}$ events from atmospheric neutrino backgrounds. It summarises the methods used in producing simulations and performing CNN training, and then analysing the network’s performance by investigating event topologies. Finally, the network’s performance is analysed in the context of signal selection efficiency and background rejection rate, and the propagation of these to an equivalent DUNE sensitivity to the lifetime of free $n - \bar{n}$ oscillation.

7.1 Atmospheric neutrino backgrounds

Due to DUNE’s extreme depth of 1.5 km [9], the rate at which charged particles produced by cosmic ray interactions in the atmosphere will interact in the detector will be minimal. Furthermore, charged particles produced

external to the detector can be identified as crossing the detector boundary. As in all previous searches for $n - \bar{n}$ oscillation with bound neutrons, the largest expected source of background in DUNE is atmospheric neutrinos. In order to perform studies to separate signal from background, high-statistics samples of atmospheric neutrino simulations were produced.

Atmospheric neutrinos were produced using version 6.24.00 of the LArSoft simulation framework [3], using an implementation of the GENIE atmospheric neutrino event generator [58]. The atmospheric neutrino fluxes used were Honda 2015 for the DUNE far detector site at SURF [71]. Honda flux tables do not include the effects of neutrino oscillations, and are separated out into bins in energy and zenith angle θ_z , defined such that $\cos\theta_z = 1$ for downward-propagating neutrinos. The flux dependence on zenith angle is shown in Figure 7.1, as is the zenith angle dependence of the corresponding GENIE event rates.




Figure 7.1: The Honda atmospheric neutrino flux as a function of cosine zenith angle, integrated over azimuthal angle and neutrino energy (a), and the total neutrino event rate in GENIE using this flux as a function of cosine zenith angle (b). Both distributions share a similar shape, but the GENIE event rate also includes the effect of energy-dependent neutrino cross-sections.

7.1.1 Neutrino interactions

The flux of atmospheric neutrinos (shown in Figure 7.2) spans a broad neutrino energy range, over which a variety of different types of interaction can occur, as shown in Figure 7.3. In order to properly model backgrounds for this search, all of these processes must be simulated. This section discusses the different types of neutrino interaction considered and their simulation methods in GENIE.




Figure 7.2: Honda atmospheric neutrino fluxes as a function of neutrino energy.

7.1.1.1 Elastic and quasielastic scattering

In the 0.1-1 GeV energy region, elastic and quasielastic neutrino interactions dominate. The final state of the initial interaction itself (shown in Figure 7.4) is a clean topology; in charged current interactions, the final state contains

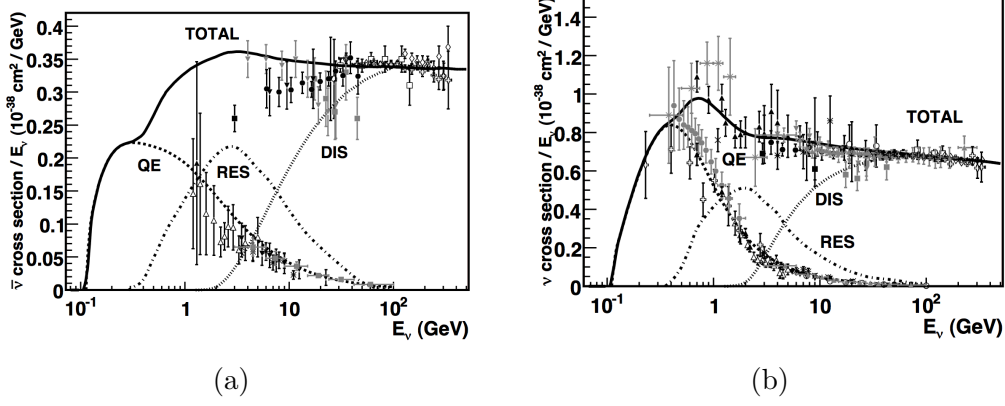


Figure 7.3: Total charged current cross-sections for (a) neutrinos and (b) antineutrinos. These figures originate from [72], which provides a comprehensive summary of all data sources used.




Figure 7.4: Feynman diagrams for (a) charged current quasielastic and (b) neutral current elastic neutrino scattering processes [73].

a lepton and a nucleon, while for neutral current events the neutrino escapes without detection, and only a nucleon is observed. Final state interactions can produce variations in final state, adding pions and nucleons as interaction byproducts reinteract.

Charged current quasielastic scattering cross-sections are modelled in GENIE using an implementation of the Llewellyn-Smith model [74]. In this model, the cross section is parameterised as a function of several scalar and vector form factors — of these, two have been measured using electron scattering, and one is constrained using partially conserved axial current. The remaining axial form factor is assumed to have dipole shape,

$$F_A(Q^2) = \frac{g_A}{\left(1 + \frac{Q^2}{M_A^2}\right)^2}, \quad (7.1)$$

where g_A is an effective axial coupling constant measured in nuclear beta decay [72], and M_A is an axial mass term for which GENIE's default value of 0.99 GeV/c² is used for this work. An additional cross-section suppression is applied for larger nuclei by requiring that the exiting nucleon's momentum is greater than the Fermi momentum for that nucleus [58].

Neutral current elastic cross-sections are modelled in GENIE with an implementation of the Ahrens model [75]. This model is similar to the Llewellyn-Smith model discussed above, but uses a different form for the axial form factor,

$$G_A(Q^2) = \frac{1}{2} \frac{g_A}{\left(1 + \frac{Q^2}{M_A^2}\right)^2} (1 + \eta), \quad (7.2)$$

where η is an adjustable parameter to account for isoscalar contributions to the axial form factor [58]. The same suppression for larger nuclei is performed as for the Llewellyn-Smith model.

7.1.1.2 Resonant production

Once there is enough energy available in the centre-of-mass frame, a neutrino interaction can excite the struck nucleon to produce a baryonic resonance.

The de-excitation of this resonance then produces additional final state particles. For instance, the charged current process $\nu_\mu n \rightarrow \mu^- \Delta^{++}(1232)$ is followed by the decay $\Delta^{++} \rightarrow p\pi^+$, leading to a total final state of $\nu_\mu n \rightarrow \mu^- p\pi^+$.

In the lowest-energy resonance region, dominated by Δ baryon resonances, final states will typically contain a single additional pion. At higher energies, final states can contain multiple pions, kaons or other final states [72].

The cross-section of resonant neutrino events is modelled in GENIE using the Rein-Sehgal model [76], using a resonance axial mass $m_A^{\text{res}} = 1.12 \text{ GeV}/c^2$ [58].

7.1.1.3 Deep inelastic scattering

In the few-GeV neutrino energy region, inelastic scattering interactions begin to occur [72]. Rather than interacting with a nucleon, the neutrino interacts with a single quark inside a nucleon; the quark is ejected from the nucleon, and produces a hadronic shower due to quark confinement. Above 10 GeV, deep inelastic scattering (DIS) events dominate, and higher-multiplicity final states are common. Since the neutrino can scatter from either a valence or sea quark, these events are capable of producing final states containing more exotic particles. Neutral current DIS events are expected to be the dominant background to $n - \bar{n}$ oscillation, as they can produce events with high energies and multiplicities without a final state lepton.

In GENIE, DIS cross-sections are calculated using effective leading order calculations made by Bodek and Yang [77]. They add a scaling factor and Q^2 corrections to GRV98 parton distribution functions [78] in order to calculate neutrino DIS cross-sections.

7.1.1.4 Coherent pion production

Neutrinos can also interact coherently with the entire nucleus, producing a pion via either charged ($\nu A \rightarrow l\pi^+ A$) or neutral ($\nu A \rightarrow \nu\pi^0 A$) currents. Coherent interactions occur at very low Q^2 , producing a forward-going pion

and a nucleus that recoils with low energy without breaking up.

Coherent pion production is modelled in GENIE with a Rein-Sehgal model [79], using an axial mass of $m_A = 1.0$ GeV [58].

7.1.1.5 Neutrino oscillations

In recent decades, it was discovered that neutrinos can undergo a change in flavour as they propagate [39]. In order to correctly model our backgrounds, neutrinos changing flavour in-flight between production in the atmosphere and detection in the DUNE far detector must be simulated. A short summary of neutrino oscillation phenomenology is provided in Appendix B.

The GENIE `FlavorMixer` class included in GENIE is able to change neutrino flavour during event simulation, but three-flavour oscillation is not implemented. As part of this work, a modified version of the Prob3++ neutrino oscillation framework [80,81] was implemented within GENIE, allowing for neutrino flavour oscillation to be performed after GENIE has performed ray-casting and selected a neutrino energy and direction. Prob3++ models neutrino oscillation in the atmosphere, and also accounts for matter effects in an Earth treated as several layers of uniform matter density. Although the publicly available Prob3++ assumes an on-surface detector, a modified version that accounts for underground detectors is used. These modifications were performed by IceCube collaborators [82], and are not yet publicly available.

For each bin in zenith and azimuth angle, Honda atmospheric fluxes are provided accompanied with a probability distribution for neutrino production height in the atmosphere. Since only the direction — not the production height — is modelled by GENIE, this probability distribution is used in conjunction with a Monte Carlo simulation method to select a production height when calculating neutrino path length in order to compute neutrino oscillation probability. The neutrino oscillation parameters used, shown in Table 7.1, are the late-2016 NuFit values [83].

Table 7.1: NuFit normal mass hierarchy oscillation parameters used for oscillating atmospheric neutrino flux (from [83]).

Oscillation parameter	Value
δm_{21}^2	$7.50 \times 10^{-5} \text{ eV}^2$
δm_{32}^2	$2.52 \times 10^{-3} \text{ eV}^2$
$\sin^2 \theta_{12}$	0.306
$\sin^2 \theta_{23}$	0.441
$\sin^2 \theta_{13}$	0.022
δ_{cp}	0

7.2 Monte Carlo production

Signal and background Monte Carlo events were produced using version 6.24.00 of the LArSoft framework [3], with DUNE-specific simulation tools from the `dunetpc` repository. Two separate samples of signal and background events — one for network training, the other for testing — were produced, each consisting of 200,000 events.

The geometry used for these simulations is a reduced version of a DUNE 10 kt single phase module. Since simulating 200 full drift volumes for a single process that occurs over only one or a handful of these drift volumes is an inefficient use of computational resources, instead a reduced $1 \times 2 \times 6$ geometry is used. This consists of 12 wrapped anode planes, each with a drift volume and cathode on both sides, stacked two high and six deep such that the anode wire planes of all 12 modules are aligned along a single plane. This reduced geometry is 7.3 m wide, 12.1 m high and 13.9 m deep. A simple diagram is provided in Figure 7.5.

High-statistics samples were produced by parallelising MC production across many computing nodes using Fermilab grid computing resources. The `project.py` tool was used for job submission and bookkeeping. Simulation is carried out by several successive modules, all of which are described below.




Figure 7.5: Layout of DUNE $1 \times 2 \times 6$ geometry, shown from the side (a) and the front (b). This geometry is stacked two modules high and six modules deep, where each module contains an anode plane in the centre and a drift region on either side.

7.2.1 Event generation

In this stage, the final state products of the initial $n - \bar{n}$ or atmospheric neutrino interaction are produced. Signal events were simulated using the GENIE $n - \bar{n}$ event generator described in Chapter 5, while background event were simulated using the GENIE atmospheric neutrino event generator discussed in Section 7.1.

The vertex of simulated $n - \bar{n}$ events is randomly selected homogeneously inside the active volume of the TPC. The range of simulated vertex positions for atmospheric neutrino interactions is wider, homogeneously selected from the range of the full LArTPC cryostat, which extends an additional 50 cm in all directions except for the negative z direction, in which it extends 3 m beyond boundary of the active volume.

7.2.2 Physics simulation

The byproducts of the initial particle interaction are propagated through the TPC. This is performed using a LArSoft implementation of GEANT4 [84], which consists of several steps: first, all particles produced in the previous

generator stage are propagated through the detector, until they either range out or escape. A separate stage is then performed for simulating ionisation electrons, producing `simchannel` objects that simulate the number of electrons collected on each TPC wire.

Finally, an ‘MC reco’ stage is performed. This stage considers all MC truth particles produced by GEANT inside the detector, and determines which would be visible in the detector. For all visible particles, it then produces a geometric object (either a 3D line for a track, or a cone for an EM shower) associated with that particle, which represents an idealised version of a reconstructed particle.

For a particle to be considered visible, it must be inside the detector’s fiducial volume, it must be a particle type that can be detected by the TPC, and it must be above a certain detection energy threshold. For instance, a π^0 meson does not have an associated MC reco object, since it does not leave a signature in the detector, but if it decays into two γ showers, each of these produces an associated MC shower object.

7.2.3 Detector simulation

There are several stages to detector simulation in LArSoft. First, `simchannel` objects are converted into a waveform format. Noise on waveforms is then simulated, assuming an exponential frequency spectrum for noise with a high-frequency cutoff. A pedestal value is added to the waveform.

Due to the large scale of the DUNE detector, storing every TPC waveform would take up a prohibitively large, amount of mostly redundant data. Therefore a form of zero suppression is performed, and any wire waveforms without any time ticks over 10 ADC after deconvolution are discarded. Since subsequent ROI-finding involves applying a 20 ADC threshold, as described below, this zero suppression involves minimal information loss in terms of the produced image.

7.2.4 Wire reconstruction

Having simulated the detector output, including simulating TPC noise, wire reconstruction is then performed in order to reduce this noise. This reconstruction is designed to be applied identically to data and simulations.

Due to the difference in waveform polarity between different TPC wire planes, a deconvolution stage is applied to DUNE wire waveforms in order to produce unipolar waveforms. During this stage, other wire processing stages such as noise reduction can be applied, but are not performed for the purposes of this analysis. These processed wires are stored as `recob::Wire` objects.

7.3 Image preparation

Before events can be processed by a CNN, some preparatory steps must first be taken. The `recob::Wire` objects from the reconstruction stage are separated out into different APAs; for the purposes of these studies, only a single APA is considered at a time. Usually only wires from a single APA remain after zero suppression, or several APAs if the event crosses multiple modules. For each APA, the ADC counts after deconvolution for all time ticks on all wires across all three planes are added together, and the APA with the most activity is selected by maximising this parameter.

A region of interest is then selected in each of the three wire planes. A 20 ADC threshold is defined, and a rectangular region of interest is found according to the first and last time tick and wire that exceed this threshold. The region of interest is defined as this region plus a small surrounding margin region with a size of 10 pixels on all sides in the final image, corresponding to either 10 wires and 40 time ticks or 20 wires and 80 time ticks either side, depending on the level of downsampling performed.

Due to memory limitations on the size of images provided to the network, the ROI images must be downsampled to a size of 600×600 pixels. Since there are $\mathcal{O}(4000)$ time ticks per event, and only $\mathcal{O}(1000)$ wires per plane, the image is automatically downsampled in the time direction by taking the




Figure 7.6: Image processed by network for the collection plane (top), and zoomed into the interesting region (bottom), for a well-classified signal $n - \bar{n}$ event. This event was given a score of 0.99999968 by the network.




Figure 7.7: Image processed by network for the collection plane (top), and zoomed into the interesting region (bottom), for a poorly-classified (ie. classified as background) signal $n - \bar{n}$ event. This event was given a score of 7.7×10^{-7} by the network.




Figure 7.8: Image processed by network for the collection plane (top), and zoomed into the interesting region (bottom), for a poorly-classified (ie. classified as signal) background atmospheric neutrino event. This event was given a score of 0.99999988 by the network.




Figure 7.9: Image processed by network for the collection plane (top), and zoomed into the interesting region (bottom), for a well-classified background atmospheric neutrino event. This event was given a score of 0 by the network.

average of every four ADC counts in the time dimension. If the image is still larger than 600×600 pixels, the image is then downsampled again in both dimensions, by averaging ADC counts across time ticks and wires. This downsampled image is then embedded in an empty image of 600×600 pixel size, to ensure a uniform image size for all images processed by the network.

As previously mentioned, the region of interest for a CNN image is selected using a 20 ADC threshold. In some events there is zero or minimal activity inside the detector, preventing the ability to apply this threshold and select an ROI. Since these events would produce effectively empty images, which do not contribute to CNN training and cannot be efficiently classified, these events are instead omitted entirely, leading to an inefficiency in image production.

This inefficiency is considerably larger for the atmospheric neutrino sample than the $n - \bar{n}$ sample. Since $n - \bar{n}$ events are simulated inside the fiducial volume only, empty events require that all final state particles produced are either undetectable in LArTPCs (ie. neutrons), or below LArTPC energy detection thresholds. Such topologies are extremely rare for $n - \bar{n}$ events (of the 200,000 test images produced, only 9 events did not produce an image) and so the efficiency is effectively 100%.

In contrast, the image production inefficiency for atmospheric neutrinos in DUNE is approximately 35%. Investigating MC truth information for events for which no image is produced reveals that this inefficiency has two sources. Firstly, due to the fact that events are simulated in the region surrounding the detector, there are some events for which the neutrino interaction occurs entirely outside the TPC fiducial volume. This effect is visible in all spatial dimensions, and is shown for the z component of vertex position in Figure 7.10; the range of the detector in z is -0.8 cm to 1393.4 cm. The second source is a lack of visible energy — many interactions do occur inside the TPC, but produce little or no detectable activity. The total visible energy of atmospheric neutrino events is shown in Figure 7.11, and events for which no image is produced exhibit a clear correlation with lower total visible energy.






Figure 7.10: The z component of atmospheric neutrino event vertex, for events that produced no image (red) and events that did produce an image (blue).



(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.

(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).



(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.

(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).

Figure 7.11: Total event visible energy distributions for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).

7.4 Network training

Once labelled input images had been produced, they were used to train a convolutional neural network. The network was trained using a gradient descent method, as defined in Section 6.3.4. For a more broad overview of CNN concepts such as network training, see Chapter 6.

Network training was performed using an nVidia GTX TITAN X graphics card. The VGG16 network architecture was used (see Section 6.4.2), training on labelled samples of approximately 50,000 signal and background events each. The typical training time for an optimised network is several days.

Several networks were trained at different values of learning rate — 0.001, 0.0003, 0.0001 and 0.000001 — and the optimal learning rate was found to be 0.0003. In addition, networks were trained using different network depths — both *collection plane only* and *3-plane* versions. The network training parameters for the optimal configuration are shown in Table 7.2.

Table 7.2: The solver parameters used during network training. The learning rate and policy are defined in Section 6.3.5, as are the γ and power variables. Weight decay is a regularisation term that prevents the CNN from relying too heavily on any single network weight. The RMS decay is defined in Equation 6.8. Maximum iterations defines how long the network will train for before stopping, while the snapshot defines how often network models are saved.

Parameter	Value
Base learning rate	0.0003
Learning rate policy	“inv”
Gamma	1×10^{-5}
Power	0.75
Weight decay	0.001
RMS decay	0.99
Maximum iterations	100000
Snapshot	250
Network depth	Collection plane

The network’s training progress is output to log files containing the network’s computed loss and accuracy as a function of iteration. These metrics

are visualised and monitored during training by parsing this text file. The training loss for the parameters in Table 7.2 is shown in Figure 7.13, while the corresponding accuracy is shown in Figure 7.12.




Figure 7.12: CNN training accuracy monitored as a function of training iteration. The accuracy is the proportion of images correctly classified by the network, with the training sample shown in blue and red, and the accuracy on an independent test sample shown in green. The optimal network model at 21,000 iterations is indicated by the black dotted line.

7.5 Network performance

7.5.1 Network tests

When a CNN is trained, the intention is that it will learn from features associated with the labelled events it is provided — the differences in topology between $n - \bar{n}$ events and atmospheric neutrino events. The network weights learned in the very first network layer are shown in Figure 7.14 as an example.

One challenge of network training is *overtraining* — given a finite training set, the network will learn on a combination of features associated with a




Figure 7.13: CNN training loss monitored as a function of training iteration. The loss is the parameter the CNN attempts to minimise as it learns, defined in Section 6.3.3. The loss is shown for the training sample in blue and red, and for an independent test sample in green. The optimal network model at 21,000 iterations is indicated by the black dotted line.




Figure 7.14: The network weights for filters in the `conv1_1` layer of a network trained on DUNE $n - \bar{n}$ and atmospheric neutrino events. 64 individual filters are shown, each with a size of 3×3 pixels. Convolution filters are explained in more detail in Section 6.2.1.

type of event and features specific to the events it has been provided with. The former describes the network’s ability to identify any event of that type, and should be maximised; by contrast, the latter should be minimised, as it incorporates features that are not representative of the broader event type.

To account for this effect, the effectiveness of the network is tested using separate Monte Carlo samples — identical to the training samples, but generated independently. Since these events share the same high-level event features as the training sample, but differ in the specifics of each individual event, an overtrained network will perform worse on this test sample. The effects of overtraining are mitigated by testing each snapshot of the network and selecting the model with the lowest loss. The accuracy and loss for the test sample are shown in Figures 7.12 and 7.13.

A trained network was scored using high-statistics samples of signal and background events. 200,000 simulated events were generated of $n - \bar{n}$ and atmospheric neutrino events each. Due to the inefficiency in ROI image production, 199,991 $n - \bar{n}$ ROI images and 132,076 atmospheric neutrino ROI images were produced.

The network assigns a softmax score to each event, giving them scores for the signal and background labels (as described in Section 6.3.2). Since there are only two labels, and the softmax scores by definition sum to unity, we use the signal label score for both signal and background events. Signal-like events are classified with a score closer to 1, while background-like events are scored closer to 0.

The CNN score distribution for the test samples is shown in Figure 7.15. We place a discriminator cut on this score, and count the number of signal and background events that pass. By taking the ratio of passed events to total event rates, we calculate a signal selection efficiency and background rejection rate. Varying this cut and calculating the efficiency and background rate produces an efficiency vs background curve, which is shown in Figure 7.16.

Atmospheric neutrino event rates as a function of CNN discriminator cut are shown in Table 7.3. The majority of neutral current elastic interactions do not produce an image — the ROI image production rate is $\sim 30\%$, compared to $70 - 80\%$ for all other interaction types.




Figure 7.15: Score distributions after CNN training for $n - \bar{n}$ (blue) and atmospheric neutrino (red) events from independent test samples.




Figure 7.16: CNN selection efficiency vs background rate curve. Each point on this curve corresponds to a given cut on CNN score, for which the efficiency and background rate are computed.

Table 7.3: Rate of atmospheric neutrino events produced in GENIE using Honda fluxes, and surviving CNN score discriminator cuts. Percentages are efficiencies relative to the total number of events. Neutral current (NC) deep inelastic scattering (DIS) events are the most commonly mis-identified event type, followed by charged current (CC) resonant pion production, CC DIS and CC quasielastic ν_τ events.

Type		Number of events					
		Total	w/ ROI image	%	CNN score > 0.99	%	CNN score > 0.99995
ν_e	CC QE	31321	23087	73.7	73	0.2	0
	NC elastic	11178	3692	33.0	0	0.0	0
	CC res+coh	9113	6984	76.6	100	1.1	2
	NC res+coh	3340	2432	72.8	21	0.6	0
	CC DIS	6348	5064	79.8	73	1.1	1
	NC DIS	2006	1534	76.5	58	2.9	0
ν_μ	CC QE	32527	23780	73.1	1	0.0	0
	NC elastic	14205	4834	34.0	1	0.0	0
	CC res+coh	11861	9128	77.0	27	0.2	0
	NC res+coh	4787	3450	72.1	23	0.5	0
	CC DIS	11301	9140	80.9	91	0.8	0
	NC DIS	3582	2780	77.6	63	1.8	0
ν_τ	CC QE	110	91	82.7	2	1.8	0
	NC elastic	8431	2805	33.3	0	0.0	0
	CC res+coh	171	133	77.8	2	1.2	0
	NC res+coh	2021	1489	73.7	18	0.9	0
	CC DIS	286	242	84.6	3	1.0	0
	NC DIS	1325	1064	80.3	28	2.1	0
$\bar{\nu}_e$	CC QE	6143	4516	73.5	19	0.3	0
	NC elastic	4345	1428	32.9	0	0.0	0
	CC res+coh	81	69	85.2	1	1.2	0
	NC res+coh	1306	963	73.7	7	0.5	0
	CC DIS	1374	1115	81.1	16	1.2	1
	NC DIS	586	442	75.4	15	2.6	0
$\bar{\nu}_\mu$	CC QE	7716	5568	72.2	0	0.0	0
	NC elastic	5957	1984	33.3	0	0.0	0
	CC res+coh	112	90	80.4	0	0.0	0
	NC res+coh	2152	1569	72.9	11	0.5	0
	CC DIS	2925	2346	80.2	15	0.5	0
	NC DIS	1214	945	77.8	23	1.9	2
$\bar{\nu}_\tau$	CC QE	49	34	69.4	0	0.0	0
	NC elastic	3596	1124	31.3	0	0.0	0
	CC res+coh	57	47	82.5	2	3.5	0
	NC res+coh	869	639	73.5	3	0.3	0
	CC DIS	86	75	87.2	0	0.0	0
	NC DIS	396	312	78.8	3	0.8	0

Neutral current deep inelastic scattering events are the most commonly misclassified event type, with 2 – 3% of events classified with a score of > 0.99 . The network very efficiently rejects ν_e and ν_μ CCQE events, but classifies 2% of ν_τ CCQE events with a score > 0.99 . Resonant, coherent and charged current DIS events are generally misclassified at the $\sim 1\%$ level, with the exception of ν_μ and CC resonant and coherent events which the CNN misclassifies only 0.2% of the time.

With the strongest CNN score discriminator of 0.99995 applied, only six background events remain from the 130,000 event sample — two ν_μ CC resonant events, one ν_μ CC DIS event, one $\bar{\nu}_e$ CC DIS event and two $\bar{\nu}_\mu$ NC DIS events. This corresponds to a background rejection rate of 99.997%. By comparison, 27,815 simulated signal events survived this cut out of the 199,991 event sample.

7.5.2 MC truth feature identification

For a given image classified by the CNN, we retain the event number and match the event to its corresponding Monte Carlo truth information. Access to this information allows for correlations to be drawn between network-selected events and the physics characteristics of the event.

In order to understand the network’s characteristics, the terms **signal-like** and **background-like** are defined. A **signal-like** event is defined as any event with a CNN score higher than 0.99, while a **background-like** event is defined as any event with a CNN score lower than 0.01.

Histograms were produced for various physical quantities, both for *signal-like* signal events and *all* signal events. The ratio of these histograms was then taken, since if the network is able to more efficiently select events in a certain region of parameter space, then the corresponding ratio plot will demonstrate a ratio closer to 1 in this region. Similar ratio plots were produced using the ratio of *background-like* background events to *all* background events.

In signal events, the only particles produced by the GENIE event generator at MC truth level are nucleons and pions. Due to the broad range of neutrino interaction types, the range of final state particles in the background

sample is far more diverse.

Kinematic variables were calculated using MC reco objects, as defined in Section 7.2.2. The advantage of using MC reco objects, rather than MC truth information, is that it allows consideration of only the *visible* portion of the event, discarding any invisible particles such as neutrons and low-energy protons. Since the CNN is obviously unable to learn from invisible particles, this allows us to concentrate solely on the subset of event kinematics the network is sensitive to.


Event kinematics are calculated by considering the energy and momentum of final state particles. Final state particles are defined as:

- Any visible object with an associated MC truth generator-level primary.
- Any visible child particle of an invisible generator-level primary.

The majority of final state particles satisfy the first condition; the second condition is included primarily to account for π^0 mesons, which are themselves invisible but decay into a $\gamma\gamma$ or e^+e^- pair that can be detected.


A particle’s momentum is taken as the MC reco object’s momentum at its first trajectory point. Particle energy is taken in the same manner, with the additional caveat that for a proton, rest mass is subtracted in quadrature — since the majority of protons in the final state will have been ejected from the nucleus during final state interactions, their rest mass should be removed from the calculation of event energy. This removal of proton rest mass is propagated through to event invariant mass, which is calculated using sums over the event’s momentum and energy.

The network’s resolving power as a function of total event kinematic variables — visible energy, invariant mass and net momentum — are shown in Figures 7.11, 7.17 and 7.18, respectively. The network’s ability to correctly classify signal and background events exhibits a weak dependence on total event kinematics, particularly visible energy and invariant mass. The network exhibits some dependence on net momentum, as both signal and background events are classified more effectively at higher net momentum.



(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.


(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).




(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.

(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).


Figure 7.17: Event invariant mass distributions for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).




(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.



(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).




(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.




(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).

Figure 7.18: Total event net momentum distributions for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).



(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.

(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).



(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.

(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).

Figure 7.19: x component of event net momentum for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).

The x component of total event momentum is shown in Figure 7.19. Similar behaviour is exhibited in this distribution to the total event net momentum — events with a higher net momentum are marginally more likely to be correctly classified by the network. The network is not using event net momentum as a discriminator between signal and background, as is the case in a cut-based analysis. Instead, it more effectively classifies both signal and background events at higher net momenta. Net momentum distributions in y and z exhibit the same behaviour as the distribution in x .


The network exhibits a stronger dependence on particle multiplicities than it does on event kinematics, as shown in Figures 7.20, 7.21 and 7.22. The network more efficiently selects signal events with higher multiplicities of charged pions and protons, and with lower multiplicities of neutral pions — in other words, the network exhibits a preference for tracks over EM showers. Since EM showers consist of erratic charge deposition contained within a cone, while tracks consist of a straight line of deposited charge, it stands to reason that the pattern recognition approach of a CNN would more efficiently select the latter.

The weak dependence of CNN performance on high-level physics features demonstrates the power of the CNN approach. The information used by the CNN to select events is complementary to that of traditional reconstruction techniques, which use such physics features to perform an event selection.


7.6 Discussion

7.6.1 APA stitching


In this study, images are produced from whichever DUNE APA contains the most activity, due to the challenge presented in stitching images from multiple APAs together. In doing so, much useful information is discarded. For instance, a neutrino interaction containing a long muon track would be easily distinguished from any $n - \bar{n}$ event due to its extent in space, but limiting images to a single APA truncates the muon track and limits the information available to the network.




(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.



(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).




(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.




(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).


Figure 7.20: Charged pion multiplicity distributions for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).




(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.



(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).




(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.




(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).


Figure 7.21: Neutral pion multiplicity distributions for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).




(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.



(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).



(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.



(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).

Figure 7.22: Proton multiplicity distributions for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).


Stitching event displays from multiple APAs is relatively simple in the collection view, where wires are vertical. Stitching wires in induction planes with diagonal wires is extremely non-trivial, especially considering wires are wrapped around both sides of the collection plane(s).

One suggested method is to stitch images across multiple APAs in the collection view, while retaining the single-APA images in the two induction views. This approach provides the network with a full view of the event in one plane, while also providing a close-up view of the interaction's most interesting region in the remaining planes. Exploring this approach in future studies is strongly advocated.


7.6.2 Effect of containment

Although no consideration for event containment within the selected APA is applied during image production, the effects of containment on network performance were explored. The shortest distance from event vertex to detector module boundary was calculated, and network performance was evaluated as a function of this variable, as shown in Figure 7.23. Although the network is trained and tested on both contained and uncontained events, it exhibits a slight decrease in efficiency for signal events close to a detector module boundary. This behaviour is not observed for background events.


Since the DUNE detector is modular, events are only included in this calculation if the detector module containing the vertex and the module used to produce an image were the same. Due to the larger average spatial extent of atmospheric neutrino events, they are more likely to cross multiple APAs, and so the chances of the image being produced for a subsequent module are higher. Of the events for which an image was produced, that image contained the vertex in 96% of $n - \bar{n}$ events and 80% of atmospheric neutrino events.




(a) Absolute distribution for $n - \bar{n}$ events, including well-classified events (CNN score > 0.99), poorly-classified events (CNN score < 0.99), and events for which no image was produced.



(b) Ratio of well-classified $n - \bar{n}$ events (CNN score > 0.99) to all classified events (omitting events for which no image was produced).



(c) Absolute distribution for atmospheric neutrino events, including well-classified events (CNN score < 0.01), poorly-classified events (CNN score > 0.01), and events for which no image was produced.



(d) Ratio of well-classified atmospheric neutrino events (CNN score < 0.01) to all classified events (omitting events for which no image was produced).

Figure 7.23: Shortest distance from event vertex to closest detector module boundary for $n - \bar{n}$ (top) and atmospheric neutrino (bottom) events, including absolute distributions (left) and event ratios for well-classified events (right).

7.7 Sensitivity calculation

Measurements of processes such as nucleon decay and $n - \bar{n}$ oscillation typically take the form of a lifetime limit at 90% confidence level. For such rare processes, which have never been observed and are not predicted under the Standard Model of particle physics, it is hypothesised that the process will *not* be observed, given an expected background rate and a measured event rate consistent with that expected background rate. This information is then used to set a lower limit on the process's lifetime (in the case that a significant excess above expected background is observed, a different approach must be taken) [85].

Considering two events, A and B , each of which has a distinct probability of occurring, we construct a Bayesian distribution to compute the probability of A given B ,

$$p(A|B) = \frac{p(B|A)p(A)}{p(B)} , \quad (7.3)$$

where $p(A)$ and $p(B)$ are the individual probabilities of A and B . For this kind of search, A refers to a prediction made by theory, and B to the observed data. We therefore calculate the probability of the theory being accurate, provided with a given data set, $p(\text{theory}|\text{data})$.

When measuring discrete events during a continuous measurement period, the appropriate probability distribution to use is the Poisson distribution,

$$p(r; \lambda) = \frac{e^{-\lambda} \lambda^r}{r!} , \quad (7.4)$$

where λ is the mean expected event rate, and r is the number of observed events. This distribution describes the probability of observing a given set of data given a theoretical assumption.

The number of events observed is

$$\lambda = \Gamma E \epsilon + b , \quad (7.5)$$

where Γ is the bound $n - \bar{n}$ oscillation width, E is the total DUNE exposure in neutron-years, ϵ is the signal selection efficiency and b is the expected background rate.

By substituting Equation 7.5 into Equation 7.4, and then placing that total expression into Equation 7.3, we arrive at the total probability of observing a given event rate as a function of the bound $n - \bar{n}$ oscillation width Γ ,

$$P(\Gamma|n_{\text{obs}}) = A \int \int \int \frac{e^{-(\Gamma E \epsilon + b)} (\Gamma E \epsilon + b)^{n_{\text{obs}}}}{n_{\text{obs}}!} P(\Gamma) P(E) P(\epsilon) P(b) dE d\epsilon db . \quad (7.6)$$

This equation is integrated over Gaussian priors for the uncertainty distributions for Γ , λ and ϵ . Since the DUNE detector is not yet operational, preliminary values are chosen for systematic uncertainties based on the recently published Super-Kamiokande analysis, as shown in Table 7.4. The sensitivity for the systematics-free case is also considered — in this case, the exposure, selection efficiency and background rate are held at their central values, without Gaussian priors, and a one-dimensional integral over the $n - \bar{n}$ width is performed.

Table 7.4: Preliminary systematic uncertainties used for DUNE $n - \bar{n}$ sensitivity studies. Since a full systematic uncertainty evaluation is not possible at present, these preliminary values, derived from the Super-Kamiokande $n - \bar{n}$ search, are used as placeholders.

Variable	Systematic uncertainty
Exposure	3 %
Signal selection efficiency	25 %
Background rate	25 %

The DUNE sensitivity is calculated assuming the measured event rate is equal to the expected background rate, for a given value of total detector size and run time, selection efficiency, and expected number of background events. The total integral is evaluated by integrating over Γ from $0 \rightarrow 1 \times 10^{30}$. The upper limit of integration is then iteratively reduced until the total integral

is 90% of its original value. The inverse of this width is $\tau_{n-\bar{n}}$, the bound lifetime limit at 90% confidence level.

The integration is performed using Cubature, an open-source C package for multi-dimensional integration [6]. By computing selection efficiency and background rate as a function of CNN score discriminator, these can then be propagated to sensitivities. Optimal sensitivity is achieved by strongly rejecting background events, rather than efficiently selecting signal events. The bound and free-equivalent sensitivities as a function of CNN score cut for an aggressive score cut are shown in Figures 7.26 and 7.27 respectively, calculated using Equation 2.49 and the conversion factor $0.666 \times 10^{23} \text{ s}^{-1}$ provided in Section 2.2.

At the tail end of the sensitivity vs background rate distribution, the sensitivity is approximately linear as a function of background rate. In order to prevent discontinuities in the sensitivity surface, which exist due to the discrete nature of the size of the statistical sample, a simple linear fit is performed on this distribution, as shown in Figure 7.25, smoothing out the sensitivity surface. In the longer term, this sensitivity distribution can be smoothed by producing a higher-statistics test sample. The relationship between CNN score cut and signal selection efficiency in this region is also shown in Figure 7.24.

Using an optimised cut of 0.99995, a value of 14% for signal selection efficiency and 99.997% for background rejection rate is achieved. Assuming systematic uncertainties similar to Super-Kamiokande, the sensitivity to free $n - \bar{n}$ oscillation lifetime in DUNE after ten years will be $1.6 \times 10^9 \text{ s}$. The way in which this evolves as a function of time is shown in Figure 7.28.

For reference, this sensitivity was compared to the sensitivity for an equivalent exposure calculated using signal selection efficiency and background rejection rates achieved by Super-Kamiokande, with ^{40}Ar bound-to-free conversion factor. This calculation assumed a signal selection efficiency of 12.1% and a background misidentification rate of 0.12%, where this value was extrapolated from an absolute background rate by assuming the same atmospheric neutrino rate per neutron yr exposure as used in DUNE. This free-equivalent $n - \bar{n}$ lifetime sensitivity is $3.3 \times 10^8 \text{ s}$, indicating that DUNE's




Figure 7.24: Signal selection efficiency as a function of CNN score discriminator.

improved sensitivity is due both to its improved mass and its improved resolution.




Figure 7.25: Background rate as a function of CNN score discriminator, with linear fit performed to prevent statistical fluctuations.

7.7.1 Three-plane vs one-plane training

Since network depth (ie. the number of wire plane images used per event) is a configurable parameter, a study was performed to investigate the effect of training a network using images from all three wire planes simultaneously. Network performance was found to be worse for a network trained on all three planes simultaneously. Due to a weaker resolving power, a more relaxed CNN score discriminator of 0.999 was applied, leading to a free $n - \bar{n}$ lifetime sensitivity of 1×10^9 s at 90% confidence level, as shown in Figure 7.29.

Future CNN implementations could adopt a hybrid network architecture, consisting of layers trained on each of the three wire planes independently




Figure 7.26: Bound $n - \bar{n}$ oscillation lifetime sensitivity in DUNE at 90% CL as a function of CNN score discriminator. The discontinuity occurs when the sensitivity calculation switches from using a Gaussian probability distribution to a Poisson distribution.




Figure 7.27: Free-equivalent DUNE sensitivity at 90% CL as a function of CNN score discriminator, compared with the free-equivalent 90% CL limit from Super-Kamiokande. The discontinuity occurs when the sensitivity calculation switches from using a Gaussian probability distribution to a Poisson distribution.




Figure 7.28: Free-equivalent DUNE sensitivity as a function of exposure, for a CNN score discriminator of 0.99995. The sensitivity is shown for the full far detector coming online at once, as well as the sensitivity assuming one 10 kt module coming online per year, over a period of four years. In both cases, it is assumed that all four 10 kt modules use a single phase design.

and layers trained on all three planes simultaneously. Such architectures would likely yield the best results for classification, and their adoption is advocated in futures studies.




Figure 7.29: Free-equivalent $n - \bar{n}$ lifetime sensitivity as a function of exposure for a CNN trained simultaneously on all three wire planes. The sensitivity is shown for the full far detector coming online at once, as well as the sensitivity assuming one 10 kt module coming online per year, over a period of four years. The sensitivity is calculated assuming Super-Kamiokande-like systematic uncertainties, as shown in Table 7.4.

8 $n - \bar{n}$ Oscillation Search in MicroBooNE

In the previous chapter, it was demonstrated that convolutional neural networks show great promise as an analysis tool in LArTPCs. However, these studies are based in Monte Carlo simulations, and implicitly make the assumption that LArTPC data will be perfectly modelled by simulations, an assumption that is questionable at best.

In order to develop these techniques for use in analysing actual data, it is necessary to understand not only the performance of CNNs in simulation, but also how they are applied to data. In this chapter, the studies performed in the previous chapter for DUNE simulations are repeated in the smaller MicroBooNE LArTPC for which data is currently available.

A CNN is trained on Monte Carlo simulations of simulated signal and background events in MicroBooNE. Since MicroBooNE is operating on-surface, and also significantly smaller in size than DUNE, signal and background images differ from the DUNE detector, and so this study is not a direct analogue to the DUNE case. This CNN is used to classify test samples of signal and background simulations, as in the DUNE case, but also to classify MicroBooNE data.

Due to the low exposure considered in MicroBooNE (time exposure of ~ 15 minutes and 87 tons active mass), the rate of atmospheric neutrinos is negligible compared to the large flux of cosmogenic particles. Therefore, cosmogenic interactions are the source of background considered in this study.

Significant differences between the network's classification of simulated and data background events were found, proving that CNNs are highly sen-

sitive to data-MC disagreement. These results are discussed in the context of potential future work, with several different avenues of research proposed for understanding these disagreements and taking steps to mitigate them.

8.1 Cosmogenic background simulation

As MicroBooNE is an on-surface detector, it is constantly exposed to cosmogenic particles produced in the atmosphere by high-energy cosmic rays.

MicroBooNE uses CORSIKA [86] to simulate air showers produced in the atmosphere by high-energy cosmic rays. Hadronic showers above 80 GeV are simulated using GHEISHA routines [87]; below this threshold, CORSIKA passes off hadronic shower simulation to FLUKA [88, 89].

Cosmogenic interactions in MicroBooNE are produced using a 5-component cosmic ray flux model, the constant mass composition model, which considers cosmic ray nuclei in five mass groups — p, α , N, Mg and Fe [90]. Event generation in MicroBooNE samples a previously produced database of hadronic showers generated by CORSIKA, preventing the need to rerun CORSIKA and FLUKA whenever new cosmogenic interactions are simulated.

Cosmogenic particles produced in the region surrounding the MicroBooNE detector in a 50,000 event sample are shown in Table 8.1. These events were simulated using version 6.35.00 of LArSoft [3] and CORSIKA version 7.4003. The most common particles are muons and protons, in addition to a significant contribution from electrons and positrons, photons and neutrons. Charged pions and kaons are simulated at a very low rate, and the rate at which kaons are visible in the detector is vanishingly small. The flux of antinucleons, the particles most likely to emulate an $n - \bar{n}$ topology, is so low that the chances of observing an annihilation event are remote.

8.2 Event selection

The presence of cosmogenic interactions in MicroBooNE necessitates a different approach to DUNE. In DUNE, the interaction rate is low enough that we can assume any given interaction in the detector will be the only interaction

Table 8.1: Multiplicity of detected cosmogenic primaries in a 50,000 event sample of CORSIKA simulations in MicroBooNE, corresponding to a time exposure of 110 s. Particles are generated in the region surrounding MicroBooNE and are only counted if they have an associated MC track or shower, meaning they enter the TPC active volume and have energy above the threshold for detection.

Particle type	No. visible particles	Average no. per event
μ^-	337205	6.74
μ^+	268429	5.37
e^-	26711	0.53
e^+	27308	0.55
γ	113727	2.27
π^+	83	0.0017
π^-	133	0.0027
K^+	1	0.00002
K_L^0	0	0
p	212525	4.25
\bar{p}	9	0.00018
n	22030	0.44
\bar{n}	7	0.00014

occurring in that readout window. Identifying the location of the interaction and defining an image ROI is trivial — we simply search for wires with energy deposition above the expected noise level.

However, in MicroBooNE any given readout window will contain many cosmogenic interactions in addition to the signal $n - \bar{n}$ event. This renders the technique used in DUNE for identifying an ROI via energy deposition information impractical, and so an alternative approach is necessary.

For the purposes of this study, no attempt is made to remove cosmogenic interactions, and no ROI finding is performed. Instead, the entire TPC readout is downsampled to fit into an image size the network can handle. Since the signal and background MC samples contain identical cosmogenic simulations, the network is trained to identify $n - \bar{n}$ oscillation events inside the TPC.

8.3 Monte Carlo production

Since each MicroBooNE signal window contains many cosmogenic interactions, the definition of signal and background is reframed with respect to the DUNE studies. The challenge is not rejecting background events, as in DUNE, but instead in picking out candidate signal events among the cosmogenic interactions.

Two different kinds of Monte Carlo samples were produced. The first was a standard simulation of cosmogenic interactions in MicroBooNE, simulated using CORSIKA as detailed in Section 8.1. The second was another cosmogenic sample, identical except for an additional GENIE $n - \bar{n}$ event generator stage run in addition to the CORSIKA generator. This second sample consists of MicroBooNE cosmogenic events, each of which also contain an $n - \bar{n}$ event.

A MicroBooNE truncated waveform is used, consisting of 6400 time ticks in the time range -400 to 6000. Waveforms are sampled at 2 MHz, and so each time tick is $0.5 \mu\text{s}$ in duration [35]. The $n - \bar{n}$ event vertex is randomly selected homogeneously within a region 5 cm inside the MicroBooNE active volume at time tick 0. Since a full drift readout window of 2.2 ms is smaller

than the extent in time of a readout window, the $n - \bar{n}$ vertex position is inhomogeneous in the produced image, as it cannot appear at the very beginning or end of the drift readout.

Due to the presence of cosmogenic interactions within each MicroBooNE readout window, ROI-finding cannot be performed using ADC thresholding. Instead, no ROI finding is performed at all. Images are produced using the entire MicroBooNE event-sized readout window — 8456 wires and 6400 time ticks. These images are heavily downsampled, with much image resolution being lost in the process. Images are downsampled by a factor of 4 in wire space for U and V planes, and a factor of 6 in wire space for the Y plane. All images are downsampled by a factor of 10 in time, producing images with a uniform size of 600×640 pixels.

Due to the constant flux of cosmogenic particles present in every readout window, no simulated image is empty of activity. The image production inefficiency in the DUNE studies, especially in the atmospheric neutrino background sample, is not present in these MicroBooNE studies. Example CNN images for the signal and background sample are shown in Figures 8.1 and 8.2, respectively.

8.4 Network training

A network was trained using labelled samples of 50,000 events each of signal ($n - \bar{n}$ plus cosmogenics) and background (cosmogenics only) samples. The same network architecture (Table 6.1) and solver parameters (Table 7.2) are used as for the DUNE study. The training accuracy and loss are shown in Figures 8.3 and 8.4 respectively.

8.5 Network performance

The network was tested using independent Monte Carlo samples of signal and background events. The score distributions are shown in Figure 8.5. As is the case for DUNE, the network displays impressive resolving power in




Figure 8.1: CNN Monte Carlo image of full MicroBooNE readout, containing both simulated cosmogenic interactions and a visible $n - \bar{n}$ oscillation star in the bottom centre of the image.




Figure 8.2: CNN Monte Carlo image of full MicroBooNE readout, containing only simulated cosmogenic interactions.




Figure 8.3: MicroBooNE CNN training accuracy monitored as a function of training iteration.




Figure 8.4: MicroBooNE CNN training loss monitored as a function of training iteration.

identifying events containing $n - \bar{n}$ in among the cosmogenic interactions in Monte Carlo.

As in the DUNE case, applying a CNN score cut to classified signal and background events yields a signal selection efficiency and background rejection rate, and varying this score cut produces a curve in efficiency vs background rate, shown in Figure 8.6.

Again the network demonstrates an impressive ability to select events containing signal. In this study the network is not only able to recognise $n - \bar{n}$ interactions, but also pick them out inside events containing cosmogenic interactions.




Figure 8.5: MicroBooNE CNN score distributions for Monte Carlo $n - \bar{n}$ with cosmogenics (blue) and cosmogenic-only (red) events.




Figure 8.6: MicroBooNE signal selection efficiency vs background rate curve for Monte Carlo $n - \bar{n}$ and cosmogenic events.

8.6 MicroBooNE sensitivity to $n - \bar{n}$ oscillation




Figure 8.7: Signal selection efficiency in MicroBooNE as a function of CNN score discriminator.

Using the same sensitivity calculation method detailed in Section 7.7, the $n - \bar{n}$ lifetime sensitivity was evaluated for a search in MicroBooNE cosmogenic data. The central frame of a single MicroBooNE readout window is ≈ 2.2 ms. The current statistics of beam-external triggers in MicroBooNE is approximately 380,000 events, corresponding to a live-time of around 15 minutes, with an active mass of 87 tons. Due to the small MicroBooNE exposure, as well as the large and constant flux of cosmogenic particles, MicroBooNE's $n - \bar{n}$ oscillation lifetime sensitivity is far below current experimental limits.

The signal selection efficiency as a function of CNN score discriminator is shown in Figure 8.7, while the background rate is shown in Figure 8.8.




Figure 8.8: Background rejection rate in MicroBooNE as a function of CNN score discriminator.

The $n - \bar{n}$ lifetime sensitivity was profiled as a function of CNN score discriminator, as shown in Figure 8.9. The MicroBooNE CNN exhibits weaker resolving power than that from DUNE, and so a CNN score discriminator of 0.99 is applied, corresponding to a signal selection efficiency of 45% and a background rejection rate of 99%. Assuming one hour of exposure in MicroBooNE, the free $n - \bar{n}$ lifetime sensitivity is 1.3×10^5 s using only statistical uncertainties.

Although not within the scope of this thesis, it should be noted that in addition to beam-external triggers, the zero-suppressed MicroBooNE supernova readout could be used as a source of beam-external data. This could potentially drastically increase available statistics, providing an exposure on the order of years instead of hours. However, due to the high background rate and small detector active volume, the $n - \bar{n}$ lifetime sensitivity would still be far below current experimental limits. Furthermore, the impact that data zero suppression has on CNN performance has yet to be established.

8.7 Data-MC agreement

The network was also tested using a sample of real cosmogenic data collected in MicroBooNE, with an example image shown in Figure 8.11. The `prod_extbnb_reco_neutrino2016_goodruns_v5` data set was used to produce images. This data set consists of beam-external cosmogenic events, and is slightly biased in comparison to the Monte Carlo training sample. In this data set, the MicroBooNE software trigger imposes a requirement that all events pass a 3.5 photoelectron (PE) optical threshold [91], slightly biasing images classified towards higher TPC activity. This is not expected to have a large impact on image classification.

The CNN score distribution is shown in Figure 8.12. The CNN is found to respond very differently to data events compared to Monte Carlo, misclassifying the majority of images as containing an $n - \bar{n}$ event. Given current lifetime limits on $n - \bar{n}$ oscillation, this is not realistic.

This data test highlights the necessity for further development in order for these techniques to be an effective analysis tool. The studies presented




Figure 8.9: Bound $n - \bar{n}$ lifetime sensitivity at 90% confidence level in Micro-BooNE as a function of CNN score discriminator, assuming a total exposure of 1 hour.




Figure 8.10: Free $n - \bar{n}$ lifetime sensitivity at 90% confidence level in MicroBooNE as a function of exposure, with a CNN score discriminator cut of 0.99 applied. MicroBooNE’s sensitivity with only statistical uncertainties is 1.3×10^5 s after only one hour (a factor of ~ 1000 below current best limits). The discontinuity occurs at the boundary between sensitivity calculation with a Poisson distribution at lower statistics and a Gaussian distribution at higher statistics.




Figure 8.11: CNN data image of cosmogenic interactions in MicroBooNE.




Figure 8.12: MicroBooNE CNN score distributions for Monte Carlo and data cosmogenic events.

in this thesis highlight the great potential of CNNs as a tool for identifying different event topologies — but they also highlight the need to carefully understand and account for differences between data and Monte Carlo.

The distinguishing power of CNNs can only be utilised effectively if the effect of data-Monte Carlo difference can be overcome. Since this technique shows great promise as an analysis tool, further studies are warranted into investigating and solving this challenging problem.

The fact that the network responds very differently to images from data and simulation is indicative of differences between the two. Since a CNN approach utilises different information to a more traditional reconstruction approach, simulations that agree well with data from the viewpoint of traditional reconstruction can be identified as different by a CNN.

There are several methods through which these effects can be mitigated. Since the network is sensitive to differences between simulation and data, we can either entirely eliminate all differences between data and simulation, or search for ways to understand why CNNs respond differently to data and simulation, independently of its ability to classify event types. Since a perfect modelling of data using simulation is unrealistic, the ideal solution will likely be some combination of the two.

8.7.0.1 Noise simulations

The fact that the network is sensitive to data-MC differences is itself a motivator for further study. We can train a network using nominally identical samples from data and simulation, such as the MicroBooNE cosmogenic samples discussed in this chapter. By looking in detail at network weights, we can understand *how* the network is sensitive to data-MC differences.

This technique could be used to improve the precision with which events in a LArTPC are simulated. The network's ability to resolve data from Monte Carlo could be assessed as a function of, for instance, variations in the simulation of noise, or in the flux of cosmogenic particles. Minimising the network's resolving power with respect to these quantities could lead to a more accurate modelling of data.

8.7.0.2 Wires and hits

An early attempt at CNN training in MicroBooNE [92] made use of an MCC (Monte Carlo Challenge) sample for cosmogenic background events, and a custom-generated sample of $n - \bar{n}$ signal events with cosmogenic overlay. The two samples were not consistent in their noise simulation, and the resulting trained CNN was able to separate signal and background with perfect accuracy. A custom background MC sample was then produced, with noise simulation identical to the signal sample, and the network retrained to produce the results presented above. This indicates that CNNs are highly sensitive to differences in noise simulation, which is likely the primary cause of data-MC disagreements.

Since TPC noise is a major contributor to data-MC disagreements, it may be beneficial to produce images using detector hits instead of deconvoluted wires. Since hits are produced by fitting Gaussian distributions to wire signals (as demonstrated in Figure 4.3), using hits to produce images may help to mitigate the effects of noise on network response, particularly in regions of images that contain only noise. This approach has been taken by other MicroBooNE collaborators working with CNNs .

8.7.0.3 Network improvements

In terms of mitigating the network’s sensitivity to differences between simulation and data, both low-level and high-level approaches can be envisioned. In terms of low-level features, training a network to distinguish between data and MC provides us with an insight into the types of convolutional filters the network uses to identify such differences. We can envision modifying a network to penalise it for learning on this type of information, forcing it only to learn on information which is insensitive to data-Monte Carlo differences. This would be a highly non-trivial undertaking, but could be worthwhile if such techniques are to receive widespread use as a LArTPC analysis tool.

Another potential approach to overcoming this problem, previously suggested by other MicroBooNE collaborators, is the use of Generative Adversarial Networks (GANs) [93]. Such networks utilise a generative model, which

generates new images based on a training sample of images, and set it against an adversarial network that attempts to ascertain whether an image has been generated or is part of the training set. We can envision adapting a similar framework that combines a network that classifies Monte Carlo events with another network that understands the differences between data and Monte Carlo.

Blank page

9 Conclusions

The ultimate goal of this thesis is to understand future prospects for exploring new parameter space for $n - \bar{n}$ oscillation in the context of the upcoming Deep Underground Neutrino Experiment. The motivations for this work are several, as discussed in more detail in Chapter 2; measurement of $n - \bar{n}$ oscillation would have implications for the source of matter-antimatter asymmetry in the universe and the nature of neutrino mass, in addition to being a landmark discovery in its own right.

Since its first proposal in the late 1960s, many experiments have searched for — and failed to find — $n - \bar{n}$ oscillation. Chapter 3 summarises previous experimental searches, and also offers a brief summary of the antiproton scattering data from which $n - \bar{n}$ oscillation branching ratios are derived.

If this $n - \bar{n}$ oscillation is realised in nature, it occurs with a lifetime longer than 2.7×10^8 s at 90% CL for a free neutron, or 1.9×10^{32} years for a neutron bound in oxygen. The next generation of neutrino detectors, utilising liquid argon time projection chamber technology, offer high resolution, large exposures and low background environments; the basic principles of this technology is summarised in Chapter 4.

9.1 Event generator development

Chapter 5 discusses the development of an $n - \bar{n}$ event generator as part of GENIE, later included in an official release. The event generator uses branching ratios from antiproton scattering data, and simulates nuclear effects such as Fermi momentum and final state interactions.

Some validation studies are also presented that demonstrate neglect-

ing intermediate states in nucleon-antinucleon annihilation (such as ρ resonances) is an acceptable approximation, given that such structure is largely eliminated by the profound effect that final state interactions have on the annihilation byproducts.

Several additional avenues for improvement to the event generator, including the implementation of updated branching ratios, have been outlined in this thesis. Additional modelling of intermediate states could also be a beneficial improvement, though as discussed above, it is a labour-intensive process that provides little benefit.

The development of this event generator serves a twofold purpose: firstly, it was a fundamental component of the analysis work described in this thesis, and secondly, it is this author's hope that the development and public release of an $n - \bar{n}$ event generator lays important groundwork for future work carried out by others in $n - \bar{n}$ oscillation. The public availability of this generator provides an easy simulation source for others to study $n - \bar{n}$ oscillation, and also provides the opportunity for others to update, modify and refine this event generator in the future.

9.2 Convolutional neural networks

Convolutional neural networks are a powerful tool for image classification, and are increasingly being used as an analysis tool in particle physics. The work described in this thesis involves the use of CNNs to select $n - \bar{n}$ events, and so Chapter 6 provides an overview of basic CNN operating principles, and describes techniques and tools used.

Chapter 7 discusses work carried out in using CNNs to classify and distinguish Monte Carlo simulations of $n - \bar{n}$ and atmospheric neutrino events in DUNE. Simulation methods for atmospheric neutrino events are described, as is the method for producing images for classification. Networks trained on DUNE simulations demonstrates an ability to very efficiently select $n - \bar{n}$ events while rejecting atmospheric neutrino backgrounds. Efforts are made to correlate network learning with event MC truth information — CNNs seem to prefer tracklike objects to showerlike objects. The network's abil-

ity to resolve signal from background is equivalent to a free $n - \bar{n}$ lifetime sensitivity of 1.6×10^9 s at 90% confidence level, a factor of 5 improvement over current best limits, and the first ever evaluation of DUNE's sensitivity to $n - \bar{n}$ oscillation.

The study performed in MicroBooNE, described in Chapter 8, emphasises again the power of convolutional neural networks as an analysis tool — in this context, a CNN is not only able to distinguish signal from background, but also identify signal $n - \bar{n}$ events embedded within LArTPC event displays of cosmogenic interactions, with a free $n - \bar{n}$ lifetime sensitivity of 1.3×10^5 s at 90% CL (neglecting systematics) after only one hour of exposure. However, it also demonstrates the impact of data-MC disagreement on network performance, highlighting the challenges that must be overcome in order for CNNs to be viable as an analysis tool.

9.3 Future challenges

Chapter 8 also summarises several potential avenues through which the effects of data-MC disagreements on CNN performance can be mitigated. Due to the complicated geometry of DUNE, there are undoubtedly optimisations to be made in image production, especially with regards to the wrapped wire geometry and the combination of information across multiple TPC modules.

In addition, testing a network trained on Monte Carlo on data events demonstrates that CNNs are highly sensitive to data-MC disagreements. Several suggestions are made for avenues to overcome this issue, including noise simulation improvements, alternate approaches to image production, and alternate network architectures.

The work carried out in this thesis demonstrates the potential for both the Deep Underground Neutrino Experiment and the use of convolutional neural networks for exploring new parameter space in $n - \bar{n}$ oscillation. Through detailed simulation and image production, a CNN approach is shown to have the potential to improve on current measurements by a factor 5. Given the challenges demonstrated in moving from Monte Carlo to data, strong motivation is provided for the continued development of these techniques over the coming decade.

Blank page

A Antiproton scattering data

This appendix provides the branching ratio tables from several independent measurements of at-rest antiproton annihilation. These branching ratios are discussed within the context of $n-\bar{n}$ oscillation branching ratios in Chapter 3, and are here reproduced in full for reference.

A.1 Antiproton annihilation in hydrogen

The $\bar{n}n$ annihilation branching ratios are derived from four data sets of stopping antiprotons on hydrogen. Each of these four sources is briefly summarised below.

A.1.1 Cresti *et al.*

Cresti *et al.* measured approximately 10,000 antiproton annihilation events in the Saclay 81 cm hydrogen chamber, produced using the CERN proton synchrotron [94]. The branching ratios are shown in Table A.1, and compared with Super-Kamiokande branching ratios in Table A.2. The paper describes these events as stopping antiprotons, with 93% of events being induced by antiprotons which stopped inside the detector, although no numbers are provided for the momentum range of these stopping antiprotons.

A.1.2 Baltay *et al.*

Baltay *et al.* measured approximately 45,000 antiproton annihilation events in the 30 inch Columbia-BNL hydrogen bubble chamber, produced using the Brookhaven Alternating Gradient Synchrotron [95]. The branching ratios are

Table A.1: Branching ratios for $\bar{p}p$ annihilation, as published in Armenteros and French's review of nucleon-antinucleon annihilation [31]. The data is taken from two publications of bubble chamber measurements: one from the CERN proton synchrotron by Cresti *et al.* [94], and the other from the Brookhaven Alternating Gradient Synchrotron by Baltay *et al.* [95].

Channel	Branching ratio [%]	
	Cresti <i>et al.</i>	Baltay <i>et al.</i>
Multi- π^0	—	3.20 ± 0.5
$\pi^+\pi^-$	0.33 ± 0.04	0.32 ± 0.03
$\pi^+\pi^-\pi^0$	8.8 ± 1.4	11.9 ± 1.5
$\pi^+\pi^-n\pi^0$ ($n > 1$)	—	34.5 ± 1.2
$2\pi^+2\pi^-$	$11.4 \pm$ —	$12.0^{+0.6}_{-1.6}$
$\pi^+\pi^-\omega$	5.2 ± 0.3	3.8 ± 0.04
$2\pi^+2\pi^-\pi^0$	38.7 ± 1.9	35.0 ± 5.2
$2\pi^+2\pi^-n\pi^0$ ($n > 1$)	—	21.3 ± 1.1
$3\pi^+3\pi^-$	1.7 ± 0.2	1.9 ± 0.2
$3\pi^+3\pi^-\pi^0$	3.6 ± 0.7	1.6 ± 0.3
$3\pi^+3\pi^-n\pi^0$ ($n > 1$)	—	0.3 ± 0.1

Table A.2: A comparison between Super-Kamiokande $n - \bar{n}$ branching ratios and Cresti *et al.* antiproton annihilation data collected at the CERN proton synchrotron [94], and reproduced in Armenteros and French's 1969 review [31]. Only channels for which both tables have ratios available are considered. The full tables for each experiment are provided in Tables 3.2 and A.1 respectively.

Channel	Branching ratio [%]		Ratio [SK/data]
	SK	Cresti <i>et al.</i>	
$\pi^+\pi^-$	2	0.33	6.1
$\pi^+\pi^-\pi^0$	6.5	8.8	0.74
$2\pi^+2\pi^-$	7	11.4	0.61
$\pi^+\pi^-\omega$	10	5.2	1.9
$2\pi^+2\pi^-\pi^0$	24	38.7	0.62

shown in Table A.1, and compared to Super-Kamiokande branching ratios in Table A.3. The paper describes these events as stopping antiprotons, but again an upper limit on antiproton momentum is not provided.

Table A.3: A comparison between Super-Kamiokande $n - \bar{n}$ branching ratios and Baltay *et al.* antiproton annihilation data collected at the Brookhaven Alternating Gradient Synchrotron [95], and reproduced in Armenteros and French's 1969 review [31]. Only channels for which both tables have ratios available are considered; the full tables for each experiment are provided in Tables 3.2 and A.1 respectively.

Channel	Branching ratio [%]		Ratio [SK/data]
	SK	Baltay <i>et al.</i>	
Multi- π^0	1.5	3.20	0.47
$\pi^+\pi^-$	2	0.32	6.3
$\pi^+\pi^-\pi^0$	6.5	11.9	0.55
$\pi^+\pi^-n\pi^0$ ($n > 1$)	39	34.5	1.1
$2\pi^+2\pi^-$	7	12.0	0.58
$\pi^+\pi^-\omega$	10	3.8	2.6
$2\pi^+2\pi^-\pi^0$	24	35.0	0.69
$2\pi^+2\pi^-n\pi^0$ ($n > 1$)	10	21.3	0.47

A.1.3 Pavlopoulos *et al.*

Pavlopoulos *et al.* measured antiproton annihilations in an NaI γ spectrometer at the CERN Proton Synchrotron [96]. The branching ratios are shown in Table A.4, and compared to the Super-Kamiokande branching ratios in Table A.5. These measurements were made using a beam of antiprotons with an initial momentum of 800 MeV/c, but a momentum range for stopping antiprotons is not provided.

A.1.4 Backenstoss *et al.*

Backenstoss *et al.* measured antiproton annihilations in an NaI γ spectrometer at the CERN Proton Synchrotron [97]. The branching ratios are reproduced in Table A.6, and compared to Super-Kamiokande branching ratios in Table A.7. The antiproton beam with an initial momentum of 600 MeV/c

Table A.4: Branching ratios for $\bar{p}p$ annihilation, as published by Pavlopoulos *et al.* [96]. This data was collected at the CERN Proton Synchrotron using a large NaI γ spectrometer.

Channel	Branching ratio [%]
$\pi^+\pi^-$	0.33 ± 0.04
$2\pi^+2\pi^-$	5.8 ± 0.3
$3\pi^+3\pi^-$	1.9 ± 0.2
$\pi^+\pi^-\pi^0$	7.8 ± 0.6
$2\pi^+2\pi^-\pi^0$	15.8 ± 4.3
$\pi^+\pi^-\omega$	3.8 ± 0.4
$3\pi^+3\pi^-\pi^0$	1.9 ± 0.5
$\pi^+\pi^-2\pi^0$	9.3 ± 3.0
$\pi^+\pi^-3\pi^0$	23.3 ± 3.0
$\pi^+\pi^-4\pi^0$	2.8 ± 0.7
$2\pi^+2\pi^-2\pi^0$	16.6 ± 1.0
$2\pi^+2\pi^-3\pi^0$	4.2 ± 1.0
$2\pi^0$	$0.04 \pm \text{—}$
$n\pi^0$ ($n > 2$)	3.2 ± 0.5

Table A.5: A comparison between Super-Kamiokande $n - \bar{n}$ branching ratios and Pavlopoulos *et al.* antiproton annihilation data collected at the CERN Proton Synchrotron [96]. Only channels for which both tables have ratios available are considered; the full tables for each experiment are provided in Tables 3.2 and A.4 respectively.

Channel	Branching ratio [%]		Ratio [SK/data]
	SK	Pavlopoulos <i>et al.</i>	
$\pi^+\pi^-$	2	0.33	6.1
$2\pi^+2\pi^-$	7	5.8	1.2
$\pi^+\pi^-\pi^0$	6.5	7.8	0.83
$2\pi^+2\pi^-\pi^0$	24	15.8	1.5
$\pi^+\pi^-\omega$	10	3.8	2.6
$\pi^+\pi^-2\pi^0$	11	9.3	1.2
$\pi^+\pi^-3\pi^0$	28	23.3	1.2
$2\pi^+2\pi^-2\pi^0$	10	16.6	0.6
$2\pi^0$	1.5	0.04	37.4

was slowed in a graphite target, and stopped in a 25 cm liquid hydrogen target; no momentum range is provided for the stopping antiprotons.

Since the branching ratios presented in this reference are used for Monte Carlo simulations, rather than measurements from data, this paper does not provide uncertainties on branching ratios.

Table A.6: Branching ratios for $\bar{p}p$ annihilation, as published by Backenstoss *et al.* [97]. These branching ratios were used for Monte Carlo simulations at the CERN Proton Synchrotron using an NaI γ spectrometer.

Channel	Branching ratio [%]
$\pi^+\pi^-\pi^0$	6.4
$\pi^+\pi^-2\pi^0$	9.3
$\pi^+\pi^-3\pi^0$	23.3
$\pi^+\pi^-4\pi^0$	2.8
$\pi^+\pi^-\omega$	3.8
$2\pi^+2\pi^-\pi^0$	13.7
$2\pi^+2\pi^-2\pi^0$	16.6
$2\pi^+2\pi^-3\pi^0$	4.2
$3\pi^+3\pi^-\pi^0$	1.3
$n\pi^0$	3.2
$\pi^+\pi^-$	0.4
$2\pi^+2\pi^-$	6.9
$3\pi^+3\pi^-$	2.1

A.2 Antiproton annihilation in deuterium

Only one source for $\bar{p}n$ annihilation is considered: Bettini *et al.* measured $\sim 20,000$ antiproton annihilation events using a deuterium bubble chamber at the CERN Proton Synchrotron [27]. The measured branching ratios are reproduced in Table A.8, and compared to Super-Kamiokande in Table 3.4. Only events with a spectator proton momentum of $p < 250$ MeV/c are considered.

Table A.7: A comparison between Super-Kamiokande $n - \bar{n}$ branching ratios and Backenstoss *et al.* antiproton annihilation data collected at the CERN Proton Synchrotron [97]. Only channels for which both tables have ratios available are considered; the full tables for each experiment are provided in Tables 3.2 and A.6 respectively.

Channel	Branching ratio [%]		Ratio [SK/data]
	SK	Backenstoss <i>et al.</i>	
$\pi^+\pi^-\pi^0$	6.5	6.4	1.0
$\pi^+\pi^-2\pi^0$	11	9.3	1.2
$\pi^+\pi^-3\pi^0$	28	23.3	1.2
$\pi^+\pi^-\omega$	10	3.8	2.6
$2\pi^+2\pi^-\pi^0$	24	13.7	1.8
$2\pi^+2\pi^-2\pi^0$	10	16.6	0.6
$n\pi^0$	1.5	3.2	0.47
$\pi^+\pi^-$	2	0.4	5.0
$2\pi^+2\pi^-$	7	6.9	1.0

Table A.8: Branching ratios for $\bar{p}n$ annihilation, as published by Bettini *et al.* [27] and reproduced in Armenteros & French's review of antiproton annihilation data [31]. This data was collected at the CERN Proton Synchrotron using a deuterium bubble chamber. Here the charge of π^\pm is left as originally measured, and has not been inverted.

Channel	Branching ratio [%]
$\pi^-n\pi^0$ ($n > 0$)	16.4 ± 0.5
$\pi^-\pi^0$	$0.7 \pm \text{—}$
$2\pi^-\pi^+n\pi^0$ ($n \geq 0$)	59.7 ± 1.2
$2\pi^-\pi^+$	1.57 ± 0.21
$2\pi^-\pi^+\pi^0$	21.8 ± 2.2
$3\pi^-2\pi^+n\pi^0$ ($n \geq 0$)	23.4 ± 0.7
$4\pi^-3\pi^+n\pi^0$ ($n \geq 0$)	0.39 ± 0.07

B Neutrino oscillations

Neutrino oscillations were first proposed by Bruno Pontecorvo in 1958 [98], and first definitively measured at the Sudbury Neutrino Observatory in 2002 [99]. This discovery has serious implications for Standard Model neutrinos: they are not massless, as previously believed, but possess a small but non-zero mass. Three distinct neutrino mass states exist, and mixing occurs between mass and flavour states [39]. Neutrinos interact via a definitive flavour state, and propagate in a definitive mass state. Mixing between mass and flavour states is described using matrix elements,

$$\nu_{\alpha L} = \sum_{k=1}^N U_{\alpha k} \nu_{kL} , \quad (\text{B.1})$$

where U is a unitary matrix, and the subscripts α and k refer to flavour and mass eigenstates respectively — Greek indices are used for flavour states and Roman indices for mass states.

Assuming that no sterile neutrinos exist [100], and there are only three massive neutrinos, this mixing matrix can be parameterised in terms of three mixing angles θ_{12} , θ_{13} and θ_{23} and one CP-violating phase δ_{CP} . The full mixing matrix is

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} , \quad (\text{B.2})$$

first proposed by Maki, Nakagawa and Sakata [101] based on Pontecorvo's

prediction, referred to as the PMNS matrix.

The flavour oscillation probability of a function of time can be expressed as

$$P(\nu_\alpha \rightarrow \nu_\beta) = \langle \nu_\beta | \nu_\alpha(t) \rangle = \sum_{k=1}^3 U_{\alpha k}^* U_{\beta k} e^{-iE_k t} . \quad (\text{B.3})$$

The PMNS matrix describes three-flavour neutrino oscillations in a vacuum, but does not account for the effect of the constant density of electrons in matter. While all three neutrino flavours can interact via neutral current interactions, only electron-flavoured neutrinos (and antineutrinos) can interact via charged current interactions. An additional potential is experienced,

$$\mathcal{H}_F = U \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{\Delta m_{21}^2}{2E} & 0 \\ 0 & 0 & \frac{\Delta m_{31}^2}{2E} \end{pmatrix} U^\dagger + \begin{pmatrix} V_{CC} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (\text{B.4})$$

where $V_{CC} = \pm\sqrt{2}G_F n_e$, where G_F is the Fermi coupling constant and n_e is the electron density in the medium.

Neutrino oscillations are included in the simulation of atmospheric neutrinos as a background process in the search performed in Chapter 7. A full simulation of three-flavour oscillations is performed, including matter effects when neutrinos propagate through the Earth.

References

- [1] K. Abe *et al.* (Super-Kamiokande Collaboration), Search for $n - \bar{n}$ oscillation in Super-Kamiokande, Phys. Rev. D **91**, 072006 (2015), arXiv: 1109.4227.
- [2] J. Gustafson, A Search for Baryon Number Violation by Two Units at the Super-Kamiokande Detector, PhD thesis, 2016.
- [3] E. D. Church, LArSoft: A Software Package for Liquid Argon Time Projection Drift Chambers, arXiv:1311.6774 (2013).
- [4] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, Caffe: Convolutional Architecture for Fast Feature Embedding, arXiv:1408.5093 (2014).
- [5] R. Acciarri *et al.* (MicroBooNE Collaboration), Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber, J. Instrum. **12**, P03011 (2017), arXiv: 1611.05531.
- [6] S. G. Johnson, Multi-dimensional adaptive integration (cubature) in C, 2017, URL: <https://github.com/stevengj/cubature> (last accessed 25 Sept 2017).
- [7] D. G. Phillips II *et al.*, Neutron-Antineutron Oscillations: Theoretical Status and Experimental Prospects, arXiv:1410.1100 (2014).
- [8] P. Langacker, Grand unified theories and proton decay, Phys. Rep. **72**, 185 (1981).

- [9] R. Acciarri *et al.* (DUNE Collaboration), Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 1: The LBNF and DUNE Projects, arXiv:1601.05471 (2016).
- [10] R. N. Mohapatra, Neutron-Anti-Neutron Oscillation: Theory and Phenomenology, J. Phys. G **36**, 104006 (2009), arXiv: 0902.0834.
- [11] C. Cohen-Tannoudji, B. Diu, and F. Laloe, Quantum Mechanics Vol. 1 (Wiley, 1977).
- [12] E. Friedman and A. Gal, Realistic calculations of nuclear disappearance lifetimes induced by neutron-antineutron oscillations, Phys. Rev. D **78**, 016002 (2008), arXiv: 0803.3696.
- [13] K. G. Chetyrkin, M. V. Kazarnovsky, V. A. Kuzmin, and M. E. Shaposhnikov, On the possibility of an experimental search for n – \bar{n} oscillations, Phys. Lett. B **99**, 358 (1981).
- [14] A. D. Sakharov, Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe, Sov. Phys.-JETP **5**, 24 (1967).
- [15] V. Kuzmin, CP-noninvariance and Baryon Asymmetry of the Universe, Pis'ma Zh. Eksp. Teor. Fiz. **12**, 335 (1970).
- [16] R. N. Mohapatra and R. E. Marshak, Local $B - L$ Symmetry of Electroweak Interactions, Majorana Neutrinos, and Neutron Oscillations, Phys. Rev. Lett. **44**, 1316 (1980).
- [17] G. Fidecaro *et al.*, Experimental search for neutron-antineutron transitions with free neutrons, Phys. Lett. B **156**, 122 (1985).
- [18] M. Baldo-Ceolin *et al.*, A new experimental limit on neutron-antineutron oscillations, Z. Phys. C **63**, 409 (1994).
- [19] G. Bressi *et al.*, Final results of a search for free neutron-antineutron oscillations, Nuovo Cimento A **103**, 731 (1990).

-
- [20] M. J. Frost, The NNbar Experiment at the European Spallation Source, arXiv:1607.07271 (2016).
- [21] M. L. Cherry, K. Lande, C. K. Lee, R. I. Steinberg, and B. Cleveland, Experimental Test of Baryon Conservation: A New Limit on Neutron-Antineutron Oscillations in Oxygen, Phys. Rev. Lett. **50**, 1354 (1983).
- [22] T. W. Jones *et al.* (Irvine-Michigan-Brookhaven Collaboration), Search for $n - \bar{n}$ Oscillation in Oxygen, Phys. Rev. Lett. **52**, 720 (1984).
- [23] M. Takita *et al.* (Kamiokande Collaboration), Search for neutron-antineutron oscillation in ^{16}O nuclei, Phys. Rev. D **34**, 902 (1986).
- [24] C. Berger *et al.* (Frèjus Collaboration), Search for neutron-antineutron oscillations in the Frejus detector, Phys. Lett. B **240**, 237 (1990).
- [25] J. Chung *et al.* (Soudan 2 Collaboration), Search for neutron-antineutron oscillations using multiprong events in Soudan 2, Phys. Rev. D **66**, 032004 (2002).
- [26] B. Aharmim *et al.* (SNO Collaboration), The search for neutron-antineutron oscillations at the Sudbury Neutrino Observatory, Phys. Rev. D **96**, 092005 (2017), arXiv:1705.00696.
- [27] A. Bettini *et al.*, Annihilation Into Pions of the $\bar{p}n$ System from Antiprotons at Rest in Deuterium, Nuovo Cimento **47**, 642 (1967).
- [28] C. Amsler *et al.* (Crystal Barrel Collaboration), Annihilation at rest of antiprotons and protons into neutral particles, Nucl. Phys. A **720**, 357 (2003).
- [29] E. Klempt, C. Batty, and J.-M. Richard, The antinucleon nucleon interaction at low energy: Annihilation dynamics, Phys. Rep. **413**, 197 (2005).
- [30] T. Bressani and A. Filippi, Antineutron physics, Phys. Rep. **383**, 213 (2003).

- [31] R. Armenteros and B. French, Antinucleon-Nucleon Interactions Vol. 4, p. 237, 1969.
- [32] A. Bettini *et al.* (ICARUS Collaboration), The ICARUS liquid argon TPC: a complete imaging device for particle physics, Nucl. Instr. Meth. Phys. Res. **315**, 223 (1992).
- [33] R. Guenette, The ArgoNeuT experiment, arXiv:1110.0443 (2011).
- [34] J. Paley *et al.* (LArIAT Collaboration), LArIAT: Liquid Argon In A Testbeam, arXiv:1406.5560 (2014), arXiv: 1406.5560.
- [35] R. Acciarri *et al.* (MicroBooNE Collaboration), Design and Construction of the MicroBooNE Detector, arXiv:1612.05824 (2016).
- [36] R. Acciarri *et al.* (SBND Collaboration), A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam, arXiv:1503.01520 (2015).
- [37] J. Joshi and X. Qian, Signal Processing in the MicroBooNE LArTPC, arXiv:1511.00317 (2015).
- [38] B. J. P. Jones, Sterile Neutrinos in Cold Climates, PhD thesis, 2015.
- [39] C. Patrignani *et al.*, Review of Particle Physics, Chin. Phys. **C40**, 100001 (2016).
- [40] B. Baller, Liquid Argon TPC Signal Formation, Signal Processing and Hit Reconstruction, arXiv:1703.04024 (2017).
- [41] J. S. Marshall and M. A. Thomson, The Pandora Particle Flow Algorithm, arXiv:1308.4537 (2013).
- [42] M. Diwan, M. Potekhin, B. Viren, X. Qian, and C. Zhang, A novel method for event reconstruction in Liquid Argon Time Projection Chambers, J. Phys. Conf. Ser **762**, 012033 (2016).

-
- [43] C. Adams *et al.* (LBNE Collaboration), The Long-Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe, arXiv:1307.7335 (2013).
 - [44] R. Acciarri *et al.* (ArgoNeuT Collaboration), A study of electron recombination using highly ionizing particles in the ArgoNeuT Liquid Argon TPC, J. Instrum. **8**, P08005 (2013), arXiv: 1306.1712.
 - [45] R. Acciarri *et al.* (ArgoNeuT Collaboration), The detection of back-to-back proton pairs in Charged-Current neutrino interactions with the ArgoNeuT detector in the NuMI low energy beam line, Phys. Rev. D **90**, 012008 (2014), arXiv: 1405.4261.
 - [46] M. Wallbank, Automated Reconstruction, Signal Processing and Particle Identification in DUNE, 2017, URL: <https://indico.cern.ch/event/466934/contributions/2590688/> (last accessed 25 Sept 2017).
 - [47] Y.-T. Tsai, MicroBooNE and its Cross Section Measurement, arXiv:1705.07800 (2017).
 - [48] A. Ankowski *et al.*, Energy reconstruction of electromagnetic showers from π^0 decays with the Icarus T600 liquid argon TPC, Acta Phys. Pol. B **41**, 103 (2010).
 - [49] R. Acciarri *et al.* (MicroBooNE Collaboration), Study Towards an Event Selection for Neutral Current Inclusive Single π^0 Production in MicroBooNE, MICROBOONE-NOTE-1006-PUB (2016).
 - [50] C. Rubbia, The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors, CERN-EP-INT-77-08 (1977).
 - [51] V. Radeka and W. J. Willis, Liquid Argon Ionization Chambers as Total Absorption Detectors, Nucl. Instrum. Methods **120**, 221 (1974).
 - [52] P. Benetti *et al.* (ICARUS Collaboration), A three-ton liquid argon time projection chamber, Nucl. Instr. Meth. Phys. Res. **332**, 395 (1993).

- [53] S. Amerio *et al.* (ICARUS Collaboration), Design, construction and tests of the ICARUS T600 detector, Nucl. Instr. Meth. Phys. Res. **527**, 329 (2004).
- [54] C. Anderson *et al.* (ArgoNeuT Collaboration), The ArgoNeuT Detector in the NuMI Low-Energy beam line at Fermilab, J. Instrum. **7**, P10019 (2012), arXiv: 1205.6747.
- [55] J. M. Conrad, W. C. Louis, and M. H. Shaevitz, The LSND and Mini-BooNE Oscillation Searches at High Δm , Annu. Rev. Nucl. Part. Sci. **63**, 45 (2013), arXiv: 1306.6494.
- [56] G. Carugno, B. Dainese, F. Pietropaolo, and F. Ptohos, Electron lifetime detector for liquid argon, Nucl. Instr. Meth. Phys. Res. **292**, 580 (1990).
- [57] R. Acciarri *et al.* (DUNE Collaboration), Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report, Volume 4 The DUNE Detectors at LBNF, arXiv:1601.02984 (2016).
- [58] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, The GENIE Neutrino Monte Carlo Generator: Physics and User Manual, arXiv:1510.05494 (2015).
- [59] C. Andreopoulos, GENIE releases, 2017, URL: `releases.genie-mc.org` (last accessed 25 Sept 2017).
- [60] Particle Data Group Monte Carlo Particle Numbering Scheme, 2007, URL: <http://pdg.lbl.gov/2007/reviews/montecarloorpp.pdf> (last accessed 25 Sept 2017).
- [61] R. D. Woods and D. S. Saxon, Diffuse Surface Optical Model for Nucleon-Nuclei Scattering, Phys. Rev. **95**, 577 (1954).
- [62] A. Bodek and J. L. Ritchie, Fermi-motion effects in deep-inelastic lepton scattering from nuclear targets, Phys. Rev. D **23**, 1070 (1981).

- [63] A. Bodek, M. E. Christy, and B. Coopersmith, Effective Spectral Function for Quasielastic Scattering on Nuclei, *Eur. Phys. J. C* **74**, 3091 (2014), arXiv: 1405.0583.
- [64] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).
- [65] K. O'Shea and R. Nash, An Introduction to Convolutional Neural Networks, arXiv:1511.08458 (2015).
- [66] K. He, X. Zhang, S. Ren, and J. Sun, Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification, arXiv:1502.01852 (2015).
- [67] G. Gordon, D. Dunson, and M. Dudík, editors, *Deep Sparse Rectifier Neural Networks*, , Proceedings of Machine Learning Research Vol. 15, PMLR, 2011.
- [68] S. Ruder, An overview of gradient descent optimization algorithms, arXiv:1609.04747 (2016).
- [69] Google Protocol Buffer format, 2017, URL: <https://developers.google.com/protocol-buffers/docs/reference/overview> (last accessed 22 Dec 2017).
- [70] LArCV: Liquid Argon Computer Vision, 2016, URL: <https://github.com/LArbys/LArCV> (last accessed 25 Sept 2017).
- [71] M. Honda, M. S. Athar, T. Kajita, K. Kasahara, and S. Midorikawa, Atmospheric neutrino flux calculation using the NRLMSISE00 atmospheric model, *Phys. Rev. D* **92**, 023004 (2015), arXiv: 1502.03916.
- [72] J. A. Formaggio and G. P. Zeller, From eV to EeV: Neutrino Cross Sections Across Energy Scales, *Rev. Mod. Phys.* **84**, 1307 (2012), arXiv: 1305.7513.
- [73] J. Ellis, TikZ-Feynman: Feynman diagrams with TikZ, *Comput. Phys. Commun.* **210**, 103 (2017), arXiv: 1601.05437.

- [74] C. H. Llewellyn Smith, Neutrino Reactions at Accelerator Energies, *Phys. Rep.* **3**, 261 (1972).
- [75] L. A. Ahrens *et al.*, Measurement of neutrino-proton and antineutrino-proton elastic scattering, *Phys. Rev. D* **35**, 785 (1987).
- [76] D. Rein and L. M. Sehgal, Neutrino-excitation of baryon resonances and single pion production, *Ann. Phys. (N. Y.)* **133**, 79 (1981).
- [77] A. Bodek and U. K. Yang, Higher twist, ξ_w scaling, and effective LO PDFs for lepton scattering in the few GeV region, *J. Phys. G* **29**, 1899 (2003).
- [78] M. Glück, E. Reya, and A. Vogt, Dynamical parton distributions of the proton and small-x physics, *Z. Phys. C* **67**, 433 (1995).
- [79] D. Rein and L. M. Sehgal, Coherent π^0 production in neutrino reactions, *Nucl. Phys. B* **223**, 29 (1983).
- [80] Prob3++ neutrino oscillation framework, 2012, URL: <http://webhome.phy.duke.edu/~raw22/public/Prob3++> (last accessed 25 Sept 2017).
- [81] V. Barger, K. Whisnant, S. Pakvasa, and R. J. N. Phillips, Matter effects on three-neutrino oscillations, *Phys. Rev. D* **22**, 2718 (1980).
- [82] M. G. Aartsen *et al.* (IceCube Collaboration), The IceCube Neutrino Observatory: Instrumentation and Online Systems, *Journal of Instrumentation* **12**, P03012 (2017), arXiv: 1612.05093.
- [83] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz, Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity, *J. High Energy Phys.* **01**, 087 (2017), arXiv: 1611.01514.
- [84] S. Agostinelli *et al.*, Geant4: A simulation toolkit, *Nucl. Instr. Meth. Phys. Res.* **506**, 250 (2003).

-
- [85] R. Barlow, Statistics: a guide to the use of statistical methods in the physical sciences, Manchester Physics Series (Wiley, 1989).
- [86] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw, CORSIKA: a Monte Carlo code to simulate extensive air showers., 1998.
- [87] H. Fesefeldt, The simulation of hadronic showers, Report PITHA-85/02, RWTH Aachen (1985).
- [88] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, FLUKA: A multi-particle transport code, Report No: CERN-2005-010 (2005).
- [89] T. T. Böhlen, F. Cerutti, M. P. W. Chin, A. Fassò, A. Ferrari, P. G. Ortega, A. Mairani, P. R. Sala, G. Smirnov, and V. Vlachoudis, The FLUKA Code: Developments and Challenges for High Energy and Medical Applications, Nuclear Data Sheets **120**, 211 (2014).
- [90] R. Acciarri *et al.* (MicroBooNE Collaboration), A Comparison of Monte-Carlo Simulations and Data from MicroBooNE, MICROBOONE-NOTE-1014-PUB (2016).
- [91] R. Acciarri *et al.*, (MicroBooNE Collaboration), Software Trigger Documentation, 2016, URL: <https://microboone-docdb.fnal.gov/cgi-bin/private/ShowDocument?docid=5205> (last accessed 25 Sept 2017).
- [92] J. Hewes, Searching for rare processes using CNNs in MicroBooNE, 2017, URL: <https://microboone-docdb.fnal.gov/cgi-bin/private/ShowDocument?docid=10484> (last accessed 25 Sept 2017).
- [93] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative Adversarial Networks, arXiv:1406.2661 (2014).
- [94] M. Cresti *et al.*, Protonium Annihilations Into Pions, Proceedings from the International Conference on Elementary Particles, Sienna **1**, 263 (1963).

- [95] C. Baltay, P. Franzini, G. Lütjens, J. C. Severiens, D. Tycko, and D. Zanello, Annihilations of Antiprotons at Rest in Hydrogen. V. Multipion Annihilations, *Phys. Rev.* **145**, 1103 (1966).
- [96] P. Pavlopoulos *et al.*, Possible evidence for narrow bound states related to the $p\bar{p}$ system, in *AIP Conference Proceedings* Vol. 41, pp. 340–356, AIP Publishing, 1978.
- [97] G. Backenstoss *et al.*, Proton-antiproton annihilations at rest into $\pi^0\omega$, $\pi^0\eta$, $\pi^0\gamma$, $\pi^0\pi^0$, and $\pi^0\eta'$, *Nucl. Phys. B* **228**, 424 (1983).
- [98] B. Pontecorvo, Neutrino experiments and the problem of conservation of leptonic charge, *Sov. Phys.-JETP* **26**, 984 (1958).
- [99] Q. R. Ahmad *et al.* (SNO Collaboration), Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory, *Phys. Rev. Lett.* **89**, 011301 (2002).
- [100] J. M. Conrad and M. H. Shaevitz, Sterile Neutrinos: An Introduction to Experiments, *arXiv:1609.07803* (2016).
- [101] Z. Maki, M. Nakagawa, and S. Sakata, Remarks on the Unified Model of Elementary Particles, *Progr. Theor. Exp. Phys* **28**, 870 (1962).