FERMILAB-THESIS-2022-29

MANCHESTER 1824
The University of Manchester

Light and Dark in Liquid Argon Time Projection Chamber Neutrino Detectors

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

2022

Patrick J. Green
School of Natural Sciences
Department of Physics and Astronomy

Contents

Word Count: 64577

Abstract

Liquid Argon Time Projection Chamber (LArTPC) neutrino detectors have emerged as a detector technology poised to perform measurements of neutrino interactions with unprecedented precision and as a result answer some of the largest open questions in neutrino physics in the coming decade. This thesis describes methods developed to tackle the computational challenges faced as LArTPC detectors increase in scale and complexity moving towards the multi-kiloton DUNE detectors. These include a new, approximated, model that enables rapid simulation of scintillation light in very large scale detectors, as well as the first demonstration of running the LArSoft software framework on a high performance computer. The high precision of LArTPC detectors designed for next-generation neutrino measurements enables them to also search for Beyond the Standard Model physics produced in high-energy proton–fixed-target collisions in neutrino beams. This thesis presents searches for two dark-sector models performed with the ArgoNeuT experiment: Heavy Neutral Leptons and Heavy QCD Axions. Between them, these models can provide solutions for various unresolved puzzles including neutrino mass generation, the baryon asymmetry of the universe, dark matter and the strong CP problem. In both cases, the dark-sector particles could be produced in the NuMI neutrino beam and can then decay to a pair of oppositely charged muons observable in ArgoNeuT and the downstream MINOS near detector. Both measurements required the development of novel experimental selection techniques and enabled new constraints to be set on the existence of these particles in previously unexplored parameter-space. These searches are both the first of their kind in LArTPC neutrino detectors and pave the way for searches at future neutrino facilities.

Declaration of originality

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 other institute of learning.

Copyright statement

Acknowledgements

First and foremost I would like to thank my supervisor Andrzej, without whom none of the work presented in this thesis would have been possible. Your guidance has been invaluable throughout my PhD and I am grateful for the many opportunities you have provided, allowing me to contribute towards a plethora of interesting projects with different collaborators across multiple experiments – even if the quantity was sometimes overwhelming! I would also like to give special thanks to Ornella, Diego, David and Corey for their extensive guidance and support on the various projects I have worked on; and to Kevin for providing a theorists perspective to an experimental analysis and answering my many questions.

To the many students, postdocs and academics of the Manchester and Edinburgh particle physics departments and at Fermilab, thank you for making me feel welcome throughout my entire PhD. To the members of the ArgoNeuT, SBND, MicroBooNE and DUNE collaborations, it has been a pleasure to work alongside you all and I look forward to continuing to do so in the future.

To my family, thank you for your endless support throughout this journey – all the way from showing me the wonders of science as a child to supporting my somewhat circuitous path to studying particle physics. I am also extremely grateful for allowing me to move home from the US at the outset of COVID-19 in the midst of lock-down and for providing the place I could continue to work from during all of the pandemic upheaval uninterrupted. Completing my PhD during this challenging time would have been all the more difficult without this.

And finally to my partner Beth, thank you for putting up with all of the ridiculous demands of completing a PhD – not least moving to another continent for months at a time. I am forever grateful for your selfless support throughout and for persistently reminding me not to just work all of the time. I could not have done this without you.

Chapter 1

Introduction

Liquid Argon Time Projection Chamber (LArTPC) neutrino detectors have in recent years emerged as a detector technology promising measurements of neutrino interactions with unprecedented precision. They are rapidly increasing in scale and complexity, moving towards the multi-kiloton-scale Deep Underground Neutrino Experiment (DUNE) [1]. Over the coming decade, experiments using LArTPC detectors are poised to answer some of the largest open questions in neutrino physics. DUNE will perform precision long-baseline neutrino oscillation measurements in order to determine the neutrino mass hierarchy and the scale of charge-parity (CP) violation in the neutrino sector [2], while the Short-Baseline Neutrino (SBN) program will seek to resolve a variety of long standing anomalies observed at short baselines which could be interpreted as hints of a fourth, yet undiscovered, neutrino state [3]. LArTPC detectors can also be used to probe physics beyond the standard model produced in high-energy proton–fixed-target collisions in neutrino beams. This thesis will present searches for two well-motivated dark-sector models using the ArgoNeuT experiment: Heavy Neutral Leptons and Heavy QCD Axions. The first of these models introduces one or more heavy right-handed sterile neutrinos, commonly referred to as Heavy Neutral Leptons. Extending the standard model in this way can provide solutions to various unresolved puzzles including neutrino mass generation, the baryon asymmetry of the universe (via leptogenesis) and dark matter. The second of these models is an extension of a model introducing a new mechanism that can resolve the strong CP problem. This mechanism gives rise to a new light neutral boson that couples with gluons, the QCD Axion. Constraints on these light axions, however, necessitate significant fine-tuning in order to maintain the solution to the strong CP problem. Heavy QCD Axions arise from an alternative to the original model that enables heavier axions to avoid these constraints. In addition to these two dark-sector searches, this thesis will present two models and methods that provide solutions to the simulation and computational challenges arising from the increase in scale and complexity of LArTPC detectors. First, a new

model to perform fast simulation of the transport of scintillation light in large scale LArTPC detectors will be presented. This allows fast light simulation to be performed at the scale of the DUNE far detectors. Second, a method to run LArTPC simulation and reconstruction on high performance computers will be presented. This will be demonstrated using the Theta high performance computer at Argonne National Laboratory.

This thesis will be structured as follows. Chapter 2 will provide a brief introduction to standard model neutrinos and an overview of the two dark-sector models searched for in this work: Heavy Neutral Leptons and Heavy QCD Axions. In both cases the model motivation, existing constraints, production in neutrino beams and searches in neutrino detectors will be discussed. Chapter 3 will give an overview of the LArTPC detector technology along with a description of the detectors relevant for the models and analyses presented in this thesis: ArgoNeuT and the MINOS near detector, the Short Baseline Near Detector (SBND) and the DUNE far detectors. Chapter 4 will describe the Neutrinos at the Main Injector (NuMI) neutrino beam in the context of the ArgoNeuT experiment. Chapter 5 will then provide an overview of simulation and reconstruction in LArTPC detectors. Chapter 6 will present the semi-analytical model for fast simulation of the transport of scintillation light in large scale LArTPC detectors. It will also discuss the current application of this model in SBND and the DUNE far detectors, where it is now the standard light simulation method used. The model presented in this chapter has been published in European Physical Journal C [4]. Chapter 7 will present a method to perform LArTPC simulation and reconstruction on the Theta high performance computer. A paper describing this method is in preparation. Chapter 8 will present a search for Heavy Neutral Leptons decaying with the signature $N \to \nu\mu^+\mu^-$ using the ArgoNeuT detector. This is the first search for Heavy Neutral Leptons decaying with a di-muon signature in a LArTPC detector. The results of this search exclude a significant region of previously unexplored parameter space for tau-coupled Heavy Neutral Leptons. The search presented in this chapter has been published in Physical Review Letters [5]. Chapter 9 will present a search for Heavy QCD Axions decaying with the signature $a \to \mu^+\mu^-$ again using the ArgoNeuT detector. This is the first search for Heavy QCD Axions in a LArTPC detector. The results of this search exclude a significant region of previously unexplored parameter space for Heavy QCD Axions. A paper describing this search is available as a preprint [6] and has been submitted to a journal.

The author's contributions to the models and searches presented in this thesis are as follows. In Chapter 6 the author contributed to the refinement of the direct light models, was the primary developer of the reflected light models and contributed extensively to the development of the xenon-doped argon light models. The author

was solely responsible for the implementation of all of these models into the LArSoft software framework. The author was also responsible for the applications of the reflected and xenon-doped argon light models in SBND and the DUNE horizontal drift detector, and was the sole developer of the extended model for the DUNE vertical drift detector. In Chapter 7 the author was responsible for the development, optimisation and benchmarking of the workflows on Theta. Technical support was provided throughout from the Argonne Leadership Computing Facility and from the Balsam development team. Finally, both of the dark-sector searches presented in Chapters 8 and 9 were performed in collaboration with phenomenologists. The author was responsible for the experimental parts of these analyses.

Chapter 2

Neutrinos, Heavy Neutral Leptons and Heavy QCD Axions

This chapter will first provide an introduction to standard model neutrinos. It will then discuss two extensions to the standard model: Heavy Neutral Leptons and Heavy QCD Axions. In each case discussing the motivation, existing experimental constraints and how they can be searched for in neutrino experiments. Searches for Heavy Neutral Leptons and Heavy QCD Axions using the ArgoNeuT experiment will be the subject of Chapters 8 and 9, respectively.

2.1 Standard Model neutrinos

2.1.1 Neutrinos and neutrino mixing

Neutrinos are neutral light fermions that interact only via the weak force. They were postulated by Pauli in 1930 as a possible solution to the β-decay problem [7]. The electrons emitted from β-decays were expected to have a discrete energy corresponding to the energy difference between the initial and final nuclear states. However, they were instead observed to have a continuous energy spectrum in apparent violation of conservation of energy given the two-body nature of the decay. Pauli proposed the existence of a new light neutral particle, later named a neutrino, that would be produced during β-decays in addition to the electron. The sum of the energies of the neutrino and the electron would equal the expected energy difference between nuclear states, preserving conservation of energy. The theory of the weak interaction devised by Fermi in 1934 then led to a full theoretical explanation of β-decay including the production of electron anti-neutrinos [8].

Early calculations of the neutrino-nucleon interaction cross section led to the belief that, as a result of them interacting only via the weak force, the neutrino would be impossible to observe experimentally [9]. It was not until significantly

later, after the advent of nuclear reactors in the 1940s and 1950s, that experimental detection became possible. The reactors provided a continuous high-intensity source of electron anti-neutrinos from β-decays compensating for the small interaction cross section. The first observation of the electron anti-neutrino was achieved in 1956 by Reines and Cowan using the Savannah River reactor [10, 11].

The detection of a second flavour of lepton, the muon, in 1937 [12] led to the prediction of the existence of a second flavour of neutrino associated with this lepton, the muon neutrino. This second flavour of neutrino was later detected by Lederman, Schwartz and Steinberger in 1962 [13]. Finally, the discovery of the tau lepton in 1975 [14] led to the prediction of the third neutrino flavour, the tau neutrino, which was first observed in 2000 by the DONUT collaboration [15]. The total number of weakly-interacting standard model neutrinos is constrained to these three flavours as a result of precision measurements of the Z boson decay width [16].

Mixing between neutrinos and anti-neutrinos was proposed by Pontecorvo in 1957 [17], analogous to the mixing between quarks in neutral Kaons [18]. With the discovery of a second type of neutrino, this was expanded to include mixing between the different neutrino flavours by Maki, Nakagawa and Sakata in 1962 [19] and Pontecorvo in 1967 [20]. After the discovery of the tau neutrino, this was further expanded to encompass mixing between all three flavours. In the 1960s measurements of solar neutrinos performed by the Homestake experiment [21] observed a factor of three fewer solar neutrinos than expected by theoretical predictions [22], in what became known as the solar neutrino problem. Pontecorvo proposed that oscillation between the different neutrino flavours could occur as a result of their mixing [20], which could then provide a solution to the solar neutrino problem [23]. Since the Homestake experiment was only sensitive to electron neutrinos, oscillation of the emitted electron neutrinos to other flavours as they propagated from the Sun to the Earth could account for the observed discrepancy. The existence of neutrino oscillations was confirmed in the late 1990s and early 2000s by measurements of solar and atmospheric neutrinos performed by the Sudbury Neutrino Observatory [24, 25] and the Super-Kamiokande experiment [26].

2.1.2 Neutrino mixing formalism

Neutrino mixing, and hence neutrino oscillation, is formalised expressing the neutrino weak eigenstates (flavours), ν_α ($\alpha = e, \mu, \tau$), as a mixture of the neutrino mass eigenstates, ν_i ($i = 1, 2, 3$),

$$|\nu_\alpha\rangle = \sum_{i=1}^{3} U^*_{\alpha i} |\nu_i\rangle \,, \tag{2.1}$$

where $U_{\alpha i}$ are elements of the 3×3 Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix[1],

$$U_{PNMS} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}. \tag{2.2}$$

The parameters of U_{PNMS} can be expressed in terms of three mixing angles θ_{12}, θ_{13} and θ_{23}, along with a charge-parity (CP) violating phase δ_{CP} accounting for the difference in the oscillations of neutrinos and anti-neutrinos,

$$U_{PNMS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.3}$$

where $c_{ij} \equiv \cos(\theta_{ij})$ and $s_{ij} \equiv \sin(\theta_{ij})$. A full discussion of neutrino mixing and oscillations can be found in various reviews provided as part of Reference [27].

An often convenient approximation is to consider neutrino mixing and oscillations between two flavours. The mixing matrix in Equation 2.2 reduces to 2×2 and the probability of oscillation from one flavour to another can be expressed as [27]

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\theta) \sin^2\left(1.27 \cdot \frac{\Delta m^2\,[\mathrm{eV}^2] \cdot L\,[\mathrm{km}]}{E\,[\mathrm{GeV}]}\right), \tag{2.4}$$

where θ is the mixing angle, Δm^2 is the mass splitting between the two mass eigenstates, L is the distance the neutrino has propagated and E is the neutrino energy. In order for the oscillation probability $P(\nu_\alpha \to \nu_\beta)$ to be non-zero, the mass splitting between the states must be non-zero. This therefore implies that in the three-flavour mixing scenario at least two of the neutrinos must have non-zero mass.

In experiments different regions of the neutrino oscillation parameter space can be probed by either varying the neutrino energy or the distance between the source and the detector. For a given neutrino energy, oscillations driven by larger mass splittings can be probed at shorter distances and smaller splittings at longer distances. A summary of the energies and distances probed by a variety of different neutrino oscillation searches is shown in Figure 2.1. Oscillations driven by two different neutrino mass splittings, the so-called atmospheric splitting Δm^2_{32} and solar splitting Δm^2_{21}, are shown by the dashed lines. Neutrino beam experiments can typically be grouped into two categories: short-baseline searches over distances $\mathcal{O}(100\,\mathrm{m})$ to $\mathcal{O}(1\,\mathrm{km})$, and long-baseline searches over distances as large as $\mathcal{O}(1000\,\mathrm{km})$. Short-baseline searches such as the Short-Baseline Neutrino (SBN) program [3] are sensitive to larger mass

[1] $U^*_{\alpha i}$ denotes the complex conjugate of $U_{\alpha i}$. For anti-neutrinos, $U_{\alpha i}$ is used instead of $U^*_{\alpha i}$ in Equation 2.1.

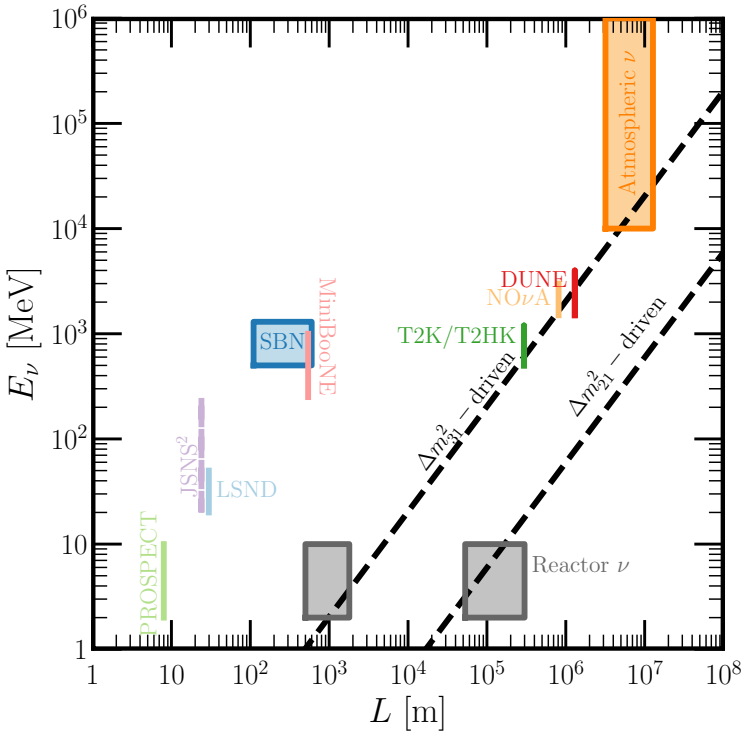

Figure 2.1: Summary of the energies and distances probed by a variety of different neutrino oscillation searches. Figure from Reference [34].

splittings. In particular they are designed to probe $\Delta m^2 \sim 1\,\text{eV}^2$ hinted at in various experimental anomalies [28–31]. Long-baseline experiments such as DUNE [1, 2], NOvA [32] and T2K [33] are sensitive to smaller mass splittings. They are primarily designed to perform precision measurements of oscillations driven by the Δm^2_{32} mass splitting. The regions of phase space explored by atmospheric and reactor neutrino experiments are also shown.

2.1.3 Neutrino masses and mixing parameters

The observation of neutrino oscillation implies that neutrinos have non-zero mass. The neutrino masses are experimentally constrained to be very small. Recent experimental measurements of the absolute electron anti-neutrino mass performed by the KATRIN experiment place an upper limit of $m_{\bar{\nu}_e} < 0.8\,\text{eV}$ [35]. Cosmological measurements provide stronger constraints, requiring the sum of the three neutrino masses to be less than approximately 0.2 eV [36].

Neutrino oscillation measurements probing the neutrino mixing are not sensitive directly to the neutrino mass scale, but can determine the mass splittings. Atmospheric neutrino oscillations are sensitive to Δm^2_{32} and the mixing angle θ_{23} through measurements of muon neutrino disappearance. The Super-Kamiokande experiment found $\Delta m^2_{32} \sim 2 \times 10^{-3}\,\text{eV}^2$ and $\theta_{23} \sim 45°$ [26, 27]. Solar neutrino oscillations are

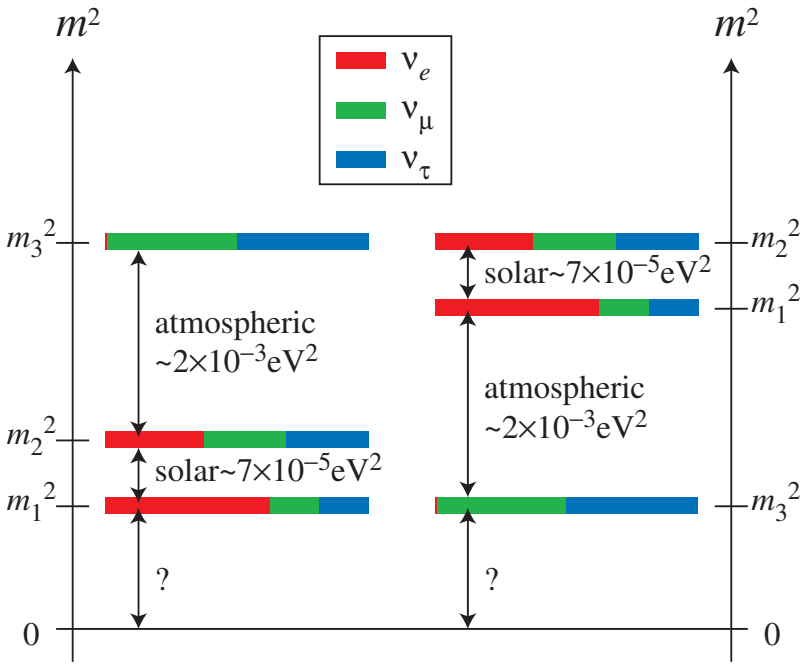

Figure 2.2: Illustration of the two possible neutrino mass orderings: normal ordering (left) and inverted ordering (right). The mass squared splittings from solar and atmospheric neutrino oscillation measurements are also shown. The absolute mass scale is not known, but is constrained to be sub-eV. Figure from Reference [44].

sensitive to Δm^2_{21} and θ_{12} through measurements of the fraction of solar neutrinos reaching the Earth with electron flavour. Measurements performed by the Sudbury Neutrino Observatory (SNO) found $\Delta m^2_{21} \sim 7.5 \times 10^{-5}\,\text{eV}^2$ and $\theta_{12} \sim 33^\circ$ [24, 25, 27]. Finally, the mixing angle θ_{13} can be probed using reactor experiments. The Double Chooz, Daya Bay and RENO experiments found $\theta_{13} \sim 9^\circ$ [27, 37–39].

In solar neutrino oscillation measurements the high matter density of the Sun impacts the oscillation rate of electron neutrinos as they propagate from the core to the surface. This is known as the Mikheyev-Smirnov-Wolfenstein effect [40, 41]. Accounting for this, the best fit to solar and reactor measurements performed by the SNO [24, 25, 42] and KamLAND [43] experiments show that Δm^2_{21} is positive and hence that the ν_2 state is heavier than the ν_1 state. The sign of Δm^2_{32}, however, is not currently known. This leads to two possible neutrino mass orderings: *normal ordering* where ν_3 is the heaviest state and *inverted ordering* where ν_3 is the lightest state. These scenarios are illustrated in Figure 2.2, along with the mass splittings from solar and atmospheric neutrino oscillation measurements.

The final parameter governing neutrino oscillations is the CP-violating phase δ_{CP}. The value δ_{CP} is not known, however recent measurements performed by T2K [45] indicate that it is non-zero. The next generation of beam neutrino experiments, DUNE [1, 2] and Hyper-Kamiokande [46], will aim to determine the value of δ_{CP} along with the neutrino mass ordering.

The existence of neutrino oscillations requires that neutrinos have non-zero mass, however the standard model does not provide a mechanism for the generation of these masses. There are two proposed mechanisms for neutrino mass generation depending on whether neutrinos are Dirac [47] or Majorana [48] particles. Both mechanisms require the introduction of new physics beyond the standard model. To gain mass through the Dirac mechanism, Dirac particles and anti-particles are required to exist with both left-handed and right-handed chiral symmetry. Unlike the other standard model fermions, however, neutrinos (anti-neutrinos) are observed to have only left-handed (right-handed) chirality [49]. Therefore, if neutrinos are Dirac particles, the introduction of additional right-handed neutrino states, often called sterile neutrinos, is required [50]. Alternatively, neutrinos could gain their mass through the Majorana mechanism. If neutrinos are Majorana particles, the distinction between neutrinos and anti-neutrinos vanishes [48]. They are then able to gain mass without requiring the introduction of additional right-handed states. This also introduces processes forbidden in the standard model that could be observed experimentally, most notably neutrinoless double-β decay [51]. This process has not yet been observed and stringent constraints have been placed on its lifetime from a variety of searches in recent years including those performed by the GERDA [52], KamLAND-Zen [53] and CUORE [54] experiments.

2.2 Heavy Neutral Leptons

Heavy Neutral Leptons are a proposed extension to the standard model that could be produced in neutrino beams and searched for in neutrino detectors. A search for Heavy Neutral Leptons decaying with an $N \to \nu\mu^+\mu^-$ signature using the ArgoNeuT experiment will be the focus of Chapter 8.

2.2.1 Motivation

As discussed in Section 2.1.3, the existence of neutrino oscillation requires the standard model neutrinos to have non-zero mass and the standard model lacks a mechanism to generate this. One solution is the introduction of one or more right-handed sterile neutrinos [50, 55], which would allow neutrinos to gain mass through the Dirac mechanism [47]. Sterile neutrinos are neutral right-handed gauge singlet fermions that, unlike the standard model neutrinos, do not interact through the weak force. Instead, they can couple to the standard model neutrinos via an extended neutrino mixing matrix. In some models, they can also interact via a new "right-handed" weak force that is weaker than the standard model weak force [56]. As a result, these new right-handed neutrinos are commonly referred to as sterile neutrinos whereas the

standard model neutrinos are referred to as active neutrinos. Heavy sterile neutrinos are also often referred to as Heavy Neutral Leptons. The right-handed neutrino states enable neutrino mass generation by the Dirac mechanism. Furthermore, the mixing of heavy right-handed neutrinos with the standard model neutrinos can be used to explain the observed small neutrino masses via one of a variety of different so-called seesaw mechanisms [57–63].

The existence of sterile neutrinos, across a broad range of different mass scales, could provide solutions for a variety of unresolved puzzles and experimental anomalies. At the small mass scale, $\mathcal{O}(1)$ eV or less, light sterile neutrinos could provide a solution to a series of long-standing experimental anomalies in short-baseline neutrino data observed by the LSND [28] and MiniBooNE [29–31] experiments. At these scales, oscillation between active and sterile neutrinos is possible. The mass splittings between these states would be larger than those between the active neutrinos, therefore experiments at short baselines would be especially sensitive to their effects [55]. The Short-Baseline Neutrino program at Fermilab is being constructed to resolve these anomalies [3, 64, 65].

At slightly larger mass scales, sterile neutrinos could become candidates for dark matter. Of particular interest are sterile neutrinos with masses $\mathcal{O}(10)$ keV that could form a viable warm dark matter candidate [66, 67]. These are motivated by the observation of an unidentified 3.55 keV emission line that could potentially originate from the decay of keV scale dark matter [68]. Sterile neutrinos of this type could be produced via two different mechanisms. If existing in thermal equilibrium in the early universe via interactions with a new right-handed W boson, where the interaction cross section is suppressed relative to standard weak interaction, they could freeze-out in a manner analogous to active neutrino freeze-out. For this mechanism to generate the observed abundance of dark matter, a second heavier sterile neutrino is also needed with mass $\mathcal{O}(1)$ GeV in order to suppress the rate of production [56]. Alternatively, in the absence of a new interaction, sterile neutrinos of this type could be produced via active-sterile neutrino oscillation while the active neutrinos are in thermal equilibrium in the early universe [69]. In this second scenario, the introduction of a lepton asymmetry in the early universe would be required. This would lead to resonances in the rate of active-sterile neutrino oscillation, allowing a smaller active-sterile mixing angle that is consistent with current constraints [70–73]. These models are strongly constrained by existing astrophysical measurements [74, 75], but have not yet been ruled out.

At higher mass scales, $\mathcal{O}(1)$ GeV or larger, sterile neutrinos can provide a mechanism for baryogenesis, a process producing the observed matter-antimatter asymmetry in the universe, via so-called leptogenesis [76–78]. In this process a lepton asymmetry is generated from decays of the heavy sterile neutrinos. This can then be

converted to a baryon asymmetry as a result of the conservation of $B-L$ in the early universe, where B is the baryon number and L is the lepton number. At even larger scales sterile neutrinos can account for the observed small active neutrino masses via seesaw mechanisms, as previously mentioned. Unlike light sterile neutrinos, heavy sterile neutrinos cannot oscillate with the standard model neutrinos as a result of their much larger mass [79]. Instead, they can decay to standard model particles through their mixing with the standard model neutrinos.

Models have also been constructed that attempt to provide solutions to several of the unresolved puzzles via the introduction of multiple sterile neutrinos at different mass scales. A prominent example is the neutrino Minimal Standard Model (νMSM) [80, 81]. This introduces three right-handed sterile neutrinos, analogous to the three standard model neutrinos. The smallest of these is at the $\mathcal{O}(10)$ keV scale providing a warm dark matter candidate, and then the other two are required to have near-degenerate masses at the $\mathcal{O}(1)$ GeV scale. This model is simultaneously capable of providing mechanisms for neutrino mass generation (and hence neutrino oscillations), dark matter and the baryon asymmetry of the universe [50]. However, it does not provide the lower $\mathcal{O}(1)$ eV scale sterile neutrino needed to resolve the short-baseline neutrino oscillation anomalies.

2.2.2 Experimental searches

Searches for sterile neutrinos (or Heavy Neutral Leptons) have been performed across a vast range of different mass scales. These include neutrino oscillation analyses looking for light sterile neutrinos, searches for heavier sterile neutrinos produced in colliders or neutrino beams, along with constraints from astrophysics and cosmology. Although many models favour or require the introduction of multiple right-handed neutrinos, experimental searches commonly simplify these to include only a single new particle. The mixing between this new sterile neutrino (or Heavy Neutral Lepton), N, and the standard model neutrinos can then be described using one or more new non-zero mixing angles $|U_{eN}|^2$, $\left|U_{\mu N}\right|^2$, and $|U_{\tau N}|^2$. These then form an extended 4×4 version of neutrino mixing matrix shown in Equation 2.2. It can also be assumed, again for simplicity, that only one of these new mixing angles is non-zero at a time. This reduces the sterile-active mixing to a two flavour model, analogous to that described in Equation 2.4[2], allowing easier evaluation of the expected experimental sensitivity. However, it should be noted that some recent studies have suggested that this may be a detrimental oversimplification [82].

A summary of the existing experimental constraints on sterile neutrinos with

[2]This analogy only applicable for light sterile neutrinos, heavier sterile neutrinos cannot oscillate due to loss of coherence [79].

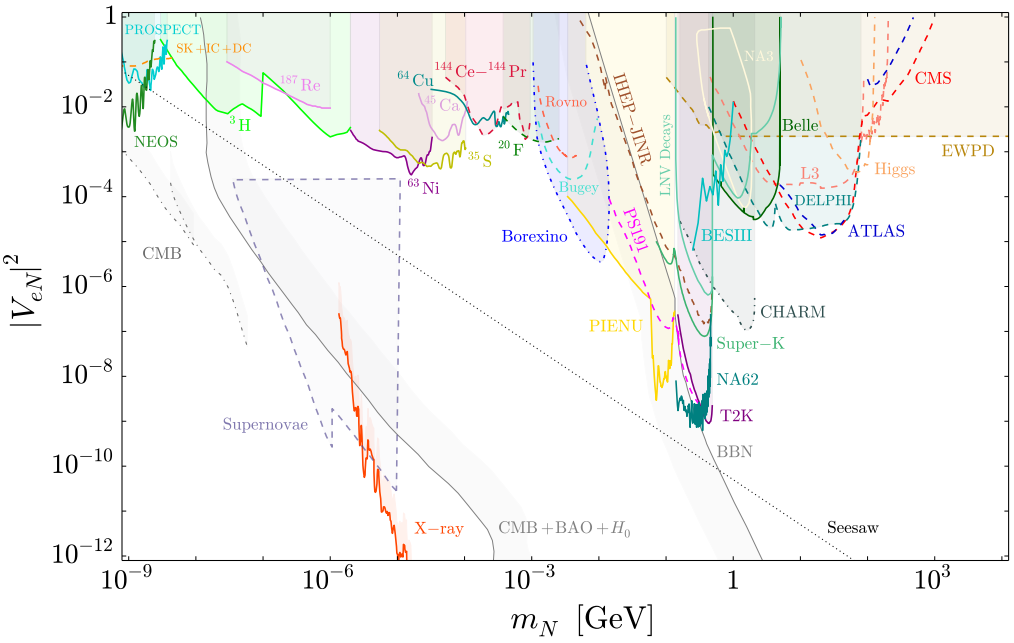

Figure 2.3: Summary of current constraints on the sterile neutrino mixing angle $|U_{eN}|^2$ for sterile neutrino masses, m_N, between 1 eV and 10 TeV. A seesaw relation $|U_{eN}|^2 = m_\nu/m_N$ is also shown, where $m_\nu = 0.5$ eV. Figure from Reference [83].

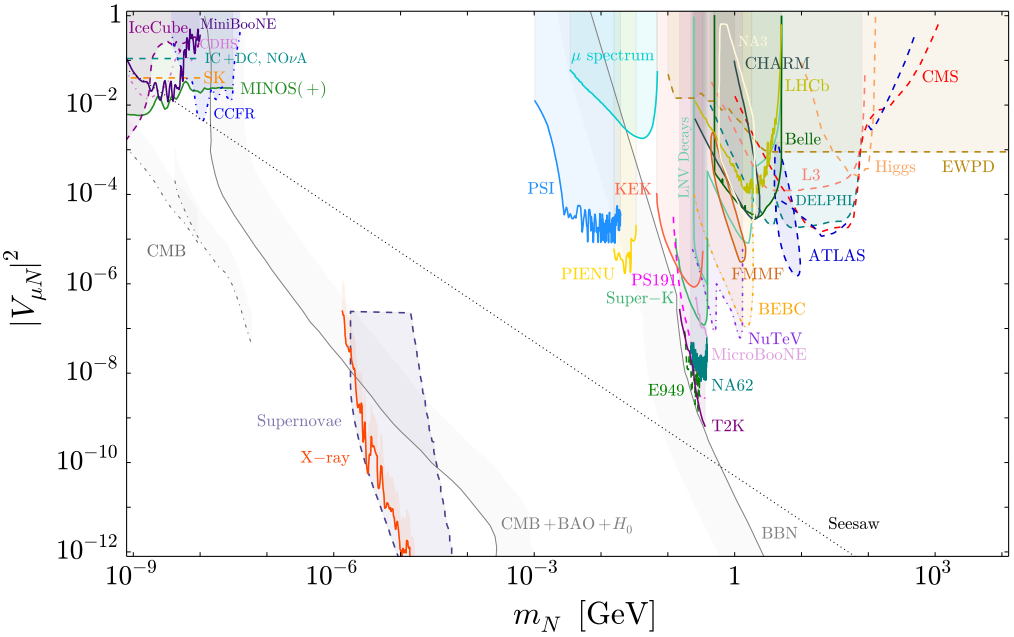

Figure 2.4: Summary of current constraints on the sterile neutrino mixing angle $|U_{\mu N}|^2$ for sterile neutrino masses, m_N, between 1 eV and 10 TeV. A seesaw relation $|U_{\mu N}|^2 = m_\nu/m_N$ is also shown, where $m_\nu = 0.5$ eV. Figure from Reference [83].

masses ranging from 1 eV to 10 TeV scale is shown in Figures 2.3, 2.4 and 2.5. These are split into the constraints on the three different mixing angles with the standard

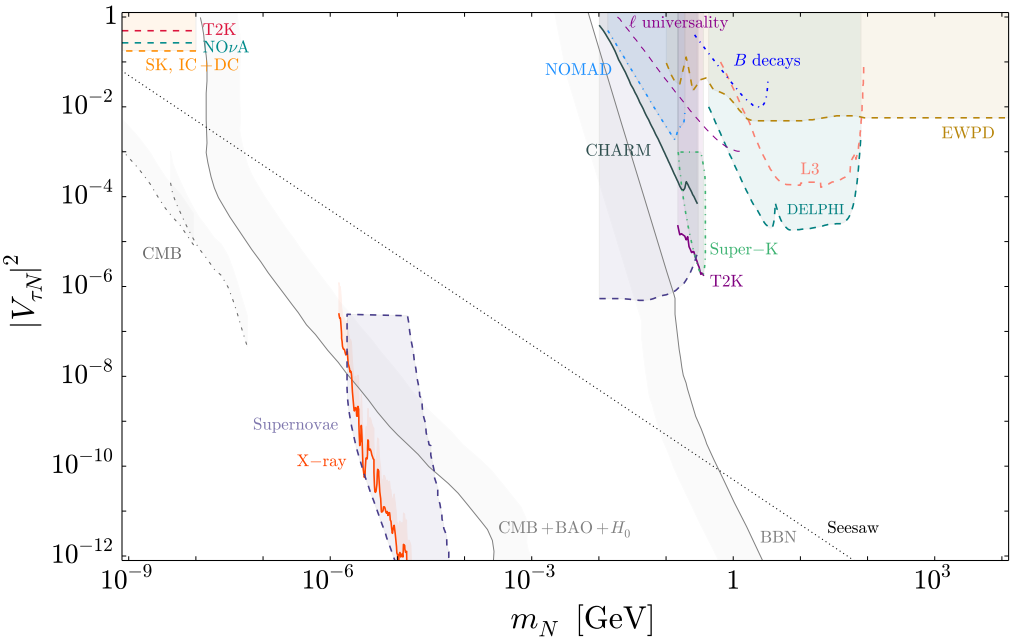

Figure 2.5: Summary of current constraints on the sterile neutrino mixing angle $|U_{\tau N}|^2$ for sterile neutrino masses, m_N, between 1 eV and 10 TeV. A seesaw relation $|U_{\tau N}|^2 = m_\nu/m_N$ is also shown, where $m_\nu = 0.5$ eV. Figure from Reference [83].

model neutrinos[3], $|U_{eN}|^2$, $|U_{\mu N}|^2$, and $|U_{\tau N}|^2$. In the majority of the searches shown, only a single non-zero mixing angle is considered and the other two are assumed to be zero. Other assumptions are made in some cases, however, for example an even mixture between the three mixing angles. This has been accounted for to allow approximate comparison between the resulting constraints [83]. At lower masses, $\mathcal{O}(1)$ eV, the strongest constraints originate from active-sterile neutrino oscillation measurements performed by a variety of reactor, atmospheric and beam neutrino experiments. In the $|U_{eN}|^2$ case, for mass ranges up to $\mathcal{O}(1)$ MeV significant constraints also arise from measurements of the β-decay spectrum that would be altered by the presence of sterile neutrinos. These searches can probe different mass ranges depending on the isotope used. Towards higher masses, beam-dump and collider experiments dominate. The sensitivity of these depends on the intensity and energy of the beam. Beam-dump experiments, including those in neutrino beams, have typically lower interaction energy and are therefore sensitive to lower masses, $\mathcal{O}(10)$ MeV to $\mathcal{O}(10)$ GeV. However, they are often able to probe smaller mixing angles due to their higher intensity beams. Collider experiments on the other hand can reach much higher masses up to the $\mathcal{O}(1)$ TeV scale due to the much higher interaction energies. Finally, several constraints from astrophysics and cosmology across a range of different masses are shown. Further details on all of these constraints can be found in Reference [83]. Experimental searches for sterile neutrinos have

[3]In the figures the notation $|V_{eN}|^2$, $|V_{\mu N}|^2$, and $|V_{\tau N}|^2$ is used, these are the same angles.

been performed using numerous different detector technologies sensitive to masses spanning many orders of magnitude. Nonetheless, large regions of parameter space remain unexplored.

2.2.3 Searches in beam neutrino experiments

Heavy sterile neutrinos can be produced as a result of high-energy proton–fixed-target collisions in neutrino beams. In these beams, protons are accelerated to energies $\mathcal{O}(10)\,\text{GeV}$ to $\mathcal{O}(100)\,\text{GeV}$ and then collide with a fixed target. This produces a flux of charged mesons that can then decay resulting in outgoing light standard model neutrinos, ν_α ($\alpha = e, \mu, \tau$). A heavy sterile neutrino with mass m_N and mixing angle $|U_{\alpha N}|^2$ with these light neutrinos could be produced in any process that would normally result in an outgoing ν_α, provided that it is kinematically allowed. The heavy sterile neutrino mass range that can be produced in these beams typically spans between $\mathcal{O}(10)\,\text{MeV}$ to $\mathcal{O}(1)\,\text{GeV}$. Sterile neutrinos in this mass range are commonly referred to as Heavy Neutral Leptons (HNLs), and this terminology will be used subsequently.

The HNL lifetime, τ_N, can be calculated as a function of the mixing angle $|U_{\alpha N}|^2$ considering all kinematically accessible final states [84, 85]. The decay width, Γ, is suppressed by both the weak scale, G_F^2, and by the small mixing angles $|U_{\alpha N}|^2$,

$$\Gamma(N \to X) \propto \frac{G_F^2 m_N^5 |U_{\alpha N}|^2}{192\pi^3}, \tag{2.5}$$

where X denotes various possible final states. Hence, the decay length, $c\tau_N$, can be relatively long, on the scale of $\sim$1 km,

$$c\tau_N \sim 1\,\text{km} \left(\frac{100\,\text{MeV}}{m_N}\right)^5 \frac{1}{|U_{\alpha N}|^2}. \tag{2.6}$$

As a result of this, HNLs produced in neutrino beams could propagate to downstream neutrino detectors before decaying. A fraction of them could then decay in these detectors producing an observable signature. The high intensity of neutrino beams allows neutrino experiments to have significant sensitivity to HNLs. In particular, searches using neutrino experiments at short-baselines or using the near detectors of experiments at long-baselines can set competitive constraints on the HNL mixing with the standard model light neutrinos.

The HNL decay branching fractions to various final states are shown in Figures 2.6, 2.7 and 2.8 for electron-coupled ($|U_{eN}|^2 \neq 0$), muon-coupled ($|U_{\mu N}|^2 \neq 0$), and tau-coupled scenarios ($|U_{\tau N}|^2 \neq 0$), respectively [84]. In each case the dominant observable decay channels, those that produce charged particles, are shown.

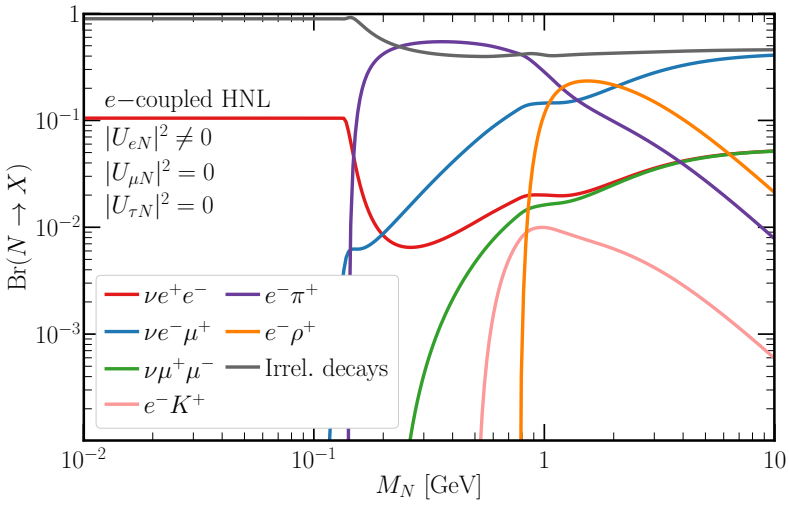

Figure 2.6: Branching fractions of electron-coupled Heavy Neutral Leptons to the dominant observable final states as a function of m_N. Figure from Reference [84].

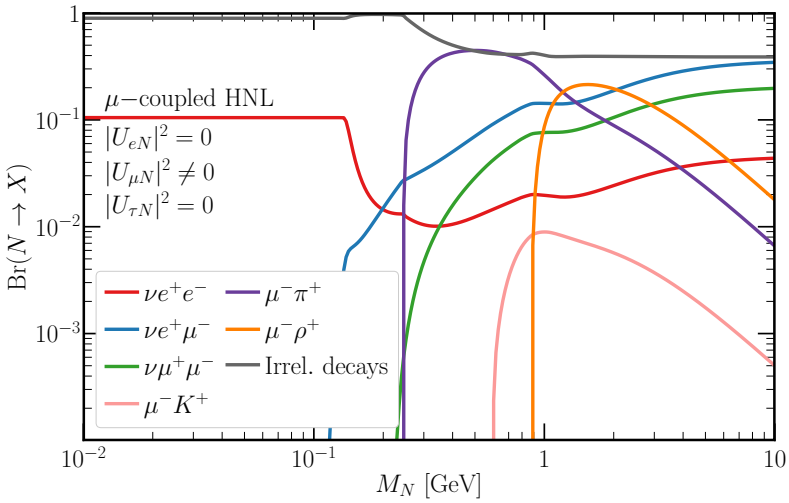

Figure 2.7: Branching fractions of muon-coupled Heavy Neutral Leptons to the dominant observable final states as a function of m_N. Figure from Reference [84].

Decays to non-observable final states, such as $N \to \nu\nu\nu$ or $N \to \nu\pi^0$, are denoted as irrelevant decays. In the mass range of interest in beam neutrino experiments, the dominant decay channels are $N \to e^-\pi^+$ in the electron-coupled scenario and $N \to \mu^-\pi^+$ in the muon-coupled scenario. Searches in neutrino experiments have therefore focused on these signatures in the past [86, 87]. For the tau-coupled scenario, however, there are far fewer possible decay channels in the mass range of interest. This is because, since $m_N < m_\tau - m_\pi$, decays including an outgoing $\tau^\pm$ are not kinematically allowed. Instead, the signatures $N \to \nu e^+e^-$ and $N \to \nu\mu^+\mu^-$ are the dominant observable decay modes. The signature $N \to \nu\mu^+\mu^-$ in particular is simple to reconstruct in neutrino detectors, easy to distinguish from backgrounds, and, in the past, has not been focused on. This signature will therefore be the focus

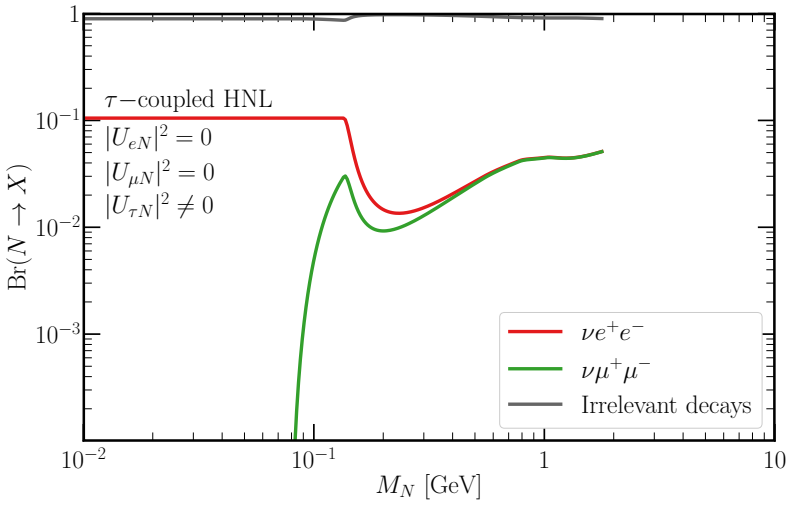

Figure 2.8: Branching fractions of tau-coupled Heavy Neutral Leptons to the dominant observable final states as a function of m_N. Figure from Reference [84].

of the search presented in this thesis.

In all three scenarios it is assumed that the HNL is a Dirac particle. It therefore conserves lepton number, restricting the possible decay channels of N and $\bar{N}$. For example, a Dirac HNL produced in the process $K^+ \to e^+N$ would have a lepton number $L = 1$. Therefore, only the decay process $N \to e^-\pi^+$ would be allowed and $N \to e^+\pi^-$ would be forbidden. If the HNL was instead a Majorana particle, the distinction between N and $\bar{N}$ disappears and the lepton number can be violated. Hence, both processes would be allowed leading to an approximately factor of two increase in the HNL decay width [84, 85]. In evaluating the production and subsequent decay of HNLs, the contributions from both N and $\bar{N}$ need to be accounted for.

2.2.4 Existing constraints at scales relevant to neutrino experiments

In the mass range of interest for HNL searches in neutrino beams various existing constraints are present from searches in beam-dump and neutrino experiments. In the electron-coupled scenario the strongest existing constraints are from the CHARM [88], PS191 [89, 90], T2K [87] and NA62 [91] experiments. These are shown in Figure 2.9 at 90% confidence level for HNL masses between 200 MeV and 2 GeV. At lower masses strong constraints are present probing mixing angles around $|U_{eN}|^2 \sim 10^{-9}$. These become significantly weaker, however, for higher masses beyond the threshold where HNL production via $K^\pm$ decay is kinematically allowed, reducing to around $|U_{eN}|^2 \sim 10^{-6}$.

In the muon-coupled scenario, the strongest existing constraints are from the

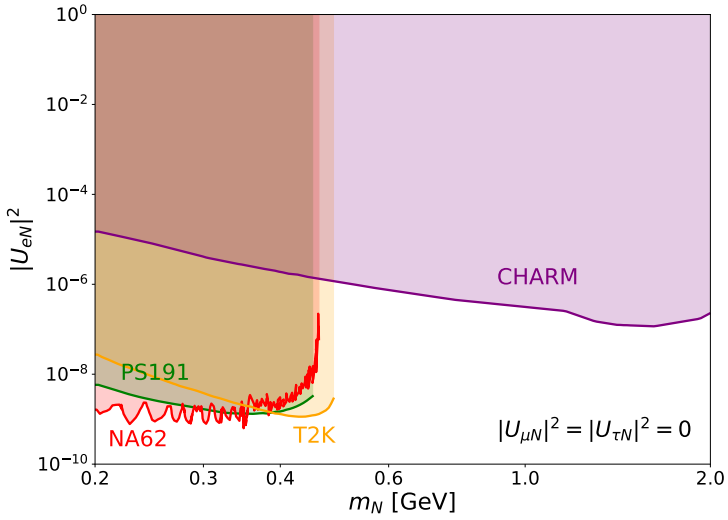

Figure 2.9: Existing constraints on electron-coupled Heavy Neutral Leptons at 90% confidence level. Constraints from the CHARM (purple), PS191 (green), T2K (orange) and NA62 (red) experiments are shown.

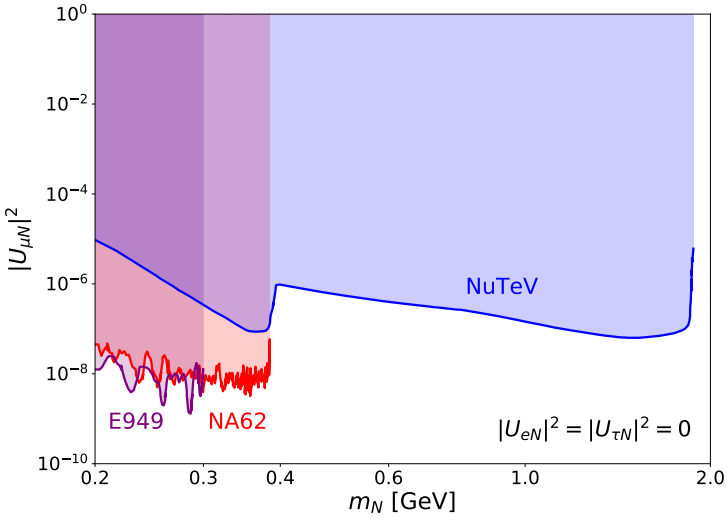

Figure 2.10: Existing constraints on muon-coupled Heavy Neutral Leptons at 90% confidence level. Constraints from the E949 (purple), NA62 (red) and NuTeV (blue) experiments are shown.

E949 [92], NA62 [93] and NuTeV [94] experiments. These are shown in Figure 2.10 at 90% confidence level for HNL masses between 200 MeV and 2 GeV. Additionally, there are constraints from the MicroBooNE experiment [86, 95]. However, they are weaker than other existing limits in the mass range of interest. Similar to the electron-coupled case, there are strong constraints around $\left|U_{\mu N}\right|^2 \sim 10^{-8}$ at lower masses. However, at higher masses where production via $K^{\pm}$ decay is no longer

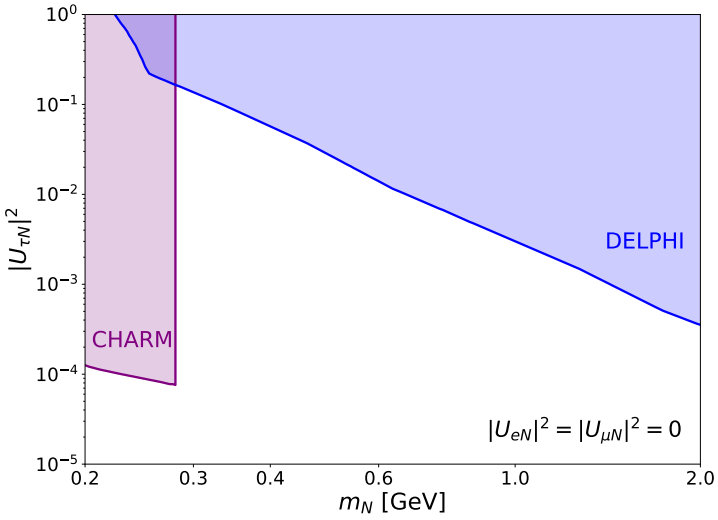

Figure 2.11: Existing constraints on tau-coupled Heavy Neutral Leptons at 90% confidence level. Constraints from the CHARM (purple) and DELPHI (blue) experiments are shown.

possible the constraints again become weaker, around $\left|U_{\mu N}\right|^2 \sim 10^{-6}$.

Finally, in the tau-coupled scenario the strongest existing constraints are from the CHARM [96] and DELPHI [97] experiments. These are shown in Figure 2.11 at 90% confidence level for HNL masses between 200 MeV and 2 GeV. The existing constraints in this scenario are significantly weaker than the electron-coupled and muon-coupled cases: around $|U_{\tau N}|^2 \sim 10^{-4}$ up to 280 MeV and ranging from $|U_{\tau N}|^2 \sim 10^{-1}$ to $|U_{\tau N}|^2 \sim 10^{-3}$ above this. For this reason, the tau-coupled scenario will be the focus of the search presented in this thesis.

2.3 Heavy QCD Axions

Heavy QCD Axions are another proposed extension to the standard model that could be produced in neutrino beams and searched for in neutrino detectors. A search for Heavy QCD Axions decaying with an $a \to \mu^+\mu^-$ signature using the ArgoNeuT experiment will be the focus of Chapter 9.

2.3.1 Motivation

The Quantum Chromodynamic (QCD) axion was originally proposed in the Peccei-Quinn mechanism [98, 99] to address the strong Charge-Parity (CP) problem [100]. This problem arises from the apparent contradiction between experimental measurements of the CP violating angle, $\bar{\theta}$, and theoretical expectations. In the standard model CP violation gives rise to a neutron electric dipole moment. Non-observation

of this experimentally [101–104] constrains $\bar{\theta}$ to be less that $\mathcal{O}(10^{-10})$. Conversely, from the theoretical perspective, there is no reason for it not to be $\mathcal{O}(1)$ [100]. While $\bar{\theta}$ could take any value between 0 and 2π and hence this constraint is not incompatible with QCD, it requires significant fine-tuning of the model and is hence disfavoured on grounds of naturalness. The Peccei-Quinn mechanism introduces a new field, via the Peccei-Quinn symmetry [98, 99], that can minimise $\bar{\theta}$. This field results in a new particle, the QCD axion, a [105, 106]. This particle is a neutral boson with a non-zero mass. It couples to gluons and the resulting strong dynamics generates a potential [107, 108],

$$V \sim -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4 m_u m_d}{(m_u + m_d)^2} \sin^2\left(\frac{a(x)}{2f_a} + \frac{\bar{\theta}}{2}\right)}, \tag{2.7}$$

where m_π is the pion mass, m_u is the up quark mass, m_d is the down quark mass, $f_\pi = 93\,\text{MeV}$ is the pion decay constant and f_a is the axion decay constant. The QCD axion dynamically relaxes to the minimum of this potential, at $\langle a \rangle = -\bar{\theta} f_a$, resulting in the effective $\bar{\theta}$ parameter becoming vanishingly small and the strong CP problem being resolved.

The existence of QCD axions would have strong implications for cosmology and astrophysics, and as a result of this various astrophysical measurements constrain their properties [109–116]. In particular, the QCD axion decay constant, f_a, is required to be large: at least $f_a \sim 10^9\,\text{GeV}$. The QCD axion mass, m_a, then becomes small, governed by the relation [107]

$$m_a \simeq 5.7\,\text{meV}\,\frac{10^9\,\text{GeV}}{f_a}. \tag{2.8}$$

These light QCD axions can form a potential candidate for cold dark matter [117–119]. In these models the axion dark matter is produced in the early universe as a result of oscillations of the axion field about the minimum of its potential. In order to be consistent with the observed abundance of dark matter along with the various astrophysical and cosmological constraints, the mass of these axions is constrained to be on the scale of $\mathcal{O}(1)\,\mu\text{eV}$ to $\mathcal{O}(1)\,\text{meV}$ [120].

The simplest implementations of this mechanism, however, suffer from the so-called axion quality problem [121–123]. Planck scale suppressed Peccei-Quinn symmetry breaking operators can drastically alter the QCD axion potential. These contributions can shift the QCD axion to a minimum away from the one required to solve the CP problem, $\langle a \rangle \neq -\bar{\theta} f_a$. Since f_a is required to be large, the axion mass becomes small as shown in Equation 2.8. This light mass corresponds to a very flat potential that is especially susceptible to the Planck-suppressed contribu-

tions. To solve the strong CP problem a large number of these operators have to be forbidden[4] [124]. This again requires significant fine-tuning of the model and is hence disfavoured on grounds of naturalness.

Heavy QCD Axions that have a much larger mass, but with a similar coupling to gluons, can significantly relax the axion quality problem. By incorporating additional axion interactions beyond QCD, such as those proposed in References [125–130], the relation between m_a and f_a in Equation 2.8 no longer needs to hold. In particular, m_a can be much larger while f_a can independently be much smaller. A larger m_a makes the potential less flat and thus more resistant to the Planck-suppressed contributions, while a smaller f_a reduces the impact of them [108]. Finally, astrophysical constraints are also significantly weaker for $m_a > 10\,\text{MeV}$. The Heavy QCD Axion can therefore resolve the strong CP problem without succumbing to the axion quality problem, providing significant motivation for these models. Heavy QCD Axions can also provide explanations for various astrophysical and cosmological phenomena such as dark matter [131], the baryon asymmetry of the universe through so-called axiogenesis [132], and inflation [133]. The larger m_a opens up various decay channels involving standard model particles and smaller f_a enhances the interaction strengths, enabling searches for these Heavy QCD Axions in beam-dump and collider experiments.

2.3.2 Experimental searches

The majority of existing axion searches have focused on the ultra-light mass axions motivated by dark matter. The constraints on these types of axions predominately originate from astrophysical measurements, including solar axions and significant constraints from the SN1987A supernova, along with cosmological constraints [109–116]. Laboratory based searches have also been performed using the "light shining through a wall" (LSW) technique [134]. In these experiments, photons from a laser are directed towards an opaque wall passing through a strong magnetic field. Axions could be produced from the photons and would then pass through the wall while the remaining photons would be blocked. The axions could then convert back into photons beyond the wall and be detected. The strongest existing constraint using this technique originates from the OSQAR [135] experiment. The constraint on the axion-photon coupling, $g_{A\gamma\gamma}$, is shown in Figure 2.12. The constraint from the Any Light Particle Search (ALPS) LSW experiment is also shown [136]. The

[4]Specifically the operators are proportional to $f_a^n/M_{\text{Pl}}^{n-4}$ with $n > 4$, where M_{Pl} is the Planck mass. At larger n the impact of these contributions becomes increasingly smaller. The particular number of terms that need to be excluded depends on the scale of f_a, see for example numerical evaluations in Reference [124]. In this scenario, where $f_a \sim 10^9\,\text{GeV}$, all terms with $n \leq 9$ have to be forbidden.

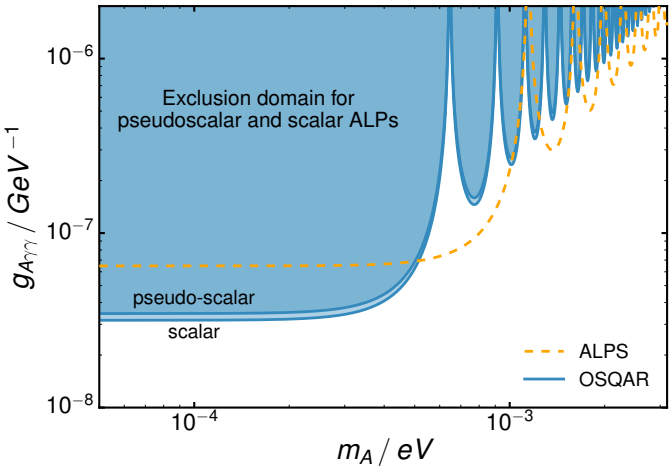

Figure 2.12: Constraints on the light QCD axion coupling with photons, $g_{A\gamma\gamma}$, from the OSQAR and ALPS LSW experiments. Figure from Reference [135].

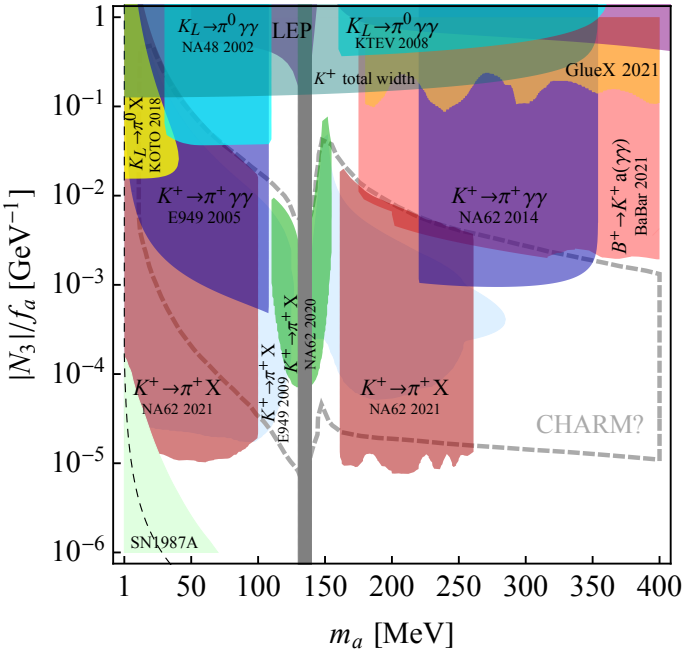

Figure 2.13: Constraints on the Heavy QCD Axion decay constant, f_a, for axions that couple with gluons only from various rare decay searches. The grey band corresponds to the π^0 mass near which perturbative treatment axion production, as performed here, is unreliable. Figure from Reference [138].

upcoming ALPS-II experiment is expected to surpass this [137]. A review of the existing experimental constraints on light axions from astrophysical, cosmological and laboratory searches can be found in Reference [120].

Constraints on Heavy QCD Axions primarily arise from searches for rare decay modes of standard model mesons in beam-dump and collider experiments. In the presence of an axion, new decay modes of these mesons would be possible. Therefore, searches for these rare decay modes place significant constraints on the axion

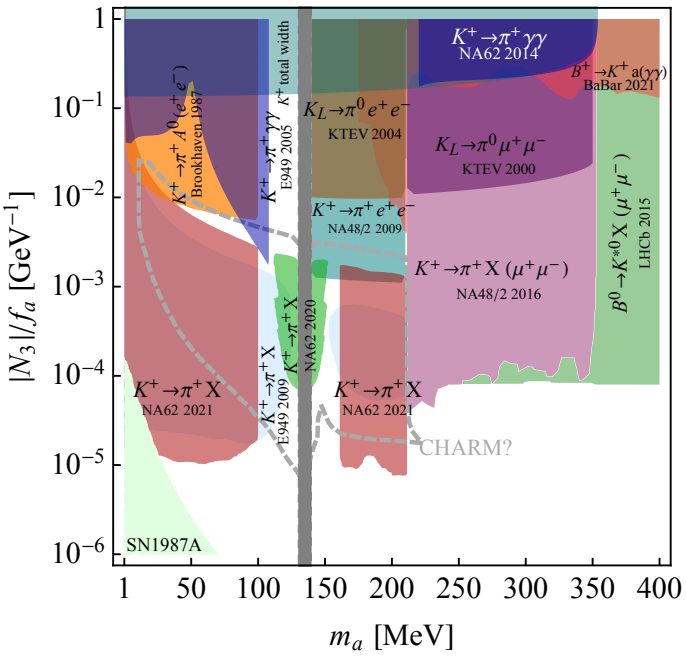

Figure 2.14: Constraints on the Heavy QCD Axion decay constant, f_a, for axions that couple with gluons and standard model leptons from various rare decay searches. The grey band corresponds to the π^0 mass near which perturbative treatment axion production, as performed here, is unreliable. Figure from Reference [138].

coupling strengths. The constraints are model-specific depending on the particular couplings with the standard model particles that are introduced [138]. Constraints on the Heavy QCD Axion decay constant, f_a, from rare decay searches are shown in Figures 2.13 and 2.14 for models where the axions couple with gluons only and with gluons and standard model leptons only, respectively[5]. Significant constraints arise from searches performed by the E949 [139], NA48/2 [140, 141], NA62 [142, 143], KTeV [144, 145], BaBar [146, 147] and LHCb [148] experiments. The constraints from the SN1987A supernova are also shown [109–112]. At lower masses, beam-dump experiments including NA62 and NA48/2 dominate. Towards higher masses beyond the region of sensitivity of these experiments, constraints set by collider experiments including BaBar and LHCb dominate. The constraints from BaBar and LHCb also extend to higher masses than shown in Figures 2.13 and 2.14. A review of the current constraints on various Heavy QCD Axions models can be found in Reference [138].

[5]In these figures the coefficient N_3 of the coupling of the axions to gluons is stated explicitly. More commonly this coefficient is absorbed into the definition of the axion decay constant, $f_a/N_3 \to f_a$. This is the case in Equations 2.7 and 2.8 and will be the case subsequently.

2.3.3 Searches in beam neutrino experiments

The Heavy QCD Axions must couple to gluons in order to solve the strong CP problem. The axions can also couple with the other gauge bosons of the standard model. These considerations motivate the following couplings,

$$\mathcal{L}_{\text{gauge}} = \frac{c_3\alpha_3}{8\pi f_a} a G\tilde{G} + \frac{c_2\alpha_2}{8\pi f_a} a W\tilde{W} + \frac{c_1\alpha_1}{8\pi f_a} a B\tilde{B}. \tag{2.9}$$

Here $G\tilde{G} \equiv (1/2)\epsilon^{\mu\nu\rho\sigma} G_{\rho\sigma} G_{\mu\nu}$ is given in terms of the gluon field strength operator $G_{\mu\nu}$. Couplings to $SU(2)$ and $U(1)$ gauge fields, $W\tilde{W}$ and $B\tilde{B}$, are defined analogously. The coefficients $\alpha_i = g_i^2/(4\pi)$ are given in terms of the three gauge couplings g_i defined at the scale m_a. Two different scenarios are commonly considered: so-called co-dominance, where coefficients $c_3 = c_2 = c_1 = 1$; and so-called gluon dominance, where $c_1, c_2 \ll c_3$ [108, 138]. The co-dominance scenario will be considered hereafter. It should be noted, however, that the model considered in this thesis primarily depends on the $aG\tilde{G}$ coupling and therefore would lead to similar results the gluon dominance scenario. The Heavy QCD Axions can also couple to the standard model fermions via the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) model [149, 150]. Coupling to both quarks and leptons can appear, however in the analysis presented in this thesis axion coupling to standard model leptons only is considered. This enables focusing on a parameter space that is complementary to the multitude of rare meson decay and flavour changing searches [138] and is theoretically well motivated. Therefore, the coupling

$$\mathcal{L}_{\text{lepton}} = \sum_{\ell=e,\mu,\tau} \frac{\partial_\mu a}{2f_a} \left(c_{V\ell} \bar{\ell}\gamma^\mu \ell + c_{A\ell} \bar{\ell}\gamma^\mu\gamma_5 \ell \right), \tag{2.10}$$

is additionally considered. Here $c_{V\ell}$ and $c_{A\ell}$ control the flavor universal vector and axial coupling of the axion to standard model charged leptons. The axions also couple to standard model neutrinos, however the impact of this is negligible for the analysis presented in this thesis.

As a result of the gluon coupling, the Heavy QCD Axions mix with the standard model neutral pseudoscalar mesons. Production of axions could therefore occur through this mixing during high-energy proton–fixed-target collisions in neutrino beams. In these beams, Heavy QCD Axions with masses $\mathcal{O}(100)\,\text{MeV}$ to $\mathcal{O}(1)\,\text{GeV}$ could be produced. These could then decay into signatures containing charged particles observable in a neutrino detector. Given the axion couplings in Equations 2.9 and 2.10, the relevant decay modes in this mass range are into pairs of photons, electrons and muons along with some hadronic states. The relative contributions of these decay modes depends on the axion-lepton coupling, c_ℓ. In general this is a

function of $c_1, c_2, c_3, c_{V\ell}$ and $c_{A\ell}$, but can be approximated as $c_\ell \approx c_{A\ell}$. To determine the axion decay width and hence evaluate the rate of axion decays at a neutrino detector, the parameter c_ℓ needs to be chosen. Two example theory benchmarks are evaluated: $c_\ell \approx c_{A\ell} = 1/36$ and $c_\ell \approx c_{A\ell} = 1/100$. The motivation for these choices are two-fold. Firstly, from the theoretical perspective, a value of c_ℓ somewhat smaller than $c_i = 1$ is motivated because a suppressed leptonic coupling can naturally emerge in models where axions directly couple to some new heavy leptons, which in turn mix with standard model leptons, giving $c_\ell \propto \theta^2_{\text{mix}} \ll 1$. An example of a model of this type can be found in Reference [151]. This constraint is also required in order to be able to treat the axion–pseudoscalar mixings perturbatively. Secondly, from the perspective of experimental axion searches using neutrino detectors, a smaller $c_\ell < c_i$ allows the axion to be sufficiently long-lived to reach a detector before decaying while not suppressing its production via the gluon coupling.

Figure 2.15 shows a summary of the contributions to the axion decay width from various channels for $0.1 \leq m_a \leq 1\,\text{GeV}$, with $f_a \sim 10\,\text{TeV}$. These include the dominant hadronic modes $a \to 3\pi$, $a \to \eta\pi\pi$ and $a \to \pi\pi\gamma$ [152], along with decays to pairs of photons, electrons and muons. The decay widths are shown for the two benchmark scenarios, $c_\ell = 1/36$ (left) and $c_\ell = 1/100$ (right). The axion–pseudoscalar mixings cannot be treated perturbatively close to the standard model pseudoscalar masses m_π, m_η and $m_{\eta'}$ (denoted by the grey bands). These regions of phase-space are therefore difficult to reliably evaluate and are excluded from subsequent analysis. Above $m_a = 2m_\mu$, the decay to muon pairs is dominant except for when the axion mass is close to the standard model pseudoscalar masses. In these regions the diphoton and hadronic modes become significant. Below $m_a = 2m_\mu$, the muon and hadronic decay modes are no longer kinematically allowed. In this region, the diphoton decay mode instead dominates. Across the full mass range, the decay into pairs of electrons is sub-dominant. The axion decay width into pairs of leptons is given by,

$$\Gamma_{a\to ll} = \frac{c_\ell^2 m_a m_l^2}{8\pi f_a^2}\sqrt{1 - \frac{4m_l^2}{m_a^2}}, \tag{2.11}$$

where $l = e, \mu$. Decays into electrons are therefore suppressed relative to muons by a factor m_e^2/m_μ^2 and hence have minimal impact on the total decay width.

The $a \to \mu^+\mu^-$ signature dominates across the majority of the phase space that can be probed by beam neutrino experiments. Additionally, it is a very clean signature in neutrino detectors that is simple to reconstruct and distinguish from backgrounds. This signature will therefore be the focus of the search presented in this thesis. Figure 2.16 shows the behaviour of the axion lifetime (top) and branching ratio to $\mu^+\mu^-$ (bottom) for the two benchmark scenarios. As before, the regions

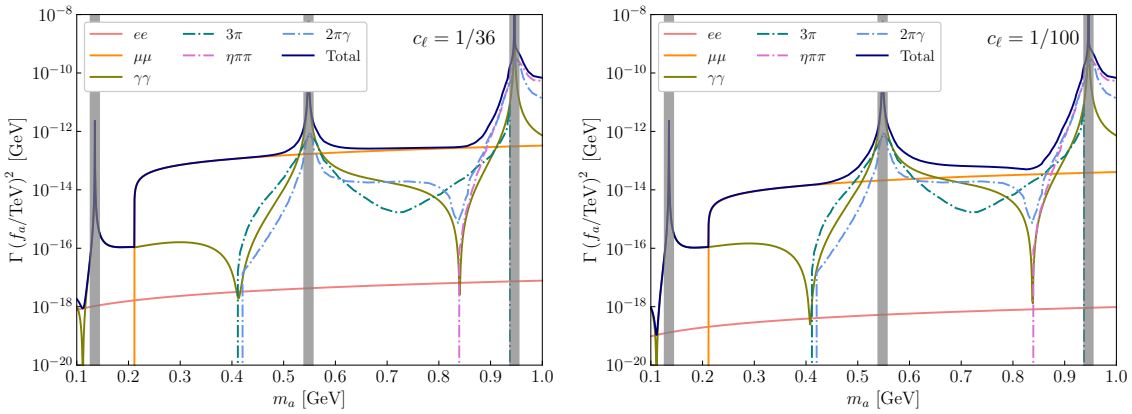

Figure 2.15: Contribution to the Heavy QCD Axion decay width from various modes in the mass range of interest for neutrino experiments. Figure from Reference [153].

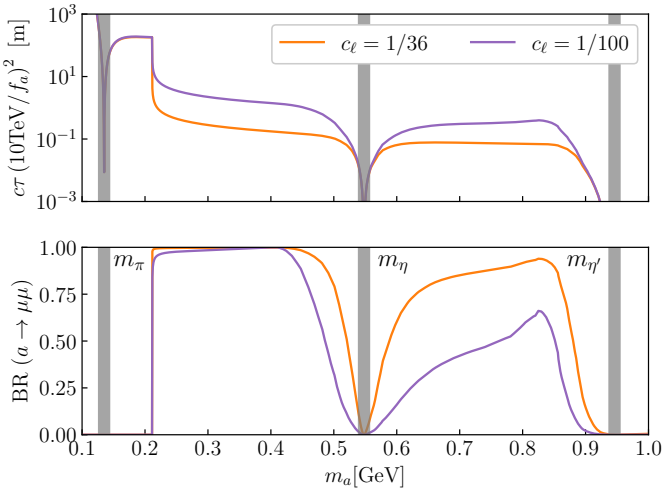

Figure 2.16: Lifetime (top) and $a \to \mu^+\mu^-$ branching ratio (bottom) of the Heavy QCD Axion for the two benchmark scenarios. Figure from Reference [153].

close to the standard model pseudoscalar masses are shown by the grey bands and are excluded. Neutrino detectors at distances of ∼100 m to ∼1 km from the axion production therefore have the potential to be sensitive to this model. In particular, searches using neutrino experiments at short-baselines or using the near detectors of experiments at long-baselines could set competitive constraints.

2.3.4 Existing constraints at scales relevant to neutrino experiments

There are no existing constraints from dedicated searches for the axion model considered in this thesis in the mass range of interest for beam neutrino experiments. However, as discussed in Section 2.3.2, the presence of an axion would enable new decay modes for standard model mesons. Searches for these rare decay modes therefore

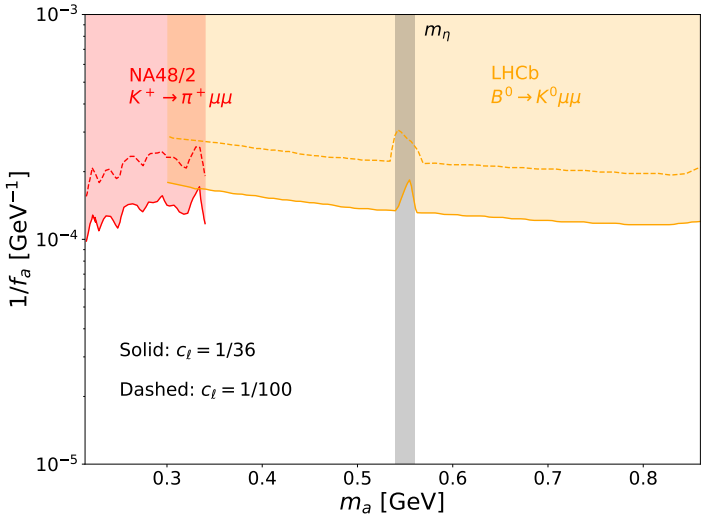

Figure 2.17: Existing constraints on the Heavy QCD Axions at 95% confidence level from NA48/2 and LHCb. The re-cast constraints for $c_\ell = 1/36$ and $c_\ell = 1/100$ are shown by the solid and dashed contours, respectively. The grey band shows the η meson mass, near which the perturbative treatment of the axion production is unreliable.

place important constraints on the parameter space of interest [138]. In the mass range that can be probed in neutrino experiments, significant constraints on the model described in Section 2.3.3 arise from searches performed by the NA48/2 [141], NA62 [143], KTeV [144], BaBar [146] and LHCb [148] experiments. The strongest of these are from the measurements of $K^\pm \to \pi\mu\mu$ by the NA48/2 experiment [141] and of $B^0 \to K^*(892)^0\mu\mu$ by the LHCb experiment [148]. For these cases, the bounds are re-cast as appropriate for the axion model considered in Section 2.3.3. The resulting constraints at 95% confidence level are shown in Figure 2.17 for the two different benchmark scenarios, $c_\ell = 1/36$ (solid lines) and $c_\ell = 1/100$ (dashed lines). The constraints from these experiments are predominantly restricted to short axion lifetimes since the axion must decay promptly after being produced in order to be observed before exiting these fixed-target or forward-spectrometer experiments. The phase space excluded by these searches is therefore different to that which can be probed by beam neutrino experiments. In order to be observed in neutrino detectors the axions are required to propagate $\mathcal{O}(1)$ km from the beam target to the detector before decaying and therefore must be significantly longer lived. These experiments would therefore be sensitive smaller $1/f_a$ than the existing constraints, but would not be sensitive to the large $1/f_a$ regions probed by NA48/2 and LHCb.

Chapter 3

Liquid Argon Time Projection Chamber neutrino detectors

Liquid Argon Time Projection Chambers (LArTPCs) are high precision neutrino detectors first proposed in 1977 by Rubbia [154]. They are an evolution of the Time Projection Chamber (TPC) concept developed by Nygren [155, 156] and the liquid argon ionisation chamber detectors developed by Willis and Radeka [157] in the late 1960s and early 1970s. This chapter will provide a general overview of the LArTPC detector technology along with detailed descriptions of the LArTPC detectors that will be relevant for the work presented in this thesis: ArgoNeuT, the Short-Baseline Near Detector (SBND) and the far detectors of the Deep Underground Neutrino Experiment (DUNE).

3.1 Operational principle

Liquid argon (LAr) provides an ideal medium for neutrino detectors: it is relatively dense ($1.41\,\mathrm{g/cm^3}$) and is chemically inert allowing ionisation charge resulting from interacting particles to be drifted over distances of multiple metres. It is also an excellent scintillator and is transparent to its own scintillation light, providing a prompt signal enabling precise interaction timing. Additionally, argon becomes liquid at approximately 87 K meaning it can be cooled on a large scale using liquid nitrogen which is readily available. These properties, combined with its abundance in the atmosphere and as a result relatively low price, make LAr an excellent choice for large scale neutrino detectors [154], enabling the construction of modules as large as multiple kilotons [1]. Figure 3.1 shows a diagram of a LArTPC detector illustrating the basic operational principle. Interacting neutrinos produce secondary charged particles that ionise and excite the LAr. Ionisation electrons are then drifted by an applied electric field towards multiple planes of sense wires with different orientations

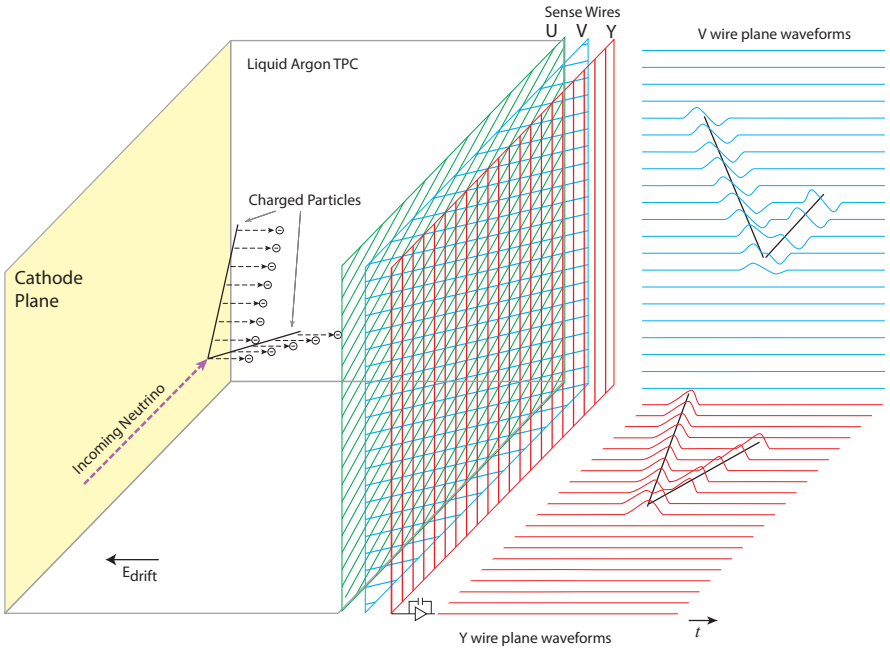

Figure 3.1: Diagram of a LArTPC detector illustrating the basic operational principle. Diagram from Reference [158].

located at the anode of the detector. The interacting secondary charged particles also produce scintillation light. This is detected by photon detectors, such as photomultiplier tubes (PMTs), typically located behind the planes of sense wires. The scintillation light signal is detected much faster than the drifted ionisation charge – $\mathcal{O}$(ns) compared with $\mathcal{O}$(ms). It therefore allows triggering of the detector readout and the interaction time to be determined. The sense wire planes then provide multiple two-dimensional projections of the drifted ionisation charge in space and time, which can be combined enabling high-resolution three-dimensional imaging of neutrino interactions.

3.2 Particle interactions in liquid argon

Charged particles traversing LArTPC detectors produce tracks and showers of ionisation electrons that are drifted towards the sense wires. Neutral particles, such as neutrinos, cannot be detected directly, but can instead be indirectly observed via secondary charged particles resulting from their interactions with the argon.

3.2.1 Tracks: muons, pions and protons

Charged particles such as muons, pions and protons form tracks of ionisation in the detector. An example of a candidate muon neutrino interaction producing track-like energy depositions in the ArgoNeuT detector is shown in Figure 3.2. The figure shows the energy deposited on one of the planes of sense wires, where the y-axis is the time the drifted charge arrives, the x-axis is the wire index and the colour scale represents the amount of charge detected. The interaction produces two track-like energy depositions in the detector: a long, lightly-ionising, muon track and a short, more highly-ionising, proton track.

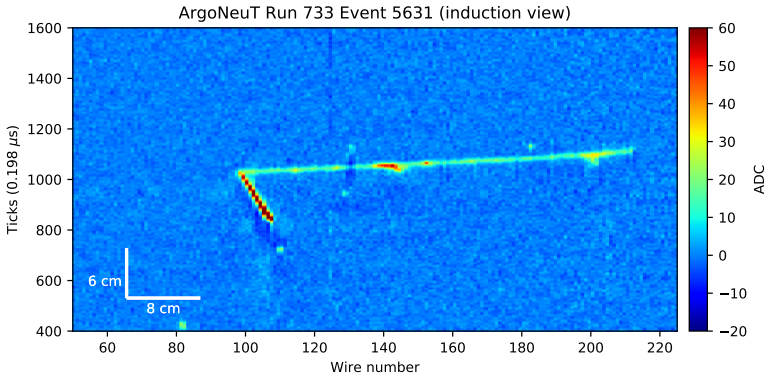

Figure 3.2: Event display showing a candidate muon neutrino interaction in the ArgoNeuT LArTPC detector. The interaction produces two track-like energy depositions in the detector: a long, lightly-ionising, muon track and a short, more highly-ionising, proton track.

For particles with energies typical of those resulting from neutrino interactions in LArTPC detectors, the average amount of energy deposited per unit length along the resulting tracks, dE/dx, is modelled by the Bethe-Bloch equation [27]. Figure 3.3 shows the Bethe-Bloch curves for muons, pions and protons in several different media for moderately relativistic momenta: $0.1 < \beta\gamma < 10000$, where $\beta = v/c$ and γ is the Lorentz factor $\gamma = 1/\sqrt{1-\beta^2}$. The secondary x-axes show the corresponding momentum for each different particle type, which varies due to the differing particle masses. The energy loss depends on both particle type and the detector medium. The average energy loss reaches a minimum at moderate $\beta\gamma$. Particles in this region are referred to as Minimally Ionising Particles (MIPs). At lower $\beta\gamma$, the energy loss rapidly increases. Therefore, as particles slow they become increasingly ionising resulting in a large amount of energy being deposited towards the end of tracks formed by stopping particles. Such a topology is referred to as the Bragg peak.

For accelerator neutrino interactions in LArTPCs, the energy of the resulting particles typically ranges from around a hundred MeV up to several GeV. Therefore, muons and pions are predominantly produced in the minimally ionising region

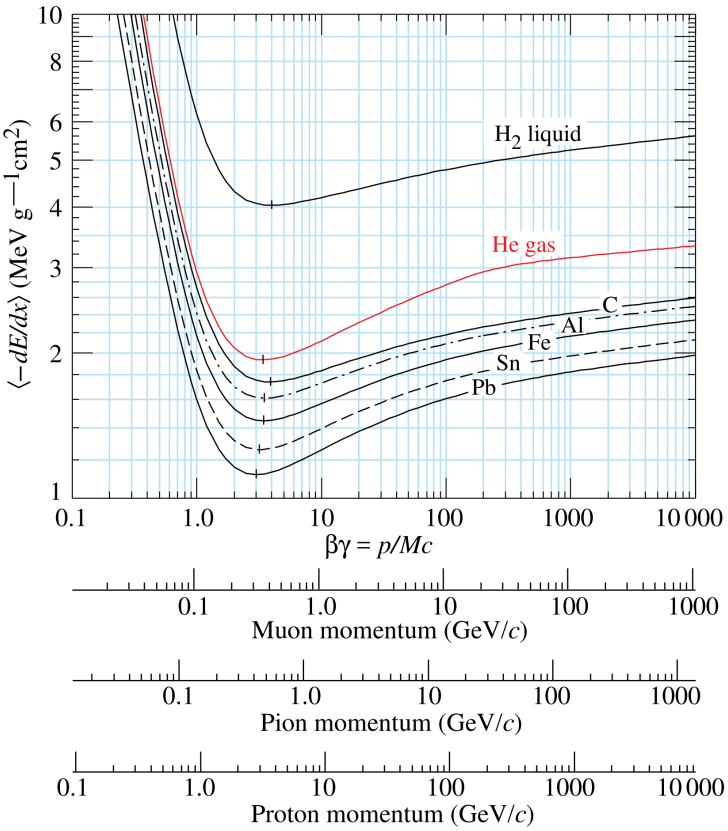

Figure 3.3: Bethe-Bloch curves showing the mean energy loss for muons, pions and protons as a function of momentum for several different media. Different x-axes are shown for the momentum of each different particle type. Figure from Reference [27].

whereas protons are often produced in the steeply rising part of the curve and are more highly ionising. For minimally ionising particles in liquid argon, the average dE/dx is approximately 2 MeV/cm [159]. The energy loss along the paths of the ionising particles is, however, a stochastic process, following a Landau distribution [160]. Particularly large energy transfers can result in electrons being emitted with sufficient energy to produce further ionisation, referred to as knock-on electrons or δ-rays [161]. Examples of this occurring can be seen in several places along the lightly ionising muon track in Figure 3.2 resulting in more highly-ionising segments, appearing as yellow and/or red. Depending on the energy of the emitted δ-ray, it may manifest as only a couple of sense wires seeing a higher dE/dx or, at higher energies, a clearly visible separate track covering many sense wires.

3.2.2 Showers: electrons and photons

Primary electrons and photons form showers of ionisation in the detector. Electrons predominantly lose energy via radiative emission of photons and photons via pair production of electrons and positrons. Provided there is sufficient energy for these

processes to occur (∼ few MeV), they can happen recursively and result in an electromagnetic cascade as illustrated in Figure 3.4.

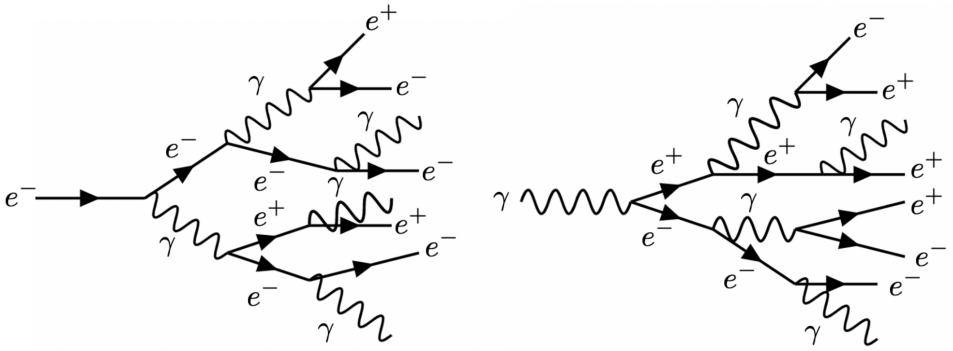

Figure 3.4: Illustration of an electromagnetic cascade initiated by an electron (left) and a photon (right). Figure from Reference [162].

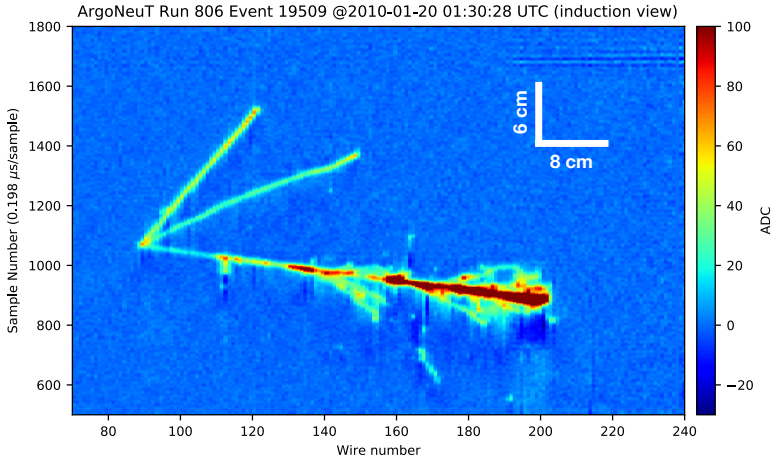

Figure 3.5: Event display showing a candidate electron neutrino interaction in the ArgoNeuT LArTPC detector. The interaction produces an electromagnetic shower from an electron, along with two additional track-like energy depositions potentially from charged pions. Figure from Reference [163].

Electromagnetic cascades are seen in LArTPC detectors as showers of ionisation charge. An example of a candidate electron neutrino interaction in the ArgoNeuT detector is shown in Figure 3.5. An electromagnetic shower is visible originating from a primary electron produced in the neutrino interaction. The interaction also produces two track-like energy depositions, potentially originating from charged pions.

3.3 Ionisation signal

3.3.1 Production

Interacting charged particles ionise the liquid argon, producing e^- and Ar^+ pairs. In liquid argon, the energy required to produce each ionisation pair is 23.6 eV [164, 165]. Therefore, ∼40000 ionisation pairs can be produced for each MeV of energy deposited in the detector, if all of this energy were to go into ionising the liquid argon. In practice, a portion of this energy would also excite the argon resulting in the production scintillation light.

In the absence of an electric field the e^- and Ar^+ pairs recombine resulting in neutral argon atoms and the production of scintillation light. In the presence of an electric field, however, a subset of the ionisation pairs separate before recombination can occur, drifting the e^- towards the planes of sense wires and the Ar^+ away from them. Therefore, the amount of charge that can be detected depends on the strength of the electric field applied [167]. The amount of recombination that occurs also depends on the ionisation density. For more highly ionising particles, such as protons, the density of the produced e^- and Ar^+ is higher resulting in an increased chance for them to recombine [168, 169]. Figure 3.6 shows the dependence between the ionisation charge observed per unit length, dQ/dx, and the energy deposited in the detector per unit length, dE/dx, for different electric field strengths. The relation between the dQ/dx and the dE/dx is described using either the semi-empirical Birks

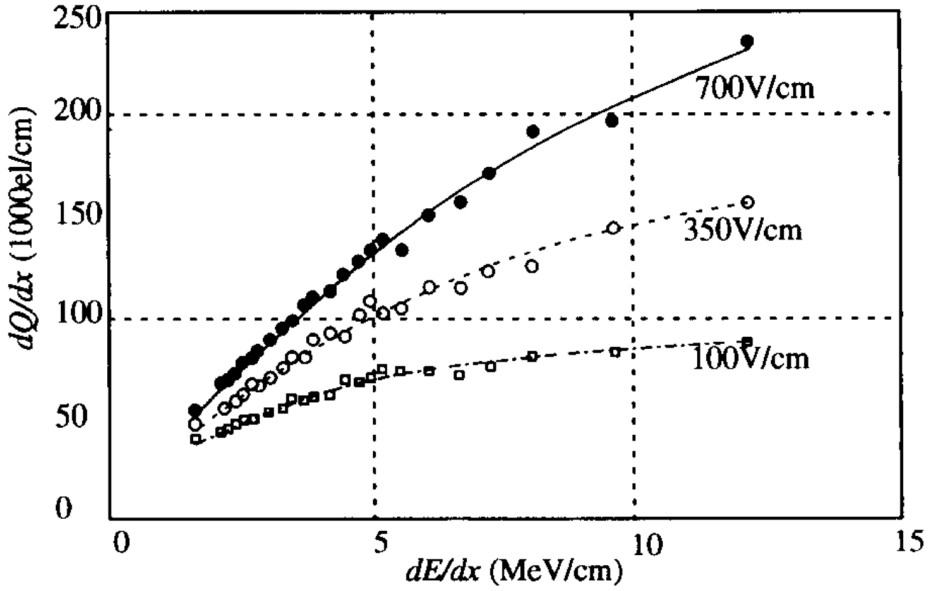

Figure 3.6: Dependence between the ionisation charge observed, dQ/dx, and the energy deposited in the detector, dE/dx, for different electric field strengths. Figure from Reference [166].

model [170] or the modified Box model [169, 171].

3.3.2 Propagation

The drift velocity of the ionisation electrons depends on the applied electric field [165]. At a typical LArTPC electric field of 500 V/cm, the drift velocity is approximately 1.6 mm/μs [172, 173]. Therefore, ionisation electrons can take $\mathcal{O}$(ms) to reach the sense wires in large-scale LArTPCs with drift distances of several metres. The drift can be hindered by the presence of electronegative impurities in the liquid argon, such as oxygen, water and carbon-dioxide, that can attach ionisation electrons suppressing the signal before it reaches the sense-wire planes [174, 175]. The effect on the strength of the observed signal can be modelled as an exponential suppression defined in terms of an electron lifetime. Through sophisticated argon purification systems modern LArTPCs have been able to achieve electron lifetimes as high as 100 ms, corresponding to contamination levels less than 5 ppt oxygen equivalent [176].

During propagation, diffusion of the ionisation electrons can occur resulting in the signal spreading out in space. Ionisation occurring close to the cathode of the detector is most strongly affected, since it has furthest to propagate. The diffusion is non-isotropic due to the presence of the electric field and is typically characterised in terms of its longitudinal and transverse components relative to the drift direction. Longitudinal diffusion results in a broader signal in time seen on individual sense wires [177], whereas transverse diffusion could result in a point charge being spread across multiple sense wires [178].

The drift velocity of the Ar^+ ions is much slower than for the ionisation electrons ($\sim$5 orders of magnitude) [179]. The Ar^+ ions can therefore accumulate in the detector as they are drifting towards the cathode, resulting in non-negligible variations in the electric field. These can distort the propagation of the ionisation electrons, impacting the signals seen by the sense wires. Variations in the electric field also impact the amount of recombination that occurs, altering the number of ionisation electrons that are detected. This is referred to as the space charge effect, and is especially significant for LArTPC detectors located on the surface that are subject to a high flux of cosmic-ray-induced tracks and showers resulting in a large number of Ar^+ ions being produced [180].

3.3.3 Detection

The drifted ionisation electrons produce signals on multiple planes of sense wires located at the anode of the detector. Figure 3.7 shows an illustration of the signals seen on the wire planes. The drifted ionisation charge induces bipolar signals on

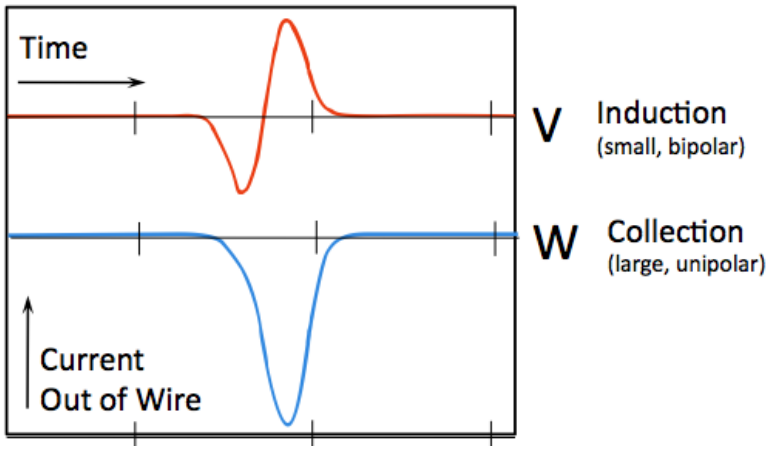

Figure 3.7: Illustration of the signals seen on the induction and collection plane wires. The drifted ionisation charge induces a bipolar signal on the wires of the inner induction plane(s) and is then collected on the wires of the outermost plane resulting in a large unipolar signal. Figure from Reference [172].

the inner plane(s) as it passes through them. These planes are therefore commonly referred to as the induction planes. The charge is then collected on the final plane resulting in a larger unipolar signal. This plane is therefore commonly referred to as the collection plane. In this example, an additional inner-most non-instrumented wire plane is present shielding the induction plane. This prevents a signal being induced as the ionisation electrons are drifted towards the induction plane until the charge passes the shield plane, resulting in a cleaner bipolar signal. In the absence of the shield plane, the signal becomes much more asymmetric with the first part of the signal spread across a much longer time period. Some LArTPC detectors [172, 176, 181, 182] use these shield planes to simplify and improve the ionisation signal reconstruction, however others [158, 183, 184] do not use them in part due to the expense and complexity of manufacturing the additional non-instrumented wire plane.

In order to allow three-dimensional imaging of interacting particles, a minimum of two sense wire planes are required. However, most modern LArTPC detectors use three instrumented planes [158, 176, 182–184]. This allows accurate three-dimensional imaging to be performed even if the interacting particle is travelling parallel to the wires in one of the planes.

Finally, the amount of charge induced or collected on the wires is correlated with the amount of energy deposited in the detector. This allows calorimetry to be performed, provided that recombination and propagation effects are accounted for.

3.4 Scintillation light signal

In addition to the primary ionisation signal, liquid argon is an excellent scintillator [185]. In LArTPC neutrino detectors the scintillation light signal has primarily been used for triggering and rejecting cosmic backgrounds [158, 183]. In experiments searching for dark matter, however, it has been employed much more extensively for energy reconstruction and background rejection [186–188]. Future LArTPC neutrino detectors could make use of the scintillation light in a similar way enabling improved time, calorimetry and position resolution [181, 189, 190].

3.4.1 Production

Liquid argon scintillation light originates from the de-excitation of argon dimer states formed as a result of interacting charged particles. The excited dimers are produced in two different processes due to either excitation or ionisation of argon atoms. In the first of these processes, excited argon atoms can couple with neighbouring ground-state argon atoms to form excited dimers that then de-excite by emitting photons,

$$\mathrm{Ar}^* + \mathrm{Ar} \rightarrow \mathrm{Ar}_2^* \rightarrow 2\,\mathrm{Ar} + \gamma. \tag{3.1}$$

In the second, argon ions combine with non-ionised neighbouring argon atoms to form ionised argon dimers,

$$\mathrm{Ar}^+ + \mathrm{Ar} \rightarrow \mathrm{Ar_2}^+. \tag{3.2}$$

These can then recombine with ionisation electrons forming excited dimers that then similarly de-excite by emitting photons,

$$\mathrm{Ar_2}^+ + e^- \rightarrow \mathrm{Ar}_2^* \rightarrow 2\,\mathrm{Ar} + \gamma. \tag{3.3}$$

Both processes emit vacuum ultra-violet (VUV) photons that have a wavelength spectrum peaking at 128 nm with a width of approximately 10 nm [191, 192].

The argon dimer de-excitation occurs from two states, a singlet $^1\Sigma$ and a triplet $^3\Sigma$. The two components produce scintillation light of approximately the same wavelength, but de-excite at significantly different rates. The time distribution of the emitted photons of each component is shown in Figure 3.8. The two components have lifetimes of $\tau_1 \sim 6 - 7\,\mathrm{ns}$ and $\tau_3 \sim 1.5 - 1.6\,\mu\mathrm{s}$ [193]. These are referred to as the fast (or prompt) and slow (or late) components, respectively. More recent measurements suggest a shorter triplet state lifetime of $\tau_3 \sim 1.3\,\mu\mathrm{s}$, once the impact of the emission time of wavelength-shifting materials used in previous measurements is accounted for [194].

In the absence of an electric field the scintillation yield of liquid argon is ap-

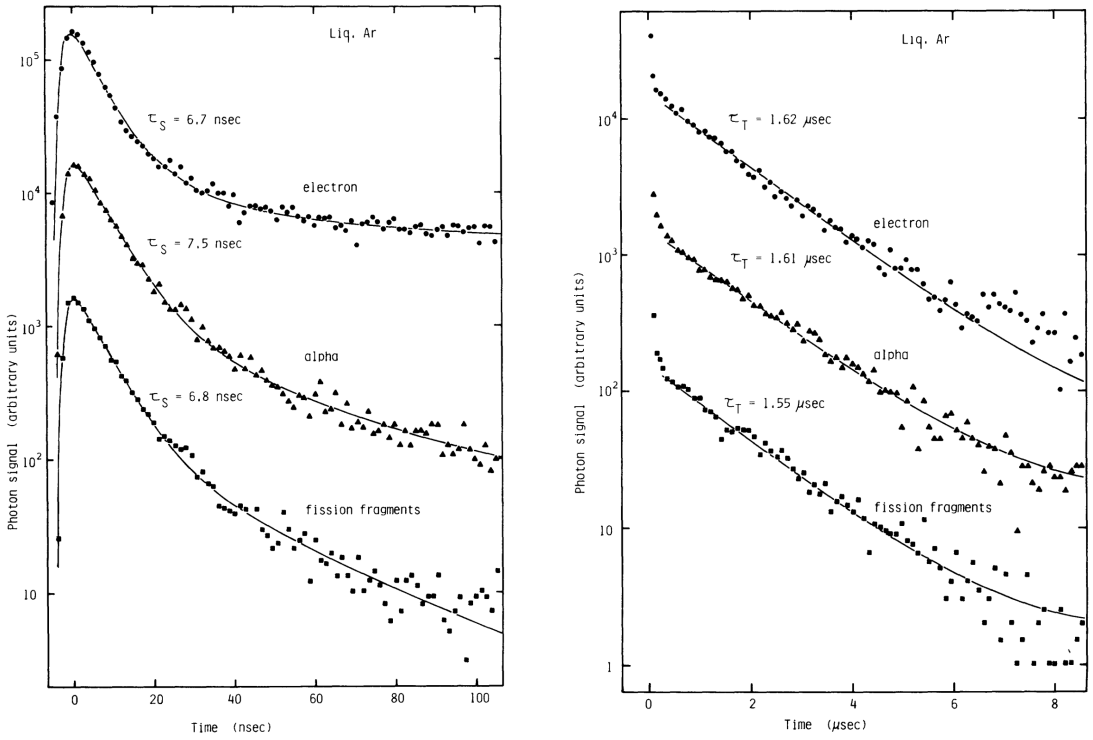

Figure 3.8: De-excitation time distributions of the two components of argon scintillation light from the singlet (left) and triplet (right) excited states of the argon dimer. In each case they are measured using several different particles species. Figures from Reference [193].

proximately 40000 photons per MeV of energy deposited in the detector [185]. The presence of an electric field, however, suppresses the recombination process in Equations 3.2 and 3.3 since a fraction of the free electrons are drifted before recombination can occur [167]. At an electric field of 500 V/cm, a typical field strength used in LArTPCs, the scintillation yield decreases to approximately 20000 photons/MeV [195]. The ionisation density created by an interacting particle also impacts the recombination and hence the amount of light produced. Additionally, at high ionisation density a quenching effect occurs [196, 197]. Therefore, the scintillation yield varies for different particle species [195]. The ratio between the amount of light emitted in the prompt and late components also depends on the ionisation density, and can therefore be used as a method of particle identification [198].

The presence of contaminants in the argon, such as nitrogen [199] or oxygen [174], can cause quenching absorbing the energy of the excited dimer without emitting light. The de-excitation time distributions, particularly of the slow component, can also be affected by these contaminants.

Finally, in addition to the scintillation light, charged particles propagating through LArTPCs produce Cherenkov radiation. The amount of Cherenkov light produced is sub-dominant to the scintillation light and therefore is typically disregarded [200].

However, recent studies have suggested that Cherenkov light may be significant for specific interaction topologies [201] and its impact is currently being explored.

3.4.2 Propagation

Liquid argon is transparent to the scintillation light emitted by the argon dimers and it is able to propagate over long distances. The mean free path of emitted photons is primarily affected by Rayleigh scattering and absorption due to the presence of contaminants. The Rayleigh scattering does not change the number of propagating photons, but alters the path that they take. Depending on the position in the detector and the distance travelled, this can be either detrimental or beneficial for the probability of light arriving at photon detectors. Light that undergoes scattering before arriving at a photon detector travels further than light propagating on a direct path. This leads to a non-trivial distribution of photon arrival times and can result in the apparent lengthening particularly of the prompt component of the scintillation light. The travel time of the scintillation light in large scale LArTPC detectors can therefore range from a few to several tens of nanoseconds. Various measurements and theoretical predictions have reported the Rayleigh scattering length, λ_{RS}, in liquid argon to be between values from around 50 cm [202–204] to as long as 110 cm [205]. The most recently reported measurement found a value of $\lambda_{RS} \sim 100$ cm [206], which has been adopted in simulation by a majority of current LArTPC experiments.

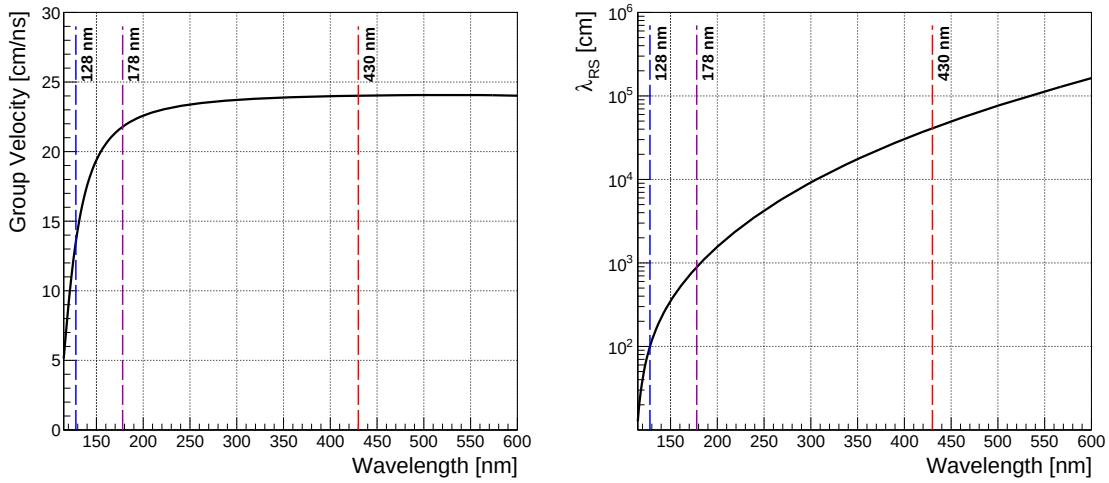

Figure 3.9: Group velocity (left) and Rayleigh scattering length, λ_{RS}, (right) as a function of the wavelength of the photons in liquid argon. Both spectra have been evaluated using the constraint added by the group velocity measurement in Reference [206]. The lines at 128 nm, 178 nm and 430 nm correspond to the wavelength peaks of scintillation emission in liquid argon, xenon-doped liquid argon and from re-emission by the wavelength-shifting compound TPB, respectively. Spectra from Reference [207].

Figure 3.9 shows the group velocity (left) and Rayleigh scattering length (right) as a function of the wavelength of the photons in liquid argon. Both spectra have been evaluated using the constraint added by the group velocity measurement in Reference [206]. Due to the high refractive index of liquid argon at VUV wavelengths [200] the group velocity of photons emitted at 128 nm is about two times slower than that of light at visible wavelengths.

Argon scintillation light can also be absorbed by contaminants with a high cross-section for VUV photons such as nitrogen [208], methane [209] and other elements that have been observed in commercial argon [203]. The total absorption can be modelled as an exponential suppression of the number of photons as a function of the distance travelled with the absorption length as a parameter. The impact that absorption has on the amount of observed light is affected by the Rayleigh scattering length. For shorter λ_{RS} photons may travel on longer, more indirect, paths increasing the possibility of their absorption before reaching a photon detector.

3.4.3 Detection

Detection of the scintillation light in LArTPC detectors requires photon detectors (PDs) that are able to operate at liquid argon temperatures (∼87 K). The primary photon-detector technology used in past and current LArTPC detectors is the Cryogenic Photomultiplier tube (PMT). These have reported Quantum Efficiencies (QE) of up to 30% [210–212]. PMTs, however, are typically large making them difficult to fit efficiently in detectors often resulting in significant non-instrumented regions of the argon volume. Current and future detectors are increasingly shifting towards using Silicon Photomulipltiers (SiPMs). These provide numerous advantages including smaller size, lower power consumption, excellent noise performance at cryogenic temperatures and high QE up to 40% [213]. In LArTPC detectors, light collection using SiPMs is typically enhanced using light-guide bars [214] or light-traps, such as the ARAPUCA [215] or X-ARAPUCA [216] devices.

The 128 nm VUV scintillation photons are predominantly absorbed by the materials used to shield the sensitive area of PDs, typically glass or plastic, resulting in them being challenging to detect. This is usually mitigated by coating the PDs with a wavelength-shifting compound that converts the 128 nm light to the visible spectrum where it is more easily observable [217]. The direction of the re-emitted light is random, hence wavelength-shifter (WLS) coated PDs suffer a ∼50% decrease in efficiency due to light emitted away from the sensitive surface. The most commonly used wavelength-shifting compound in LArTPC detectors is tetraphenyl-butadiene (TPB) [194], although alternatives such as polyethylene naphthalate (PEN) [218] have been proposed.

The wavelength-shifting compounds can also impact the photon detection time since the emission of visible light is not instantaneous. In the case of TPB, multiple time components have been observed. The majority of the light is emitted promptly on a timescale of a few nanoseconds, but a sub-dominant component could have a decay time as long as several microseconds [194]. The electronics and data acquisition chains in LArTPC detectors are usually designed with a resolution on a timescale ranging between 1-2 ns and 100 ns in most recent LArTPC neutrino experiments [158, 181, 182, 219]. The scintillation light signal therefore provides the most precise interaction timing information available in these detectors.

In LArTPC neutrino detectors, the PDs are usually placed behind the planes of sense wires [158, 181–183]. These and other components of the detector can introduce a further decrease in the number of detected photons due to shadowing effects.

3.4.4 Enhancements to the scintillation light signal

Since scintillation light is emitted isotropically, only a subset of the light propagates towards the PDs located at the detector anode. The majority of materials used to construct LArTPC detectors absorb VUV photons, therefore the remainder of the light is typically lost. Additionally, the ∼100 cm Rayleigh scattering length significantly impedes propagation over large distances reducing the total amount of scintillation light detected from regions of LArTPCs far from the PDs. This is especially problematic for large LArTPCs with drift distances of the order of several metres for which the light yield from events close to the cathode may be very low. Two approaches to enhance the scintillation light signal in large-scale LArTPCs have been proposed: xenon-doping, which alters the properties of the emitted light aiding propagation over larger distances [182, 220, 221]; and the addition of wavelength-shifting reflector foils, which allow light that would otherwise be absorbed by the walls of the detector to be recovered [181, 182, 201]. These approaches are not mutually exclusive, and could be used in combination.

Xenon doping

One of the challenges with light collection in large scale LArTPCs is the relatively short Rayleigh scattering length of light in liquid argon (∼100 cm) compared with the distances the light must propagate to reach the PDs, which can be of order several metres. Doping liquid argon with small quantities of xenon (∼10-100 ppm) can significantly alter the properties of the scintillation light, shifting a subset of the emitted photons from 128 nm to 178 nm [220, 222].

In xenon-doped argon, Xe atoms can interact with excited argon dimers, formed

via either of the processes shown in Equations 3.1 and 3.3, to produce excited $ArXe^*$ states,

$$Ar_2^* + Xe \rightarrow ArXe^* + Ar. \tag{3.4}$$

These can then interact with other nearby Xe atoms to produce excited xenon dimers,

$$ArXe^* + Xe \rightarrow Xe_2^* + Ar, \tag{3.5}$$

that can then de-excite emitting photons,

$$Xe_2^* \rightarrow 2\,Xe + \gamma. \tag{3.6}$$

This process results in vacuum ultra-violet photons with a narrow wavelength spectrum peaking at approximately 178 nm [191, 223]. It occurs on a faster time-scale than the argon triplet lifetime τ_3, effectively converting the slow component of the argon light to xenon wavelength emission. This mechanism has been demonstrated to occur with high efficiency even for relatively small quantities of xenon dopant ∼20 ppm [221].

The xenon-wavelength light has several key advantages in large scale LArTPC detectors. The Rayleigh scattering length of 178 nm light in liquid argon is significantly longer than at 128 nm: approximately $\lambda_{RS} \sim 900$ cm, as shown in Figure 3.9. The impact of this on the light yield, the total number of photons detected by the PDs per MeV of energy deposited, is shown in Figure 3.10 (left) for an example LArTPC detector – the DUNE horizontal-drift detector, described in Section 3.6.3. The light yield is shown as a function of the position of the deposited energy in the drift direction, x, where the PDs are located behind the anode at $x \sim 0$ cm and the cathode is located at $x \sim 360$ cm. The blue points show the light yield in pure argon, which is highly non-uniform across the detector volume due to a combination of the geometric acceptance of the PDs and the impact of Rayleigh scattering. The purple points show the light yield in xenon-doped argon, assuming 100% of the slow component of the argon light is converted to xenon wavelength. At short distances Rayleigh scattering boosts the light yield due to light being scattered into the PDs from the surrounding argon. Therefore, the longer Rayleigh scattering length of the xenon-doped argon results in a relative decrease in the light yield close to the PDs. Conversely at larger distances more light is lost than gained due to Rayleigh scattering during propagation, suppressing the light yield. Therefore, the longer scattering length of xenon-doped argon results in a relative increase in the light yield far from the PDs. The reduction in the light yield in the region close to the anode does not negatively impact the uses of the light since the yield remains high, at around ∼22 PE/MeV; but, crucially, the light yield approximately doubles from

~3 PE/MeV to ~6 PE/MeV close to the cathode where it is lowest.

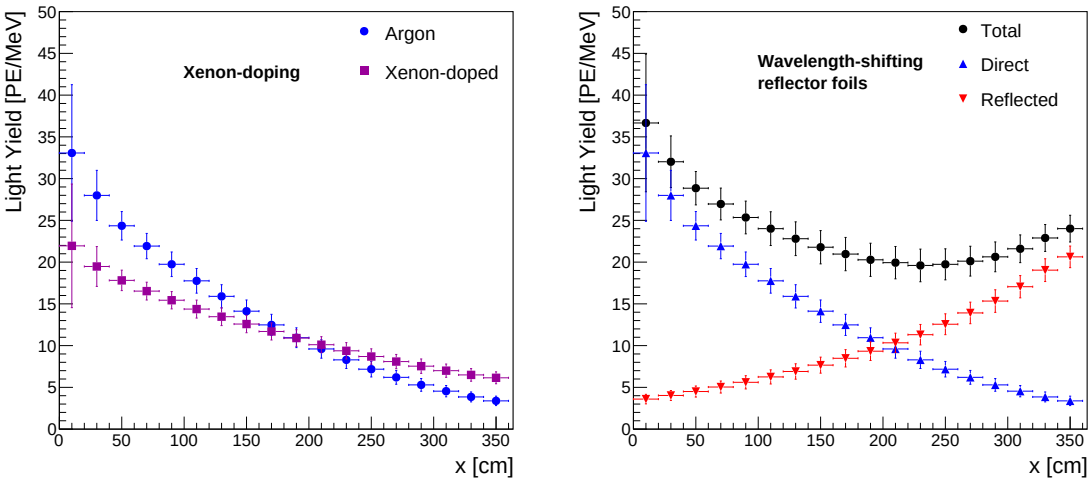

Figure 3.10: Light yield in the DUNE horizontal drift detector as a function of the position in the drift direction (x) for two different scenarios. Left: impact of xenon-doping on the light yield. Right: impact of wavelength-shifting reflector foils on the light yield, subdivided into the contributions from the direct and reflected light. In each case the uncertainties represent the variation in the light yield in the y-z plane due to the locations of the PDs. In both cases, the predictions are made using the semi-analytical light simulation model described in Chapter 6.

In addition to the advantages in light propagation, xenon-doping of liquid argon also results in a higher scintillation yield [222] as a result of the lower average energy required to produce xenon wavelength photons [185, 223]. Furthermore, PDs exist that are able to achieve high efficiency detection of xenon wavelength photons without requiring the light to be shifted to visible wavelengths. Examples include quartz window PMTs [224] and SiPMs [225]. This removes the ~50% loss in efficiency from using a wavelength-shifting compound, significantly improving the observed light yield. However, such PDs are typically more expensive to produce making the cost-benefit trade-off of using these detectors non-trivial.

The time composition of the emitted light is also altered as a result of the xenon doping. Similar to argon, the de-excitation of the xenon dimer occurs from a singlet and a triplet state. Both states de-excite rapidly with lifetimes $\tau_{1,Xe} \sim 4.3\,\text{ns}$ and $\tau_{3,Xe} \sim 22\,\text{ns}$, respectively [220, 223]. Therefore a larger fraction of the light is detected promptly which can be advantageous in some scenarios, such as for triggering purposes. It can also mitigate against the effects of quenching due to contaminants, which predominantly impacts the slow component of argon light. The time distribution of the observed light can, however, be significantly more complex due to the combination of argon and xenon emission, depending on the concentration of xenon present [221].

Xenon-doped argon is currently being explored as the baseline design for the DUNE vertical drift module [226], and as a possible enhancement for the DUNE horizontal drift modules [182]. Recent tests performed by the ProtoDUNE experiments have demonstrated the benefits of xenon-doping in large scale LArTPC detectors [221, 227].

Wavelength-shifting reflector foils

The majority of materials used to construct LArTPC detectors absorb VUV photons causing them to be lost. A method to recover them used primarily in direct dark matter search experiments is to cover the walls of the detector with a highly reflective material, such as Enhanced Specular Reflector (ESR) foils [186] or PTFE [187], coated with a wavelength-shifting compound. These surfaces then become passive elements of the light detection system, wavelength-shifting and then reflecting incident VUV photons back towards the PDs. Neutrino experiments have begun exploring the method of installing WLS-coated reflector foils on the detector cathode. Examples include the LArIAT experiment (field-cage and cathode) [181] and SBND [201], both using TPB. Implementing this solution has also been proposed for the DUNE detectors [182]. Wavelength-shifting in this way can impact the amount of light detected since the efficiency of WLS-coated PDs could be different at visible wavelengths. The photon propagation is also affected since the group velocity and Rayleigh scattering length in LAr are significantly different for photons at visible wavelengths (approximately $\sim$430 nm in the case of TPB [228]), as shown in Figure 3.9.

Figure 3.10 (right) shows the impact of wavelength-shifting reflector foils on the light yield as a function of the position of the deposited energy in the drift direction, x, for an example LArTPC detector – as before, the DUNE horizontal drift detector. The black points show the total light yield, which is subdivided into the contributions from the direct VUV and reflected visible light in blue and red, respectively. The two components are complimentary: close to the anode (and the PDs) the majority of the light is from the direct component, and close to the cathode (and the foils) the majority of the light is from the reflected component. The combination results in greatly improved light yield and light yield uniformity throughout the detector volume – with a yield of $\sim$20 PE/MeV at worst, compared with $\sim$3 PE/MeV without the reflected component. However, the reflected light is typically much more diffuse than the direct light: the initial VUV light is emitted isotropically and undergoes Rayleigh scattering during propagation to the reflector foils, then the wavelength-shifted light is re-emitted approximately isotropically and has to propagate across the full width of the detector before reaching the PDs. This results in the reflected light typically being spread across a greater number of PDs than the direct light,

smearing positional information and potentially making it more challenging to use in analyses.

TPB-coated wavelength-shifting reflector foils are typically produced using low-temperature evaporation to coat the reflector material. The method is described in detail for the SBND detector in Reference [201]. This method, however, is labour-intensive and difficult to scale to larger detectors, such as DUNE, primarily due to the vacuum required during the evaporation process. Alternative coating techniques such as solvent-based coating or embedding TPB in a polymer have been found to perform significantly worse [217]. A potential alternative is PEN, which is available in large polymeric sheets that could be used to laminate the reflector foils allowing production on a much larger scale [218]. However, PEN has been measured to convert 128 nm light to visible wavelengths with approximately 50% efficiency relative to TPB [229] reducing the efficacy of the reflector foils.

3.4.5 Generic model for the behaviour of scintillation light photons

Incorporating all the effects described in the previous sections, the number of photons detected by a given photon detector (PD) from an energy deposit ΔE at position (d, θ) can be expressed as:

$$\begin{aligned} D_\gamma = &\Delta E \times S_\gamma(\mathscr{E}) \times Q_q \times Q_{abs}(d) \times Q_{det} \times \\ &Q_{trans}(\theta) \times P(d, \theta) \times T(d, \theta), \end{aligned} \tag{3.7}$$

where the scintillation yield $S_\gamma(\mathscr{E})$ is the number of photons emitted per unit of deposited energy at an electric field $\mathscr{E}$. The position (d, θ) is defined in terms of the distance, d, between the energy deposit and the PD, and the offset angle, θ, between the energy deposit and the normal to the PD surface. The Q_x parameters account for the various effects during propagation and detection that reduce the number of photons observed: Q_q is the quenching at emission, Q_{det} is the PD efficiency, $Q_{abs}(d)$ is the loss due to absorption effects and depends on d, and $Q_{trans}(\theta)$ represents the loss of transmission due to shadowing effects and depends on θ. All of these parameters have values in the range from 0 to 1. $P(d, \theta)$ is the geometric coverage of the PD and $T(d, \theta)$ signifies other transport effects such as Rayleigh scattering, both of which depend on d and θ. Many of the parameters in Equation 3.7 are independent of each other and their order of application during modelling does not have an impact. The parameters P and T, however, are interconnected and, as a result, more complicated to model. The parameters Q_{abs} and Q_{trans} can also be impacted by T.

The time at which a photon is detected by a particular PD can be obtained by summing together the independent components resulting from the different processes that photons undergo,

$$t_\gamma = t_E + t_t(d, \theta) + t_{WLS} + t_{det}. \tag{3.8}$$

The emission time, t_E, represents the argon dimer de-excitation times accounting for any quenching effects, $t_t(d, \theta)$ is the transport time resulting from the different paths the photons can take to arrive at the PD, t_{WLS} is the time resulting from the intrinsic relaxation time of the wavelength-shifting compound and t_{det} represents the PD and electronics response time.

This framework can also be adapted to describe the enhancements to the scintillation light signal discussed in Section 3.4.4. Equations 3.7 and 3.8 can be used to model the contributions from each sub-component of the light treating them independently: the light at argon and xenon wavelengths in the case of xenon-doping; and the direct and reflected light in the case of wavelength-shifting reflector foils. For each sub-component the various parameters of Equations 3.7 and 3.8 may be different. Additionally, in the case of wavelength-shifting reflector foils, a further scaling factor $Q_r = Q_{WLS} \times Q_{foil}$ is required to model the WLS efficiency, Q_{WLS}, of the foils and their reflectivity, Q_{foil}.

3.5 Charge-light anti-correlation

The amount of ionisation charge and scintillation light produced by interacting charged particles both depend on the strength of the electric field applied, as discussed in Sections 3.3 and 3.4. For the charge the presence of an electric field results in a subset of the ionisation electrons being separated from Ar^+ ions before recombination can occur, allowing them to be drifted towards the sense-wire planes. The amount of ionisation charge that can be detected in a LArTPC detector therefore increases in the presence of a stronger electric field. Conversely, scintillation light is produced during the recombination process. The presence of an electric field reduces the rate of recombination, hence reducing the amount scintillation light. The ionisation charge and scintillation light signals are therefore anti-correlated [167]. Figure 3.11 shows the ionisation charge (Q) and scintillation light (L) anti-correlation as a function of the applied electric field for several noble elements – argon, xenon and krypton. LArTPC detectors are typically designed with an electric field of around 500 V/cm [1, 158, 172, 181, 183], approximately where the amount of energy transferred to ionisation electrons and scintillation light is equal.

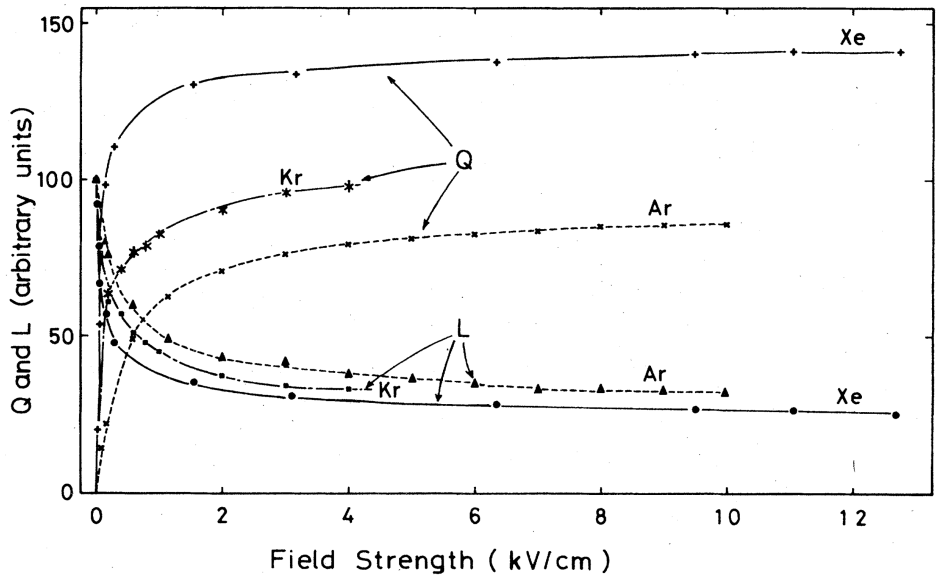

Figure 3.11: Charge (Q) and light (L) anti-correlation for argon, xenon and krypton as function of the applied electric field. Figure from Reference [167].

3.6 Overview of selected LArTPC detectors

Several LArTPC neutrino detectors have been built in the last two decades, and several more are currently under construction or have been proposed. In this section the detectors that are relevant for the work presented in this thesis will be described: ArgoNeuT (and the MINOS near detector), SBND, and the DUNE horizontal and vertical drift far detectors.

3.6.1 ArgoNeuT and the MINOS near detector

ArgoNeuT was a 0.24 ton LArTPC neutrino detector located along the 120 GeV Neutrinos at the Main Injector (NuMI) beam [230] at Fermilab. Its primary physics run lasted approximately six months from September 2009 - February 2010, collecting 1.335×10^{20} protons-on-target (POT) with the NuMI beam in so-called low-energy mode ($\langle E_\nu \rangle \sim 4\,\text{GeV}$). Of this, 0.085×10^{20} POT was in neutrino-mode (forward horn current) and 1.25×10^{20} POT in anti-neutrino mode (reverse horn current). In total, ArgoNeuT collected approximately 900 ν interactions running in neutrino-mode and approximately 4000 ν and 3500 $\bar{\nu}$ interactions running in anti-neutrino mode [172]. ArgoNeuT was the first LArTPC to run in a neutrino beam in the US and was primarily designed as a test experiment in the phased path towards the construction of future multi-kiloton scale LArTPCs. ArgoNeuT's data-set, however, has been used to produce numerous significant physics results. These include multiple first ν-argon cross-section measurements, for both ν_μ [231–235] and ν_e [163];

and beyond the standard model searches, including the first search for millicharged particles in a LArTPC [236].

The MINOS near detector (MINOS-ND) was located immediately downstream of the ArgoNeuT detector. ArgoNeuT used the MINOS-ND as a muon spectrometer, matching tracks exiting ArgoNeuT with those reconstructed downstream in the MINOS-ND. This section will give an overview of both the ArgoNeuT and the MINOS-ND detectors. More detailed descriptions can be found in References [172] and [237], respectively.

The ArgoNeuT detector

The ArgoNeuT detector was located 100 m underground in the MINOS near detector hall, 1033 m downstream of the NuMI beam target and immediately upstream of the MINOS-ND. Figure 3.12 (left) shows a photograph of the ArgoNeuT detector (orange box) with the MINOS-ND visible behind it (blue structure). Figure 3.12 (right) shows a photograph of the ArgoNeuT time projection chamber (TPC), prior to it being inserted into the cryostat (seen behind it). The cryostat was then housed in the orange box seen in the previous image.

Figure 3.12: Left: photograph of the ArgoNeuT detector structure (orange box), with the MINOS-ND behind it (blue structure). Right: photograph of the ArgoNeuT TPC prior to it being inserted into the cryostat, seen behind it. Images from Reference [172].

A diagram of the ArgoNeuT TPC is shown in Figure 3.13. The instrumented volume of the detector was $47 \times 40 \times 90\,\mathrm{cm}^3$ (horizontal (drift), vertical, beam direction), encompassing approximately 170 L of liquid argon. An electric field of

E_d = 481 V/cm was applied between the anode and the cathode, drifting electrons from interacting charged particles towards three planes of sense wires. The first (inner-most) plane contained 225 wires, orientated vertically with 4 mm wire-spacing. This plane was not instrumented, but instead shaped the electric field near the wire-planes and shielded the first instrumented plane from induced signals as the charge was drifted. The second (middle) and third (outer-most) wire planes both contained 240 wires, orientated at $\pm 60^{\circ}$ to the horizontal and with 4 mm wire spacing. For both instrumented planes, the signals were digitised at a rate of 5 MHz by the readout electronics [172].

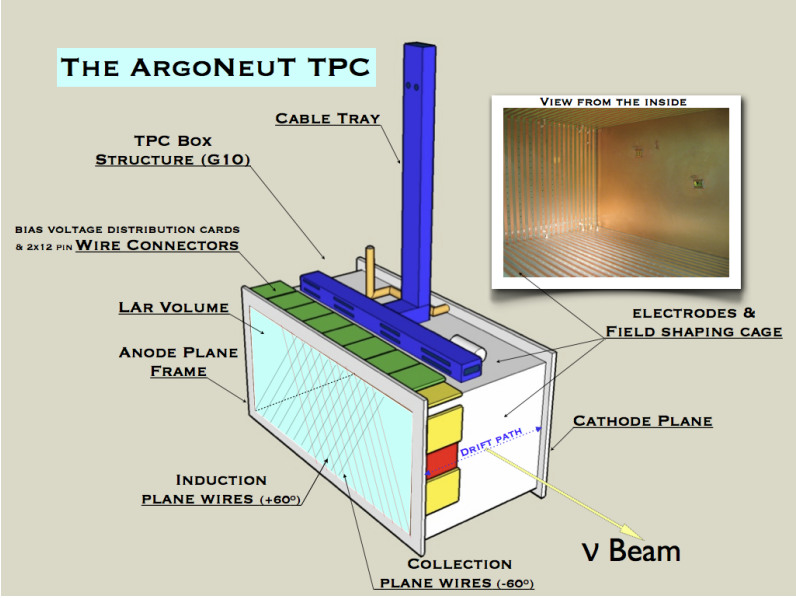

Figure 3.13: Diagram of the ArgoNeuT TPC showing the detector structure and the two instrumented wire planes. The inset shows a photograph of the inside of the detector where the field-shaping cage and cathode are visible. Diagram and image from Reference [172].

Unlike more modern LArTPCs, ArgoNeuT did not have a photon detection system and therefore did record use the scintillation light signals. Instead, the most precise timing information available was provided by the NuMI beam spill signal and ArgoNeuT's readout was triggered for every beam spill [172].

The MINOS near detector

The MINOS near and far detectors were magnetised sampling calorimeters consisting of alternating steel plates and plastic scintillator strips. Only the MINOS near detector is used by ArgoNeuT and will be described here. Figure 3.14 shows a photograph of the MINOS-ND in the MINOS near detector hall prior to ArgoNeuT's installation. The MINOS-ND consisted of 282 squashed-hexagonal steel planes each 3.8 m high, 4.8 m wide and 2.54 cm thick. The instrumented planes were affixed with

1 cm thick plastic scintillator strips and each steel-scintillator plane was separated by a 2.54 cm air gap. The total length of the MINOS-ND was 16.6 m [237].

Figure 3.14: Photograph of the MINOS-ND in the MINOS near detector hall at Fermilab prior to ArgoNeuT's installation. The squashed-hexagonal shaped steel planes that make up the detector can be seen, along with the magnetic coil entering the upstream face. Image from Reference [238].

Figure 3.15 (left) shows a top-down diagram of the MINOS-ND. The detector was divided into two main regions: the upstream calorimeter region consisting of 120 planes each of which was either fully or partially instrumented and the downstream muon spectrometer region consisting of 162 planes, of which each 5^{th} plane was instrumented. Figure 3.15 (right) shows a front-on cross-sectional diagram of the MINOS-ND. The beam (black cross) was centred to the left of the magnetic coil (black square), in the centre of the fully instrumented calorimeter region of the detector (shaded grey). The ArgoNeuT detector was aligned with the beam centre. The shape of the partial scintillator planes (shaded grey) and full scintillator planes (dashed lines) can be seen overlaid on the squashed-hexagonal steel planes.

The scintillator planes were arranged in four configurations each shown in Figure 3.16: partial and full views; and alternating between "U" and "V" orientations. To cover the different required shapes, there were 16 different types of readout modules each containing either 14, 20, or 28 scintillator strips. Throughout the entire detector, every 5^{th} steel plane was instrumented with a full view configuration, alternating between U and V orientations. Then for the upstream calorimeter region of the detector, the four planes between each full view were instrumented with a partial view configuration, again alternating between the U and V orientations [237]. Each scintillator plane provides a two-dimensional view of an interaction, and pairs of planes can be used to perform three-dimensional reconstruction

The MINOS detectors were magnetised, allowing the momentum and charge of

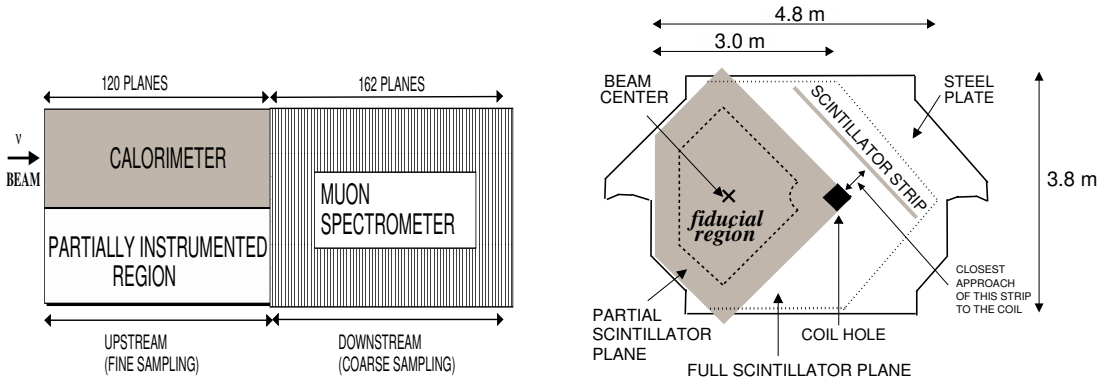

Figure 3.15: Diagrams of the MINOS-ND. Left: top-down view showing the calorimeter and muon spectrometer regions of the detector. Right: front-on cross-sectional view showing the location of the beam centre and the coil. ArgoNeuT was aligned with the beam centre. Diagrams from Reference [239].

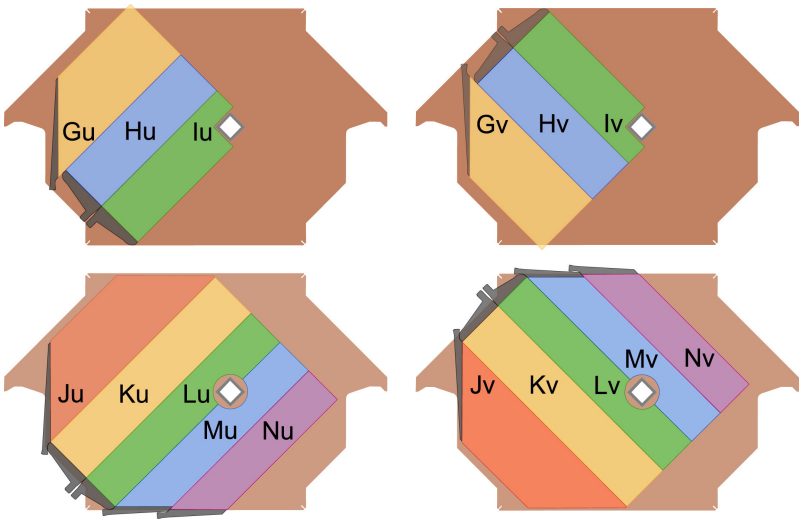

Figure 3.16: Diagrams of the four scintillator strip configurations in the MINOS-ND overlaid on the squashed-hexagonal shape of the steel planes. The partially instrumented planes are shown on the top and fully instrumented planes shown on the bottom. In each case the orientation of the strips alternates between the "U" planes on the left and the "V" planes on the right. Diagrams from Reference [237].

interacting particles to be determined by their curvature. The MINOS-ND had a toroidal magnetic field produced using a 40 kA-turns current-carrying coil running the full length of the detector. The coil location is shown in Figure 3.15 (right) (black square). It was off-set by 55.8 cm to the right of the centre of the detector and provided an average magnetic field strength of 1.28 T across the typical fiducial region used in MINOS-ND analyses (shown by the dashed lines on the partial scintillator plane) [237].

3.6.2 Short-Baseline Near Detector

The Short-Baseline Near Detector (SBND) is the near detector of the Short-Baseline Neutrino (SBN) program at Fermilab [3]. The SBN program consists of three LArTPC neutrino detectors located along the 8 GeV Booster Neutrino Beam (BNB) [240]: SBND as the near detector, MicroBooNE [158], and a refurbished ICARUS [183] as the far detector. The SBN program will perform precision short-baseline neutrino oscillation measurements, in particular searching for evidence of eV-scale sterile neutrinos. These light sterile neutrinos could resolve a series of long-standing anomalies seen in short baseline neutrino oscillation data observed by the LSND [28] and MiniBooNE [29–31] experiments. Although, the most popular explanations for these anomalies have recently been disfavoured by the MicroBooNE experiment [64, 65]. In addition, the SBND detector will collect the highest statistics sample of neutrino interactions in argon to date, observing approximately 2 million neutrinos per year. This will allow it to perform high precision ν-argon cross section measurements that will be critical for the future SBN and DUNE oscillation measurements. It will also have the potential to explore a wide range of BSM physics models [3, 241, 242].

SBND is a 112 ton LArTPC currently under construction at Fermilab that will be located along the BNB beam line 110 m downstream of the target. A diagram

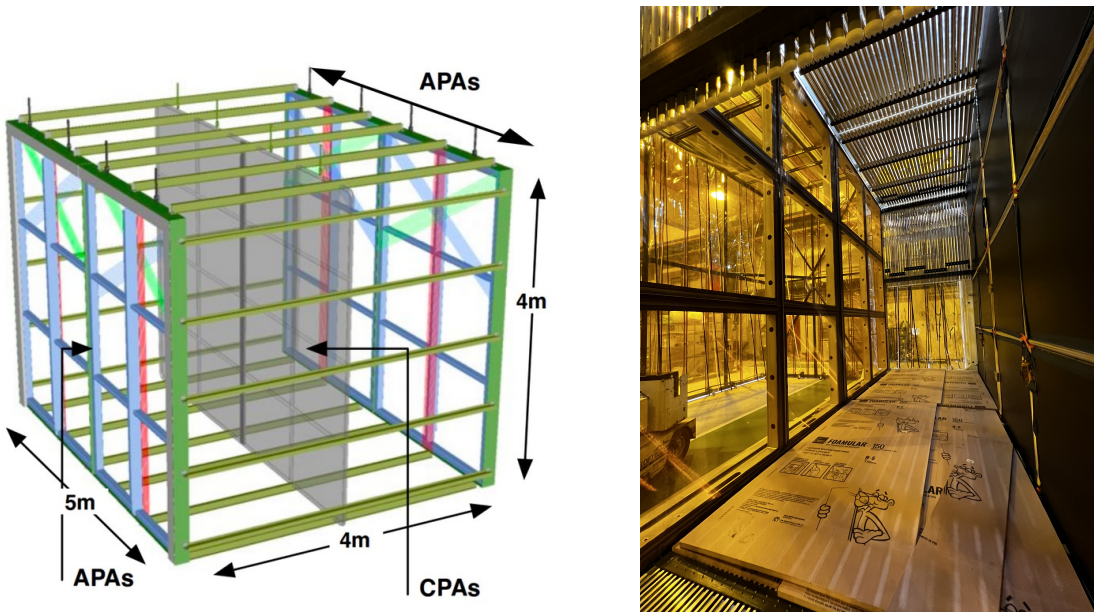

Figure 3.17: Left: diagram of the SBND detector showing the anode plane assembly (APA) planes of each drift volume and the central shared cathode plane assembly (CPA) plane. Right: photograph of the SBND detector currently being assembled at Fermilab. One of the APA planes is visible on the left with all three wire planes present, and the central CPA plane is visible on the right with the reflector foils in place with protective covers. Part of the field cage is visible at the top. Diagram from Reference [184]. Photograph from Reference [243].

of the detector is shown in Figure 3.17 (left) showing the two anode plane assembly (APA) planes and the central cathode plane assembly (CPA) plane. Figure 3.17 (right) shows a photograph of the detector being assembled. One of the APA planes is visible on the left, and the central CPA plane on the right. The detector will consist of two drift volumes, each $2 \times 4 \times 5\,\mathrm{m}^3$ (horizontal (drift), vertical, beam direction) and held at an electric field of $E_d = 500\,\mathrm{V/cm}$. The electric field will drift the ionisation charge towards the two APA planes, each of which have three planes of sense wires – two induction and one collection. The two inner induction planes,"U" and "V", are orientated at $\pm 60^\circ$ to the vertical and the final collection plane, "Y", is vertical. The wire spacing in all three planes is 3 mm and each plane is separated by 3 mm [184].

The primary photon detection system (PDS) consists of 120 8-inch PMTs: 96 TPB-coated and sensitive to both VUV and visible light; and 24 uncoated, sensitive to visible light only. The PMTs will have a timing resolution of 1-2 ns, allowing precise determination of interaction times. In addition, there will be a secondary PDS consisting of 192 X-ARAPUCA light trap modules [216] 50% of which will be sensitive to VUV light only and 50% to visible light only. In both cases, the X-ARAPUCA modules will have sensitive window areas of $19.8 \times 7.6\,\mathrm{cm}^2$. The X-ARAPUCAs will serve as a prototype for the DUNE PDS, and their mechanism of operation is described in more detail in Section 3.6.3. The PDS is subdivided into 24 PDS boxes that slot into the APA frames behind the sense wires, with 12 boxes per

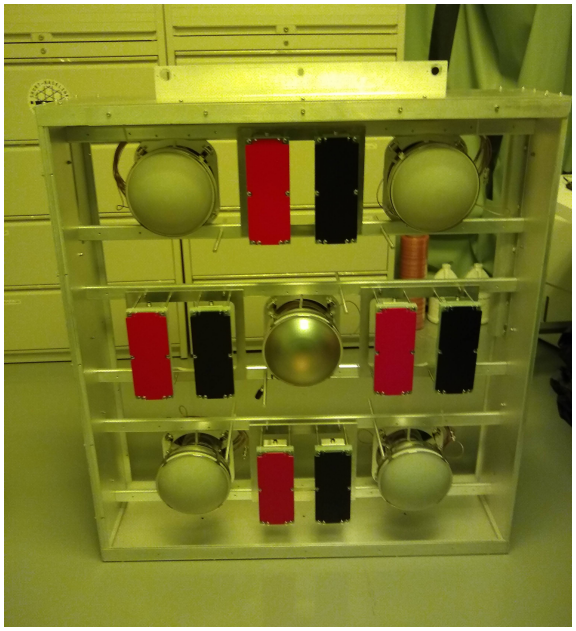

Figure 3.18: Photograph of an assembled SBND photon detection system box. Each box contains 5 PMTs, 4 TPB-coated and 1 uncoated, and 8 X-ARAPUCAs, 4 VUV sensitive (black covers) and 4 visible sensitive (red covers). Photograph from Reference [244].

APA plane. Figure 3.18 shows a photograph of an assembled PDS box. Each box contains 5 PMTs, 4 TPB-coated and 1 un-coated, and 8 X-ARAPUCA modules, half VUV sensitive and half visible sensitive. Additionally, SBND will make use of passive wavelength-shifter coated reflector foils, as described in Section 3.4.4, to boost the light response yield and uniformity. The reflector foils are coated with TPB on both sides and are inserted into the CPA frames, providing approximately 80% coverage of the cathode plane for both TPCs [201].

SBND will be located on the surface and will therefore experience a large flux of particles resulting from cosmic-ray showers. To mitigate their impact, the SBND detector will have a Cosmic-Ray Tagger (CRT) system [245]. The CRT panels consist of strips of plastic scintillator with wavelength-shifting optical fibres guiding the light towards SiPMs. They will cover the top, bottom and all four sides of the detector allowing background cosmic-ray muons to be tagged and rejected in analyses. The CRT also allows samples of muons to be selected which may be useful for calibrations.

At the time of writing, SBND is being assembled at Fermilab. Construction is expected to be finished late 2022, with data taking commencing in 2023 following the completion of installation and commissioning.

3.6.3 Deep Underground Neutrino Experiment

The Deep Underground Neutrino Experiment (DUNE) is a future multi-kiloton-scale long-baseline LArTPC neutrino experiment. It will consist of a near detector at Fermilab and a far detector located 1300 km downstream at the Sanford Underground Research Facility (SURF). The DUNE detectors will be exposed to a broad-band neutrino beam from the 60-120 GeV Main Injector with the PIP-II [246] upgrade. This will deliver the highest intensity neutrino beam to date. The primary physics goal of DUNE will be to perform long-baseline neutrino oscillation measurements, measuring the neutrino mixing parameters with unprecedented precision. This will allow it to probe matter-anti-matter asymmetries by measuring the charge-parity violating phase, δ_{CP}, in the neutrino sector and to determine the neutrino mass hierarchy [2, 247]. In addition to its oscillation programme, the DUNE detectors will be able to perform precision measurements of astrophysical neutrinos including solar neutrinos [248] and supernova bursts [249]. Furthermore, the DUNE detectors will be sensitive to a wide range of BSM physics models [250], including for example Heavy Neutral Leptons [85] and Heavy QCD Axions [108]. A detailed overview of the DUNE physics programme can be found in Reference [2]. The DUNE far detector will consist of four 10 kt instrumented volume modules. The first and third modules will follow the single-phase horizontal drift (DUNE-HD) design [182]. The

second module will follow the single-phase vertical drift (DUNE-VD) design [251], which is an evolution of the previous dual-phase design [252]. The DUNE-HD and DUNE-VD detectors will be described in greater detail below, with a particular emphasis on the photon detection systems. The design of the fourth and final module has not yet been decided. The DUNE near detector will consist of a multi-detector design with a modular LArTPC (ND-LAr), a magnetised high-pressure gaseous argon TPC (ND-GAr) and a magnetised beam monitor [253]. At the time of writing, the DUNE detectors and beam-line are currently under construction, with the first module expected to begin taking data in the mid to late 2020s. The subsequent modules are expected to follow in the late 2020s and early 2030s, with the full detector taking data throughout the 2030s. Large-scale prototype detectors have also been constructed for the single-phase horizontal drift design (protoDUNE-SP) [176] and the dual-phase design (protoDUNE-DP) [252] which started taking data in 2018 and 2019, respectively. Further prototypes are planned for the single-phase vertical drift design [251].

The DUNE horizontal drift module

The DUNE horizontal drift (DUNE-HD) module will have an instrumented volume of $14.0 \times 12.0 \times 58.2\,\mathrm{m}^3$ (horizontal (drift), vertical, beam direction), with a total fiducial mass of approximately 10 kt. This will be divided into four 3.5 m drift volumes by five rows of alternating APAs and CPAs running along the length of the detector, as illustrated in Figure 3.19 (left). Each drift volume will be held at an electric field of $E_d = 500\,\mathrm{V/cm}$, drifting ionisation charge towards rows of APAs. Each APA will have 4 planes of sense wires: the innermost plane "G" will be orientated vertically and serve as an non-instrumented shielding plane; then there will be two induction planes "U" and "V" orientated at $\pm 35.7^{\circ}$to the vertical; and finally an outermost collection plane, "X", that will also be oriented vertically. In all four planes, the wires will have a spacing of 4 mm and each wire plane will be separated by approximately 5 mm. Figure 3.19 (right) shows a photograph of a completed APA with all four wire planes at the Daresbury APA production factory in the UK. Note that the APA is on its side, relative to its orientation in the detector. In the full module there will be three rows of APAs each 2 APAs high (12 m) and 25 APAs deep (58.2 m), for a total of 150 APAs. The APAs on each side of the detector collect ionisation charge from a single drift volume while the APAs in the centre collect charge from drift volumes on each side. This is achieved by the wires being wound around both sides of the APA frames [182].

In order to maximise the active volume of the detector, the photon detection system (PDS) is restricted to the inactive space between the innermost wire planes inside the APA frames. This heavily limits the space available for the photon detec-

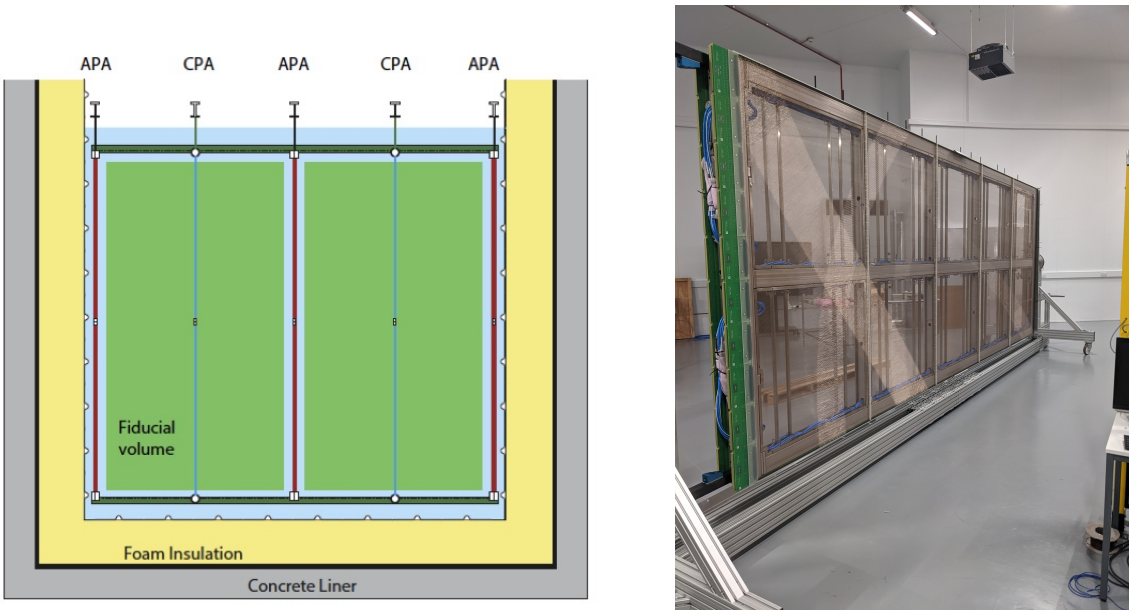

Figure 3.19: Left: diagram of the DUNE-HD module showing an end-on view of the rows of APAs and CPAs that form the four drift volumes. Each row will be 2 APAs high (12 m) and 25 APAs deep (58.2 m). Right: photograph of one of the completed DUNE-HD APAs at the Daresbury APA production factory in the UK. The APA is on its side. Diagram from Reference [182].

tors (PDs). Additionally, the PDs in the central APA must be sensitive to both sides to allow light collection from the two drift volumes. The PDS consists of 10 light collection modules per APA, each with dimensions $2.3 \times 11.8 \times 209.7\,\mathrm{cm}^3$, that slot into the APA frames as illustrated in Figure 3.20 (left). Each PDS module consists of 24 X-ARAPUCA [216] PDs with $7.3 \times 9.8\,\mathrm{cm}^2$ sensitive windows, grouped into four supercells. In the full horizontal drift module there will be a total of 1500 PDS modules, 500 double-sided and 1000 single-sided [182].

The ARAPUCA PD traps incident photons using a dichroic filter coated on each side with two different wavelength shifting compounds, as illustrated in Figure 3.21. The wavelength-shifting compound on the outer surface shifts the incident photons to a wavelength less than the cut-off of the filter such that they pass through unimpeded. Then the wavelength-shifting compound on the inner surface shifts the photons to a wavelength above the cut-off such that they cannot pass back through the filter, trapping the photons within the ARAPUCA box. The trapped photons are then reflected by the dichroic filter surface or the highly reflective walls of the ARAPUCA box until they are eventually detected by SiPMs located at the edges or are lost to absorption [215]. This enables the effective sensitive area of the SiPMs to be amplified while keeping the light collection modules within the restrictive dimension requirements of the DUNE detectors.

The X-ARAPUCA is an extension of the original ARAPUCA design combining

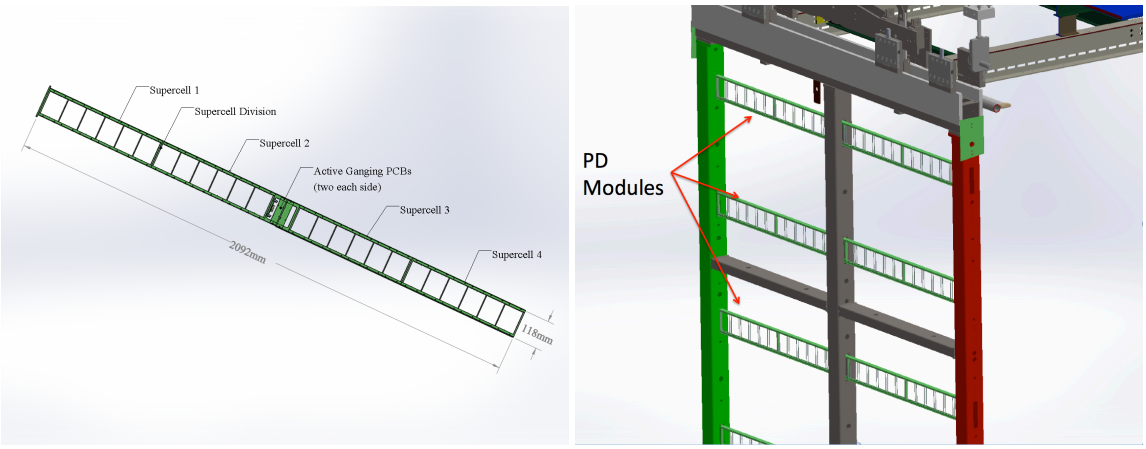

Figure 3.20: Diagrams of the DUNE-HD drift photon detection system modules. Left: a single module containing 24 X-ARAPUCA PDs grouped into four supercells. Right: three modules slotted into an APA frame. In total, there will be 10 modules per APA. Diagrams from Reference [182].

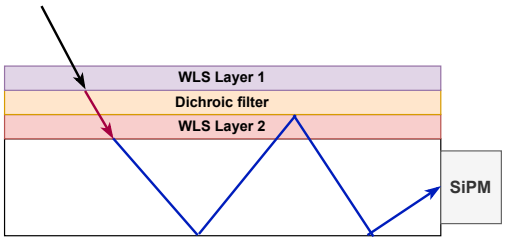

Figure 3.21: Illustration of the operational principle of an ARAPUCA photon detector.

the light trapping dichroic filter with a light guide to increase the detection efficiency, as illustrated in Figure 3.22. In the X-ARAPUCA design, the inner wavelength-shifting coating of the dichroic filter is replaced by a wavelength-shifting light guide. Photons incident at an angle greater than the critical angle undergo total internal reflection and travel along the light guide until they are detected by the SiPMs. The remaining photons pass through the light guide and are reflected off the dichroic filter surface or the highly reflective walls until being detected or lost to absorption similar to the original ARAPUCA design [216]. The X-ARAPUCA detectors can be either single- or double-sided by constructing them with either one or two windows and WLS-coated dichroic filters, with the remaining surfaces in each scenario made up of a highly reflective material.

There are two proposed extensions to the DUNE-HD photon detection system: the installation of wavelength-shifting reflector foils on the cathode of the detector, analogous to those in SBND; and doping the liquid argon with small quantities of liquid xenon. In both cases, these would boost the light yield, especially in

Figure 3.22: Illustration of the operational principle of an X-ARAPUCA photon detector showing the two mechanisms by which incident photons can be trapped. Figure from Reference [216].

the regions far from the photon detectors that see the least light. However, both extensions would result in a significant cost increase and are therefore not included in the baseline design [182].

The DUNE vertical drift module

The DUNE vertical drift (DUNE-VD) module will have an instrumented volume of $13.5 \times 13.0 \times 60.0\,\mathrm{m}^3$ (horizontal, vertical (drift), beam direction), with a total fiducial mass of approximately 10 kt. Unlike the previously described detectors, the DUNE-VD module will drift ionisation electrons vertically (up or down) towards the charge readout. The detector will be divided into two 6.5 m drift volumes, as illustrated in Figure 3.23. Each drift volume will be held at an electric field of $E_d = 450\,\mathrm{V/cm}$, drifting ionisation electrons towards charge-readout planes (CRPs) located at the top and bottom of the detector. The charge-readout in the vertical drift module will consist of printed circuit board (PCB) large electron multiplier (LEM) planes instead of the traditional sense wires. These are an evolution of the gas electron multiplier (GEM) planes developed for the DUNE dual-phase module design [254]. They consist of standard PCBs covered with small, sub-millimetre, holes. The PCBs have multiple electrical layers, consisting of planes of copper strips with different orientations. Drifted ionisation electrons passing through the holes induce charge on the inner electrical layers, resulting in bi-polar signals analogous to induction plane sense wires. The final electrical layer then collects the charge, resulting in uni-polar signals analogous to collection plane sense wires. LEM-based charge readout planes are able to achieve similar charge collection efficiency as traditional sense wires [255], while being much cheaper to manufacture. The DUNE-VD CRPs will consist of two stacked perforated PCBs with layers of copper strips on each face with different orientations. The innermost plane will serve as shield plane, similar to the innermost wire plane in the DUNE-HD design. Then there will be two induction planes and a collection plane, providing three different views of interacting particles. The DUNE-VD detector has several key advantages over the DUNE-HD

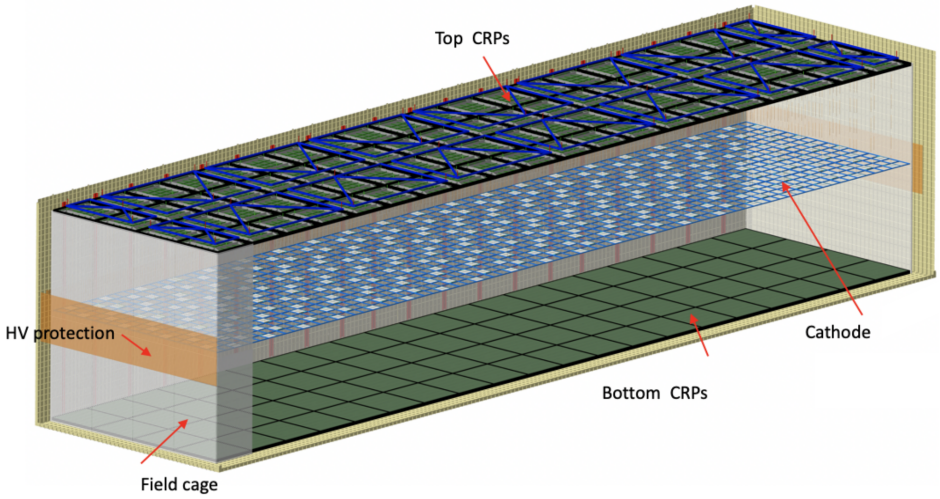

Figure 3.23: Diagram of the DUNE-VD module showing showing the two vertical drift volumes with a central cathode and charge-readout planes (CRPs) at the top and bottom. Diagram from Reference [251].

design, most significantly increased cost-effectiveness due to the reduced complexity of manufacturing the LEM-based readout compared with sense wires. The design also increases the active instrumented argon volume, while fitting within the same external cryostat size as the DUNE-HD. It achieves this by reducing the amount of dead-argon through the reduced number of drift volumes and more compact charge readout. The design comes with additional challenges, however, especially in the development of a photon detection system since, unlike the sense-wire-based designs, the photon detectors cannot be located behind the charge readout since the CRPs are fully opaque [251].

The DUNE-VD photon detection system (PDS) will use X-ARAPUCA [216] photon detectors (PDs), similar to those employed in the DUNE-HD design. Each DUNE-VD X-ARAPUCA module will have a sensitive light collection area of approximately $60 \times 60\,\mathrm{cm}^2$, much larger than those used the DUNE-HD design allowing a greater amount of light to be collected. Since the PDs cannot be located behind the CRPs, the baseline PDS design instead consists of PDs located on both the detector cathode and the side walls of the membrane cryostat, as illustrated in Figure 3.24. On the cathode, there will be a total of 320 evenly spaced double-sided X-ARAPUCA modules that will be sensitive to both drift volumes. The cathode X-ARAPUCA modules must not affect the cathode high-voltage, and hence have to be electrically insulated from ground. This poses a significant challenge in delivering power to and reading out the SiPMs. The planned solution is to adopt a non-conductive power-over-fibre and readout-over-fibre system [256]. This is an established technology, but has not previously been used at liquid-argon cryogenic

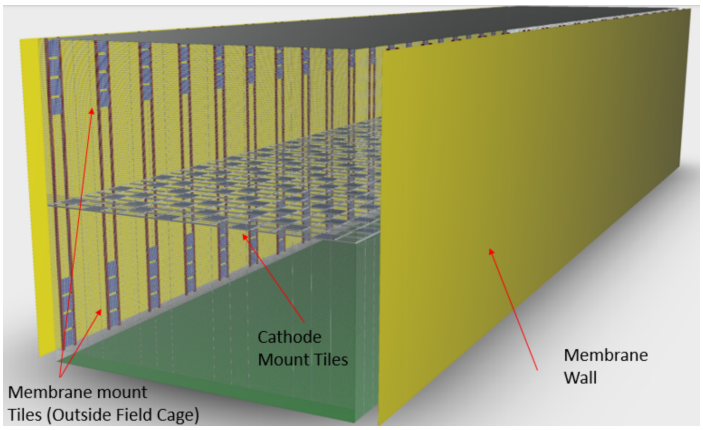

Figure 3.24: Diagram of the baseline design of the DUNE-VD module photon detection system showing the cathode and side-wall mounted (lateral) X-ARAPUCA PDs. Diagram from Reference [251].

temperatures. Development is on-going to establish the viability of this. In addition, there will be 320 single-sided X-ARAPUCA modules on the membrane cryostat walls of the long sides of the detector. In each drift volume they will be mounted in the region furthest from the cathode, as shown in Figure 3.24, where the light yield of the cathode PDs would be lowest. These detectors will be referred to as the lateral PDs subsequently. The lateral PDs will be located outside of the field-cage and hence do not need to be electrically isolated, avoiding the power delivery challenges of the cathode PDs. In the region directly in front of the lateral PDs, the field cage will have a slimmer profile allowing greater transmission of the light. As a backup design, if the cathode PDs are found to not be viable, the lateral PDs could be extended along the full width of the side walls [226, 251].

As a result of the larger drift volumes in the DUNE-VD design, there will be regions of the detector where the light has to propagate significantly further to reach the nearest PD compared with the DUNE-HD design. Xenon-doping of the liquid argon will be used compensate for this [251]. The xenon-wavelength light propagates better over the larger distances required, as described in Section 3.4.4. In addition, it is proposed that the CRP structures at the anode of the detector could be made of a material that is highly reflective to xenon wavelength light [251]. These could have a similar impact to the presence of wavelength-shifting reflector foils as described in Section 3.4.4, reflecting scintillation light incident on the anode back towards the PDs. However, the coverage of the reflective material would be limited reducing the overall efficiency of this significantly compared with full reflector foils.

Chapter 4

The NuMI neutrino beam at ArgoNeuT

This chapter will provide an overview of the Neutrinos at the Main Injector (NuMI) neutrino beam at the Fermi National Accelerator Laboratory (FNAL), with particular focus on its interaction with the ArgoNeuT experiment. Neutrino production in the beam along with the production of Heavy Neutral Leptons and Heavy QCD Axions will be described.

4.1 Neutrino production in the NuMI beam

The Neutrinos at the Main Injector (NuMI) beam [230] was originally designed to provide a flux of high-energy neutrinos for the MINOS long-baseline oscillation experiment [237]. It was then later used by a variety of other experiments including on-axis MINERvA [257], ArgoNeuT [172] and MINOS+ [258] and off-axis NOvA [259], MicroBooNE [158] and ICARUS [3, 183]. In the NuMI beam, protons are accelerated to 120 GeV in three stages. First H^- ions are accelerated to 400 MeV using a linear accelerator (Linac). They are then stripped of their electrons and injected into the Booster synchrotron where they are accelerated to 8 GeV. Finally, the protons are sent to the Main Injector synchrotron where they are accelerated to 120 GeV. Figure 4.1 shows an aerial view of the NuMI beam-line at FNAL illustrating the positions of the Linac, the Booster and Main Injector synchrotrons, and the approximate location of the ArgoNeuT detector in the MINOS near detector hall.

Figure 4.2 shows a diagram of the NuMI beam facility showing the target, magnetic focusing horns, decay pipe and hadron absorber. The 120 GeV protons from the Main Injector are incident on an approximately 1 m long carbon target producing a flux of charged mesons, primarily $\pi^{\pm}$ and $K^{\pm}$ along with a small subset of $D^{\pm}$ and $D_s^{\pm}$. These then pass through two magnetic focusing horns that allow

Figure 4.1: Aerial view of the NuMI neutrino beam at the Fermi National Accelerator Laboratory. The approximate locations of the Linac, the Booster and Main Injector synchrotrons and the ArgoNeuT detector are shown. Image adapted from Reference [260].

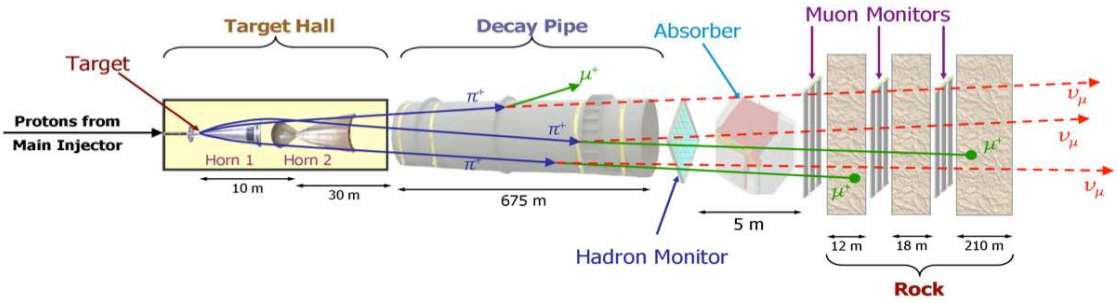

Figure 4.2: Diagram of the NuMI beam facility showing the target, magnetic focusing horns, decay pipe and hadron absorber. Diagram not to scale. Figure from Reference [230].

mesons of a particular charge to be isolated depending on whether a positive or negative horn current is applied. Positively charged mesons are focused in Forward Horn Current (FHC) mode and negatively charged mesons in Reverse Horn Current (RHC) mode[1]. The selected mesons then travel down a 675 m decay pipe where they decay to neutrinos, for example through the processes $\pi^+ \to \nu_\mu \mu^+$ and $\pi^- \to \bar{\nu}_\mu \mu^-$, producing a flux of predominantly muon neutrinos in FHC mode and muon anti-neutrinos in RHC mode. A subset of the $K^\pm$ mesons may also decay to electron neutrinos and anti-neutrinos resulting in electron neutrino contamination on the scale of a few percent. Finally, at the end of the decay pipe there is a steel

[1]The meson charge selection by the focusing horns is imperfect, and as a result there is a subset of wrong-sign contamination present. This is especially significant in the RHC mode since the majority of the initial mesons produced are positively charged.

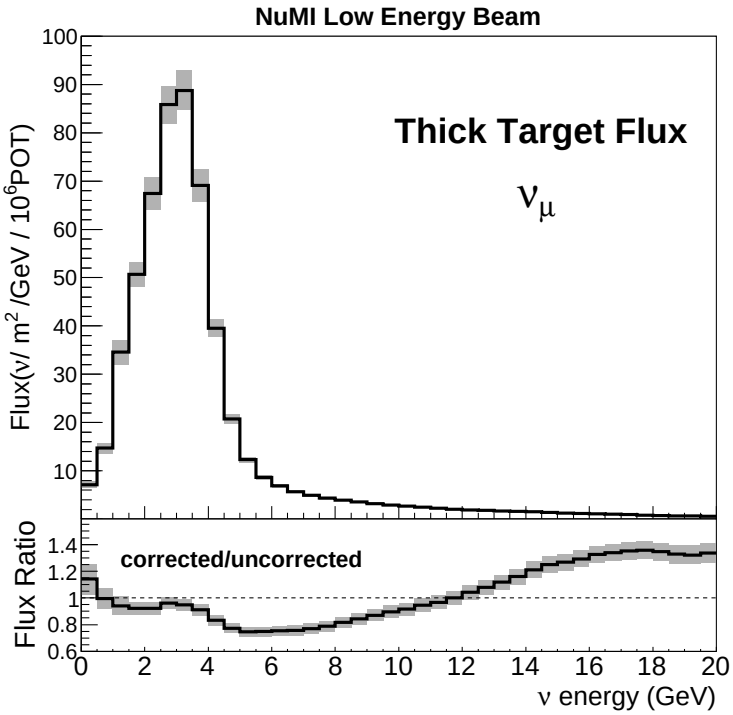

Figure 4.3: Energy distribution of the neutrinos produced in the low energy configuration of the NuMI beam. Figure from Reference [261].

and concrete hadron absorber. Any primary protons that did not interact in the target are absorbed there, along with mesons that have not decayed. There are then several hundred metres of rock between the absorber and the downstream detectors. The muons produced in the beam are stopped within this rock, leaving a pure beam of neutrinos by the time it arrives in the MINOS near detector hall [230].

The positions of the magnetic horns and target in the NuMI beam facility were designed to be configurable allowing mesons of different energies to be preferentially focused. This allows the average energy of the resulting neutrinos to be tuned. Three configurations were designed with separations between the two horns of 10 m, 23 m and 37 m corresponding to low, medium and high energy neutrinos, respectively [230]. The NuMI beam was configured in low energy mode through ArgoNeuT's data taking. The energy distribution of the neutrinos produced in this configuration is shown in Figure 4.3, ranging between approximately 1 and 10 GeV with an average energy $\langle E_\nu \rangle \sim 4\,\text{GeV}$ [261].

The ArgoNeuT detector was located along the NuMI beam-line, approximately 1033 m downstream of the beam target and 318 m downstream of the hadron absorber. For the majority of ArgoNeuT's physics run, the NuMI beam was in RHC (anti-neutrino) mode [172].

4.2 Heavy Neutral Lepton production

Heavy Neutral Leptons (HNLs) could potentially be produced in the 120 GeV proton–fixed-target collisions in the NuMI beam. As discussed in Section 2.2.3, a HNL, N, with mass m_N and mixing angle $|U_{\alpha N}|^2$ with the light standard model neutrinos ν_α ($\alpha = e, \mu, \tau$) could be produced in any process that would normally result in an outgoing ν_α, provided that it is kinematically allowed. The HNLs can therefore be produced from the decays of $\pi^\pm$, $K^\pm$ and $D^\pm_{(s)}$ mesons produced in the proton–fixed-target collisions. The production rate of N in these processes is proportional to the mixing angle squared, $|U_{\alpha N}|^2$. Three scenarios can be considered: electron-coupled, where $|U_{eN}|^2 \neq 0$; muon-coupled, where $\left|U_{\mu N}\right|^2 \neq 0$; and tau-coupled $|U_{\tau N}|^2 \neq 0$. For simplicity in all three cases the other two mixing angles are assumed to be zero.

In electron-coupled and muon-coupled scenarios the dominant HNL production channels in the NuMI beam are [84, 85]:

$$\pi^\pm \to \ell^\pm_\alpha N \qquad \text{and} \qquad K^\pm \to \ell^\pm_\alpha N, \tag{4.1}$$

where in each case ℓ_α is the corresponding lepton $e^\pm$ or $\mu^\pm$. Decays of $\pi^\pm$ mesons allow production of HNLs with masses up to $m_N = m_\pi - m_{l_\alpha}$. Therefore, for the electron-coupled case this mechanism allows production up to $m_N \sim 140\,\text{MeV}$, but for the muon-coupled case it is restricted to a maximum of $m_N \sim 30\,\text{MeV}$ as a result of the larger muon mass. Analogously, decays of $K^\pm$ mesons allow production of HNLs with masses up to $m_N = m_K - m_{l_\alpha}$. This corresponds approximately to $m_N \sim 495\,\text{MeV}$ for the electron-coupled case and $m_N \sim 390\,\text{MeV}$ for the muon-coupled case. In the NuMI beam, a non-negligible flux of charmed $D^\pm$ and $D^\pm_s$ mesons is also produced. This allows significantly higher mass HNLs to be produced through the channel [84, 85]:

$$D^\pm_{(s)} \to \ell^\pm_\alpha N, \tag{4.2}$$

where, as before, ℓ_α is the corresponding lepton $e^\pm$ or $\mu^\pm$. In this case, production of HNLs with masses up to $m_N = m_{D_{(s)}} - m_{l_\alpha}$ is possible. For the higher mass $D^\pm_s$, this allows production up to a maximum of $m_N \sim 1965\,\text{MeV}$ for the electron-coupled case and $m_N \sim 1860\,\text{MeV}$ for the muon-coupled case. However, the rate of $D^\pm_{(s)}$ meson production in the NuMI beam is far lower than $\pi^\pm$ and $K^\pm$ mesons.

In the tau-coupled scenario, production through decays of $\pi^\pm$ and $K^\pm$ is kinematically forbidden due to the need to produce an outgoing $\tau^\pm$ lepton. Production through decays of $D^\pm_{(s)}$ mesons remains possible, however is restricted to relatively low HNL masses: $m_N \sim 90\,\text{MeV}$ for $D^\pm$ decays and $m_N \sim 190\,\text{MeV}$ for $D^\pm_s$ decays. Instead, tau-coupled HNL production in the NuMI beam predominantly occurs from

decays of secondary $\tau^{\pm}$ leptons originating from decays of $D^{\pm}_{(s)}$ mesons [84, 85]:

$$D^{\pm}_{(s)} \to \tau^{\pm}\nu_{\tau} \qquad \text{then} \qquad \tau^{\pm} \to X^{\pm}N, \tag{4.3}$$

where $X^{\pm}$ represents standard model mesons such as $\pi^{\pm}$ or $K^{\pm}$. This mechanism allows the production of tau-coupled HNLs with masses up to approximately $m_N = m_{\tau} - m_{\pi} \sim 1635\,\text{MeV}$. HNLs could also be produced in decays of heavier $B^{\pm}$ mesons. However, at the energy of the NuMI beam the production rate of these is negligible.

As motivated in Section 2.2.3, the HNL decay signature $N \to \nu\mu^{+}\mu^{-}$ will be considered in the search presented in this thesis. ArgoNeuT is uniquely sensitive to this signature due to the presence of the downstream MINOS near detector, as described in Section 3.6.1. In order for this decay to be kinematically allowed, the HNLs must have a mass $m_N > 2m_{\mu} \sim 210\,\text{MeV}$. This excludes production from $\pi^{\pm}$, as well as from $D^{\pm}_{(s)}$ in the tau-coupled scenario. Direct production is therefore only possible via $K^{\pm}$ and $D^{\pm}_{(s)}$ decays in the electron-coupled and muon-coupled scenarios, while in the tau-coupled scenario only production from secondary $\tau^{\pm}$ decays is relevant. In the case of $K^{\pm}$ mesons, focusing by the magnetic horns in the NuMI beam has a non-negligible impact on the resulting HNL flux. The impact is negligible for $D^{\pm}_{(s)}$ mesons and $\tau^{\pm}$ leptons, however, as a result of their very short lifetimes.

In the NuMI beam approximately 13% of the primary 120 GeV protons pass through the target and reach the downstream hadron absorber [230]. Production of HNLs from interactions in the absorber could also occur via the same processes. However, the relative rates may be different due to the absence of the magnetic focusing horns and decay pipe, particularly for production via $K^{\pm}$ decays. The NuMI absorber is 715 m closer to the ArgoNeuT detector than the target. As a result of this, the geometric acceptance of ArgoNeuT is significantly larger ($\sim$ factor of 10) for the absorber-produced HNLs compared with the target-produced ones. Therefore, despite only a small fraction of the protons reaching the absorber it is the source of a relatively large fraction of the HNLs incident on ArgoNeuT. Additionally, the different distance between the absorber and ArgoNeuT allows a different region of the HNL phase-space to be probed compared with production in the target. In particular allowing sensitivity to HNLs with shorter lifetimes.

4.3 Heavy QCD Axion production

Similar to HNLs, Heavy QCD Axions could also potentially be produced in the 120 GeV proton–fixed-target collisions in the NuMI beam. As discussed in Section 2.3.3, Heavy QCD Axions can be produced through mixing with the neutral

standard model pseudoscalar mesons. A flux of π^0, η and η' are produced in the proton–fixed-target collisions in the NuMI beam. These mix with the axions, a, with mixing angles θ_{aM}, where $M = \pi, \eta, \eta'$, given by [108, 262, 263]:

$$\begin{aligned}
\theta_{a\pi} &= \frac{1}{6}\frac{f_\pi}{f_a}\frac{m_a^2}{m_a^2 - m_\pi^2}, \\
\theta_{a\eta} &= \frac{1}{\sqrt{6}}\frac{f_\pi}{f_a}\frac{m_a^2 - 4m_\pi^2/9}{m_a^2 - m_\eta^2}, \\
\theta_{a\eta'} &= \frac{1}{2\sqrt{3}}\frac{f_\pi}{f_a}\frac{m_a^2 - 16m_\pi^2/9}{m_a^2 - m_{\eta'}^2},
\end{aligned} \tag{4.4}$$

where $f_\pi = 93\,\text{MeV}$ is the pion decay constant. These relations hold true provided that $\theta_{aM} \ll 1$, and therefore excludes the parameter space for which $m_a \sim m_M$. The number of Heavy QCD Axions, N_a, produced in the beam then becomes,

$$N_a = N_\pi \times |\theta_{a\pi}|^2 + N_\eta \times |\theta_{a\eta}|^2 + N_{\eta'} \times |\theta_{a\eta'}|^2, \tag{4.5}$$

where N_π, N_η and $N_{\eta'}$ are the number of π^0, η and η' mesons, respectively. Production through mixing with heavier neutral pseudoscalar mesons would be negligible at the energy of the NuMI beam. Other modes of axion production, such as meson decays, are subdominant unless for example additional flavour-violating couplings are introduced [108].

Analogous to the Heavy Neutral Leptons described in Section 4.2, Heavy QCD Axions could also be produced from the approximately 13% of the primary 120 GeV protons that reach the NuMI hadron absorber. Since production occurs through the mixing with the neutral pseudoscalar mesons, rather than from their decay, the production rates at the target and the absorber are approximately the same despite the absence of the focusing horns and decay pipe. As before, the relative proximity of the absorber to ArgoNeuT compared with the target results in a significantly larger geometric acceptance and allows shorter axion lifetimes to be probed.

Chapter 5

Simulation and reconstruction in LArTPC experiments

Simulation and reconstruction in LArTPC neutrino detectors is largely similar between experiments as a result of using the common LArSoft software package [264]. This is a C++ based software suite built around the Art particle physics event processing framework [265] and the ROOT data processing toolkit [266]. Along with numerous dedicated algorithms and tools, it incorporates a variety of particle physics simulation packages such as the GENIE [267] and CORSIKA [268] event generators and the Geant4 particle propagation simulation toolkit [269]. This chapter will provide an overview of simulation and reconstruction in LArTPC detectors, with particular focus on the techniques used by the ArgoNeuT experiment. It will follow the structure of the LArSoft simulation: beginning with event generation, followed by particle propagation, then charge and light simulation, the detector response simulation, and finally reconstruction.

5.1 Event generation

The first stage of Monte Carlo (MC) simulation in LArTPC detectors is the generation of primary particle interactions and the determination of the position and momenta of the resulting secondary particles. This is performed using a variety of different event-generator software libraries. These are either directly integrated within the LArSoft suite or are used externally with the results then input into LArSoft. The kinematics of the resulting particles then provide the initial state of the subsequent simulation. Several different generators will be described in this section that are relevant for the analyses presented in this thesis.

5.1.1 Simulation of the NuMI beam in ArgoNeuT

Simulation of beam-neutrino interactions in ArgoNeuT is performed using the GENIE neutrino generator [267] integrated in LArSoft. Modelling of the neutrino flux provided by the MINERvA collaboration [261] is used, simulated using the G4NuMI Geant4-based simulation of the NuMI beam-line. The GENIE generator uses the expected neutrino energies and flux obtained from the beam simulation to determine the rate and position of neutrino interactions in the detector, accounting for the cross-sections of the different materials present. It then samples the possible neutrino interaction modes, depending on the neutrino energy, and determines the kinematics of the resulting secondary particles. In ArgoNeuT, GENIE is used to model interactions occurring in the liquid argon volume, the cryostat and the surrounding detector structure.

In addition to interactions occurring within the detector, neutrinos may interact with the various materials upstream of the detector along the beam-line. These interactions produce a significant flux of beam-induced through-going muons in ArgoNeuT [159]. These types of interactions are commonly referred to as dirt neutrino interactions. In ArgoNeuT, they are modelled using a data-driven approach. The rate, position, and momenta of the muons passing through or near to ArgoNeuT can be characterised using the downstream MINOS-ND. In each simulated ArgoNeuT beam event, a random MINOS-ND data event is selected. The kinematics of any muons that are present are extracted, and the expected start positions are calculated by projecting backwards along the muon trajectory. Any muons that would pass through or near to ArgoNeuT are then overlaid onto the GENIE event. This approach has been found to effectively model the observed neutrino-induced through-going muon rate [231, 232, 270].

5.1.2 Simulation of cosmic-ray showers in SBND

The SBND experiment will be located on the surface and therefore will be exposed to a large number of particle showers resulting from cosmic-ray interactions in the atmosphere. These are simulated using the CORSIKA cosmic-ray shower generator [268]. CORSIKA models high-energy primary cosmic-ray protons, light nuclei up to iron and other particles incident on the atmosphere. These are tracked until they interact, and the development of the resulting shower of particles is then simulated. Cosmic-ray interactions produce a significant flux of muons (and to a lesser degree other particles) that can propagate through the structures surrounding the detector and produce signals in the argon. Multiple cosmic-ray-induced muons will be present in every charge read-out window, forming a significant background that requires removal in analyses [271]. The showers produced by cosmic-ray interac-

tions are computationally challenging to model due to the large number of resulting particles. Therefore, in LArSoft a database of pre-generated CORSIKA showers is sampled for each event in place of the full simulation. The database provides the kinematics the particles that would be incident on SBND, which are then simulated subsequently as starting from a plane approximately 20 m above the detector.

An alternative approach, taken by detectors such as MicroBooNE [158], is to overlay cosmic-ray data taken between beam spills onto simulated events. This helps to remove any uncertainties related to both modelling by CORSIKA and detector-specific effects [272, 273]. An approach similar to this may be adopted by SBND in the future.

5.1.3 Simulation of Heavy Neutral Leptons and Heavy QCD Axions

Unlike modelling of the neutrino beam and cosmic-rays, Beyond the Standard Model (BSM) generators typically are not integrated directly into LArSoft. Instead the rate, position and kinematics of the BSM particles and their subsequent interactions or decays are typically generated externally. They are then provided as input to the later stages of the LArSoft simulation. Attempts have been made at creating relatively generic BSM generators covering a range of models [274, 275]. However, more commonly custom generators are produced by the phenomenologists developing the particular model. This will be the case for the analyses presented in this thesis.

As described in Sections 4.2 and 4.3, Heavy Neutral Leptons (HNLs) can be produced from decays of $K^{\pm}$, $D^{\pm}$ and $D_s^{\pm}$ mesons, along with the decays of secondary $\tau^{\pm}$ leptons, and Heavy QCD Axions can be produced via mixing with π^0, η and η' mesons. In each case, these particles are produced in the NuMI beam during high-energy proton–fixed-target interactions. An illustration of the production, propagation and subsequent decay of HNLs and Heavy QCD Axions is shown in Figure 5.1. Production occurring from protons interacting in both the NuMI beam target (87% of protons) and in the downstream hadron absorber (13% of protons) is evaluated.

Production in the NuMI beam target

Production of $K^{\pm}$ mesons in the NuMI beam is modelled using the G4NuMI simulation package. Decay-in-flight of the $K^{\pm}$ mesons is allowed, taking into account focusing by the magnetic horns of the beam. Production of $D_{(s)}^{\pm}$ mesons, and their subsequent decay into $\tau^{\pm}$ leptons, is not, however, currently incorporated in the G4NuMI simulation. Instead, the production is simulated using PYTHIA8 [276,

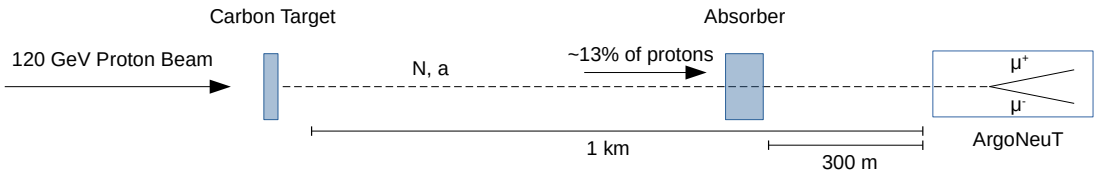

Figure 5.1: Illustration of Heavy Neutral Lepton, N, and Heavy QCD Axion, a, production in the NuMI beam, subsequent propagation to the ArgoNeuT detector and decay into an observable signature. Diagram not to scale.

277] following the approach described in Reference [85]. The production rates are determined by colliding 120 GeV protons with stationary protons and neutrons modelling a fixed target. Table 5.1 summarises the production rates of the different mesons in the NuMI beam per incident proton-on-target. The production rate of $\tau^{\pm}$ leptons arising from $D^{\pm}_{(s)}$ decays is also shown. The rates shown assume that all of the incident protons interact in the fixed target. In the NuMI beam, therefore, these rates are scaled down by a factor of 0.87 for production at the target (with the remainder of the protons interacting at the absorber). Both $D^{\pm}_{(s)}$ mesons and $\tau^{\pm}$ leptons decay promptly and as a result the impact of focusing from the horns is negligible. This is therefore not considered in the simulation.

	K	D	D_s	τ
P^+/POT	0.54	1.2×10^{-5}	3.3×10^{-6}	2.1×10^{-7}
P^-/POT	0.24	1.9×10^{-5}	4.6×10^{-6}	3.0×10^{-7}

Table 5.1: Production rates of positive and negative particles, $P^{\pm}$, in the NuMI beam per incident proton-on-target (POT) [85].

The production rate of HNLs, N, from decays of $K^{\pm}$, $D^{\pm}_{(s)}$ and $\tau^{\pm}$ is proportional to mixing angle with the light neutrinos, ν_{α} ($\alpha = e, \mu, \tau$), squared: $|U_{\alpha N}|^2$. For simplicity, it is assumed that only one angle $|U_{\alpha N}|^2$ is non-zero at a time. To be conservative, the HNLs are assumed to be Dirac particles throughout. In the electron-coupled, $|U_{eN}|^2 \neq 0$, and muon-coupled, $|U_{\mu N}|^2 \neq 0$, scenarios two production channels are included, $K^{\pm} \to \ell^{\pm}_{\alpha} N$ and $D^{\pm}_{(s)} \to \ell^{\pm}_{\alpha} N$, as described in Section 4.2. The kinematics of the resulting HNLs are simulated following the approach described in References [84, 85]. In the tau-coupled scenario, $|U_{\tau N}|^2 \neq 0$, the HNL flux is generated by simulating the decays $\tau^{\pm} \to NX$, where X is assumed to consist of SM particles, as described in Section 4.2. The kinematics are simulated by assuming $m_X = m_{\pi^{\pm}}$ and that the branching ratio of this new decay is $\mathrm{Br}(\tau^{\pm} \to NX^{\pm}) = 0.9\,|U_{\tau N}|^2\,K(m_N)$. The function $K(m_N)$ parameterises the shrinking of this branching ratio as m_N grows larger than $\sim$1 GeV [50, 84]. The factor of 0.9 accounts for the $\sim$10% of $\tau^{\pm}$ decays into heavier final states that would not allow for the production of a HNL in the mass range of interest. This is a conser-

vative assumption as long as $m_N \lesssim 1$ GeV, above which the pre-factor multiplying $|U_{\tau N}|^2$ reduces further.

Analogously, the production rates and resulting kinematics of the π^0, η and η' mesons are simulated using Pythia8 [276, 277] following the approach described in Reference [108]. The production rates per incident proton-on-target (POT) are summarised in Table 5.2. As before, the rates shown assume that all of the incident protons interact in the fixed target. These are therefore scaled by a factor of 0.87 for production at the NuMI beam target.

	π^0	η	η'
P/POT	2.89	0.33	0.03

Table 5.2: Production rates of π^0, η and η' (P), in the NuMI beam per incident proton-on-target (POT) [108].

The rate of axion, a, production from mixing with the π^0, η and η' mesons is then proportional to the axion–meson mixing angles squared $|\theta_{aM}|^2$, where $M = \pi, \eta, \eta'$, as defined in Equations 4.4 and 4.5. Since production occurs through the mixing with the neutral mesons, rather than from their decay, there is no impact from the focusing horns or other components of the NuMI beam facility. These are therefore not simulated.

Production in the NuMI beam absorber

Approximately 13% of the primary 120 GeV protons from the NuMI beam pass through the target and reach the downstream absorber [230]. These protons lose some energy and diverge slightly during propagation. A stand-alone simulation of the protons reaching the absorber is performed with Geant4 [269]. The simulation models the initial extent of the beam spot at the target, the graphite target, helium decay pipe, ambient air, upstream and downstream iron decay pipe caps, and the iron and aluminium absorber. The kinetic energy, T_p, and X-Y position distributions of the protons once they reach the absorber are shown in Figure 5.2. The majority of the protons reach the absorber with energies close to the initial 120 GeV, although a subset of the protons lose a non-negligible fraction of their energy. The proton beam also continues to be highly focused, with the majority of the protons remaining within ±20 cm of the beam centre.

Heavy Neutral Lepton and Heavy QCD Axion production is considered from protons with $T_p \geq 115$ GeV. This includes approximately 75% of the protons that reach the absorber, corresponding to approximately 10% of the total POT. The 5 GeV energy difference of these protons has a negligible impact on the production rates. For these protons, the $D^{\pm}_{(s)}$ and $\tau^{\pm}$ production rates in Table 5.1 and the π^0,

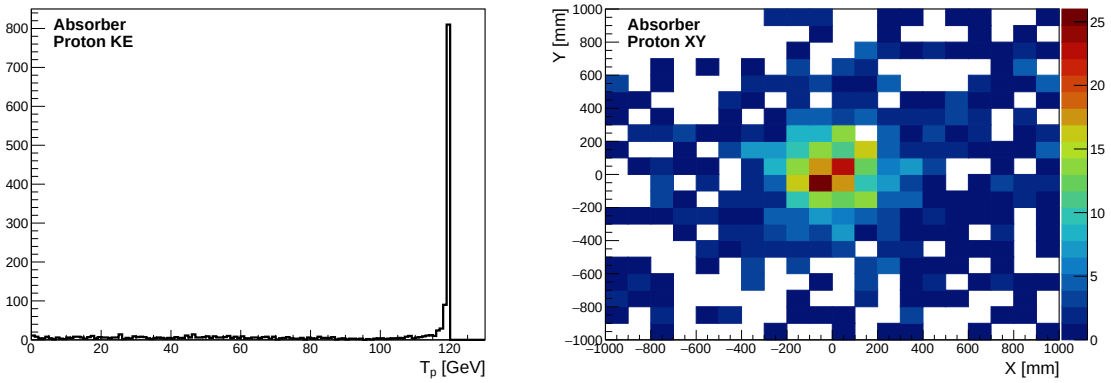

Figure 5.2: Kinetic energy (left) and X-Y position relative to the centre of the beam (right) of NuMI beam protons once they reach the hadron absorber. Simulation from Reference [278].

η and η' production rates in Table 5.2 apply. Production of $K^{\pm}$ mesons would also be similar, however the resulting propagation and decay of these mesons would be significantly different due to the absence of the focusing horns.

Heavy Neutral Lepton production in the absorber is considered to be primarily from $D^{\pm}_{(s)}$ and $\tau^{\pm}$ decays. Production from $K^{\pm}$ decays is negligible relative to production in the target as a result of the lack of focusing from the horns. The flux and kinematics of the HNLs, N, are determined using the same methods that are used for production in the target. The broader distribution of the protons reaching the absorber is incorporated in the resulting distribution of HNLs, however the impact of this is small. Since $D^{\pm}_{(s)}$ and $\tau^{\pm}$ decay promptly, the kinematics of N produced in the target and the absorber are approximately the same despite the absence of the focusing horns and the decay pipe.

Analogously, Heavy QCD Axion production in the absorber is modelled using the same methods that are used for production in the target, taking into account the slightly broader distribution of the protons. As this occurs through the mixing with the neutral pseudoscalar mesons at the point they are produced, rather than from their decay, the production rates at the target and the absorber, along with the subsequent axion kinematics, are approximately the same despite the absence of the focusing horns and decay pipe.

Propagation to ArgoNeuT and subsequent decay

The geometric acceptance of ArgoNeuT as viewed from the beam target and the hadron absorber is evaluated. Heavy Neutral Leptons and Heavy QCD Axions that would not intersect with ArgoNeuT, or the region of the upstream cavern that will be considered, are removed. As discussed in Sections 4.2 and 4.3, the geometric acceptance of ArgoNeuT is significantly larger ($\sim$ factor of 10) for production occurring at

the absorber compared with the target due to the relative proximity to the detector. Additionally this proximity allows different regions of the respective phase-spaces to be probed, in particular allowing sensitivity to different Heavy Neutral Lepton and Heavy QCD Axion lifetimes.

The rate of Heavy Neutral Lepton decays $N \to \nu\mu^+\mu^-$ is determined using the HNL lifetime and branching ratios, described in Section 2.2.3, of each scenario considered. The kinematics of the resulting decay products are then determined following the approach described in References [84, 85]. Analogously, the rate of Heavy QCD Axion decays $a \to \mu^+\mu^-$ is calculated considering the axion lifetime and branching ratios, discussed in Section 2.3.3, and the kinematics of the resulting muons determined following the approach described in Reference [108]. In both cases, the positions and kinematics of the decay products are converted into the HEPEVT standard format [279] allowing interfacing with LArSoft. They are then provided as the initial state for the subsequent LArSoft simulation.

5.2 Simulation

Once the interaction and the resulting particles have been generated, their propagation in the detector is modelled using Geant4 and the resulting charge and light signals produced are simulated. Finally the detector response to these signals is determined, modelling the signature that would be observed in data. Each of these stages is performed using the LArSoft suite itself or using tools, such as Geant4, integrated within LArSoft.

5.2.1 Particle propagation

The propagation of particles through the detector and surrounding materials is performed using Geant4. Each particle is independently stepped along its trajectory, with a step-size determined based on the physics processes available to the particular particle species. At each step, the possible interactions with the argon or other materials are evaluated. The energy deposited in the step is determined, and the resulting trajectory of the particle calculated. This process then repeats until the particle reaches an energy threshold significantly below the detector sensitivity, typically around 10-100 keV, where subsequent energy deposits would not be detectable and hence negligible. Any secondary particles produced from interactions occurring during the propagation are saved and later simulated in the same way, keeping track of their parent particle(s). The energy depositions at each step along each simulated particle trajectory are then used to determine the amount of ionisation charge and scintillation light produced, accounting for the charge-light anti-correlation as

a function of the electric field described in Section 3.5.

5.2.2 Charge simulation

Once the number of ionisation electrons produced at each step is determined, their drift to the sense wires is then modelled. Propagating each electron individually, however, would be prohibitively slow and is unnecessary since they typically follow a relatively simple path along the drift field direction. Instead the ionisation electrons are grouped into clouds, where the spatial scale of these clouds is sufficiently small to avoid discrete jumps in the amount of charge reaching the wires. The clouds of electrons are projected to the anode plane accounting for the electron lifetime, smearing due to transverse and longitudinal diffusion, and distortions in the electric field, and hence the electron trajectories, from space-charge. Each of these effects are described in Section 3.3.2. The propagation of the ionisation charge close to the wire planes and the resulting signals observed by the sense wires is more complex, however. Figure 5.3 shows an illustration of the electron drift paths as they reach the planes of sense wires. In this example, four wire planes are present: a shield plane, two induction planes and a collection plane. This corresponds to the DUNE horizontal drift detector design, discussed in Section 3.6.3. The points represent the individual wires in each plane. The drifted charge passes the shield plane and

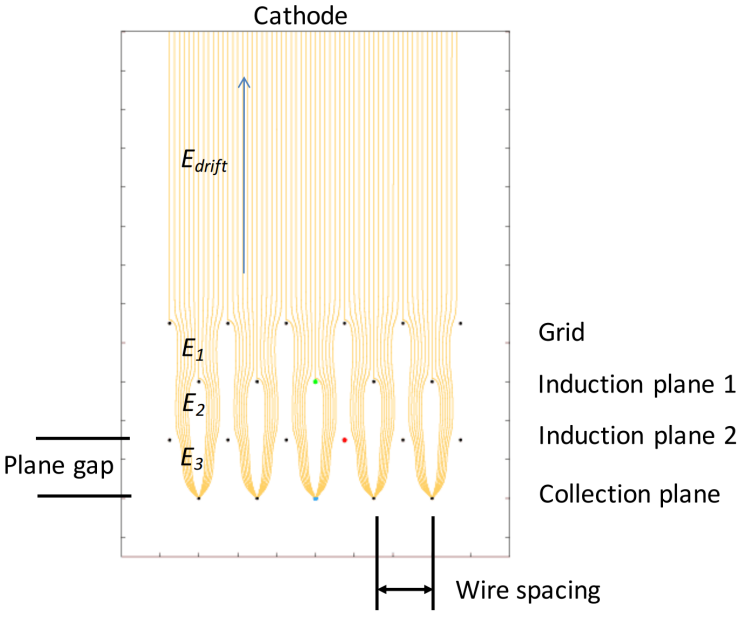

Figure 5.3: Illustration of the electron drift paths as they reach the planes of sense wires. Four wire planes are present: a shield plane, two induction planes and a collection plane. The points represent the individual wires in each plane. Figure from Reference [280].

between the two induction planes, inducing a signal on the wires, before being collected on the final plane.

Modelling of the ionisation charge close to the wire planes can be performed using, for example, the Garfield software package [281]. Due to the drifted electrons inducing signals on multiple nearby wires, the response can become quite complex. Typically in simulation, and later reconstruction, the response of the wires is modelled using a simplified fit to the prediction from Garfield that is calibrated using

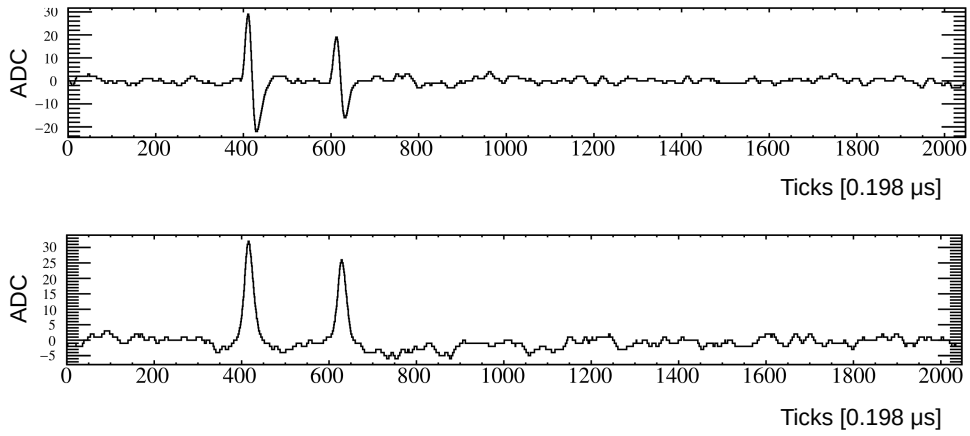

Figure 5.4: Example simulated signals on an induction plane wire (top) and a collection plane wire (bottom) from two muon tracks in the ArgoNeuT detector. Bipolar signals are seen on the induction plane, and unipolar signals on the collection plane.

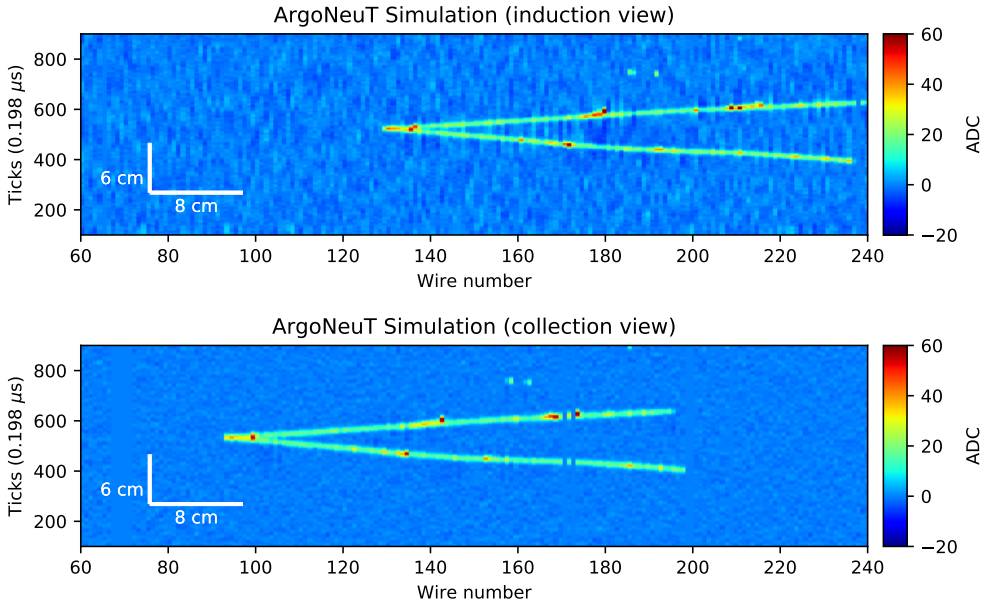

Figure 5.5: Event display showing simulated signals across the full induction (top) and collection (bottom) wire planes from two muon tracks in the ArgoNeuT detector.

the response observed in data [280]. Additionally, the expected electronics response and noise is incorporated. Figure 5.4 shows an example of the simulated signals on an individual induction plane wire (top) and collection plane wire (bottom) in the ArgoNeuT detector resulting from a pair of simulated muon tracks. In each case two peaks can be seen, originating from clouds of drifted ionisation electrons from the two muon tracks at different positions in the detector in the drift direction. The simulated noise on the wires is also visible. Figure 5.5 shows the signals seen across the full induction (top) and collection (bottom) wire planes from the pair of simulated muons. The signal on each wire (x-axis) over time (y-axis) is shown, where the colour is proportional to the amount of charge observed. The two muon tracks originating from a common vertex are clearly visible.

5.2.3 Light simulation

Similar to the ionisation electrons, once the number of photons produced at each step is determined their propagation to the photon detectors is modelled. Unlike the ionisation electrons that propagate along the electric field lines, photons can take significantly more complex paths to reach the photon detectors due to reflections and processes such as Rayleigh scattering, as described in Section 3.4.2. The propagation of the photons can be modelled using Geant4, calculating the path taken by each photon individually until they either reach a photon detector or are absorbed. This allows the number of photons arriving at each photon detector and their arrival time distribution to be determined. This approach will be referred to subsequently as *full optical simulation*. However, due to the combination of the high scintillation yield of LAr ($\sim$20,000 photons/MeV at 500 V/cm) and the high energy of interactions of interest from beam neutrinos ($\sim$GeV), full optical simulation becomes prohibitively slow for use in large-scale LArTPC detectors. To mitigate this, *fast optical simulation* techniques are employed. These are typically trained using the full optical simulation, then are used in place of it subsequently.

A common method of fast optical simulation is the optical look-up library [282]. Figure 5.6 shows a diagram illustrating the optical look-up library approach. The detector volume is divided in voxels, three-dimensional pixels, and then a large energy deposition is simulated within each voxel using the full optical simulation. The probability of photons from each voxel reaching each photon detector, the visibility of that photon detector, is calculated and stored in a look-up library. This look-up library can then be used in place of the full optical simulation subsequently to calculate the number of photons arriving on each photon detector from each voxel. The computationally challenging full optical simulation therefore only needs to be performed once for each voxel within the detector and the subsequent simulation is

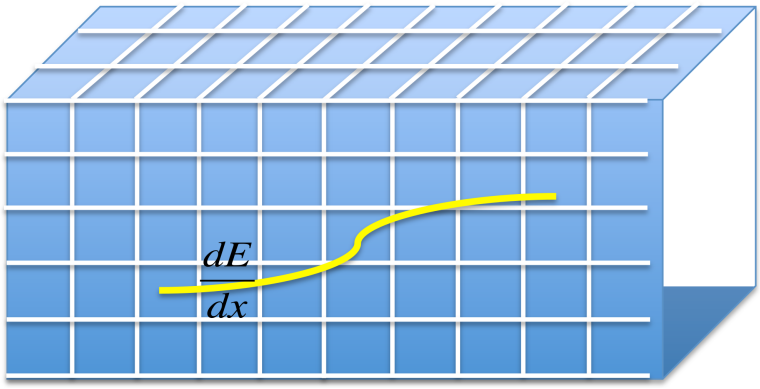

Figure 5.6: Illustration of the optical look-up library fast optical simulation approach. The visibility of each photon detector from each voxel is stored in a library and used in place of full optical simulation. Figure from Reference [207].

much faster.

The optical look-up library approach has been successfully employed in smaller LArTPC detectors, such as MicroBooNE [158], however it scales poorly as detectors increase in size. In larger detectors a larger number of voxels are required and, since larger detectors typically have more photon detectors, a larger number of visibilities must be stored per voxel. In very large detectors, such as the DUNE far detectors, or detectors with a large number of photon detectors, such as SBND, optical look-up libraries become highly memory-intensive when used in simulation. This is problematic for the typical grid-computing based MC production used by LArTPC experiments due to limitations on the amount of memory available to each grid node. Additionally, optical look-up libraries are only designed to calculate the number of photons arriving on each photon detector. They do not model the photon arrival time distribution, which can be significantly smeared out due to propagation effects as discussed in Section 3.4.2. In detectors aiming to achieve precise timing resolutions these effects are also important to model.

To mitigate the shortcomings of the look-up library method an alternative fast optical simulation model, the semi-analytical model [4], has been developed. This uses a geometric approach combined with parameterisations to calculate both the number of photons observed on each photon detector and their arrival time distribution. It is able to scale to detectors as large as the DUNE far detectors without the memory-related issues of the optical look-up library method. The semi-analytical model will be the focus of Chapter 6. It is currently the standard fast optical simulation approach used in SBND and the DUNE far detectors.

Other approaches to light simulation are also in development. One recent example is the use of generative neural networks to produce photon visibility maps analogous to those made during the optical look-up library generation [283]. The

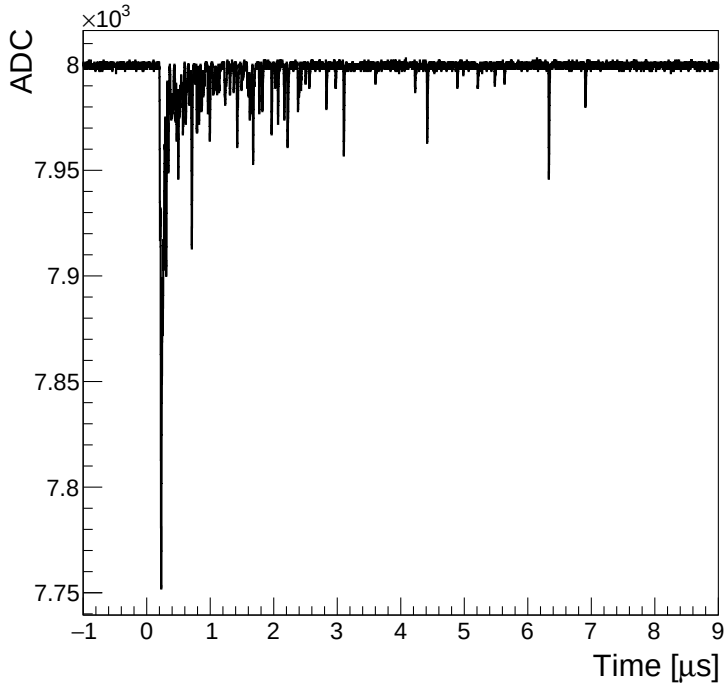

Figure 5.7: Example simulated PMT waveform from an interaction occurring in SBND. The PMT baseline in this example is at 8000 ADC.

neural network is trained using Geant4 simulation similar to the library, but then produces the visibility maps at run time avoiding the memory issues. However, in its current implementation the visibility map inference is significantly slower than alternative methods and only provides a prediction of the number of photons. In the future, as large scale computing increasingly moves towards GPU-based architectures, hardware accelerated ray-tracing could provide a more accurate approach to light modelling without the need of simplified fast optical simulation models.

Once the number of photons and their arrival times on each photon detector have been predicted, the photon detector response is modelled. This is achieved by replacing each arriving photon with the expected average single photo-electron response of the photon detector. Additional effects are also simulated such as fluctuations in the gain of PMTs due to variations in the numbers of electrons produced at each dynode and the expected level of noise. An example simulated PMT waveform from an interaction occurring in SBND is shown in Figure 5.7. A large initial peak can be seen from the prompt component of the light, followed by a long tail with peaks corresponding to individual (or small numbers of) later arriving photons from the slow component.

5.3 Reconstruction

As with simulation, reconstruction in LArTPCs is performed using the LArSoft suite. The raw signals, either from data or simulation, are first processed performing deconvolution and noise removal. The charge and light signals are then reconstructed. Finally particle identification can be performed based on topological and calorimetric information. This section will give an overview of LArTPC reconstruction, with a particular focus on the methods employed by the ArgoNeuT experiment that are relevant for the analyses described in this thesis.

5.3.1 Charge reconstruction

Signal processing

The first stage of reconstruction is the processing of the raw signals seen by each sense wire to extract the expected approximately Gaussian signal shape. The time-sampled raw signal seen on a wire consists of a convolution of the ionisation charge incident on the wire with the field response and the electronics response. To extract the ionisation charge signal, deconvolution is performed using a Fast Fourier Transform of the signal and a Wiener filter [172, 280]. A frequency filter is also used in the deconvolution process, to remove high frequency coherent noise. Further data-driven noise removal methods can also be employed [284]. The implementation of the deconvolution is typically tuned in each particular LArTPC detector to preserve the most important components of the signals registered on the wires. In particular, the threshold for removal of lower frequency components can be tuned to obtain narrower signals in time. This can improve the separation of closely spaced ionisation signals, however at the expense of the precision of the calorimetric reconstruction due to the loss of signal.

In addition, processing of induction plane signals is typically performed to convert the bipolar signals to unipolar ones. This simplifies the subsequent hit finding and calorimetric reconstruction by enabling the use of the same algorithms for all of the wire planes. This is usually performed as part of the software signal processing [280], although in some instances has been performed in the readout electronics using integrating amplifiers [183].

Hits and clustering

Once the raw signals have been processed, hit finding is performed. Hit finding characterises the position and amount of charge deposited on a wire within a Region of Interest (ROI) in which the wire signal is above a pre-defined threshold. Within each ROI a variable number of Gaussian functions are fit to the signal, depending

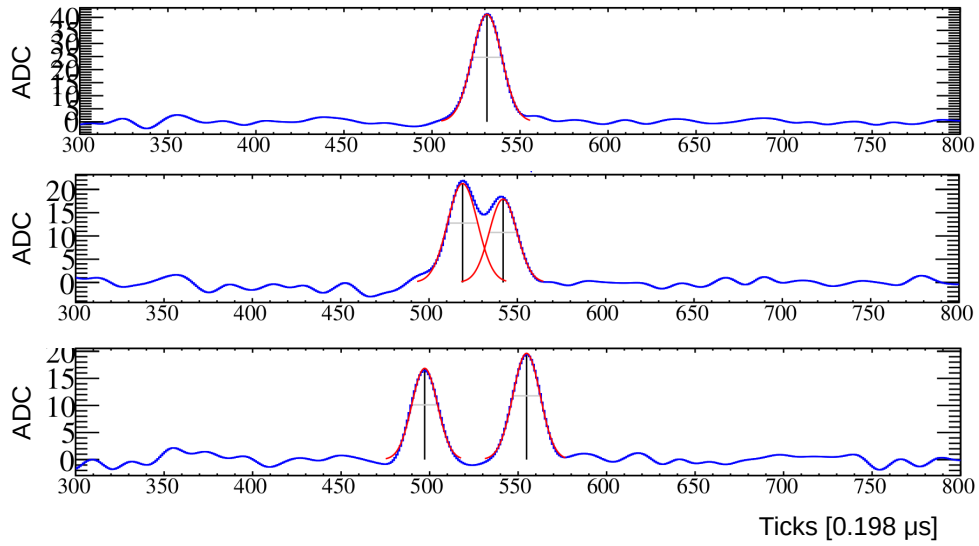

Figure 5.8: Example reconstructed hits on three different collection plane wires close to an interaction vertex in the ArgoNeuT detector. The interaction, shown in Figure 5.5, is producing two out-going approximately minimally ionising particles. Very close to the vertex (top) the signals overlap and only a single hit is reconstructed. As they begin to separate (middle) the signal becomes broader with multiple maxima and the hit finding algorithm fits multiple Gaussian functions. Further from the vertex (bottom) the signals fully separate allowing two hits to be cleanly reconstructed each fit with a single Gaussian.

on the number of local minima or maxima present. Figure 5.8 shows example reconstructed hits on three different collection plane wires close to an interaction vertex in the ArgoNeuT detector. The event shown is the same as seen in Figure 5.5. The differing number of Gaussians fit in the ROI are demonstrated as the pair of particles begin to separate.

After the hits in an event have been reconstructed, hit clustering is performed. This process groups hits together that are likely to have originated from a single interacting particle. In ArgoNeuT this is performed using the TrajCluster algorithm [235, 280], which is an evolution of the earlier LineCluster algorithm. TrajCluster attempts to group hits along trajectories in two-dimensional wire-time space on each sense-wire plane. A seed cluster is formed from a group of closely-spaced hits and the trajectory along these hits determined. Further hits found along this trajectory are then added, re-evaluating the local trajectory at each stage. This continues until a stopping condition is met such as the presence of a Bragg peak, indicating a stopping particle, or a two-dimensional interaction vertex. At each stage, the algorithm additionally attempts to match between the two-dimensional trajectories and vertices from each of the sense-wire planes to form three-dimensional clusters. An

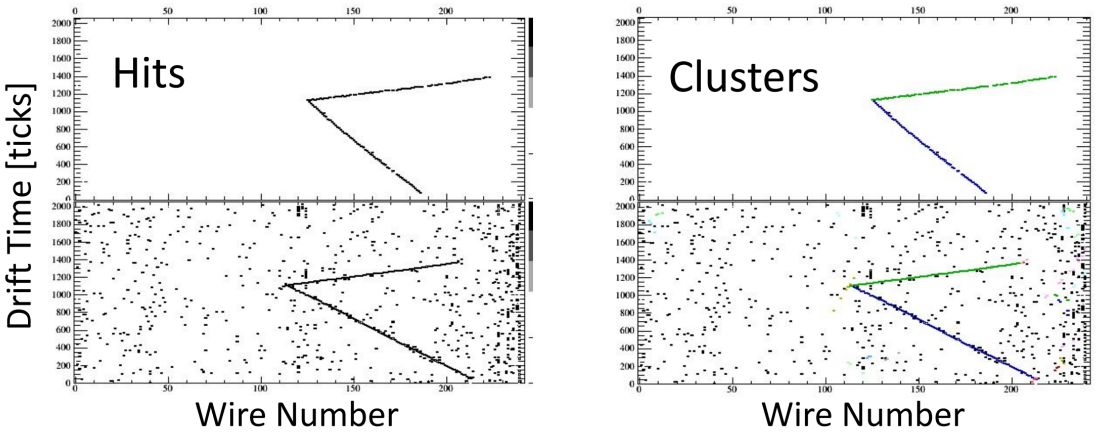

Figure 5.9: Example of the hit finding and then clustering algorithms being applied to a ν_μ candidate interaction in ArgoNeuT producing two out-going charged particles. The left figure shows the reconstructed hits on the collection (top) and induction (bottom) wire planes. The right figure shows the hits grouped into clusters that are matched between planes, denoted by the different colours.

example of the clustering algorithm being applied to a ν_μ candidate interaction in ArgoNeuT producing two out-going charged particles is shown in Figure 5.9. The left event display shows the reconstructed hits (black points), that are then grouped into two clusters on the right (green and blue points).

Tracks and showers

Next, tracks are reconstructed from the clustered hits. In ArgoNeuT this is performed using the Projection Matching Algorithm (PMA) [235, 285]. Three-dimensional tracks are constructed out of the clusters from the previous stage, matched between the sense-wire planes. The optimal position and trajectory of the track is identified by projecting it onto each available two-dimensional wire-plane view and minimising the difference between this projection and the reconstructed clusters and hits. As this is done, the clustering can also be iteratively re-optimised correcting spurious hit-to-cluster associations using the properties of the three-dimensional track. Additionally, in parallel, three dimensional vertices are identified at the start of isolated tracks or intersections between multiple tracks. These can then be used to further refine the reconstructed best matching position and trajectory of the tracks. Figure 5.10 shows the results of the PMA track reconstruction applied to the interaction previously shown in Figure 5.9. The two clusters of hits previously identified are reconstructed as two three-dimensional tracks.

Reconstruction of electromagnetic showers can also be performed. In ArgoNeuT this is achieved using the track reconstruction algorithm to determine the initial vertex and trajectory of the shower. The shower is then reconstructed around this

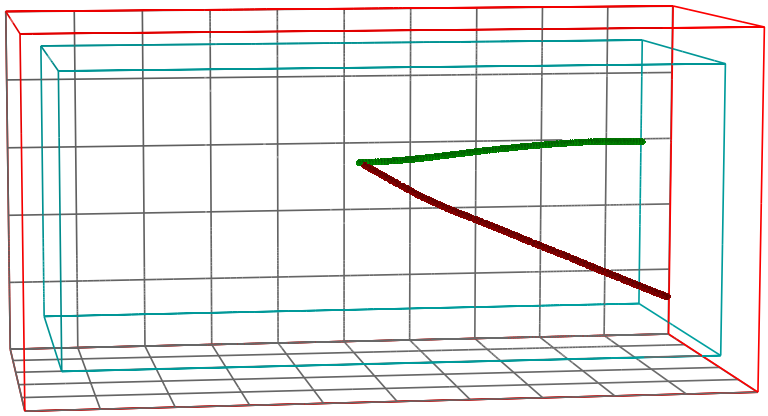

Figure 5.10: Result of the Projection Matching Algorithm three-dimensional track reconstruction applied to the ν_μ candidate interaction in ArgoNeuT previously shown in Figure 5.9. The two clusters of hits previously identified are reconstructed as two three-dimensional tracks.

initial axis by identifying nearby shower-like clusters of hits and is extended until the end of the shower is reached [163].

Calorimetry

Once the interacting particles have been reconstructed, calorimetry can be performed. The initial amount of charge per unit distance, dQ/dx, created by the interacting particle in the detector is determined by correcting the observed dQ/dx on the wire planes for losses or distortions of charge that occurred during propagation. These include accounting for the electron lifetime (impacted by charge loss due to attachment by impurities), the effects of diffusion and the impacts of space-charge. The relationship between the dQ/dx produced and the original energy deposited per unit distance, dE/dx, is then modelled using either the Birks model [170] or the modified Box model [171] accounting for recombination effects, as discussed in Section 3.3.1. In ArgoNeuT the modified Box model is used [159, 169].

Calorimetry in LArTPC detectors is typically performed using the charge information only, which provides the most accurate determination of the deposited dQ/dx. However, in modern LArTPC detectors with sophisticated photon-detection systems, the scintillation light can provide complimentary information that can enhance the performance of the calorimetric reconstruction. This has been demonstrated by the LArIAT experiment [286], which found significant improvements in the energy resolution for low energy interactions compared with the charge alone [287].

5.3.2 Light reconstruction

Signal processing

Analogous to the charge reconstruction, processing of the raw signals from the photon detectors is first performed to extract the scintillation signal from the electronics response. A similar deconvolution and noise removal process is applied, first performing a Fast Fourier Transform followed by using a Wiener filter.

Optical hits

Next, optical hit finding is performed. Similar to the ionisation charge hits, an optical hit quantifies the time and amount of light observed by a photon detector. Peaks are identified in the processed waveform and optical hits are defined when the waveform goes above a pre-defined threshold above the baseline. The optical hits then last until the waveform returns to this threshold. An example of two optical hits identified in a PMT waveform can be seen in Figure 5.11. In this case the peaks in the waveform, and hence the corresponding optical hits, are clearly separated. However, if a large number of photons arrive in a short time period, a single optical hit may contain multiple overlapping peaks. To obtain the number of incident photons in an optical hit, the signal is integrated across the full hit

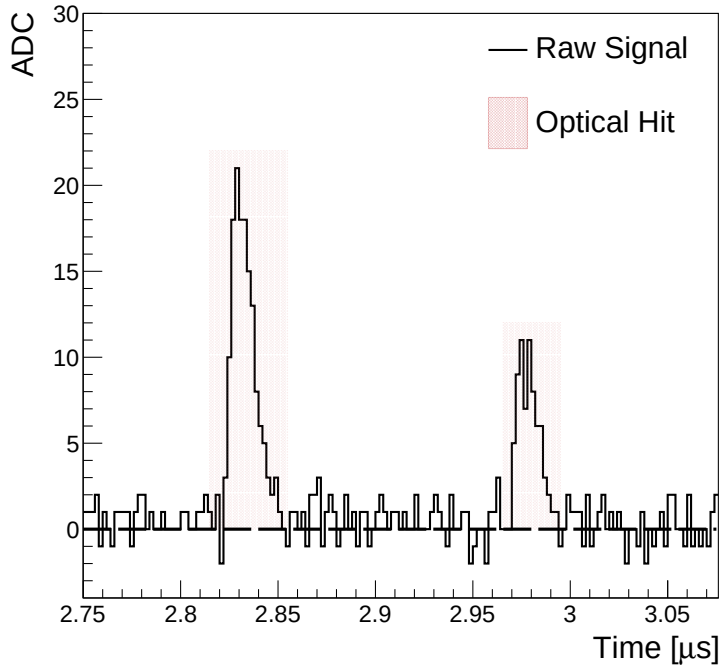

Figure 5.11: Example optical hits identified in the tail of the simulated PMT waveform shown in Figure 5.7 in SBND. The waveform has been subtracted from the initial baseline at 8000 ADC.

and divided by the expected single photo-electron response of the particular photon detector. The time of the optical hit is determined by the time of the first arriving photon, when the waveform first goes above threshold.

Flashes and flash matching

Once the optical hits have been identified for each photon detector, optical flashes are constructed. Clusters of optical hits that are in coincidence in time are identified and grouped into an optical flash allowing the total amount of light produced from an interaction in the detector to be determined. A flash is defined when a certain number of photons are observed in a particular time window. In SBND, for example, when 6 photons are observed within 10 ns. Once the initial flash time is identified, the light across all of the photon detectors is integrated over a fixed window. In SBND this window is 8 μs.

Once optical flashes have been reconstructed, they are then matched with the corresponding reconstructed ionisation signals to allow the light and charge to be used together in analysis. The ionisation tracks and showers are used to make a prediction of the number of photons incident on each photon detector by applying the fast optical simulation to each reconstructed three-dimensional segment of these objects. This forms a flash hypothesis for the interaction that can then be compared with each reconstructed flash to find the best match. The best-matching flash can then provide a precise interaction time, t_0, for the ionisation signal.

5.3.3 Particle identification

Track dE/dx

As discussed in Section 3.2.1, different species of particles produce tracks with different ionisation density. Therefore, the reconstructed energy deposited per unit length, dE/dx, along the ionisation tracks can be used to identify what type of particle produced them. Figure 5.12 shows the simulated dE/dx as a function of residual range, the distance before the particle stops, for several different particle species. The more highly ionising proton and kaon tracks can be distinguished from the less ionising muon and pion tracks. However, it is extremely difficult, if not impossible, to distinguish between muons and pions based on the dE/dx alone. The dE/dx of an example reconstructed stopping track in ArgoNeuT is also shown. This track is consistent with being a proton candidate [172].

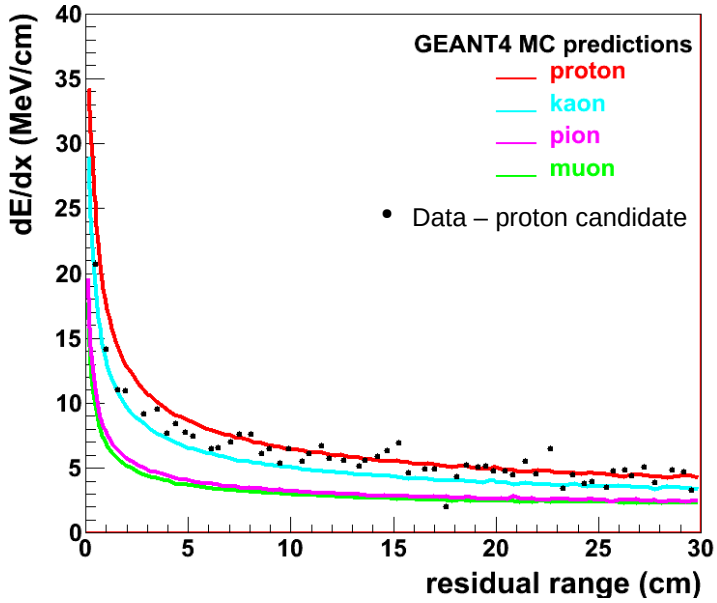

Figure 5.12: Simulated track dE/dx as a function of residual range for several different particle species in ArgoNeuT. A reconstructed proton candidate in ArgoNeuT data is also shown in the black points. Figure from Reference [172].

Electron-photon separation

Electromagnetic cascades resulting from primary electrons and photons can be distinguished based on their topology close to the interaction vertex prior to the cascade developing. Electrons initially produce a minimally ionising track, prior to losing energy via radiative emission of photons. Photons, however, are neutral and hence do not produce a signal in the detector until energy loss via pair production of electrons and positrons occurs. Therefore, for photons there is a gap between the interaction vertex and the start of the visible ionisation charge. The scale of this gap is $\mathcal{O}(10\,\text{cm})$ [288], and hence can easily be resolved in LArTPC detectors. However, this requires that additional particles are also produced in the interaction to allow the position of the interaction vertex to be determined.

An additional method of differentiating primary electron and photon electromagnetic cascades results from the fact that for photons, since an electron-positron pair is produced, the dE/dx of the initial track prior to the cascade developing is approximately double the dE/dx produced by a single electron. This allows electron and photon showers to be distinguished in the absence of vertex information or if the gap to the vertex is too small to resolve. Figure 5.13 demonstrates electron-photon separation using the reconstructed dE/dx of the track at the beginning of

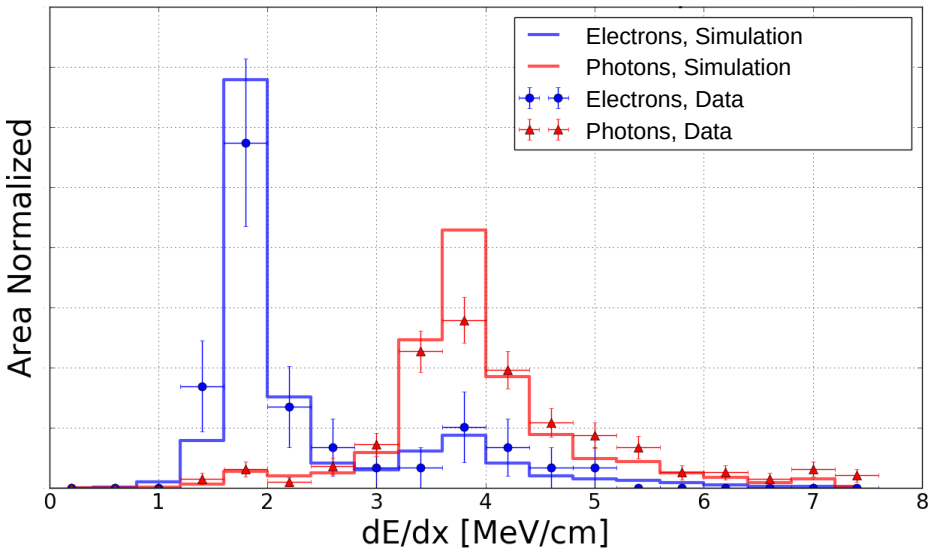

Figure 5.13: Electron-photon separation in ArgoNeuT using the reconstructed dE/dx of the track at the beginning of the electromagnetic cascade. Electrons and photons from both simulation and data are shown. Figure from Reference [288].

the electromagnetic cascade in the ArgoNeuT detector. Electrons and photons from both simulation and data are shown [288]. The electron and photon interactions can be clearly distinguished based on the track dE/dx.

Pulse-shape discrimination

The scintillation light signal can also be exploited for particle identification. As discussed in Section 3.4.1, the ratio between the prompt and late components of the light depends on the ionisation density created by the interacting particle. Therefore, the ratio between the amount of light observed in the first ∼100 ns and the total amount of light integrated over the full waveform, F_{prompt}, can be used to distinguish between particle species [198]. This technique, commonly referred to as pulse-shape discrimination, has been used extensively in liquid argon direct-detection dark matter searches [186–188]. Figure 5.14 shows an example of the technique being applied to distinguish between low energy electrons and low energy nuclear recoils in this context. Effective discrimination between these two types of interactions can be achieved. In LArTPC detectors pulse-shape discrimination is of interest for rejection of radiological backgrounds, such as highly ionising α-particles produced from radon decays, in searches for lower energy non-beam neutrinos such as solar neutrinos [248]. Since the scintillation light signal is available on a short time-frame relative to the ionisation charge, pulse-shape discrimination could be used during detector triggering to reject such backgrounds. This could allow searches for lower energy neutrino interactions than would otherwise be possible.

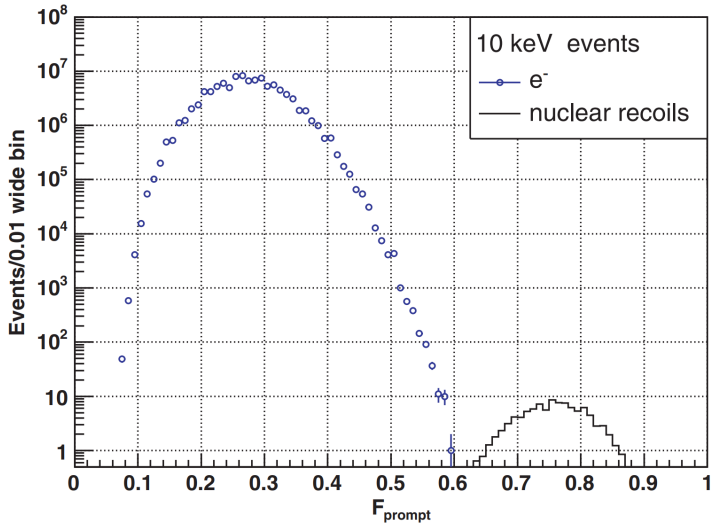

Figure 5.14: Example of pulse-shape discrimination applied to distinguish between low energy electrons and low energy nuclear recoils in direct-detection dark matter searches. Figure from Reference [198].

5.3.4 Reconstruction frameworks

The reconstruction used by ArgoNeuT described in the previous sections is relatively simple compared with larger and newer experiments. Many modern LArTPC detectors use more sophisticated reconstruction frameworks that incorporate and expand upon the techniques described previously. The most widespread of these is the Pandora reconstruction framework [289]. This is in use by the SBND, MicroBooNE, ICARUS and DUNE collaborations. Pandora uses a layered multi-algorithm pattern recognition approach to reconstruct and characterise the different particles present in LArTPC interactions. It incorporates both conventional and machine-learning based algorithms. These are designed to be generic and largely independent allowing them to be applied in different detectors and for different purposes. They can be combined together to form different reconstruction paths focusing of particular types of interactions, typically with the earliest algorithms identifying obvious topologies and then progressively increasing in complexity. For surface LArTPCs such as MicroBooNE and SBND, two optimised reconstruction paths are used: PandoraCosmic, optimised for cosmic-ray interaction identification focusing on track-like topologies; and PandoraNu, optimised for neutrino interaction identification placing emphasis on vertex identification and careful reconstruction daughter tracks and showers [290].

Various alternative reconstruction frameworks have also been developed. The Wire-Cell framework, employed in the MicroBooNE detector, takes a tomographic approach focusing on performing three-dimensional reconstruction of the deposited

ionisation charge prior to employing pattern recognition techniques [291, 292]. This contrasts with the conventional approach followed in ArgoNeuT and the Pandora framework where pattern recognition is applied based on two-dimensional information from each wire plane and then subsequently matched between planes to form a three-dimensional image. In recent analyses performed by MicroBooNE [293, 294], the Wire-Cell approach has been shown to significantly outperform the conventional techniques. The expansion of Wire-Cell reconstruction to other detectors, such as DUNE, is currently on-going. A deep-learning based reconstruction framework has also been developed, again employed in the MicroBooNE detector. This approach focuses on the use of semantic segmentation with convolutional neutral networks to identify and characterise interactions using the two-dimensional images provided by each wire plane [295–299]. This framework is currently limited to specific interaction topologies, however more comprehensive deep-learning based reconstruction is in development.

5.4 Geometry models used in simulation

5.4.1 DUNE far detectors

Horizontal drift module

Due to the size of the DUNE horizontal drift (DUNE-HD) module, described in Section 3.6.3, simulation of the full detector is extremely computationally challenging in terms of both CPU time and memory requirements. Instead for simulation studies, a geometry modelling a subset of the DUNE-HD module is used [182]. The geometry consists of the two drift volumes on either side of the central APA. It includes the full 12 m (2 APA) height of the detector, but is restricted to a depth in the beam direction of 14 m (6 APAs). This subset of DUNE-HD geometry will be referred to as the DUNE-HD 1×2×6 geometry subsequently. The DUNE-HD 1×2×6 geometry has an active volume of $7.0 \times 12.0 \times 14.0\,\mathrm{m}^3$ (horizontal (drift), vertical, beam direction) and 120 PD modules. The results of simulation in this geometry can be extrapolated to the full detector based on symmetry: the two remaining drift volumes will behave approximately the same as the two simulated; and in the beam direction the central regions of the DUNE-HD 1×2×6 approximate the majority of the detector volume that is far from the end walls, while the ends model the ends of the full detector.

Vertical drift module

Similar to the DUNE-HD module, the full DUNE vertical drift (DUNE-VD) detector geometry, described in Section 3.6.3, is not used in simulation studies due to computational constraints. Instead, subsets of the detector are used that are representative of the full geometry [251]. The geometries consist of one of the two drift volumes and, analogous to the DUNE-HD case, include the full width of the detector but are shortened in the beam direction. Two different geometries are currently in use. The first, referred to as the DUNE-VD 1×8×14 geometry, has an active volume of $13.5 \times 6.5 \times 21.0\,\mathrm{m}^3$ (horizontal, vertical (drift), beam direction) and 224 X-ARAPUCA PD modules. This geometry was used for the development of the light simulation. A second smaller geometry, referred to as the DUNE-VD 1×8×6 geometry, is now in use by default. It has an active volume of $13.5 \times 6.5 \times 9.0\,\mathrm{m}^3$ (horizontal, vertical (drift), beam direction). The use of this geometry is necessitated by memory constraints resulting from the simulation of the ionisation charge. However, it is less representative of the full detector due to the shorter length in the beam direction resulting in border effects from the end walls that are larger than they would be in reality. As the simulation of the DUNE-VD module becomes more sophisticated, the geometries used will be refined.

5.4.2 Generic geometry models for light simulation development

Custom detector geometries will be used for the development and validation of the semi-analytical fast optical simulation model, described in Chapter 6. These are defined using the GDML markup language [300] to be realistic approximations of typical LArTPC detectors. Figure 5.15 shows a schematic representation of the geometries. The main volume of liquid argon is delimited by metal walls of a cryostat[1]. The field-cage surrounding the active volume is modelled as an array of metal strips on the top, bottom, upstream and downstream walls and spaced to provide ∼30% optical transparency[2]. The photon detectors (PDs) are uniformly distributed across the open plane of the field-cage (left), referred to as the *photon-detector plane* (PD-plane). The PD sensitive windows are simulated as flat disks or rectangles facing the active volume. The cathode plane (right) is modelled as an opaque volume covered by reflector foils coated in a wavelength-shifter (WLS). The WLS used in the simulations is TPB, for which the absorption and emission spectra are taken

[1] Photons leaving the field-cage are simulated and can in principle reflect off the cryostat walls and be detected on PDs.

[2] This is a typical value in most real LArTPC detectors, with the exception of the DUNE vertical drift detector that will have significantly slimmer field cage profiles allowing greater transparency.

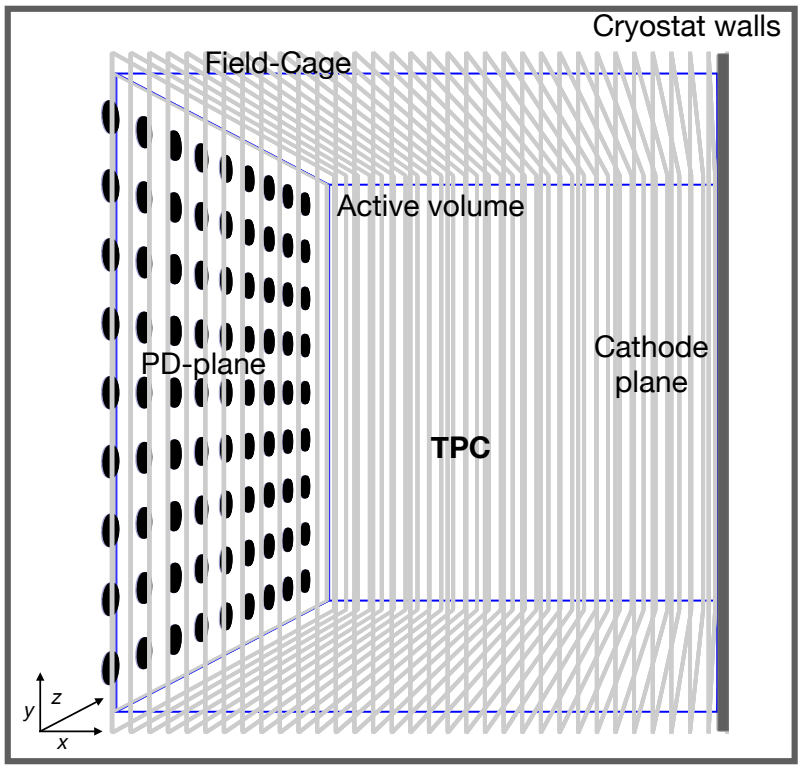

Figure 5.15: Illustration of the LArTPC detector geometries used in the simulation. In this example, the PDs are modelled as flat disks. Diagram not to scale.

Material	VUV reflectivity	Visible reflectivity
metal	25% [302]	60% [302]
reflector foils	N/A	93% [301]

Table 5.3: Material reflectivities used in the simulation. The VUV reflectivity of the reflector foils is not defined since 100% of the incident light is assumed to be converted to visible wavelengths.

from [301].

The reflectivities of the materials to VUV photons (pure scintillation emission) and visible photons (coming from the re-emission of the VUV photons absorbed by the TPB on the reflector foils) are summarised in Table 5.3. In the simulations used in this work the reflections on optical boundaries have been modelled as Lambertian on a rough surface and it is assumed that 100% of light incident on the WLS is converted to visible wavelengths.

The model is tested in two different geometries, corresponding to two different experiments employing LArTPC detectors: *SBND-like* and *subset-of-DUNE-like* (henceforth referred to as *DUNE-like*). The main properties of these two geometries are summarised in Table 5.4. These geometries are approximations of SBND, described in Section 3.6.2, and of the 1×2×6 subset of the DUNE horizontal drift detector, described in Section 5.4.1, respectively. In each case a single TPC from the detectors is modelled. In both geometries the PDs are distributed approximately

Parameter	SBND-like	DUNE-like
width [cm]	200	365
height [cm]	400	1200
length [cm]	500	1400
number of PDs	66	123
PD shape	disk	rectangle
PD size	8" diameter	$9.3 \times 9.3\,\mathrm{cm}^2$

Table 5.4: Summary of the main properties of the two geometries used.

evenly and with a PD located in the exact centre of the PD-plane. In the SBND-like geometry they are modelled as disks, representing PMTs, and in the DUNE-like geometry they are modelled as rectangles, representing X-ARAPUCAs.

5.5 ArgoNeuT's simulation of the MINOS near detector

5.5.1 MINOS-ND simulation and reconstruction

Simulation and reconstruction of particles that exit ArgoNeuT and are seen downstream in the MINOS near detector (MINOS-ND) is performed using the standard MINOS software suite [237, 239, 303] using dedicated tools provided by the MINOS collaboration for the ArgoNeuT and MINERvA experiments [172, 257]. In simulation, particles exiting the ArgoNeuT detector are propagated using the standard Geant4-based simulation, described in Section 5.2.1, to a surface corresponding to the start of the MINOS-ND. The final positions and momenta of these particles are then provided as input to the MINOS-ND simulation. The standard MINOS-ND simulation chain is run performing particle propagation with GEANT3 [304], modelling of the magnetic field using finite element analysis, and simulation of the expected detector response [239, 303].

As described in Section 3.6.1, the signals seen on the alternating orthogonal scintillator strip planes allow three dimensional reconstruction of the energy depositions created by interacting particles. A Hough transform is first used to identify linear segments of tracks [305]. These are then chained together to form longer tracks taking into account position and timing information. The curvature of the track in the magnetic field is then determined using a Kalman Filter technique to find the trajectory [306]. The momentum of the particle is reconstructed using a range based approach for contained tracks and using the curvature in the magnetic field for exiting tracks. The charge of the interacting particle is also identified using the track curvature [239, 303]. Finally, calorimetric reconstruction is performed allowing the dE/dx of the tracks to be evaluated and particle identification techniques to be

applied.

The track truth and reconstructed information output from the MINOS-ND simulation and reconstruction is then merged with the ArgoNeuT simulation, allowing the tracks to be accessed and used in subsequent ArgoNeuT reconstruction and analysis within LArSoft [172].

5.5.2 ArgoNeuT–MINOS-ND matching

Events in the ArgoNeuT and MINOS-ND data are grouped together based on common timestamps of spills from the NuMI beam [172]. Matching between tracks exiting ArgoNeuT and tracks reconstructed in the MINOS-ND is then performed based on the relative orientation and positions of the tracks in each detector. Tracks reconstructed in ArgoNeuT are projected to their expected start position and trajectory in the MINOS-ND. These are then compared with each reconstructed MINOS-ND track. Figure 5.16 shows an illustration of several different matching scenarios. Tracks are considered to match if their position and trajectory are within a radial

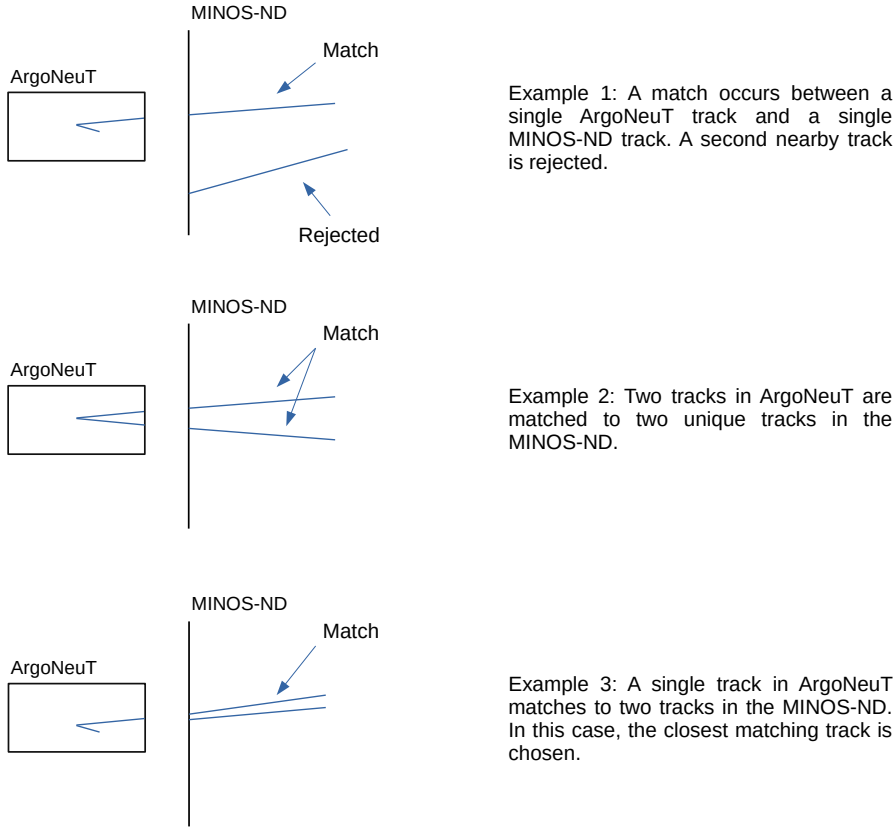

Figure 5.16: Examples of several different matching scenarios between ArgoNeuT and the MINOS-ND.

and angular tolerance: r_{diff} and θ_{diff}, respectively. The standard tolerances used in recent ArgoNeuT analyses are $r_{diff} = 12.0\,\text{cm}$ and $\theta_{diff} = 9.74^{\circ}$. These take into account uncertainty in the projection due to the angular resolution in ArgoNeuT along with the expected deflections of the particles as they propagate between ArgoNeuT and the MINOS-ND [235].

Matching between ArgoNeuT and the MINOS-ND allows further particle identification to be performed and is especially powerful for distinguishing between muons and pions that would otherwise be difficult in a LArTPC alone. In the energy range of interest for particles resulting from neutrino interactions, the majority of muons will exit ArgoNeuT and be reconstructed in the MINOS-ND. The majority of pions and protons, however, will either stop before reaching the MINOS-ND, or will only form very short tracks. Therefore, by imposing a matching requirement, muons can be identified. Additionally, the magnetic field present in the MINOS-ND allows the muon charge to be reconstructed enabling neutrino and anti-neutrino interactions to be distinguished.

Chapter 6

Modelling transport of scintillation light in large scale LArTPCs

This chapter describes the semi-analytical model for fast simulation of the transport of scintillation light in large scale LArTPC detectors. Models for predicting both the number of photons arriving on each photon detector and the distribution of their arrival times will be described for two generic LArTPC detectors, corresponding approximately to the SBND and DUNE horizontal drift detectors. The models are split into two components which describe the direct VUV light and the reflected visible light from wavelength-shifting reflective foils. The model is also extended to simulate xenon-doped liquid argon. Finally, the implementation of the model in several real LArTPC detectors will be described. The basic models describing the direct VUV and reflected visible light in liquid argon have been published in Reference [4]. Sections 6.1 to 6.4 cover the material presented there, and hence the majority of the figures and discussion are adapted from this publication.

The development of the models presented in this chapter was performed in collaboration with D. Garcia-Gamez (University of Granada) and A. M. Szelc. (University of Edinburgh). The author contributed to the refinement of the direct light models, was the primary developer of the reflected light models and contributed extensively to the development of the xenon-doped argon models. The author was solely responsible for the implementation of all of these models into the LArSoft software framework. The author was also responsible for the applications of the reflected and xenon light models in the SBND and DUNE horizontal drift detectors, and was solely responsible for the development and application of the extended model for the DUNE vertical drift detector.

6.1 Predicting the number of detected photons

This section describes the development of the model to predict the number of photons arriving at any given photon detector, based on the size of an energy deposit and the position where it occurred. Using the notation from Equation 3.7, the model described in this section is designed to estimate the geometric coverage of the photon detector, $P(d, \theta)$, and the impact from transport effects, $T(d, \theta)$. Two different detector geometries are considered: one SBND-like and one DUNE-like, each of which are described in Section 5.4.2.

6.1.1 Direct light

Geometric considerations

Scintillation light is emitted isotropically. This means that in an ideal case the number of photons arriving at a given photon detector (PD) could be calculated by simply estimating its geometric acceptance with respect to the scintillation point by determining the solid angle subtended by the sensitive window. This is a calculation that can be performed either analytically or with simple numerical integration depending on the shape of the PD. Solutions exist that cover most existing PD designs, for example for disk [307] and rectangular [308] shapes. The models described in this work can be extrapolated to any other PD shape, provided that the calculation of the solid angle subtended from a point-like source by such a shape can be determined.

This approach works in an idealised scenario: in the absence of Rayleigh scattering and reflections by the detector materials ($\lambda_{RS} \to \infty$ and all materials 100% absorptive). In this scenario, the calculation becomes purely geometric and for any given energy deposition, ΔE, at position (d, θ) the number of photons incident on a PD can be evaluated as,

$$N_{\Omega} = \Delta E \times S_{\gamma}(\mathscr{E}) \times Q_{abs}(d) \times \frac{\Omega}{4\pi}, \tag{6.1}$$

where the $S_{\gamma}(\mathscr{E})$ is the scintillation yield of LAr for a given electric field and Ω is the subtended solid angle. The position (d, θ) is defined in terms of the distance, d, between the energy deposit and the PD, and the offset angle, θ, between the energy deposit and the normal to the PD surface. Absorption effects due to contaminants can also be accounted for via $Q_{abs} = e^{-\frac{d}{\lambda_{abs}}}$, where λ_{abs} is the absorption length and d is the distance to the PD. The performance of Equation 6.1 at predicting the number of photons can be seen in Figure 6.1 (top). This shows a comparison between the number of photons hitting the PD windows predicted by Equation 6.1 and the number obtained from a full Geant4 simulation, normalised to the sensitive-

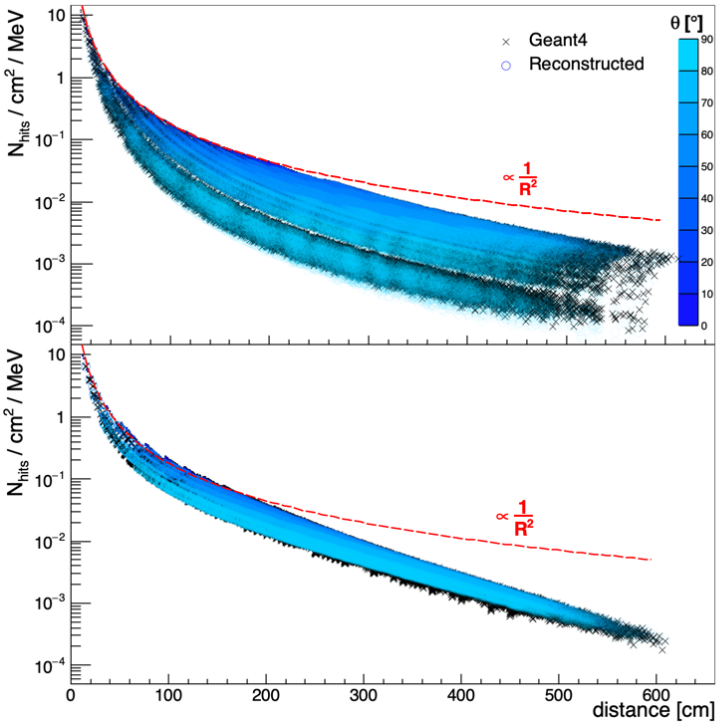

Figure 6.1: Top: Number of Geant4 tracked (black crosses) and analytically predicted (blue circles) scintillation photons arriving at the PDs per unit of deposited energy and PD sensitive-window area, in the SBND-like detector geometry. In this case Rayleigh scattering is not included and all material reflectivities are set to zero. The red dashed line represents an ideal $1/R^2$ behaviour. It diverges from the simulated points when the size of the detector excludes any further points on-axis to the PDs, at $d = 200\,\text{cm}$. The dependence on the offset-angle θ is indicated by the shades of blue of the reconstructed circle-markers. Bottom: Variation of the top scenario where Rayleigh scattering [206] is included. Rayleigh scattering strongly shapes the amount of light observed in the PDs.

window area and the energy deposited. The pure-geometric calculation agrees with the full simulation within expected Poisson fluctuations. The gradient-colours of the circles represent the offset angle, θ. A larger amount of light is observed from emission points that are closer and more on-axis to the PDs, as expected. The red dashed line indicates a perfect $1/R^2$ behaviour. Even in this simplified case, it is necessary to account for the offset angle which can change the prediction for a given distance by up to two orders of magnitude.

Corrections to the geometric approach

The simplified geometric approach breaks down once Rayleigh scattering is introduced into the simulation. The VUV scintillation photons in LAr undergo scattering during propagation with a characteristic length, λ_{RS}, that is small compared to the size of current and future LArTPC experiments. This alters the path of the major-

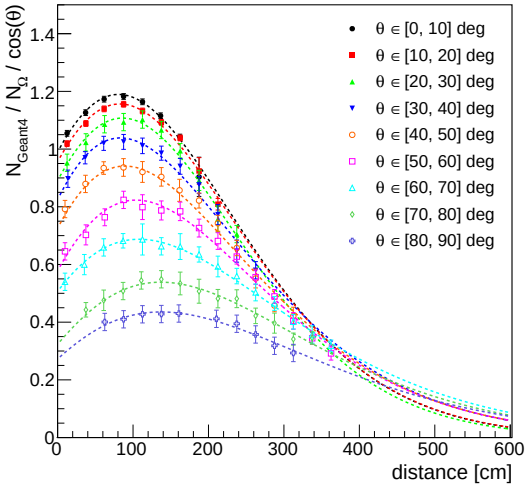

Figure 6.2: Relation between the number of Geant4 simulated hits on the PDs and the pure geometric estimation described by Equation 6.1, in the SBND-like geometry. The error bars represent the standard deviation of the distribution within each angular bin. A strong dependency is clear in both distance and offset angle. At any angle, the dependency of the ratio with distance can be accurately described by a Gaisser-Hillas function, as illustrated by the dashed curves. To avoid divergences at large offset angles it was found to be more convenient to work with the projected solid angle.

ity of the photons and consequently the number of them arriving at the PDs. Once Rayleigh scattering is included in the Geant4 simulation, the distribution of points in Figure 6.1 (top) is altered to Figure 6.1 (bottom, black crosses). The Rayleigh scattering significantly impacts the amount of light observed by the PDs and it is therefore necessary to account for it during simulation. The effect strongly depends on both the distance, d, and the offset angle, θ, of the PD relative to the light emission point.

To build corrections for the effects of Rayleigh scattering, a ratio is calculated between the number of incident photons from Geant4 simulation, N_{hits}, and the geometric estimation from Equation 6.1 projected on $\cos(\theta)$, $N_\Omega/\cos(\theta)$. For simplicity, the phase space is split into 10° bins in θ. The discretisation in θ introduces a systematic effect in the model: a more (less) sampled choice would result in a more (less) accurate correction. The choice of angular binning is a trade off between accuracy and computational time during training. The resulting ratios, shown in Figure 6.2, are smooth distributions as a function of distance and clearly separated between the different angular bins. This indicates that a parameterisation in (d, θ) incorporates the primary dependencies and consequently is sufficient to predict the number of arriving photons.

The distributions shown in Figure 6.2 can, for all angles, be accurately described

using Gaisser-Hillas (GH) functions [309], as illustrated by the dashed curves. These are defined as,

$$GH(d) = N_{\max} \left(\frac{d - d_0}{d_{\max} - d_0} \right)^{\frac{d_{\max} - d_0}{\Lambda}} e^{\frac{d_{\max} - d}{\Lambda}}, \tag{6.2}$$

where N_{max} is the maximum of the function located at a distance d_{max}, and d_0 and Λ are parameters describing the width of the distribution. The GH functions are implemented as the core of the semi-analytical model to predict the scintillation light signals in large LArTPC detectors[1]. The number of incident photons on each PD is first predicted by the solid angle that the aperture of the detector subtends, then the effects of Rayleigh scattering are accounted for via corrections to the geometric prediction. Once these corrections are applied to Equation 6.1, the model precisely predicts the number of incident photons on each PD as shown qualitatively in Figure 6.1 (bottom, blue circles). A quantitative comparison will be discussed in Section 6.4.1.

In the limit $d \to 0$, the effects of Rayleigh scattering should be negligible and the y-intercept in the corrections should correspond to the value $\cos(\theta)$ (approximately 1 for the on-axis case). The Gaisser-Hillas-like shape of the corrections suggests a behaviour of the light such that for small distances from the PD, the probability of detecting scattered photons that would otherwise escape from the detectors is larger than the fraction that is lost due to the scattering. This situation continues for larger distances until a point at which it is reversed and more photons are lost than gained. Additionally, once taking into account the $1/\cos(\theta)$ factor in Figure 6.2, it can be seen that PDs at large θ (that have a small geometric acceptance) present a higher relative probability to recover scattered photons compared with PDs located closer to on-axis. This is seen in Figure 6.1 as a significant tightening of the angular dependence when Rayleigh scattering is included (bottom) compared to the ideal case (top). These effects also result in the detector size having an impact on the required correction curves: the greater the active volume in which photons can scatter, the greater the probability that these photons will end up increasing the signal. The impact of the size of the detector under consideration can be accounted for by altering the corrections, which will be described next.

Correcting for detector size: border effects

The dependency of the derived corrections on the detector size can be treated as a border effect. These borders (the cryostat walls, field-cage and other detector components) not only delimit the active volume where photons can travel and scatter,

[1] As the Gaisser-Hillas function can be shown to be equivalent to a Gamma distribution, the latter could be used with similar results. The Gaisser-Hillas function definition was chosen in this model for mathematical convenience.

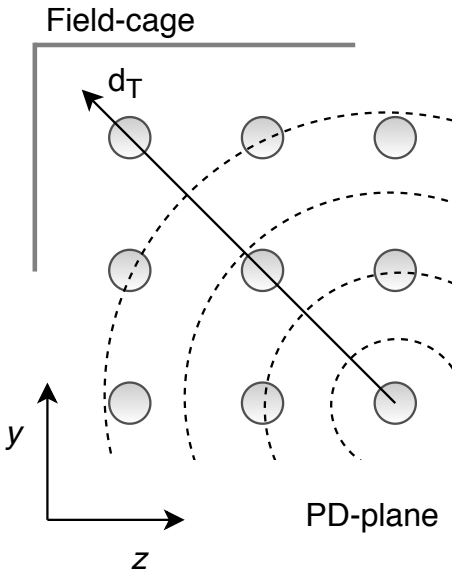

Figure 6.3: Illustration of the concentric cylinders at different radial distances, d_T, from the centre of the photon-detector plane (PD-plane). Each grey disk represents a PD. Diagram not to scale.

but also consist of surfaces that can reflect or absorb them. These effects influence the amount of light observed in the PDs and, as a consequence, the required corrections depend on the position of the scintillation light emission. To study the impact of the borders of the detector and cover different regions where the response may vary, the volume is divided into concentric cylinders at different radial distances, d_T. These range from the PD at the centre of the photon-detector plane (PD-plane) $(Y - Z)$ outwards towards the corners of the field-cage, as illustrated in Figure 6.3.

The behaviour of the parameters of the Gaisser-Hillas functions is examined as the scintillation position is varied in d_T. For simplicity, and taking advantage of the strong correlation between the d_0 and Λ parameters of the Gaisser-Hillas functions, the value of d_0 is fixed absorbing all of the d_T dependencies into the remaining three parameters. Figure 6.4 shows the results for the N_{max}, d_{max} and Λ parameters, for both of the detector geometries considered. An approximately linear dependence in d_T is seen for all of the offset angle bins. The slopes of the lines also depend on θ, increasing for the more off-axis cases. These trends can be taken into account to accurately estimate the number of scintillation photons arriving at a PD for the entire active LAr volume.

To correct for the border effects, the Gaisser-Hillas parameters in Equation 6.2

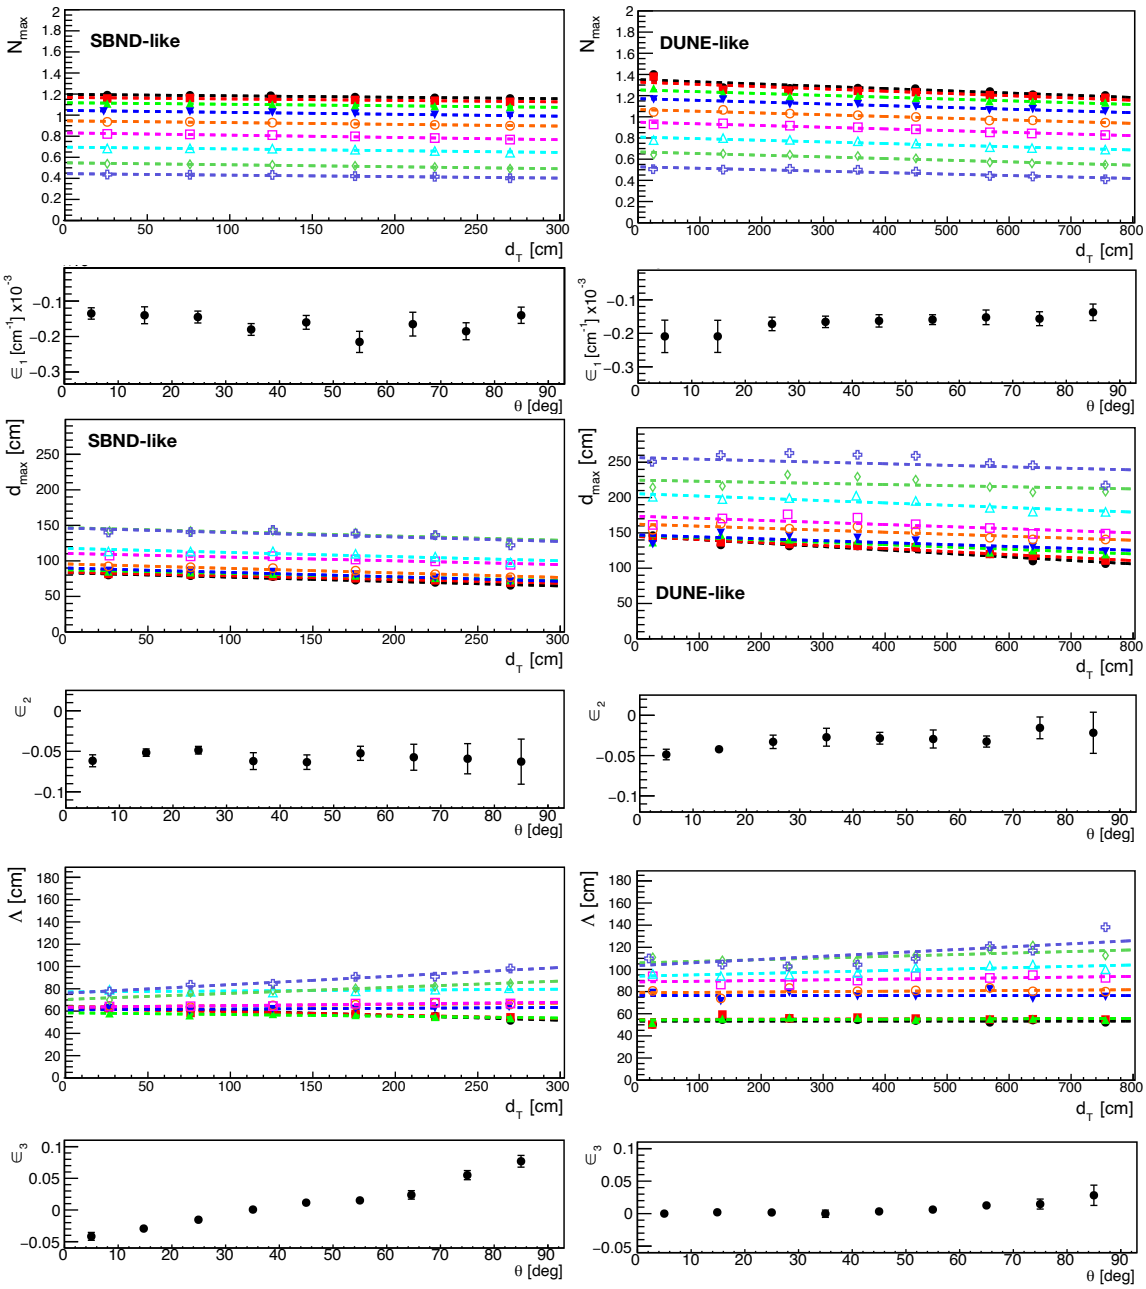

Figure 6.4: N_{max} (top), d_{max} (middle) and Λ (bottom) Gaisser-Hillas parameter dependency on distance to PD-plane centre, d_T, for the SBND-like (left) and DUNE-like (right) geometries. The different colours represent to different θ bins as shown in Figure 6.2. The lines represent the linear fit of the points. In each case, the slopes ϵ_1, ϵ_2 and ϵ_3 of the linear fits for the different offset angles are shown in the lower panels.

are redefined as:

$$\begin{aligned} N'_{\text{max}} &= N_{\text{max}} + \epsilon_1(\theta) d_T \\ d'_{\text{max}} &= d_{\text{max}} + \epsilon_2(\theta) d_T \\ \Lambda' &= \Lambda + \epsilon_3(\theta) d_T, \end{aligned} \tag{6.3}$$

where $N_{\max}$, $d_{\max}$ and Λ are the values of the parameters in the centre of the PD-plane ($d_T = 0\,\text{cm}$), and ϵ_1, ϵ_2 and ϵ_3 are the slopes of the linear corrections for each parameter respectively. This new function is referred to as GH'. To give an indication how these corrections affect the probability of photons arriving at PDs, Figure. 6.5 shows examples of correction curves for the two extreme cases, centre (top) versus corner (bottom), for both the SBND-like and DUNE-like geometries. The correction curves vary noticeably depending on both the detector size and the position of the scintillation, illustrating the impact of the border effects.

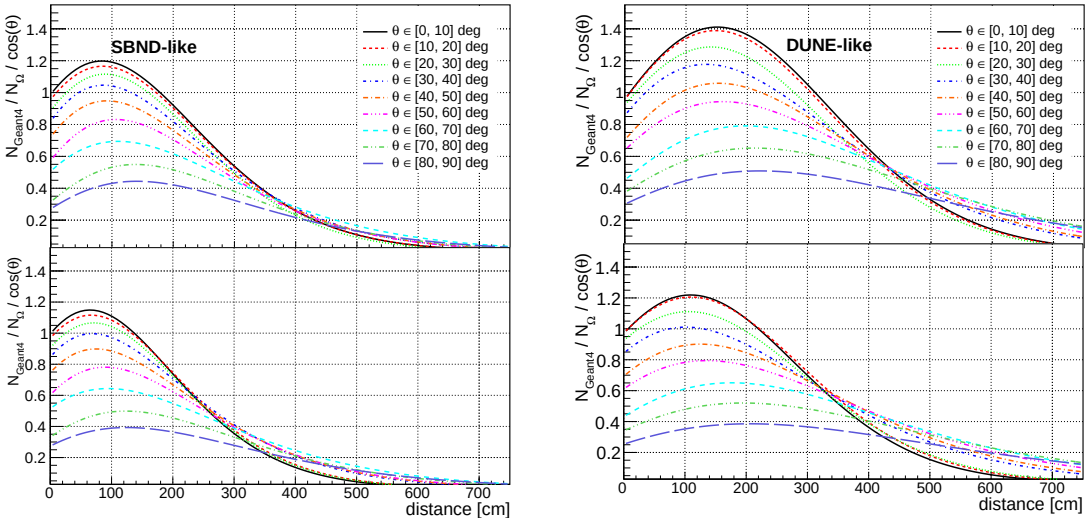

Figure 6.5: Correction curves for the two geometries under study, for scintillation in the centre of the TPC (top) and in the farthest corner (bottom).

Bringing all of the above effects together, the model is able to estimate the number of detected scintillation light photons using Equations 6.1, 6.2 and 6.3, combined as:

$$N_\gamma = N_\Omega \times GH'(d, \theta, d_T)/cos(\theta), \tag{6.4}$$

which depends only on the distance and angle between the emission point and the PD, and distance of the emission point from the centre of the detector. Here, using the notation from Equation 3.7, $N_\gamma = \Delta E \times S_\gamma \times Q_{abs}(d) \times P(d, \theta) \times T(d, \theta)$.

6.1.2 Reflected light

Basic geometric model

The number of photons arriving at the PDs in a LArTPC detector that has wavelength-shifting reflector foils on the cathode can also be predicted using a geometric approach. This requires expanding the model developed for the direct light case described in Section 6.1.1. The prediction of the wavelength-shifted and reflected *visible* light is inherently dependent on the specific detector geometry because the

distance between the reflective foils and the PDs becomes a key element of the model. Additionally, unlike for the VUV light, the wavelength-shifted light is much more likely to reflect from the borders of the detector and the field cage. Therefore, the model for wavelength-shifted light is constructed using a realistic detector geometry from the start rather than using an idealised detector.

The visible light arrives at the PDs after being re-emitted and possibly reflected by the WLS-coated reflector foils at the cathode of the detector. Therefore, the number of VUV photons incident on the reflector foils is first calculated using the solid angle that the entire cathode subtends, Ω_c. This is corrected for the effects of Rayleigh scattering using Equation 6.4. It is then assumed that these photons are re-emitted approximately isotropically after being wavelength-shifted and reflected, and that the region of the cathode in front of the scintillation in the drift direction will be the dominant source of the visible photons. The central point of this region is referred to as the *bright-spot*, and is illustrated in Figure 6.6 together with the other elements of the geometric model for the reflected light.

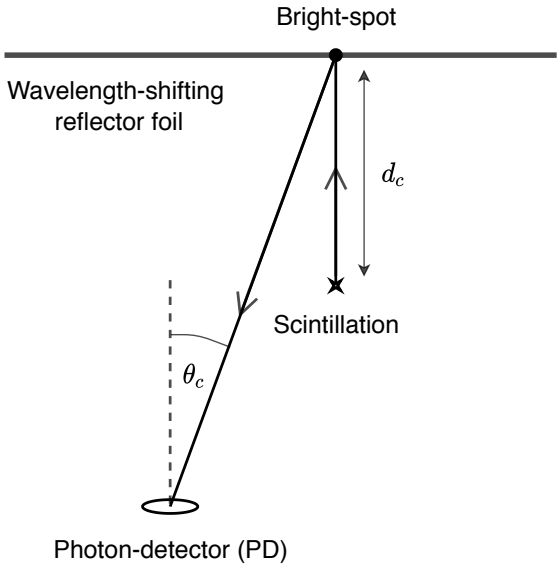

Figure 6.6: Diagram illustrating the geometric model for predicting the number of photons incident on the PDs as a result of wavelength-shifting reflector foils on the detector cathode and predicting the arrival-time distribution of these photons.

The number of photons incident on each PD can then be calculated using the solid angle subtended by the PD aperture as viewed from the bright-spot, Ω_{PD}. The geometric prediction for the number of visible photons arriving at the PDs can therefore be expressed as,

$$N_{\Omega,reflected} = N_{\gamma,direct}(\Omega_c, d_c, \theta_c, d_T) \times Q_r \times \frac{\Omega_{PD}}{2\pi}, \tag{6.5}$$

where $N_{\gamma,direct}(\Omega_c, d_c, \theta_c, d_T)$ is the prediction of the number of photons incident on

the cathode using the direct VUV light model given by Equation 6.4 and $Q_r = Q_{WLS} \times Q_{foil}$ is a scaling factor accounting for the WLS efficiency, Q_{WLS}, and the foil reflectivity, Q_{foil}. The solid angle of the PD, Ω_{PD}, is divided by 2π rather than 4π due to the presence of the highly reflective foils beneath the WLS which halves the total solid angle light is able to propagate in.

Corrections for PD position

The basic geometric model provides an initial approximation of the number of reflected photons incident on each PD. However, the assumption that the bright-spot region is dominant does not fully account for the distribution of the re-emitted wavelength-shifted photons across the whole surface of the reflective cathode. The approximation performs well for the PDs placed close to on-axis (at small θ_c) that see the majority of the light. However, the approximation worsens for the PDs located further off-axis where a larger fraction of the observed light originates from regions of the cathode opposite to the PD rather than the bright-spot. Corrective factors are therefore implemented to the basic model to account for these effects in an analogous way to the direct light model described in Section 6.1.1. Because the corrections are developed in a realistic geometry, they also account for effects of reflections of the wavelength-shifted photons from the field-cage, cryostat walls and other detector components.

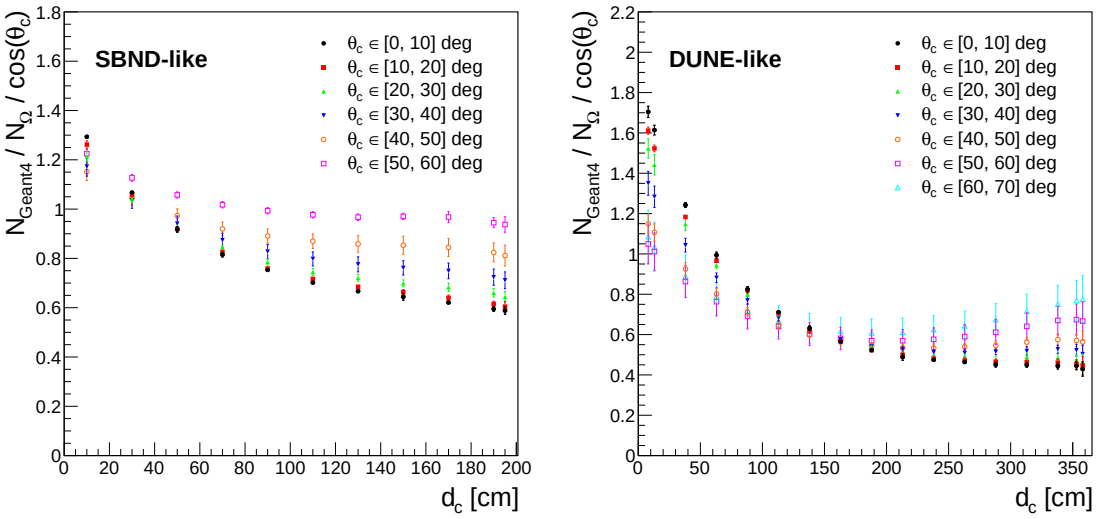

Figure 6.7: Ratio between the number of photons incident on the PDs in Geant4 simulation and the prediction from the reflected light geometric model for scintillation occurring in the central region of the SBND-like (left) and DUNE-like (right) detector geometries.

Similar to the method used for the direct light model, the required corrections are found by taking a ratio between the number of incident photons on the PDs in

Geant4 simulation and the prediction from the basic geometric model. For scintillation photons generated in the central region of a detector, the difference between the full Geant4 simulation and the model can be parameterised using only the distance between the scintillation and the bright-spot, d_c, and the offset angle between the bright spot and the normal to the PD surface, θ_c. Examples of the parameterised corrections are shown in Figure 6.7 for the SBND-like (left) and DUNE-like (right) detector geometries. In both cases, the maximum offset angle as viewed from the bright-spot is defined by the size of the detector geometry. The error bars represent the standard deviation of the distribution within each angular bin. This can be affected significantly by increases in the detector size, as can be seen comparing the DUNE-like and SBND-like cases. This is a result of the large θ_c angular bins describing larger regions of the detector volume where the variations due to reflections from the detector walls can be significant. However, the scintillation points with the highest variations, at high d_c or θ_c, account for a relatively small fraction of the total light observed. The variation could in principle be accounted for by increasing the number of angular bins used, or by using bins with different sizes.

To calculate the PD-position corrective factors the mean of the $N_{Geant4}/N_{\Omega}/cos(\theta_c)$ distributions within each angular bin is employed. Instead of using a fit to the corrective factors, a linear interpolation in d_c is used to find the exact correction for the prediction from the geometric model.

Correcting for scintillation position: border effects

In addition to the corrective factors accounting for the relative position of the PDs, further corrections are required to account for the position of the scintillation light emission inside of the detector. Light emitted closer to the detector borders will be significantly affected by their proximity: the wavelength-shifted photons can be reflected off the walls, while the VUV photons can be absorbed before they reach the WLS-coated cathode plane. These effects are again accounted for using parameterised corrective factors. Similar to the PD-position based corrections, these corrections depend on d_c and θ_c. An additional parameter is the position of the scintillation emission relative to the borders of the detector volume. Therefore, similar to the direct light model border corrections described in Section 6.1.1, sets of corrective factors are created at different distances, d_T, from the centre of the detector, as illustrated in Figure 6.3. Then, during simulation, linear interpolation is used in both d_c and d_T for the required angular bin in θ_c to calculate the exact corrective factor required.

Examples of sets of border effect corrections for the SBND-like (left) and DUNE-like (right) detector geometries are shown in Figure 6.8 for two cylinders defined by different values of d_T. As before, the corrections are taken as the ratio between the

amount of light seen in full simulation in Geant4 compared with the prediction from the geometric model. The required corrective factors become significantly larger as d_T increases and the scintillation is emitted closer to the edges and corners of the detector volume. Additionally, the angular dependence becomes more significant and larger offset angles of the PDs, as viewed from the bright-spot, become geometrically possible. The border effects are much more significant for the light reflected by wavelength-shifting foils compared with the direct light and larger corrective factors are therefore required.

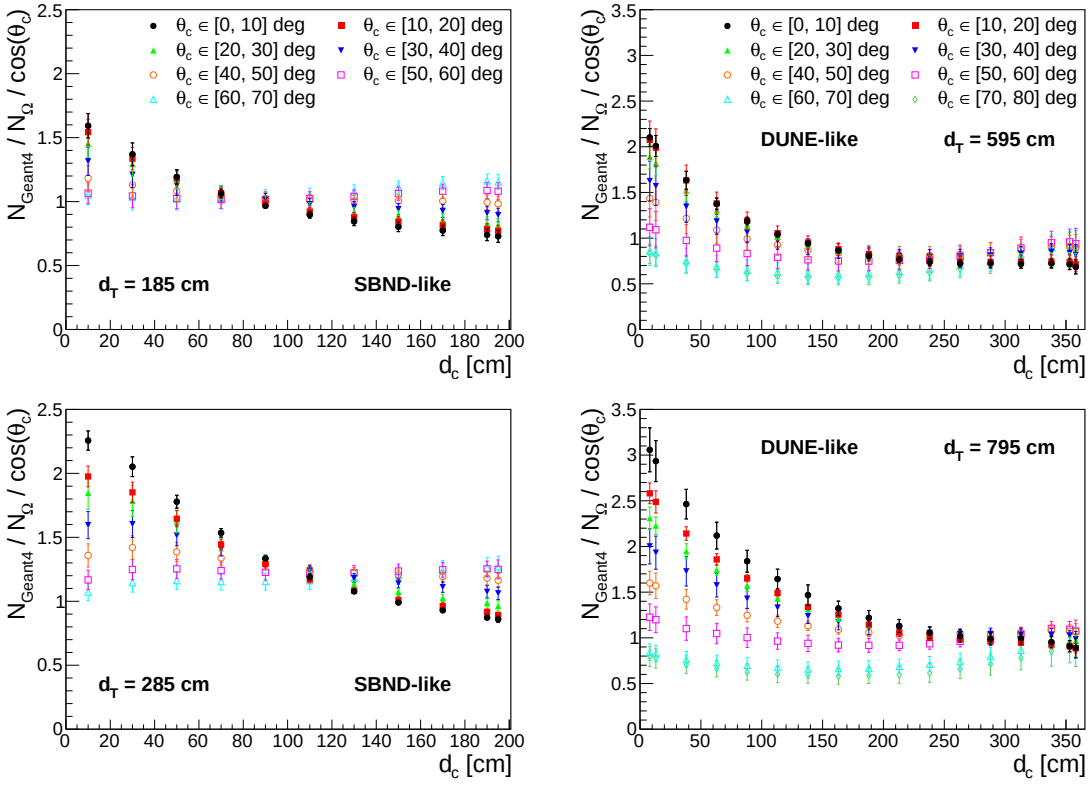

Figure 6.8: Examples of the border effect corrections required for the reflected light model in two different regions of the SBND-like (left) and DUNE-like (right) detector geometries.

Bringing the above effects together, the number of incident reflected light photons on each PD can be expressed as,

$$N_{\gamma,reflected} = N_{\Omega,reflected} \times A(d_c, \theta_c, d_T)/cos(\theta_c), \tag{6.6}$$

where $N_{\Omega,reflected}$ is the geometric prediction given by Equation 6.5 and $A(d_c, \theta_c, d_T)$ is the parameterised corrective factor accounting for PD position and border effects. This corrective factor depends only on the distance between the emission point and the bright-spot, the angle between the bright-spot and the PD, and the distance of the emission point from the centre of the detector. Here, using the notation from

Equation 3.7 for the reflected light, $N_{\gamma,reflected} = \Delta E \times S_\gamma \times Q_{abs} \times Q_r \times P(d_c, \theta_c) \times T(d_c, \theta_c)$.

6.2 Predicting the photon arrival-time distributions

The model developed in the previous section allows the prediction of the number of photons arriving at each PD, but does not address the distribution of their arrival times. As described in Sections 3.4.1 and 3.4.2, the timing of the scintillation light is dominated by the double-exponential distribution caused by the de-excitation of the two argon dimer states. However, propagation effects can impact the time distribution actually registered by the PDs. Using the notation from Equation 3.8, the model described in this section is designed to estimate the photon transport time $t_t(d, \theta)$. As with the previous section, two different detector geometries are considered: one SBND-like and one DUNE-like, each described in Section 5.4.2.

6.2.1 Direct light

The earliest arrival time of a photon on a particular PD can be predicted geometrically using the minimum distance that a photon must travel and the velocity of VUV light in LAr, shown in Figure 3.9. A geometric calculation can provide the arrival time of the fastest possible photon, but does not account for other transport effects. A typical distribution of photon arrival times due to only transport effects can be seen in Figure 6.9. The distribution shows a prompt component followed by a long diffuse tail resulting from the many different paths the photons can take as a result of Rayleigh scattering.

Empirically it is found that for essentially all combinations of emission point and PD position the distributions are of a similar nature and can be approximated by a combination of *Landau* and *Exponential* functions:

$$t_t(x) = \underbrace{N_1 \frac{1}{\xi} \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} e^{\lambda s + s \log s}\, ds}_{Landau} \; + \; \underbrace{N_2\, e^{\kappa\, x}}_{Exponential} \;, \tag{6.7}$$

where $\lambda = \frac{x-\mu}{\xi}$, with μ and ξ commonly referred as the landau *most probable value* and *width* parameters respectively, κ is the slope of the exponential and N_1 and N_2 are normalisation constants. The resulting five parameters of the *Landau + Exponential* composite that describe a given time distribution are monotonic functions of the distance between the emission point and PD, provided the incident angle is

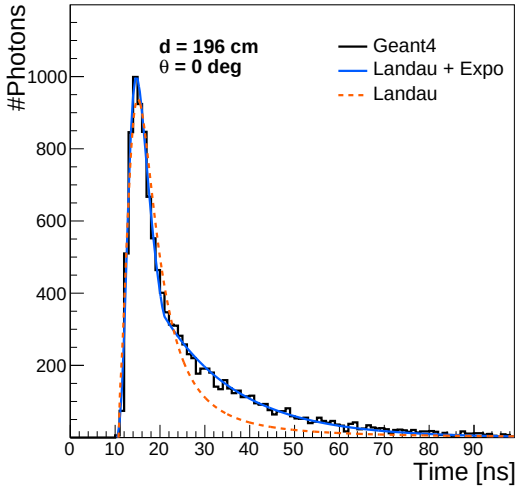

Figure 6.9: Example of the distribution of direct light photon arrival times for scintillation occurring on-axis to a PD due to only transport effects, together with the predictions from the transport time models.

accounted for. Two angular bins are used[2]: on-axis with $\theta \in [0°, 45°]$, and off-axis with $\theta \in [45°, 90°]$. Figure 6.10 shows the behaviour of the model parameters for the two angle ranges in the two geometry cases, SBND-like (blue points) and DUNE-like (black points). The spread of the parameter values depends on the detector size: in a larger detector the signals are more scattered. For simplicity and because VUV photons are predominantly absorbed by all detector materials, border effects are neglected in the model.

At larger distances the long diffuse tail of the arrival time distributions tends to disappear and the shape can be described using only a Landau distribution. A quantitative comparison of the accuracy of the two approaches as a function of the distance using the relative difference of the χ^2 of both models is shown in Figure 6.11. In both of the geometry cases a similar result is found: the Landau + Exponential model describes the shape of the signals more accurately, but at larger distances the two models perform similarly. The distance at which the two models become equivalent depends very slightly on the detector size, but for both geometries under study has a value around $d = 400\,\text{cm}$. At longer distances the simpler Landau model can be used successfully instead of the Landau + Exponential one.

During simulation the probability distribution function (PDF) of the VUV photon arrival times for each PD is constructed and then sampled using the parameters of the Landau + Exponential model. For computational reasons, a cut-off is applied at the 99^{th}-quantile when sampling the PDF of the transport time signal.

[2]To obtain greater accuracy more bins could be used, at the cost of increased computational time.

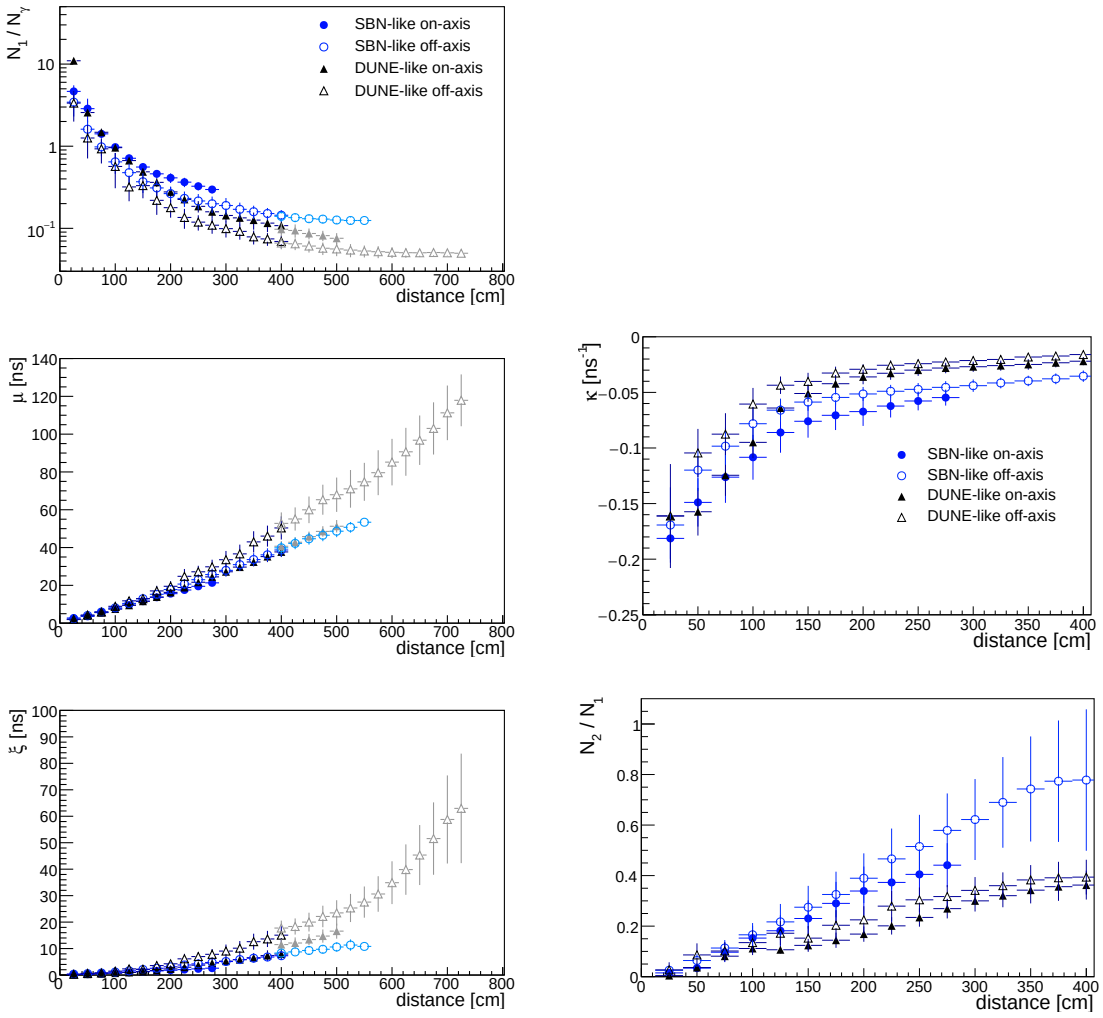

Figure 6.10: Behaviour of the Landau (left) and Exponential (right) parameters of the direct light transport time model as a function of distance between the energy deposit and PD for the DUNE-like (black points) and SBND-like (blue points) geometries. The lighter grey and blue points denote the switch to using a simple Landau instead of the Landau + Exponential (for distances larger than 400 cm).

6.2.2 Reflected light

The transport time of photons arriving at the PDs in a LArTPC detector that has wavelength-shifting reflector foils on the cathode can be modelled using a similar approach. First, a geometric prediction of the transport time of earliest arriving photons is calculated. The fastest photons are most likely to travel along the path that minimises the distance travelled at VUV wavelength, where the group velocity is slower. At visible wavelengths the photons propagate significantly faster due to the lower refractive index, as shown in Figure 3.9. Figure 6.6 shows a diagram illustrating the most likely fastest path. The emitted VUV photons travel along the shortest path from the scintillation point to the cathode. They are then are wavelength-shifted and re-emitted around the bright spot and take the shortest

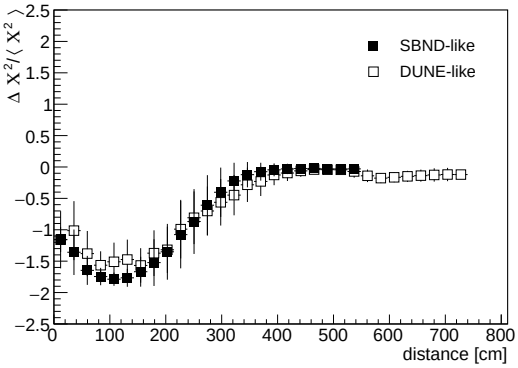

Figure 6.11: Comparison of the relative difference of the χ^2 for the two direct light transport time models: "Landau + Exponential" vs "Landau".

path to the PD. This simple model is able to predict the arrival time of the earliest photons.

The subsequent photons can be reflected from different regions of the wavelength-shifting foils and take very different paths to arrive at the PD. This results in a significantly broader distribution of their arrival times. The model describing the visible photon arrival times at the PDs is constructed in three steps. First, the direct light Landau+Exponential model, described in Section 6.2.1, is used to estimate the arrival time distribution of the VUV photons at the bright-spot on the cathode. Then the time needed for a visible photon to propagate between the bright-spot and the PD in a straight line is added. Finally, a parameterised smearing is applied to the result to account for the multitude of longer paths that can be taken. Empirically it was found that the following smearing function effectively approximates the distribution,

$$t_s = t + (t - t_f)[\exp(-\tau \ln(x)) - 1], \tag{6.8}$$

where t_s is the resulting smeared arrival time, t is the un-smeared arrival time, t_f is the fastest possible arrival time calculated geometrically, τ is a smearing factor and x is a uniformly distributed random number between 0.5 and 1. This function keeps the earliest arrival times unchanged, but increasingly smears the photons arriving later. Additionally, a maximum time cut-off is applied to avoid a non-physical long tail from the exponential distribution.

The smearing factor, τ, and the cut-off time, t_{max}, are parameterised in terms of the distance between the scintillation and the bright-spot, d_c, and the offset angle, θ_c, between the bright-spot and the PD, as shown in Figure 6.6. The cut-off time is calculated as the time needed for 99.5% of Geant4 simulated photons to arrive. The τ parameter is determined by minimising the difference between the smeared arrival time distribution calculated by the model and the distribution generated

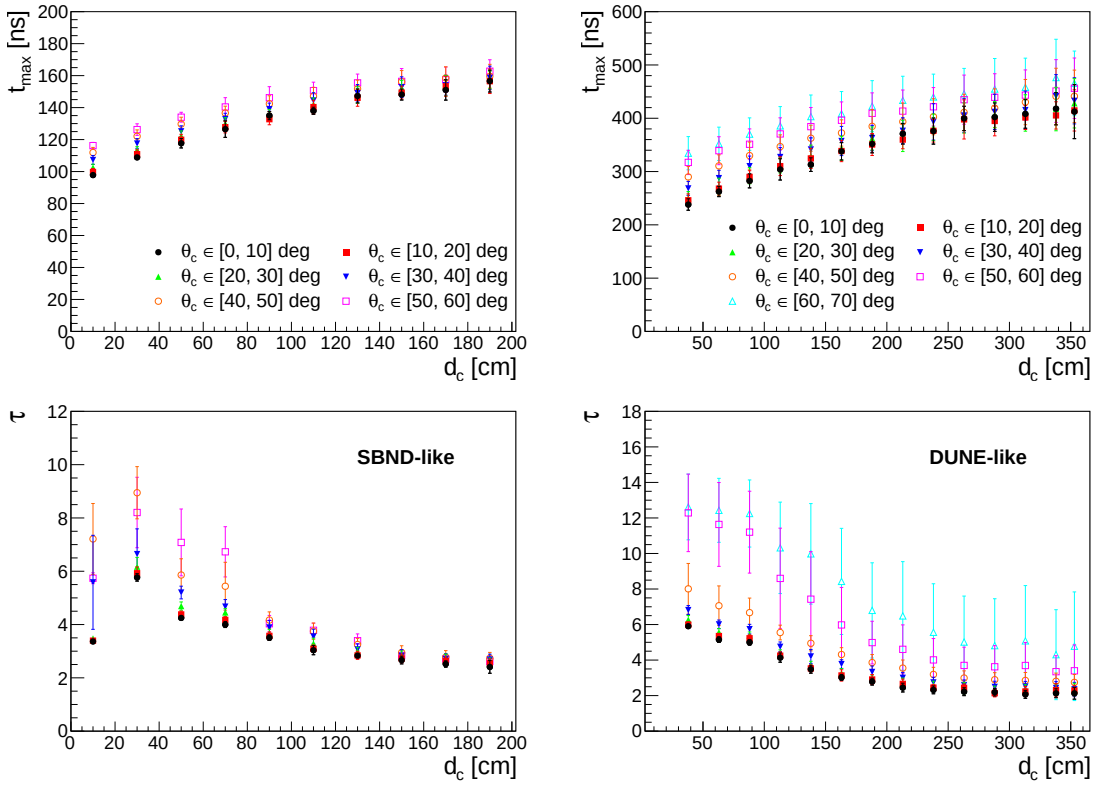

Figure 6.12: Reflected light transport time model cut-off time (top) and smearing parameter (bottom) in the central region of the SBND-like (left) and DUNE-like detector (right) geometries.

using Geant4. Unlike with the direct light transport time model, it is important to account for border effects such as reflections off the detector walls since they are highly reflective for visible wavelength photons. An approach similar to the one in Section 6.1.2 is used, creating sets of smearing parameters at different distances from the centre of the detector, d_T. These sets can then be used to calculate the smearing parameters for any location in the detector using interpolation. An example set of the parameterised cut-off times and τ smearing factors is shown for the SBND-like (left) and DUNE-like (right) geometries in Figure 6.12. It is found that the cut-off times become larger with the size of the detector. This can intuitively be explained by the longer distances the photons need to travel before reaching the PDs, including many paths where they are reflected off the detector walls. The angular dependence of the cut-off time is relatively small, with a significant overlap between bins. The τ parameter is more dependent on the angle. This effect again grows with detector size and is much more prominent for the DUNE-like case.

6.3 Training and validation samples

To train and validate the scintillation light transport models a Geant4 [269] based simulation embedded in the LArSoft software framework [264] is used. Geant4 is capable of simulating liquid argon scintillation light emission, propagation and the impact of detector boundaries. It is used as the baseline truth information for the model training. In Geant4 the optical properties of the active medium, along with all of the surrounding materials with which the photons can interact, are configurable. The optical properties of the liquid argon that are implemented are summarised in Table 6.1. Additionally, to model the impact of a typical level of contaminants present in LArTPC detectors, an absorption length of $\lambda_{abs} = 20\,\text{m}$ is applied corresponding to approximately 3 ppm of nitrogen equivalent [208].

Parameter	Type	Value
emission wavelength	spectrum	$\langle 128\,\text{nm} \rangle$ [192]
fast component decay time	number	6 ns [194]
slow component decay time	number	1300 ns [194]
refractive index	spectrum	$\langle 1.32 \rangle$ [206]
Rayleigh scattering length	spectrum	$\langle 100\,\text{cm} \rangle$ [206]

Table 6.1: Summary of the liquid argon properties used in the simulation.

Energy depositions are simulated using Geant4 at points within the active volume that cover the full phase space of distances and angles between the scintillation, the PDs and the reflector foils. To account for the impact of the borders of the detector the volume is divided into concentric cylinders at different radial distances, d_T, as illustrated in Figure 6.3. Then at each d_T, energy depositions are simulated at evenly spaced positions in the drift direction, X, covering the full drift length. In the SBND-like geometry, the simulated energy depositions are spaced in approximately 20 cm steps in the drift direction and 50 cm steps in d_T. In the DUNE-like geometry, they are spaced in approximately 25 cm steps in the drift direction and 100 cm steps in d_T. The step sizes in d_T are driven by the distances between the PDs, while those in the drift direction have been roughly chosen to cover the entire detector volumes without requiring an excessive total number of points. For each X and d_T pair five different energy deposition points are simulated. The first is chosen to be directly in front of a PD and the remaining four are placed at increasing offsets of approximately 5 cm in d_T. This allows coverage of the phase space of photons incident on the PDs from smaller angles, which otherwise would be limited due to the typically large spacing to the next nearest PD. It total following this method approximately 2500 and 3500 energy deposition points are simulated in the SBND-like and DUNE-like geometries, respectively. The models could be trained for any other detector geometry by defining simulation points in an analogous manner.

Energy depositions generating 25×10^6 photons are simulated in the SBND-like geometry and 100×10^6 photons in the DUNE-like geometry. The simulation takes approximately between 90 s and 170 s per 1×10^6 photons generated, depending on the position in the detector. The larger number of photons in the DUNE-like case was chosen to ensure the number of photons incident on the PDs is sufficient despite the larger propagation distances possible in this geometry.

Validation samples are also generated in an analogous manner to compare the predictions of the trained models against full simulation of the scintillation light in Geant4. In these samples each of the simulated points are shifted by several centimetres in a random direction to verify the performance of the models in the interpolated regions.

6.4 Validation and performance

In this section, the performance of the trained models are assessed using the validation samples described in Section 6.3. The performance is also compared with that of optical look-up libraries and the model is applied to an example realistic event.

6.4.1 Number of detected photons

Direct light

The resolution obtained with the direct light semi-analytical model as a function of d_T is shown in Figure 6.13 (left) for the SBND-like and DUNE-like geometries. An unbiased estimation of the number of VUV photons arriving to the PDs is obtained in both geometries for all values of d_T. Additionally, the global resolution is found to be better than 10%, independent of d_T. Figure 6.13 (right) shows the performance as a function of the distance between the scintillation emission and the PD. The resolution worsens slightly with distance, ranging from 5 and 15% between the closest and farthest PDs. In each case, the performance is worst at distances significantly larger than the maximum drift distance (grey line) beyond which all PDs are off-axis. These PDs, however, are a minor contribution to the overall light signal of an interacting particle and do not significantly affect the overall resolution.

Reflected light

The resolution obtained with the reflected light semi-analytical model as a function of d_T is shown in Figure 6.14 for the SBND-like and DUNE-like geometries. The model performs well throughout the entire detector volume in both cases. It has a resolution better than 10% in the SBND-like geometry and better than 15% in

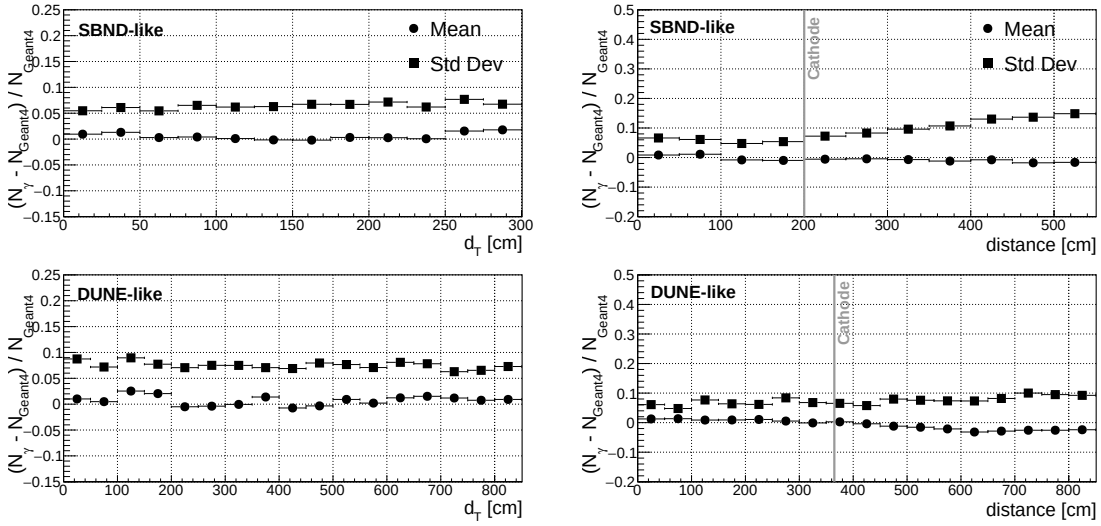

Figure 6.13: Resolution of the direct light semi-analytical model as a function of the distance from the PD-plane centre, d_T, (left) and as a function of the distance between the scintillation emission and the PD (right) for the SBND-like (top) and DUNE-like (bottom) geometries. The position of the cathode is illustrated for both geometries by the grey lines.

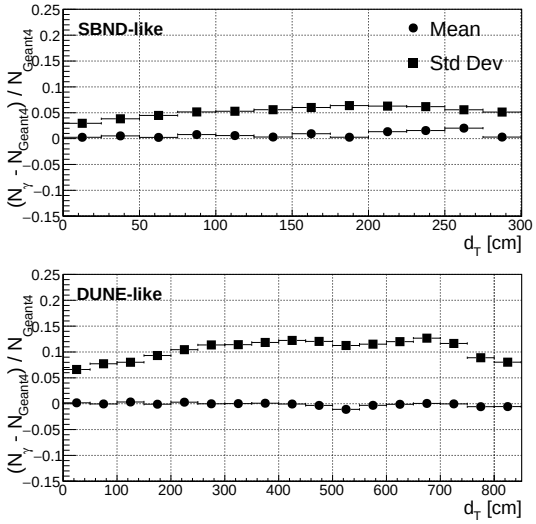

Figure 6.14: Resolution of the reflected light semi-analytical model as a function of the distance to the PD-plane centre, d_T.

the DUNE-like geometry, with minimal bias in each case. For both geometries, the resolution is best in the central region of the detector, at small d_T, where the effects of the borders are smallest. It then degrades slightly at larger d_T as the border effects become more substantial and complex. The performance of the model in the DUNE-like case is poorer than for the SBND-like case due to the larger number of possible positions within the detector and larger number of different PDs for each d_T and θ_c bin, especially at larger angles. This results in greater uncertainty and

spread in the corrective factors required, as seen in Figures 6.7 and 6.8.

6.4.2 Photon arrival-time distributions

Direct light

The performance of the direct light model at predicting the time of the earliest arriving photon, t_0, is shown in Table 6.2 for the SBND-like and DUNE-like geometries. In both cases t_0 is predicted with a resolution better than 0.5 ns and with minimal bias. This resolution is smaller than the sampling of the PD electronics in current and upcoming LArTPC detectors, as described in Section 3.4.3.

	SBND-like		DUNE-like	
model	mean	std dev	mean	std dev
Direct: Δt_0 [ns]	-0.2	0.2	0.0	0.3
Reflected: Δt_0 [ns]	0.0	0.3	0.3	0.9

Table 6.2: Resolution of the photon transport time model prediction of the earliest arriving photon time for the direct and reflected light in the SBND-like and DUNE-like geometries. In each case, $\Delta t_0 = t_{0,Geant4} - t_{0,model}$. The uncertainties on the mean and standard deviation are negligible.

An example comparison between the direct light photon transport time distribution predicted by the model and simulation in Geant4 is shown in Figure 6.15 for the SBND-like detector geometry at two difference distances and offset angles. In each case, the same number of photons is used. The distribution of the photon arrival times is accurately predicted at both distances. At the shorter distance (left), there

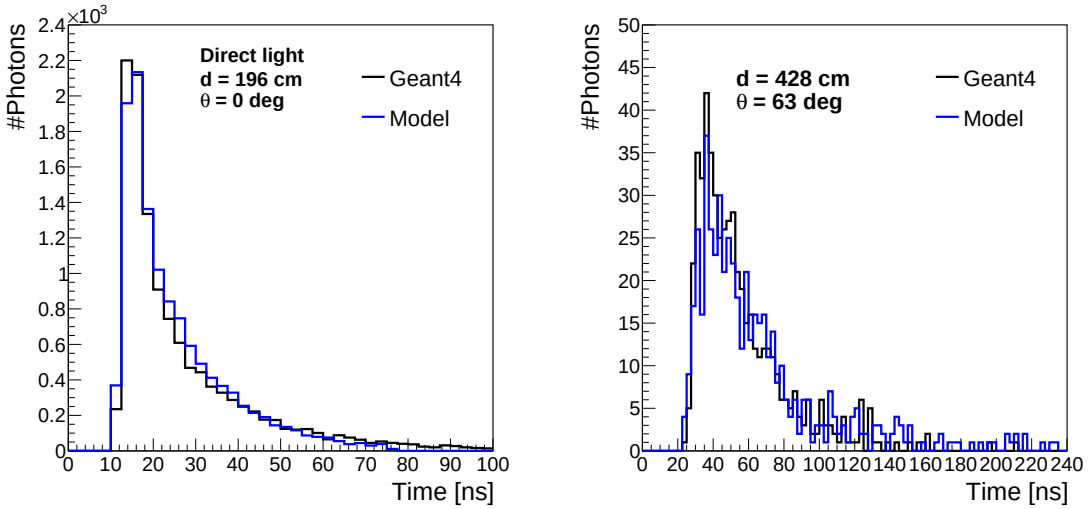

Figure 6.15: Examples of the performance of the direct light transport time model compared with simulation in Geant4 in the SBND-like detector geometry for two different distances and angles.

is a slight tail-offset between the distribution from Geant4 and the model. This is due to the example θ lying at the extreme edge of the parameterised $\theta \in [0^\circ, 45^\circ]$ angular bin. This offset could be reduced by increasing the number of angular bins used in the parameterisation or by using variable bin sizes with higher density in less linear regions. At the longer distance (right), the tail of the distribution is possibly slightly over-estimated. However, at this distance and angle very few photons reach the PDs and hence the predicted distributions are statistically limited.

Reflected light

The performance of the reflected light model at predicting time of the earliest arriving reflected photon, t_0, is shown in Table 6.2 for the SBND-like and DUNE-like geometries. In the SBND-like geometry, t_0 is predicted with a resolution better than 0.5 ns and without bias. In the DUNE-like geometry the performance is slightly worse, however the model still predicts t_0 with a resolution better than 1 ns and minimal bias. As with the direct light model, these numbers are well within the timing resolution of the PD electronics in typical LArTPC detectors.

An example comparison between the reflected light photon transport time distributions predicted by the model and simulation in Geant4 is shown in Figure 6.16 for the SBND-like detector geometry at two different distances and offset angles. The model accurately predicts the arrival time of the earliest photons and provides a reasonable approximation of their overall distribution. The model slightly underestimates in the first part of the tail of the distribution and overestimates towards the end of it. This behaviour most prominently affects off-axis PDs (as seen in the

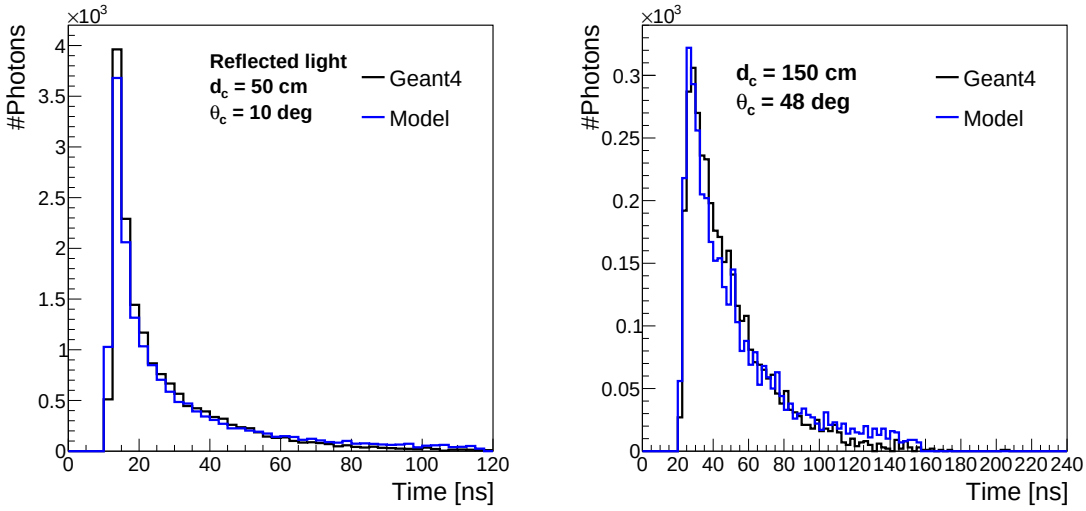

Figure 6.16: Examples of the performance of the reflected light transport time model compared with simulation in Geant4 in the SBND-like geometry for two different distances and angles.

right plot), which see substantially less light than those closer to the energy deposit, resulting in a relatively small overall impact.

6.4.3 Comparison with optical look-up libraries

An important consideration is how the performance of the semi-analytical model compares to that of the optical look-up library method commonly used in neutrino LArTPCs, as discussed in Section 5.2.3. This test is performed for both the SBND-like and DUNE-like detector geometries. To directly compare the performance, dedicated look-up libraries were generated with the same total number of photons as used to train the semi-analytical model. The configurations are summarised in Table 6.3. A uniform voxel size was used throughout the detectors and a uniform distribution of photons per voxel[3]. For completeness a "Hi-Res" version of the SBND-like look-up library was also generated to compare performance with a larger number of photons/voxel. In the "Hi-Res" case 500×10^3 photons were generated per voxel, a typical value used in recent optical look-up libraries produced for the SBND and DUNE detectors. The library generation takes approximately between 90 s and 170 s per 1×10^6 photons simulated, depending on the position in the detector.

Library	Total Photons	Photons per Voxel	Voxel Size [cm^3]	File Size[4] [MB]
SBND-like	61.4×10^9	192×10^3	5×5×5	390
DUNE-like	353.5×10^9	158×10^3	5×5×11	826
SBND-like Hi-Res	159.9×10^9	500×10^3	5×5×5	499

Table 6.3: Parameters of the look-up libraries generated to compare with the semi-analytical model. Except for the "Hi-Res" case, the total number of photons corresponds to the number of photons used to train the model.

The results of the comparison for the direct VUV light model can be seen in Figure 6.17 (left) for the two geometries under study. These plots can be compared with Figure 6.13, where the performance for the semi-analytical model is shown (note the different axes). The semi-analytical model performs significantly better than the look-up libraries in terms of both bias and standard deviation, especially at larger distances. This is at least partially a result of under-sampling of the look-up libraries, as shown by the improved performance of the "Hi-Res" library in the SBND-like case. In the DUNE-like case the fluctuations are exacerbated by the fact

[3]This is common practice in generating optical look-up libraries. However, it should be noted that varying the voxel size or the number of photons/voxel could improve the performance compared to the results shown here. This would likely require a separate optimisation process.

[4]The size of optical look-up libraries is proportional to the number of PDs. Note that the SBND-like and DUNE-like geometries used in these studies have a factor of 4 fewer PDs than the real SBND and DUNE detectors.

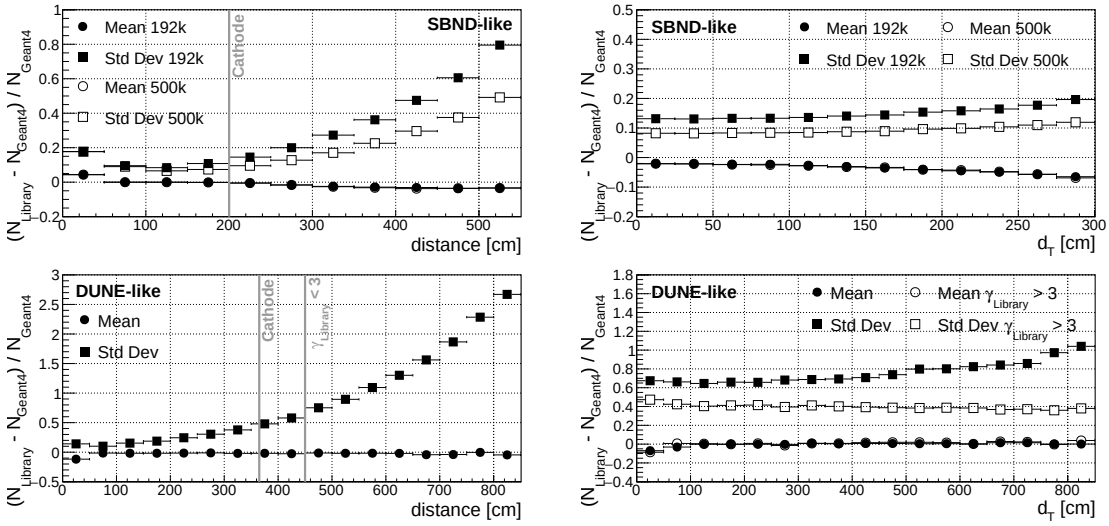

Figure 6.17: Performance of the look-up library method for the SBND-like (top) and the DUNE-like (bottom) geometries in the estimation of the number of direct light photons (left), as a function of the distance between the scintillation emission and the PD, and of the number of the reflected light photons (right), as a function of the distance from the centre of the PD-plane, d_T. The black points in each case were obtained using the same total number of photons to train the semi-analytic model. These can be compared with Figure 6.13 (direct light) and Figure 6.14 (reflected light), noting the different axes. In the top plots, the white points represent the look-up library generated with an increased total number of photons. In the bottom left plot, a vertical line indicates the distance beyond which the majority of the look-up library predictions are based on samples of less than 3 photons, which results in large fluctuations in the predictions of the library. In the bottom right plot, this effect is instead illustrated by the white points which show only voxel-PD pairs where the number of photons was larger than 3. Additionally, in both of the left plots the position of the cathode is illustrated by vertical lines.

that for distances larger than 450 cm the severe under-sampling in photons/voxel at the library generation stage causes the majority of predictions to be based on samples of less than 3 photons per voxel-PD pair. Additionally, at very short distances the look-up libraries suffer from a higher uncertainty due to the voxel size introducing discrete jumps in the predictions very close to the PDs. This second problem cannot be resolved by increasing the number of photons used per voxel, instead it requires reducing the size of the voxels or using a different approach altogether in this region.

Figure 6.17 (right) shows the performance of the generated look-up libraries in predicting the number of reflected light photons. Due to the nature of modelling the reflected light d_T is used as the variable instead of distance from the PD. These plots can be compared with Figure 6.14, where the performance for the semi-analytical model is shown. In the SBND-like geometry the look-up library method has comparable performance to the semi-analytic model, especially if the "Hi-Res" version is

used, although a small under-prediction is observed in the regions of high d_T. In the DUNE-like case the under-sampling effects are so severe that the standard deviation of the look-up library prediction is much larger than for the semi-analytic method. The effect is again caused by many voxel-PD pairs where the prediction is made based on samples of a few photons. This could be mitigated by using a significantly higher number of photons/voxel to generate the look-up library.

Overall the semi-analytical model is found to perform significantly better than look-up libraries trained using the same number of photons.

6.4.4 Simulating realistic events

In the previous sections, it has been demonstrated that the semi-analytical model effectively predicts the number of photons and their arrival times from point-like energy depositions. In simulations of particle detectors extended objects such as tracks or showers are more commonly dealt with. The model can easily simulate these kinds of events using the paradigm employed in Geant4, where particle trajectories are composed of discrete energy depositions called steps. To simulate realistic particle events the semi-analytical model can be applied to each of these steps and the results combined to obtain the simulation of the full particle trajectory. An

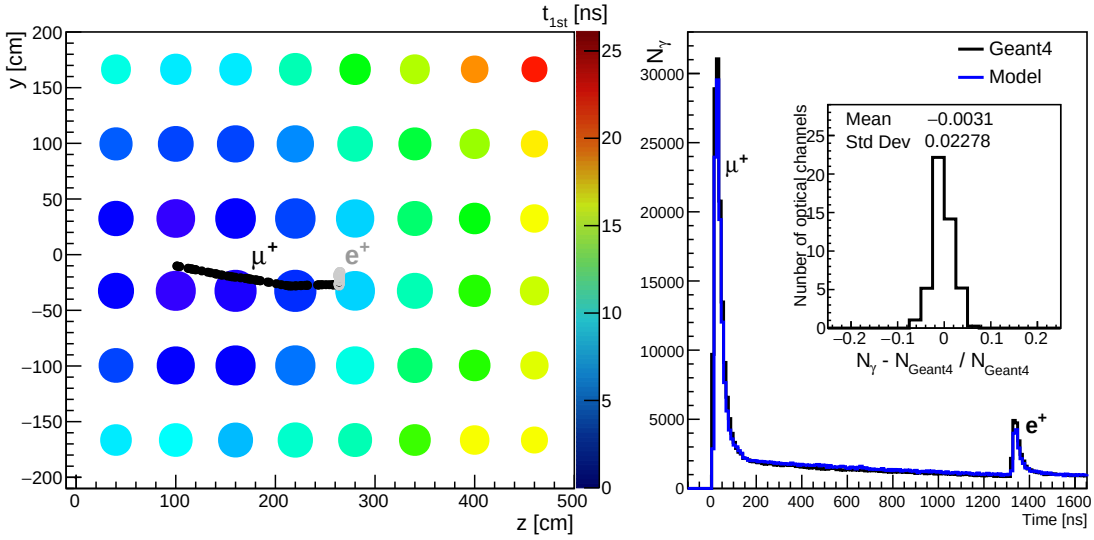

Figure 6.18: Event display of a stopping anti-muon $\mu^+ \rightarrow e^+\nu_e\overline{\nu}_\mu$ simulation. The left figure shows the charge (Geant4) and light (semi-analytic model) footprint projected on the PD-plane. Each circle represents a PD, where different colours indicate the starting time, t_{1st}, of the signals (the anti-muon is entering from the left), and the size is proportional to the number of detected photons ($\propto log_{10}(N_\gamma)$). The right figure shows the summed Geant4 signal of all of the PDs overlapped with the model prediction for comparison. Excellent agreement is seen between them, quantified by the resolution histogram of the light-model for this particular event.

example of this approach is shown in Figure 6.18. This presents the results of simulating the scintillation light originating from an anti-muon track decaying into a Michel positron $\mu^+ \rightarrow e^+\nu_e\overline{\nu}_\mu$ [310] inside the SBND-like geometry. The left figure shows the charge (Geant4) and light (semi-analytic model) footprint projected on the PD-plane. The right figure shows the summed Geant4 signal of all of the PDs overlapped with the model prediction. Excellent agreement is seen between them, quantified by the resolution histogram of the light-model for this particular event. The model prediction of the waveform observed by the PDs is also compared to that of the full Geant4 simulation. Very good agreement is seen for both the primary anti-muon scintillation peak and the secondary scintillation peak caused by the positron.

6.5 Modelling xenon-doped liquid argon

The semi-analytical model can be extended to model xenon-doped liquid argon. Doping of liquid argon with xenon is a proposed method of enhancing the scintillation light signal in large scale LArTPCs such as DUNE, as discussed in Section 3.4.4. Unlike in the previous sections, the examples for xenon-doped liquid argon will be shown using the realistic DUNE horizontal and vertical drift detector geometries, described in Section 5.4.1, for which these models were developed.

6.5.1 Predicting the number of detected photons

The xenon-doped argon can be modelled using the same approach described in Section 6.1. As described in Section 3.4.4, a subset of the photons emitted in xenon-doped argon are converted from 128 nm to 178 nm. The scintillation light signal can therefore be predicted using a combination of the argon wavelength corrections and a second set of corrections trained at xenon emission wavelength. At xenon wavelength, the Rayleigh scattering length is significantly longer: $\lambda_{RS} \sim 900$ cm, as shown in Figure 3.9. This allows the light to propagate over much larger distances before significant attenuation occurs. The corrections required to the geometric prediction of the number of photons are therefore much smaller. Figure 6.19 shows an example comparison of the direct light corrections required at argon (left) and xenon (right) wavelengths, in both cases in the DUNE vertical drift detector geometry. The dashed lines show the fitted GH curves. In the xenon case, the curves are much closer to the value of $\cos(\theta)$ indicating a smaller correction from the geometric prediction. Additionally, the dependence on distance is weaker and the curves do not sharply drop at larger distances as in the argon case, illustrating the significantly reduced attenuation of the light due to Rayleigh scattering. The corrections

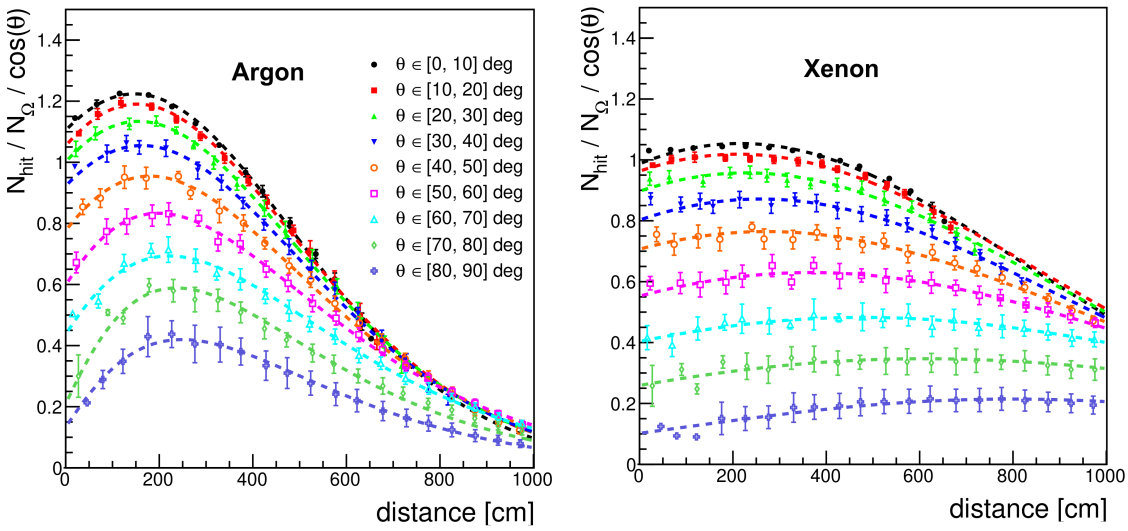

Figure 6.19: Example comparison of the direct light corrections from the geometric prediction required at argon (left) and xenon (right) wavelengths, in both cases in the DUNE vertical drift detector geometry. The dashed lines show the fitted GH curves. In the xenon case, the curves are much closer to the value of $\cos(\theta)$ indicating a smaller correction from the geometric prediction.

for detector size and border effects can then be applied using the same approach as those for argon wavelength, described in Section 6.1.1. Due to the longer Rayleigh scattering length, the impact of the decreased volume available for scattering is reduced. However, the cryostat walls and other detector components may be more reflective at xenon wavelength than at argon wavelength. Therefore, the corrections for border effects remain important.

The visible light resulting from wavelength-shifting reflector foils can also be modelled in xenon-doped liquid argon. As with the direct light, this can be modelled using the same method described in Section 6.1.2. The primary difference in this case is the amount of direct light arriving onto the reflector foils, which can be predicted using the direct light model trained at xenon wavelength emission. Once the incident light has been converted to visible wavelengths, the propagation of the light is unchanged. However, the efficiency of the wavelength-shifting compound used may be different at xenon wavelength and must be accounted for.

The performance of the direct and reflected light semi-analytical models at xenon wavelength is similar to that at argon wavelength. Figure 6.20 (left) shows the resolution of the direct light model as a function of the distance between the scintillation emission and the PD for the DUNE horizontal drift detector geometry. An unbiased estimation of the number of photons arriving at the PDs is obtained, with a resolution ranging from approximately 10-15% depending on the distance. This is comparable with the performance at argon wavelength shown in Figure 6.13. Figure 6.20 (right) shows the resolution of the reflected light model at xenon wavelength

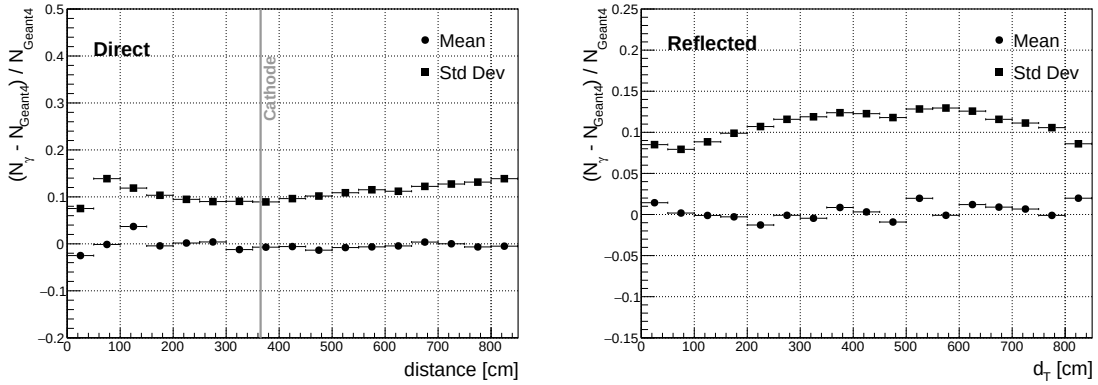

Figure 6.20: Left: resolution of the direct light model at xenon wavelength as a function of the distance between the scintillation emission and the PD. The position of the cathode is illustrated by the grey line. Right: resolution of the reflected light model at xenon wavelength as a function of the distance to the PD-plane centre, d_T. In both cases, the DUNE horizontal drift detector geometry is used.

as a function of the distance to the PD-plane centre, d_T. Similar to the direct light model, and unbiased estimation of the number of photons is obtained with a global resolution of approximately 7.5-12.5% depending on the position within the detector. Again, this is comparable with the performance at argon wavelength shown in Figure 6.14.

Different concentrations of xenon dopant can then be modelled by varying the ratio of argon wavelength and xenon wavelength light simulated. In simulation this can also be split depending on whether the argon wavelength photons would be emitted promptly from the singlet state, or slowly from the triplet state. The majority of the singlet emission would remain at argon wavelength, whereas the majority of the triplet emission would be converted to xenon wavelength as discussed in Section 3.4.4.

6.5.2 Predicting the photon arrival time distributions

The impact of transport effects on the photon arrival time distribution is also altered by xenon-doping. At xenon wavelengths, since the Rayleigh scattering length is much longer, the photons on average travel on more direct paths to the PDs than at argon wavelengths. Examples of the direct light arrival time distribution at xenon wavelengths due to transport effects are shown in Figure 6.21 (left) for three different distances. These can be compared with the argon wavelength example seen in Figure 6.9. At xenon wavelength, the resulting arrival time distributions are much narrower, with the majority of the photons from each distance arriving within a few nanoseconds. The direct light transport time can therefore be effectively approximated using a purely geometric prediction of the propagation time along the

shortest path, as described in Section 6.2.1. A small fraction of the light does, however, arrive later, especially at longer distances. To fully model this the *Landau + Exponential* parameterisation described in Section 6.2.1 could be applied, retrained at xenon wavelength. However, since the distribution is significantly sharper, an alternative approach may be required.

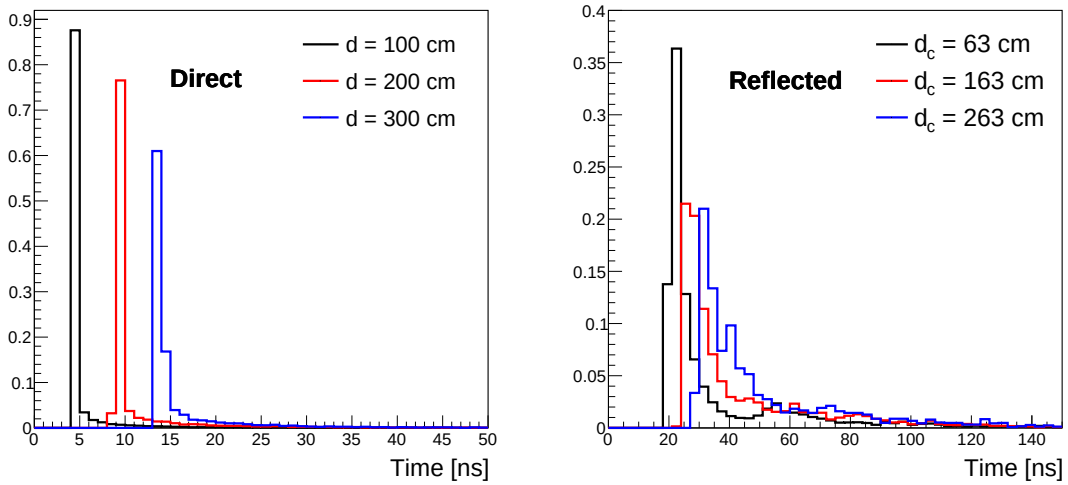

Figure 6.21: Examples of the distribution of direct (left) and reflected (right) photon arrival times due to transport effects only for xenon wavelength light. In each case the distributions are from full simulation of the light in the DUNE horizontal drift detector geometry, performed using Geant4. Note the different x-axes.

Figure 6.21 (right) shows examples of the reflected light arrival-time distribution at xenon wavelengths due to transport effects for three different distances to the cathode, d_c. The reflected light arrival time distributions are significantly broader as a result of the multitude of the paths the photons can take incident on different regions of the reflector foils. The arrival time of the earliest photons can be predicted geometrically as described in Section 6.2.2. However, this is a much poorer approximation of the arrival time distribution compared with the direct light. To accurately predict the distribution would require an approach analogous to the model described in Section 6.2.2, requiring the development of a direct light parameterisation.

At the time of writing, parameterisations for the direct and reflected light transport times at xenon wavelength have not yet been developed. Instead, the geometric only approach is in use as an approximation.

6.6 Application in real-life LArTPC detectors

The semi-analytical model is currently the standard fast optical simulation approach used in the SBND, DUNE horizontal drift and DUNE vertical drift detectors. The implementation in each of these detectors will be described in this section. The

photon transport time parameterisation is also employed in the ICARUS detector, in conjunction with an optical look-up library. Finally, the model is currently in development for the MicroBooNE detector.

To enable this, the models have been implemented into the LArSoft software suite [264] used by numerous LArTPC detector collaborations. Additionally, the specific configurations for each detector are implemented in the SBND-, DUNE- and ICARUS-specific software suites.

6.6.1 SBND

Parameter sets

The SBND photon detection system uses a combination of PMT and X-ARAPUCA PDs, as described in Section 3.6.2. In addition, wavelength-shifting reflector foils are present on the cathode of the detector. The semi-analytical model is used as the standard fast optical simulation for both the direct (VUV) and reflected (visible) light in SBND. The acceptance windows of the PDs in SBND have significantly different shapes: a dome-like surface in the case of the PMTs, and a flat rectangular surface in the case of the X-ARAPUCAs. The acceptance of both of these can be calculated geometrically using the subtended solid angle[5]. However, the corrections required for propagation and border effects are different for the PMTs and X-ARAPUCAS. This is especially significant for off-axis scintillation due to the extended three-dimensional shape of the PMTs. Therefore, separate sets of corrections are required for each PD type for both the direct and reflected light.

Unlike in the idealised SBND-like geometry, the SBND has sense-wire planes and other detector components that can shadow the PDs reducing the amount of light that reaches them. This is represented by the factor Q_{trans} in Equation 3.7. Earlier iterations of the simulation used a flat scaling factor to account for the impact of this. However, the current implementation incorporates these effects directly into the parameterisations allowing the full distance and angular dependency to be modelled. This is achieved through using a realistic geometric description of these detector components in the Geant4 simulation used to train the model.

The performance of the model at predicting the number of photons is summarised in Table 6.4 for all four scenarios, comparing against full simulation in Geant4. For the X-ARAPUCAS (flat PDs) the model performs well for both direct and reflected light, with global resolution of approximately 5-10% and minimal bias. For the PMTs (dome PDs), however, the direct light model appears to perform worse with

[5]In the PMT case, the analytical on-axis expression for the solid angle of a hemisphere is used. A parameterised correction is then applied when off-axis. This reduces the accuracy of the prediction, but significantly improves computational performance compared with the full calculation.

Scenario	Mean	Std Dev
direct, flat	±1-2%	5-10%
direct, dome	±1-3%	5-35%
reflected, flat	±1-2%	5-10%
reflected, dome	±1-2%	5-10%

Table 6.4: Summary of the performance of the semi-analytical model in SBND, compared with full simulation in Geant4. The performance varies as a function of distance and angle in each case, indicated by the ranges given.

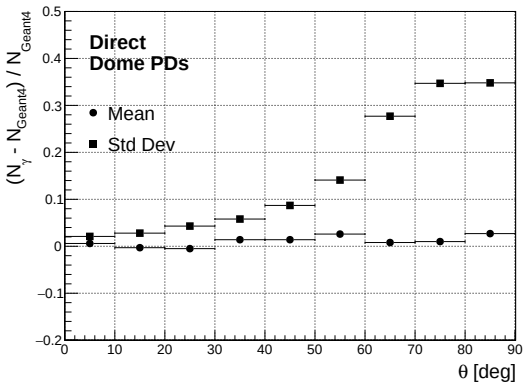

Figure 6.22: Resolution in SBND of the direct light model for PMTs (dome PDs) as a function of the offset angle between the scintillation emission and the normal to the PD surface, θ.

a resolution as poor as 35% in some regions. Figure 6.22 shows the performance of the direct light model as a function of the offset angle, θ. The model performs well, except for at large θ where the resolution becomes significantly worse. This is predominantly caused by shadowing effects where one PMT shadows adjacent PMTs that are off-axis. However, this effect only occurs in very specific regions of the detector where these large angles are possible and has a negligible effect on the overall light seen from the majority of interactions. The reflected light model has a resolution of approximately 5-10%. It is not impacted by the shadowing effects since the detected photons are always propagating from at least the drift distance away, as a result of the location of the foils, and hence the large offset angles are not possible. For both the direct and reflected light, there is minimal bias.

The photon propagation time model is also used in SBND for both the direct and reflected light. Unlike the number-of-photons prediction, separate parameters are not required for the two PD types as the impact of the different shapes was found to be negligible.

Hybrid model

The semi-analytical model is effective at predicting the amount of light incident on the PDs from within the active volume of the detector. For it to function the geometric acceptance of the PDs must provide a reasonable estimation of the number of photons incident on them, with only relatively small corrections needed to account for propagation and border effects. However, scintillation light is also produced in the non-instrumented regions of the detector external to the field-cage, such as behind the PDs. The amount of light originating from these regions is enhanced by the increased scintillation yield of the non-instrumented argon as a result of the zero (or near-zero) electric field present, see Section 3.5. In some cases a non-negligible fraction of this light can reach the PDs, often via complex paths undergoing multiple Rayleigh scatterings or reflections. The semi-analytical model is not able to effectively predict this light. In this case geometric acceptance is a poor approximation, and propagation and border effects dominate.

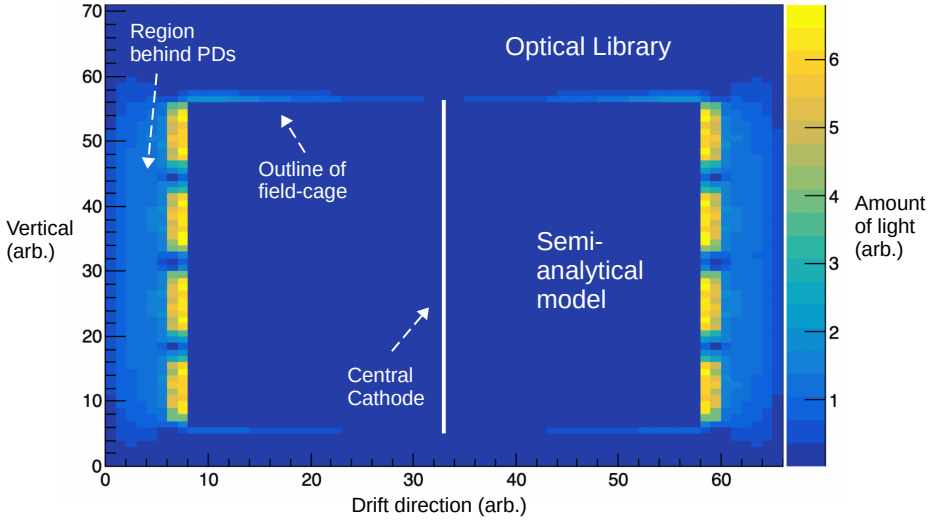

Figure 6.23: Light map of the hybrid model optical look-up library in SBND. Each bin represents a voxel in the library, and the colour scale represents the amount of light detected by the PDs from that position. The brightest regions are those immediately behind the PDs, where a non-negligible amount of light from the external argon is observed. The outline of the field-cage is also visible, indicating there is light arriving from interactions occurring just outside of it. The regions inside the field cage are modelled using the semi-analytical approach and no library is present. Note that the axes are not to scale. Figure adapted from Reference [207, 311].

To accurately model this light in SBND, a hybrid approach is employed: the semi-analytical model is used within the active volumes that make up the majority of the argon in the cryostat, and a set of slimmed-down optical look-up libraries are used for the external regions [207, 311]. This allows light from the external

regions to be predicted, while keeping the optical libraries small negating the related computational challenges of large libraries discussed in Section 5.2.3. Figure 6.23 shows a light map of the hybrid model optical libraries in SBND. Each bin represents a voxel in the library and the colour scale represents the amount of light detected by the PDs from that position. The brightest regions are those immediately behind the PDs, where a non-negligible amount of light from the external argon is observed. The outline of the field-cage is also visible, indicating there is light arriving from interactions occurring just outside of it.

Computational performance

The computational performance of the semi-analytical model in SBND has been assessed using a sample of 50 MeV electrons uniformly distributed throughout the detector volume and a sample of 5 GeV muons passing through the detector in the beam direction. The performance of the fast optical simulation using the hybrid approach is compared with the full optical simulation using Geant4, in both cases run using the standard implementations of the models within the LArSoft software suite. In simulation various scaling factors, such as the efficiency of the PDs, can be applied to the scintillation yield at the point of production prior to the photon propagation. This pre-scaling of the scintillation yield can substantially reduce the number of photons that need to be propagated, improving the computational performance. A realistic pre-scaling factor of 0.05 is used. The performance is summarised in Table 6.5. For both samples the fast optical simulation is much faster than the full optical simulation: approximately 60 times faster for the 50 MeV electrons and 70 times faster for the 5 GeV muons. The memory usage of the semi-analytical model is higher in both cases, predominantly as a result of the hybrid mode optical library and from loading additional LArSoft libraries. However, it remains relatively low.

Mode	Sample	CPU time [sec/event]	Memory [MB]
Fast optical simulation	50 MeV electrons	0.12	1800
	5 GeV muons	3.43	2060
Full optical simulation	50 MeV electrons	7.44	650
	5 GeV muons	249.70	1380

Table 6.5: Comparison between the computational performance of the fast optical simulation using the semi-analytical model and full optical simulation using Geant4 in SBND.

Light yield

Figure 6.24 shows the light yield, the number of photons detected per MeV of energy deposited, in SBND as a function of the distance to the PD-plane. The yield is sub-

divided into the contributions from the direct (VUV) and reflected (visible) light. The direct light dominates close to the anode of the detector, where the PDs are located, and the visible light dominates close to the cathode. The combination results in a high light yield throughout the detector with a minimum of approximately 70 PE/MeV.

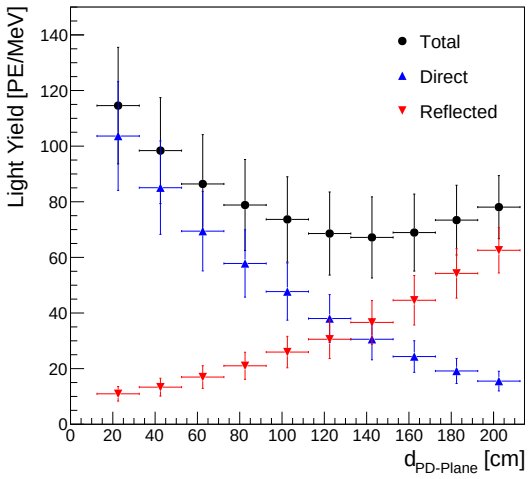

Figure 6.24: Light yield in SBND as a function of the distance to the PD-plane, $d_{PD-Plane}$, subdivided into the contributions from the direct and reflected light. The uncertainties represent the variation in the light yield in the y-z plane due to the locations of the PDs.

6.6.2 DUNE horizontal drift

Parameter sets

The DUNE horizontal drift (DUNE-HD) photon detection system consists of X-ARAPUCA PDs, as described in Section 3.6.3. Additionally, both wavelength-shifting reflector foils and xenon-doping are considered as potential extensions to the baseline design to improve the light yield. The semi-analytical model is used as the standard fast optical simulation in DUNE-HD for each of these scenarios. As described in Section 6.5, the xenon-doped liquid argon is modelled by using a combination of parameters trained at argon and and xenon wavelengths. In each case, the parameters have been trained for both the direct (VUV) and the reflected (visible) light. As with the SBND model, the impacts of shadowing from the sense wires and other detector components is incorporated directly into the parameterisations.

The performance of the model at predicting the number of photons is summarised in Table 6.6 for each set of parameters, comparing against full simulation in Geant4. In all four scenarios the DUNE-HD 1×2×6 subset geometry is modelled, as described in Section 5.4.1. The performance of each model at argon and xenon wavelengths

are similar. For the direct light, the number of photons is predicted with minimal bias and a resolution of approximately 10-20% in each case. For the reflected light, a resolution of 7.5-12.5% is obtained again with minimal bias.

Scenario	Mean	Std Dev
direct, argon	±1-3%	10-20%
direct, xenon	±1-3%	10-15%
reflected, argon	±1-2%	7.5-12.5%
reflected, xenon	±1-2%	7.5-12.5%

Table 6.6: Summary of the performance of the semi-analytical model in DUNE-HD, compared with full simulation in Geant4. The performance varies as a function of distance and angle in each case, indicated by the ranges given.

The default configuration of the xenon-doped argon simulation is for 100% of the prompt component of the argon scintillation light to remain at argon wavelength, and 100% of the slow component to be converted to xenon wavelength. This is a reasonable approximation of the impact of the levels of xenon doping proposed for DUNE-HD, $\mathcal{O}(10)$ ppm [182, 221]. The fraction of the light that is converted to xenon wavelength can be altered, however, without requiring any re-training of the semi-analytical model.

The photon propagation time model is also applied in DUNE-HD, for both the direct and reflected light. At argon wavelength the standard models are used, as described in Section 6.2. At xenon wavelength the geometric only approximation is used, as described in Section 6.5. Future iterations of the model may include a full parameterisation of the xenon wavelength propagation time, particularly if xenon-doping becomes the baseline design.

The hybrid model, described in Section 6.6.1 for SBND, is also currently in development for DUNE-HD. This will allow more realistic modelling of the non-instrumented argon regions in the detector.

Computational performance

The computational performance of the semi-analytical model in the DUNE-HD detector has been assessed using a sample of 50 MeV electrons uniformly distributed throughout the detector volume and a sample of 10 GeV muons passing through the detector in the beam direction. In both cases, the simulation is performed using the baseline DUNE-HD detector configuration without foils and without xenon-doping. Analogous to the computational performance tests in SBND, a realistic scintillation yield pre-scaling factor of 0.05 is used. The performance of the fast and full optical simulation is summarised in Table 6.7. As before, for both cases the fast optical simulation is much faster than the full optical simulation: approximately 35 times

faster for the 50 MeV electrons and 60 times faster for the 10 GeV muons. The memory usage of the semi-analytical model is higher than the full optical simulation for the electrons, predominantly as a result of loading additional LArSoft libraries, but significantly lower for the muons. The memory usage of the semi-analytical model in the DUNE-HD detector is lower than in SBND, see Table 6.5, as a result of the absence of the hybrid model library.

Mode	Sample	CPU time [sec/event]	Memory [MB]
Fast optical simulation	50 MeV electrons	0.15	1080
	10 GeV muons	7.8	1170
Full optical simulation	50 MeV electrons	5.23	640
	10 GeV muons	461.7	2550

Table 6.7: Comparison between the computational performance of the fast optical simulation using the semi-analytical model and full optical simulation using Geant4 in the DUNE-HD detector.

Comparison with optical look-up library

A full scale optical look-up library was produced for the DUNE-HD 1×2×6 geometry allowing direct comparison between the semi-analytical model and the library approach in a realistic detector geometry. The library was generated with $5\times5\times11\,\text{cm}^3$ voxels and 500×10^3 photons per voxel, matching the configuration of the most recent full library generated for DUNE-HD. The liquid argon and material properties were configured to be identical to those used to train the semi-analytical model. Additionally, in both cases the impact of shadowing from the sense wires and other detector components was incorporated. The performance of the semi-analytical model and the optical library were then both compared with the high statistics full simulation in Geant4, with 100×10^6 photons generated per comparison point. Figure 6.25 shows a comparison of the performance of the two approaches for the direct argon-wavelength light as a function of the distance between the scintillation and the PD. The semi-analytical model outperforms the look-up library at both very short and large distances. At short distances, this is due to the size of the voxels introducing discrete jumps in the predictions when very close to the PDs. At larger distances, this is due to severe under-sampling in the number of photons reaching the PDs during the library generation. Despite the higher number of photons used in the generation compared with the library discussed in Section 6.4.3, at distances large than approximately 600 cm fewer than 10 incident photons are used to make the library predictions resulting in large statistical fluctuations. However, at medium distances, up to approximately the drift length of the detector (350 cm), the library slightly out-performs the semi-analytical model.

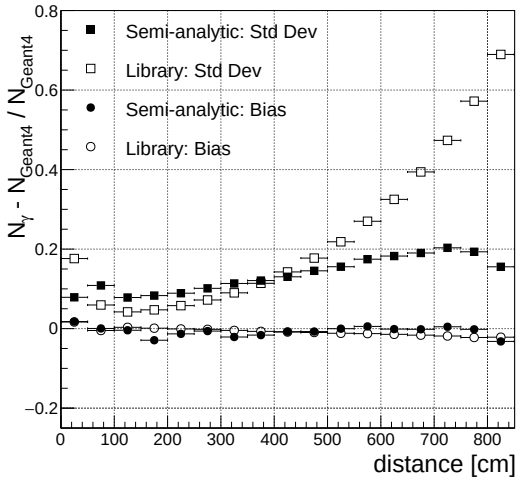

Figure 6.25: Comparison of the performance of the semi-analytical model and the optical look-up library in the DUNE-HD 1×2×6 geometry for the direct argon-wavelength light as a function of the distance between the scintillation and the PD. The white points show the results for the optical library, and the black points for the semi-analytical model.

Due to the large volume and large number of PDs in the DUNE-HD 1×2×6 geometry, the resulting optical library was very large. It required approximately 12 GB of memory to load using LArSoft, making it effectively impossible to use in grid-based production campaigns where the majority of grid nodes have only a few GBs of memory available per core. In comparison the memory footprint of the semi-analytical approach is negligible. Additionally, due to the large number of voxels required to cover the DUNE-HD 1×2×6 geometry, the Geant4 simulation required to train the optical library was extremely slow to produce – taking approximately 2-3 weeks on the available DUNE grid computing resources. This compares with a production time of 2-3 days using the same resources for the semi-analytical model. The semi-analytical model can therefore be much more flexible if multiple configurations are required.

Light yield

Figure 6.26 shows the light yield, the number of photons detected per MeV of energy deposited, in the DUNE-HD detector as a function of the distance to the PD-plane. The light yield is shown for the baseline design, without reflective foils or xenon-doping. The impact of these enhancements to the DUNE-HD photon detection system can be seen in Figure 3.10. The light yield in the DUNE-HD detector is highly non-uniform, varying between approximately 33 PE/MeV at the anode and 3 PE/MeV at the cathode.

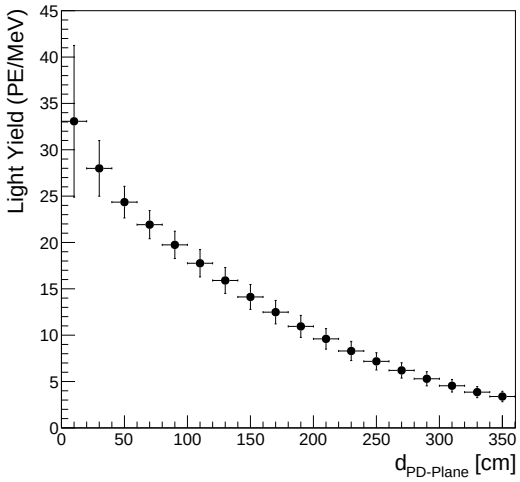

Figure 6.26: Light yield in the DUNE-HD detector as a function of the distance to the PD-plane, $d_{PD-Plane}$, for the baseline photon detection system design. The uncertainties represent the variation in the light yield in the y-z plane due to the locations of the PDs.

6.6.3 DUNE vertical drift

The DUNE vertical drift (DUNE-VD) photon detection system consists of X-ARAPUCA PDs located on both the cathode and the laterals of the detector, as described in Section 3.6.3. The baseline design also incorporates xenon-doping of the liquid argon to allow propagation of the light over larger distances.

Modelling the cathode PDs

The X-ARAPUCA PDs mounted on the cathode of the detector are modelled using the standard approaches for the argon and xenon wavelength light as described in Sections 6.1 and 6.5, respectively. At xenon wavelength, in addition to the direct light, photons may be reflected from highly reflective surfaces on the anode of the detector before arriving at the PDs as described in Section 3.6.3. This is modelled using the same approach used to predict the number of photons as a result of wavelength-shifting reflector foils, as described in Section 6.1.2. Except, instead of the photons being wavelength shifted, they are simply reflected by the anode and remain at xenon wavelength. Since the light is reflected from the anode rather than the cathode, the solid angle of the anode plane is used in making the geometric predictions and the corrections are parameterised in terms of the distance to the anode plane, d_a, instead of d_c. The average reflectivity of the anode is treated as a scaling factor in the model, analogous Q_r in Equation 6.5, allowing it to be adjusted without re-training. This will allow the impact of different materials and designs to be assessed. Since both the direct and reflected photons have the same wavelength,

they are effectively indistinguishable unlike the light from wavelength-shifting reflector foils. Therefore, the two components are summed in the simulation rather than being treated separately.

The performance of the model at predicting the number of photons arriving on the cathode PDs is summarised in Table 6.8 for both argon and xenon wavelength light, comparing against full simulation in Geant4. In each case the DUNE-VD 1×8×14 geometry is modelled, as described in Section 5.4.1. At argon wavelength, the number of photons is predicted with a resolution of approximately 10-20% and with minimal bias. At xenon wavelength, the resolution is similar at approximately 10-15%, however there are slightly larger biases in some regions – as much as 8% at large distances. This is predominantly caused by the contribution from the reflected photons, and is a result of the method used during training to distinguish between the direct and reflected components rather than a deficiency in the model. When simulating wavelength-shifting reflector foils in Geant4, the direct and reflected photons arriving at the PDs can be easily distinguished based on their wavelength allowing the two components to be precisely separated to train the models. This is no longer possible in this scenario, since the photons are not wavelength shifted. Instead, two training samples were produced: one with anode reflections and one without. The contribution from the reflections was then calculated by subtracting the photons from the simulation without for each energy-deposition-PD combination. Due to statistical fluctuations, especially at large distances, this could lead to nonphysical negative numbers of photons. These had to be removed from the training samples, potentially resulting in systematic biases in the resulting corrections. In future iterations of this model, this could be addressed by developing a technique to properly separate the two components in the training samples.

Scenario	Mean	Std Dev
argon, cathode	±1-3%	10-20%
xenon, cathode	±1-8%	10-15%

Table 6.8: Summary of the performance of the semi-analytical model in DUNE-VD for the cathode PDs, compared with full simulation in Geant4. The performance varies as a function of distance and angle in each case, indicated by the ranges given.

Modelling the lateral PDs

The lateral PDs are mounted on the sides of the detector. As a result of the asymmetry in the detector dimensions in the horizontal versus vertical directions, they see a liquid argon volume with different light propagation behaviour. The impact of propagation and border effects on the light is non-negligibly altered to that of the cathode PDs, necessitating separate sets of corrections to be produced. Fur-

thermore, the lateral PDs are located only in the half of the detector closest to the anode, resulting in an asymmetry in the impact of the borders in the drift direction. As a result of this, a strong dependency was found in the required border effect corrections with the proximity of the scintillation emission to the anode plane. To account for this, the direct component border effect corrections were altered to be parameterised in terms of the distance to the anode plane, d_a, instead of the radial distance from the centre of the detector, d_T. Figure 6.27 shows examples of the required direct light corrections for the lateral PDs at different distances from the anode plane. The required corrections vary significantly as a function of d_a, especially for scintillation emission occurring very close to the anode. Unlike the standard border effect model, described in Section 6.1.1, the variation in the GH parameters is not well modelled by a linear fit. Instead, linear interpolation in d_a is used to find the exact correction required similar to the reflected light model.

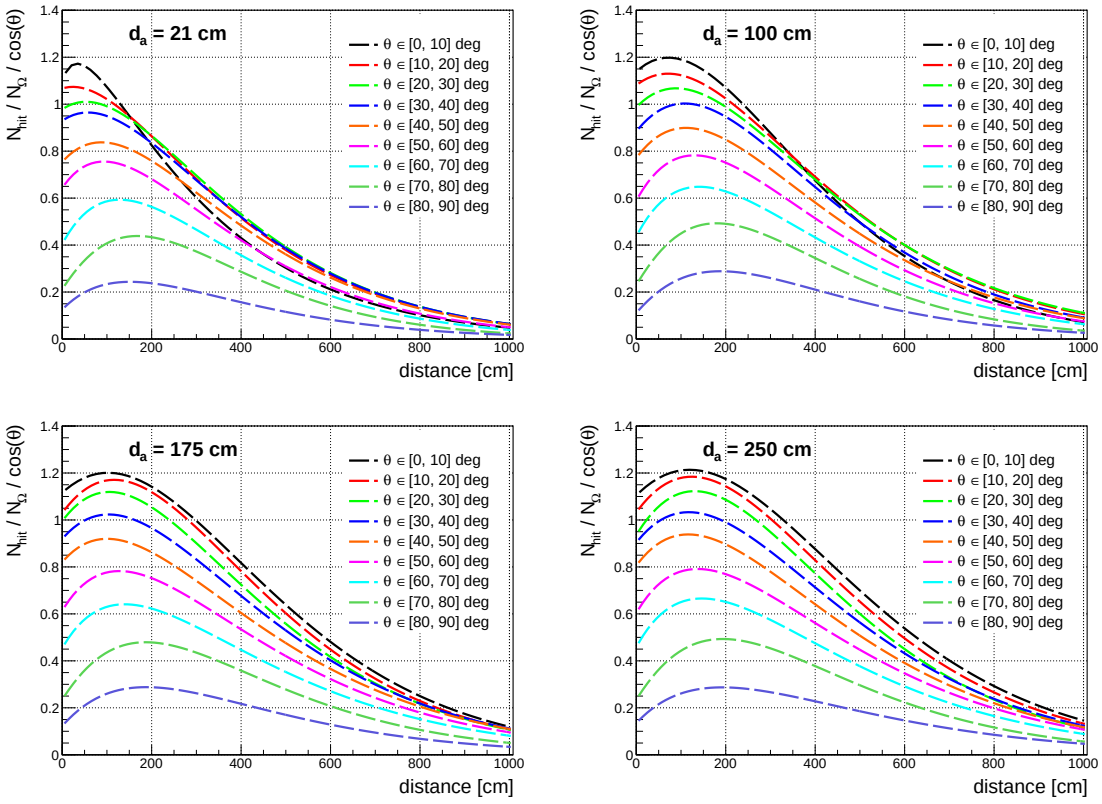

Figure 6.27: DUNE-VD lateral PD correction curves at for scintillation occurring at different distances from the anode plane, d_a. The required corrections vary significantly as a function of d_a, especially very close to the anode.

At xenon-wavelength, as with the cathode PDs, the impact of a highly reflective anode plane is modelled using an adapted version of the reflected (visible) light model. The prediction from this model is then summed with the direct light prediction to get the total number of photons incident on the PDs. Unlike the direct

light model, the d_a dependence is already incorporated into the adapted reflected light model. Therefore, the border corrections approach remains unchanged.

The lateral PDs are located external to the field cage, with approximately 50 cm of argon between them and the edge of the field cage. The shadowing effects as a result of the field cage are different depending on whether the scintillation occurs within the active volume (inside the field cage) or outside of it. In the current version of the DUNE-VD simulation, the shadowing effects are accounted for using different scaling factors depending on the position of the scintillation. In future iterations, the shadowing effect could be incorporated directly into the parameters similar to the approach taken in SBND and DUNE-HD for the sense wires. Alternatively, due to the complexity of modelling the light in this region and the potential large impact of border effects, it may be better modelled using the hybrid approach.

The performance of the model at predicting the number of photons arriving on the lateral PDs is summarised in Table 6.9 for both argon and xenon wavelength light, comparing against full simulation in Geant4. In each case the DUNE-VD 1×8×14 geometry is modelled, as described in Section 5.4.1. The performance is similar to the cathode PDs in both cases, with a resolution of approximately 10-20% at argon wavelength and 10-15% at xenon wavelength. As with the cathode PDs, there are slightly larger biases at xenon wavelength as a result of the difficulties isolating the reflected component during training of the model. Future iterations will attempt to address this, as previously discussed.

Scenario	Mean	Std Dev
argon, lateral	±1-4%	10-20%
xenon, lateral	±1-8%	10-15%

Table 6.9: Summary of the performance of the semi-analytical model in DUNE-VD for the lateral PDs, compared with full simulation in Geant4. The performance varies as a function of distance and angle in each case, indicated by the ranges given.

Propagation times

The photon propagation time model is also applied in the DUNE-VD module. At argon wavelength, the standard approach described in Section 6.2 is used. At xenon wavelength, the geometric only approximation is used as described in Section 6.5. The current implementation uses the same parameter sets that were trained for the DUNE-HD detector, which provide a reasonable approximation for initial studies. However, dedicated parameters for the DUNE-VD geometry will be needed in the future. Additionally, a full parameterisation of the propagation time distribution at xenon wavelength will be required.

6.7 Conclusions and future

In this chapter a semi-analytical model to predict the number of scintillation light photons incident on photon detectors and their arrival times was presented. The model was applied in simulations of large scale LArTPC neutrino detectors, however it could also in principle be applied to any detector constructed from materials where the Rayleigh scattering length is comparable to the volume size. Multiple scenarios were considered: argon-wavelength VUV scintillation light that propagates directly to the photon detectors, scintillation light that is reflected off wavelength-shifter coated highly-reflective foils, and xenon-wavelength scintillation light in both of the previous cases. In each scenario, the models start with a prediction from pure geometric considerations, then corrections are applied for photon transport and border effects. For the prediction of the direct argon-wavelength VUV light, a resolution better than 10% is obtained in the two different geometries studied: one SBND-like and one DUNE-like. For the reflected light, comparable performance is obtained in the smaller SBND-like detector and better than 15% resolution in the larger DUNE-like detector. For the xenon-wavelength light, similar performance was obtained for both the direct and reflected light. Furthermore, in each scenario, the prediction of the earliest photon arrival time provided by the models is within one nanosecond - better than the typical sampling used in LArTPC neutrino detectors - and the photon arrival time distribution is accurately predicted. The method is dramatically faster than the full Geant4 optical simulation ($\sim$30-70$\times$ in realistic scenarios) and outperforms the currently used look-up library method when trained with the same number of fully simulated photons. It can be applied in any large scale liquid argon detector, as well as liquid xenon or xenon-doped argon detectors, with a simple tuning of the model parameters. The model described in this chapter has been published in Reference [4].

The semi-analytical model is currently the standard fast optical simulation approach employed in the SBND, DUNE horizontal drift and DUNE vertical drift detectors. The photon transport time parameterisation is also used in the ICARUS detector, in conjunction with an optical look-up library. Additionally, the application of the model in the MicroBooNE detector is currently in development. On-going and future development will include the extension of the hybrid model (combining the semi-analytical model with scaled-down optical look-up libraries) to both of the DUNE far detectors and the development of a full parameterisation of the xenon-wavelength light propagation times. The model predictions are also currently being compared with data in the MicroBooNE detector, assessing its performance relative to the existing optical look-up library approach [312]. Additionally, proof-of-principle of data-driven correction generation has been demonstrated for the SBND

detector [311] which could negate uncertainties in the detector properties, such as the material reflectivities, in the initial Geant4 simulation used to train the model. The semi-analytical model provides an accurate and efficient fast simulation of the transport of scintillation light in large scale liquid argon detectors. It can also be employed in reconstruction algorithms, such as flash-matching, to generate predictions of the expected scintillation light signals.

Chapter 7

Simulation and data processing on high performance computers

This chapter will present the development of a method to run LArTPC simulation and data processing on high performance computers (HPCs). The method is demonstrated using the Theta HPC at the Argonne Leadership Computing Facility (ALCF) at Argonne National Laboratory (ANL). A paper describing the work presented in this chapter is in preparation.

The development of the method presented in this chapter was performed in collaboration with C. Adams (ANL) and M. Salim (ANL). The author was responsible for the development, optimisation and benchmarking of the workflows. C. Adams provided technical support throughout and M. Salim was the primary Balsam developer who implemented the various required optimisations into the Balsam workflow management software framework used to run LArSoft jobs.

7.1 Motivation

As LArTPC neutrino detectors increase in scale and complexity moving towards the multi-kiloton DUNE far detector modules, performing simulation and reconstruction is becoming increasingly computationally challenging. Additionally, more sophisticated methods of simulation and reconstruction that enhance the capabilities and uses of these detectors are continually being developed. These are typically more demanding, further increasing the computational resources needed. As a result of this the time required to process and reconstruct data, along with the generation of high-statistics Monte-Carlo (MC) simulation samples, is increasingly becoming a limiting factor in performing analyses in LArTPC neutrino experiments. In the current MicroBooNE experiment [158], processing and reconstructing the data from each year-long physics run takes multiple months. In addition to this, the high-

statistics simulation samples that are required to develop analyses can also take months to generate. This can delay the release of results and is highly inflexible, requiring the software to be frozen many months in advance to allow sufficient processing time. These issues will only get worse in future experiments. For example, the upcoming SBND experiment is expected to collect equivalent statistics to MicroBooNE's full 5 year data set every 3-4 months as a result of the proximity of the detector to the beam source [3]. This will make data processing at a rate that keeps pace with data taking extremely challenging, if not impossible, with the currently used infrastructure. High performance computing (HPC) resources could provide a solution to these problems, especially as they move towards the so-called exascale[1] computing era. There are numerous current and upcoming HPCs that could be available for use by the US-based LArTPC detector programme. These include Theta, Cori, Summit and Perlmutter; along with the upcoming exascale Aurora and Frontier.

LArTPC simulation and data processing is inherently event based, where each event can be treated individually and does not impact any other event. It is therefore trivially parallelisable by running separate events or sets of events on each available core. This is already exploited in current data processing and MC sample generation, which is primarily performed using grid-based computing resources including the FermiGrid [313] and the Open Science Grid (OSG) [314]. LArTPC production campaigns can be on the scale of many millions of events consisting of multiple different data sets or simulation samples. Each of these are split into *jobs* processing sets of events in sequence, typically $\mathcal{O}(100)$ events/job. Grid-based resources then enable processing of these in parallel at scales of typically $\mathcal{O}(1000)$ simultaneous jobs. HPC systems could expand this to a far greater scale running as many as $\mathcal{O}(100{,}000)$ simultaneous jobs, enabling each individual event to be processed in parallel fully exploiting the independent nature of LArTPC events. This has the potential to greatly reduce the length of production campaigns and to facilitate the increased computational demands of future experiments.

7.2 Theta high performance computer

The method presented in this chapter was developed and tested using the Theta HPC at the Argonne Leadership Computing Facility (ALCF) at Argonne National Laboratory (ANL). Figure 7.1 shows a photograph of Theta at ALCF. Theta is a homogeneous system based around the Intel Knights Landing (KNL) Xeon Phi processor. It consists of 4392 compute nodes, the properties of which are sum-

[1] Exascale refers to computing systems capable of processing at least 1×10^{18} double-precision floating point operations per seconds, 1 exaFLOPS.

Figure 7.1: Photograph of the Theta HPC at ALCF. Image from Reference [315].

Property	Value
CPU	Xeon Phi 7230
Cores	64
Threads per core	4
Clock speed	1.3 GHz
On-chip memory	16 GB
Regular memory	192 GB
Solid state drive	128 GB

Table 7.1: Properties of each Theta compute node [315].

marised in Table 7.1. The full system has a total of 281,088 CPU cores with a peak performance of 11.69 petaFLOPS [315]. The KNL processor clock speeds are significantly lower than typical grid nodes resulting in slower single core processing times. However, this is made up for by the orders-of-magnitude larger number of cores available. Theta was also recently expanded adding an additional 24 GPU-based nodes (Theta-GPU). However, these were not used in the simulations performed in this chapter.

7.3 Workflow design

Simulation and reconstruction in LArTPCs is performed using the LArSoft software suite [264]. An overview of the simulation and reconstruction can be found in Chapter 5. To allow production on Theta to be performed alongside conventional grid-based production, the workflow is designed to use the same releases of LArSoft as regular production campaigns with minimal modifications. To this end, the binaries provided for each LArSoft release by the Fermi National Accelerator Laboratory (FNAL) Scientific Software (SciSoft) team are used. These binaries are compiled on Scientific Linux 7 and are run on Theta using a Singularity [316] con-

tainer. The design of LArSoft incorporates the use of various software and data distribution services to facilitate grid-based production, namely the CernVM File-System (CVMFS) [317] and the Open Science Data Federation cache system (Stash Cache) [318], that are not accessible from Theta. Instead, local versions of the software binaries and required input files were transferred to Theta's central file-system and LArSoft was configured to access these directly. In addition, access to various externally hosted databases is required, typically through HTTP requests. These contain, for example, detector calibrations used during data processing. It was found that these databases could easily be overwhelmed by the potentially orders-of-magnitude more requests from jobs on Theta compared with the typical grid-based production loads. To mitigate this, local versions of the data required from these databases were also copied to Theta and LArSoft was configured access them directly. With these changes in place, LArSoft production workflows could be run on Theta and produce identical results to grid jobs.

The LArSoft software suite was historically designed to be predominantly single threaded, where each particle interaction *event* is simulated sequentially. Efforts to enable internal parallelisation of LArSoft are on-going, but will only be available in the longer term. The workflow design in conventional grid-based production campaigns is illustrated in Figure 7.2. Separate instances of LArSoft are run on each grid node, each of which processes a pre-defined set of events in sequence. This typically corresponds to the number of events in the input data or simulation file.

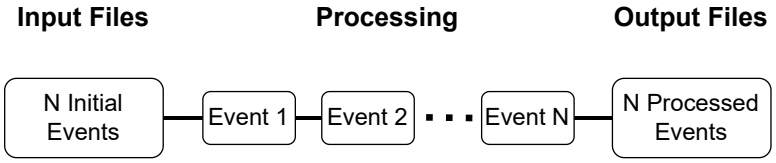

Figure 7.2: Illustration of the LArTPC data processing and simulation workflow in conventional grid-based production. Each grid node processes N events in sequence, corresponding to the number of events in the input file.

A different approach is required on Theta in order make efficient use the available compute resources, maximising the number of parallel jobs. The workflow design is illustrated in Figure 7.3. Each input data or simulation file is split into individual or small numbers of events. These are then processed in parallel running separate instances of LArSoft on each available core, cycling through as many events as possible in the available time. Finally, to avoid having many fragmented single-event files, the outputs from the individual event jobs are merged into larger files of a convenient size for subsequent processing or analysis. This removes the pre-defined job size of grid-based production and, employing this strategy, over 250,000

events could potentially be processed in parallel on Theta. This would allow high-statistics sample generation on a much shorter timescale compared with conventional production.

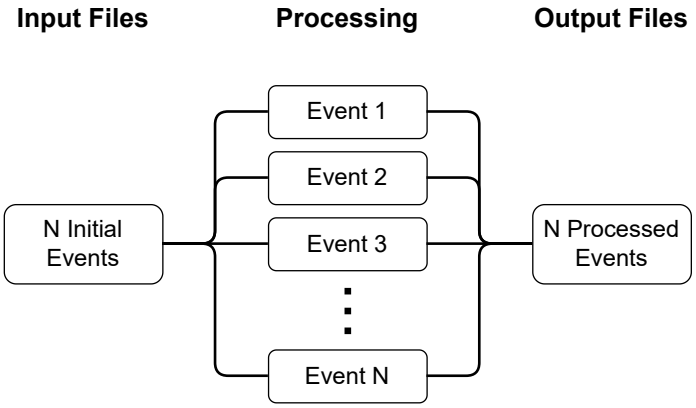

Figure 7.3: Illustration of the LArTPC data processing and simulation workflow on Theta. Each input file of N events is split into single events to be run separately on each available CPU core. They are then merged back into a single file at the end of the job, keeping track of the parentage.

To operate at such scales requires highly-efficient automated management of the workflow. This is performed by the Balsam software package [319], using the *serial-mode* configuration. This allows single node or single core applications to be run using the Balsam back-end Message Passing Interface (MPI) implementation to manage the forking of jobs to each available node or core. It is able to automate the process of splitting the input simulation or data files into individual event jobs, running each job maintaining maximal occupancy of the available cores, and merging the final output files back together keeping track of their original parent files.

7.4 Workflow optimisation

7.4.1 Initial performance

Following the approach described in Section 7.3, LArTPC simulation could be run close to out-of-the-box on Theta at very small scales. However, at larger scales severe bottlenecks were discovered. To quantify them and evaluate the workflow performance two metrics are considered: the total throughput in terms of the number of jobs completed running for one hour; and the average occupancy of the nodes with running jobs throughout this period. Figure 7.4 shows the initial job throughput running LArSoft on Theta at small scales. At a scales of 4 and 8 nodes (256 and 512

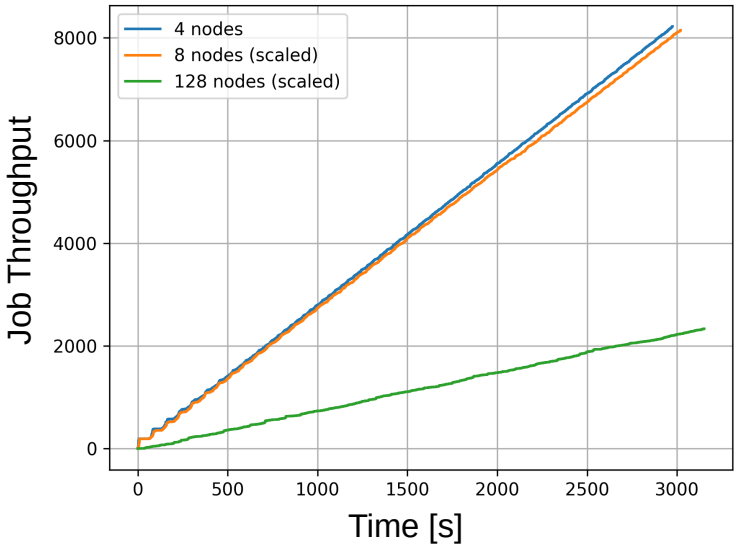

Figure 7.4: Initial job throughput running LArSoft on Theta at scales of 4, 8 and 128 nodes prior to optimisation. At 128 nodes (green), the throughput decreases to approximately 30% of optimal.

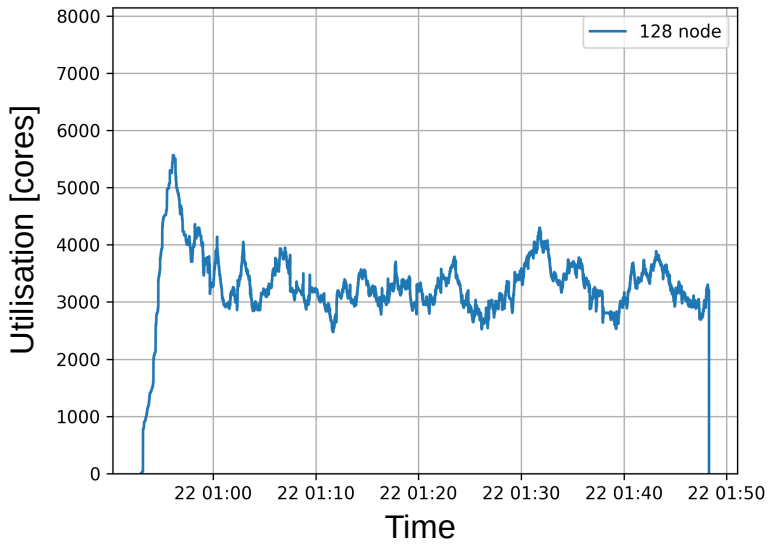

Figure 7.5: Initial core utilisation running LArSoft on Theta at 128 nodes prior to optimisation. The core utilisation is frequently below 50% of optimal (8128 cores).

cores, respectively) near optimal job throughput was achieved. However, at a scale of 128 nodes (8192 cores) this decreased dramatically with a throughput of only approximately 30% of optimal. Figure 7.5 shows the core utilisation throughout the 128 node run, which was below 50% of optimal (8192 cores) for the majority of the time. Bottlenecks were identified in two main areas: the scheduling of the jobs in Balsam and in the LArSoft input/output (I/O).

7.4.2 Balsam optimisation

Serial-mode Balsam was designed for applications running a single instance per node, where the application is then internally able to make use of the 64 cores that are available. In this scenario, a maximum of 4392 instances would be required running at the full scale of Theta. However, since LArSoft is single threaded, in this application a separate instance is required per core rather than per node. This results in a 64-fold increase in the number of sub-processes that the Balsam master process must manage. The Balsam master process has to send messages via the MPI to and from each sub-process whenever jobs start or finish. Running even at small scales in this way can rapidly outstrip the design bandwidth of the master process. For example, running at 128 nodes requires 8192 instances of LArSoft - nearly double the number simultaneous instances that would be required to run at the scale of the full system in the more conventional way. As a result of this, Balsam was unable to keep all of the available cores occupied continuously as can be seen in the low utilisation running at 128 nodes in Figure 7.5.

In order to enable efficient running at 128 nodes extensive optimisations and fixes were made in serial-mode Balsam. The optimisations with the most significant impacts are as follows:

- Several bugs were fixed that were affecting the running of more than one process per node. These were resulting in delays between jobs finishing and the core being made available for a new job.

- Non-blocking messaging and job caching were introduced. This enables the worker nodes to have a list of available jobs that they can run if the master process is busy. They can start these jobs without waiting for a response from the master process, and instead the master process can catch up when it is less busy. The same total number of messages between the master process and sub-processes are required, but spikes in the demand are smoothed out over time. This helps to avoid the master process being overloaded if, for example, a large batch of jobs finish in quick succession.

- Balsam serial-mode was re-written to use ZeroMQ messaging instead of the previously used Cray-MPI to handle process management. This was done to avoid crashes when large numbers of jobs were initialised in a short space of time as a result of incompatibilities between Cray-MPI and Theta's operating system. The crashing jobs would have to be relaunched increasing the number of jobs the master process needs to initialise, in turn increasing the fraction of jobs crashing. This cycle of crashes was leading to major issues when large batches of jobs needed to be started, especially at the beginning of a run.

Swapping to ZeroMQ resolved this issue by bypassing these incompatibilities.

- The number of logging messages written was minimised and logging was split across multiple files to reduce stress on the central file-system. The original logging was overly verbose in Balsam and scaled poorly with the increased number of processes leading to additional slowdown. The majority of logging was instead shifted to a specific debugging mode to avoid any bottleneck when running at large scale in the standard mode.

With these optimisations in place, efficient running at 128 nodes was achieved along with much greater stability. All of these optimisations were subsequently made available in a dedicated Balsam release.

To allow running at even larger scales, a new functionality was introduced that enabled the use of multiple master processes. Each master process would then manage a subset of the available nodes. This is illustrated in Figure 7.6. A central Balsam database contains information about the states of all of the jobs. Each master process then draws from this database as required. The master processes are independent and update the central database in blocks, as sets of jobs are completed, to avoid it being overloaded. They then each manage the jobs on a configurable number of nodes. It was found during testing that using 128 nodes per master process was optimal. Since the previous optimisations allow each master process to manage jobs on 128 nodes with near perfect efficiency, having multiple master processes managing this number then allows running up to the full scale of the system without further Balsam-related bottlenecks.

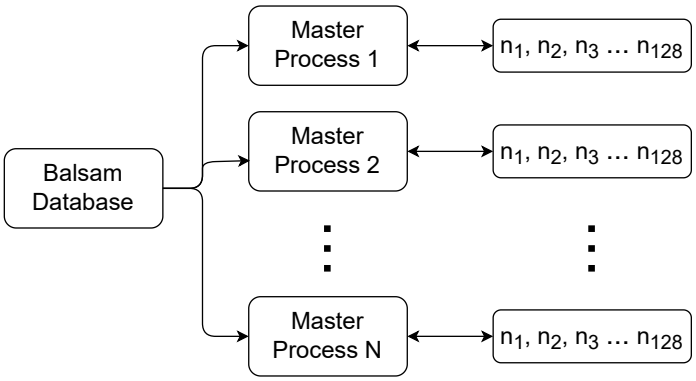

Figure 7.6: Illustration of the multiple master process structure introduced in Balsam serial-mode.

7.4.3 LArSoft workflow optimisation

The other major bottleneck identified was in the LArSoft I/O. The LArSoft binaries and required input files are stored on Theta's central Lustre file-system. Each instance of LArSoft has to read these binaries and inputs and then write any outputs to this central file-system. Since separate instances of LArSoft are run, the reads and writes to and from the central file-system are uncoordinated. This can result huge spikes in the I/O demands resulting in significant slowdown even at relatively small scales. These issues are further exacerbated as a result of LArSoft being designed to continually read and write to and from the input and output art-root event files. This results in a large number of small reads and writes, rather than the I/O being more efficiently performed in blocks, again further stressing the file-system. Despite the very high I/O bandwidth of Theta's Lustre file-system, the large number of separate uncoordinated LArSoft instances can quickly overwhelm it. The impacts of this were so severe that for runs at large scales the jobs would lock-up entirely and could even impact other unrelated jobs running on the system at the same time.

These issues cannot be resolved entirely without major structural changes in LArSoft that are beyond the scope of this work. However, they can be sufficiently mitigated through use of the solid state drives (SSDs) available as scratch space on each node. At the start of each run, a tarball containing the LArSoft binaries and any required input files can be distributed efficiently to each node using MPI-I/O. Since this is in the form of a single large ($\sim$10GB) file, so-called file striping across the Lustre file-system's multiple Object Storage Targets (OSTs) can be employed to improve transfer speeds. This allows efficient distribution of the binaries and input files without stressing the file-system. They are then unpacked on the compute nodes and each LArSoft instance uses the local version of the binaries on the node that it is running on. The copying and unpacking typically takes approximately 10 minutes at the start of the run. However, the subsequent jobs run significantly faster and, crucially, do not result in an I/O bottleneck. Additionally, to minimise the number of reads and writes to and from the central file-system, the workflow was altered to run LArSoft directly on the SSDs and to only copy the final required output files back to the central file-system at the end of each job. This allows the writing to be done in one go rather than being fragmented and avoids unnecessary files, such as intermediate stage outputs in multi-stage workflows, from being needlessly copied. Running LArSoft directly on the local SSDs also has the added benefit of marginally improving run times due to the faster I/O.

Following these changes to the workflow, efficient performance at large scales could be achieved. However, the job I/O remains a limiting factor that needs to be carefully balanced to avoid bottlenecks. As a result of this it is beneficial to run

longer jobs, for example processing several events or running multiple stages in a single job, to reduce the frequency of the transfers to the central file-system. This balance is also easier to achieve generating simulation samples since, provided the generator stage is run as part of the job, minimal input is required unlike when processing existing input files such as data.

7.5 Performance at scale

To assess the performance at scale two different benchmark applications were employed. The first was designed as a simple test of the performance of serial-mode Balsam avoiding any potential bottlenecks, such as I/O limitations, that could impact a more complex workflow. This test consisted of performing repeated matrix multiplication. The second was a realistic workflow in LArSoft performing reconstruction of cosmic-ray-induced interactions in SBND. Both applications were run with one instance per available core (64 instances per compute node), at scales between 8 and 4096 nodes (512 and 262,144 cores, respectively). In each case the performance is evaluated using the two previously-defined metrics: the total throughput in terms of the number of jobs completed running for one hour and the average occupancy of the nodes with running jobs throughout this period.

7.5.1 Benchmarking using matrix multiplication

The matrix multiplication benchmark application performs multiplication of randomly generated 1000×1000 matrices using NumPy [320]. Each job is configured to run on an individual core with a single thread and performs 1000 iterations of the matrix multiplication. The average run-time of the application on the Theta compute node cores is approximately 300 seconds. This relatively short run-time was chosen to allow multiple batches of jobs to be cycled through during a 1 hour run on Theta. This allows stress-testing serial-mode Balsam processing extremely large numbers of jobs, emulating a worst-case scenario.

Figures 7.7 and 7.8 show the performance of the matrix multiplication benchmark runs at scales between 8 and 4096 nodes. Figure 7.7 shows the utilisation of the cores over time and Figure 7.8 shows the job throughput in terms of the number of completed jobs per node over time. The small scale 8 node run (blue) provides a baseline optimal performance with negligible scaling related overhead from Balsam serial-mode. The run maintains 100% utilisation of the available cores throughout, except for small dips as sets of jobs finish and new jobs are started to replace them. The run throughput shows a step-like structure as a result of the jobs approximately starting at the same time and having the same length, and hence finishing in batches.

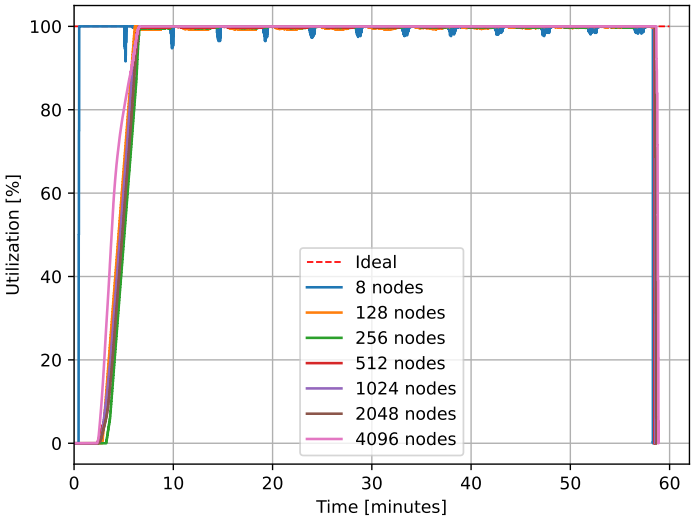

Figure 7.7: Matrix multiplication benchmark core utilisation over time for runs at scales between 8 and 4096 nodes. The core utilisation is plotted relative the start time of the run on Theta, allowing the differences in the Balsam start-up time at different scales to be seen.

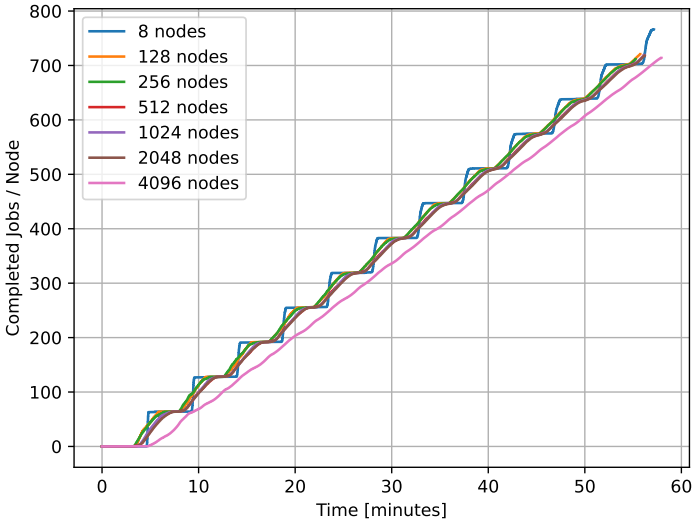

Figure 7.8: Matrix multiplication benchmark run throughput over time for runs at scales between 8 and 4096 nodes. The throughput is plotted relative to the start time of the first running job to allow easier comparison in the job completion rate at different scales.

The throughput is approximately uniform over time throughout the run. At 128 nodes (orange), there is an increased start-up time due to the larger number of jobs managed by the Balsam master process resulting in a delay before 100% core utilisation is achieved. Once started, full occupancy of the cores is maintained throughout the run. The start-up delay results in the step-like structure of the

throughput being smeared out. However, once started, Balsam is able to maintain a continuous throughput comparable to the small scale job indicating there are minimal additional scaling-related overheads. At larger scales between 256 nodes (green) and 2048 nodes (brown) the same pattern is seen, with the utilisation and throughput curves approximately overlapping throughout the runs. At these larger scales multiple master processes are used each of which manages the jobs on 128 nodes. Therefore, no additional start-up overhead is present and serial-mode Balsam is able to scale effectively without further decrease in throughput. The 4096 node run (pink) follows a slightly different pattern with a slower initial start up resulting in the throughput line being off-set from the others. Once running however, near 100% core utilisation is maintained and the rate of jobs completing is comparable to the smaller scale runs with only a small decrease in total throughput. This suggests that this run is on the threshold of hitting a bottleneck, possibly in the communication between the larger number of master processes with the central job server. This bottleneck is exacerbated by the relatively short run-time of each job, resulting in a large number of jobs starting and finishing within the 1 hour run. In practice this could be avoided by designing workflows with longer run times per job.

Scale [nodes]	Throughput [jobs]	Relative efficiency
8	6128	-
128	92,265	94.1%
256	182,164	92.9%
512	367,894	93.8%
1024	736,428	93.9%
2048	1,470,209	93.7%
4096	2,923,470	93.2%

Table 7.2: Matrix multiplication benchmark throughput and relative efficiency for runs at scales between 8 and 4096 nodes. The efficiency is quoted relative to the throughput of the small scale run at 8 nodes where any scaling related bottlenecks would be negligible.

Table 7.2 shows the total throughput the matrix multiplication benchmark runs at each scale, and their efficiency relative to the small scale 8 node run. There is an approximately 6% decrease in throughput efficiency between the 8 node and 128 node runs as a result of the longer start-up time. The impact of this is especially significant in this example application due to the short run time of the jobs allowing a full additional set of jobs to complete during the 8 nodes run compared with at 128 nodes. In an application with a longer run-time, or if running on Theta for longer than 1 hour, the impact of the start up delay would be reduced. Between 128 and 4096 nodes, although there is some small variation between individual runs ($\sim 1\%$), no further significant decrease in the efficiency is observed. In total during

the 4096 node run nearly 3 million jobs are completed in 1 hour, demonstrating the extremely high throughput achievable by the optimised serial-mode Balsam.

7.5.2 Benchmarking using LArSoft reconstruction

To benchmark the performance of serial-mode Balsam running LArSoft at scale on Theta, a typical LArTPC reconstruction workflow was run. This was performed using the SBND software framework. The SBND experiment is described in Section 3.6.2 and an overview the reconstruction performed in LArTPC detectors is given in Section 5.3. In SBND the Pandora reconstruction framework [289] is used, which is discussed in Section 5.3.4. The reconstruction was run on a sample of simulated interactions resulting from cosmic-ray showers. The sample was generated using the CORSIKA generator, described in Section 5.1.2, and the standard particle propagation and detector simulation performed as described in Section 5.2. Each job runs a separate instance of LArSoft on a single core and performs reconstruction of five simulated events. For the purposes of benchmarking, the same set of five events are used each time. The average run-time of this benchmark on the Theta compute node cores is approximately 500 seconds. As with the matrix multiplication benchmark, a relatively short run-time was chosen to allow multiple batches of jobs to be cycled through during a one hour run on Theta. In practice, a larger number of events would likely be run per job to reduce the relative impact from start-up time.

Figures 7.9 and 7.10 show the performance of the LArTPC reconstruction benchmark runs at scales between 8 and 4096 nodes. Figure 7.9 shows the utilisation of the cores over time and Figure 7.10 shows the job throughput in terms of the number of completed jobs per node over time. As with the matrix multiplication benchmark, the small scale 8 node run (blue) provides a baseline, optimal, performance with negligible scaling-related overheads from Balsam serial-mode or other factors such as I/O. At the start of each run the LArSoft binaries are distributed in a tarball to each compute node local SSD, as previously described, taking approximately 10 minutes to complete. The 8 node run then follows a similar, step-like, pattern to that of the matrix multiplication benchmark runs, maintaining near 100% core utilisation throughout. At 128 nodes (orange), there is an increased start-up time as before due to the larger number of jobs managed by the Balsam master process. Once 100% core utilisation is reached, full occupancy is maintained throughout the run and a comparable total throughput is achieved. As with the matrix multiplication runs, the step-like structure is smoothed out slightly due to the start-up delay. However, the impact of this is less noticeable due to the longer per-job run-time. Between 256 nodes (green) and 4096 nodes (pink) the initial start up times increase

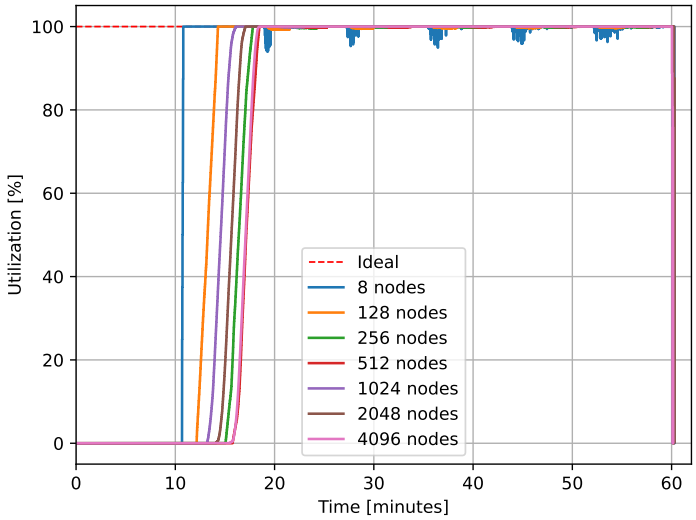

Figure 7.9: LArSoft reconstruction benchmark core utilisation for runs at scales between 8 and 4096 nodes. The core utilisation is plotted relative the start time of the run on Theta, allowing the differences in the Balsam and LArSoft start-up time at difference scales to be seen.

slightly and vary between runs. This is due to the larger number of nodes that the binaries need to be distributed to and the file-system bandwidth available which depends on the usage by other applications running on Theta at the same time. Once 100% utilisation is reached, a similar pattern to the smaller scale runs is seen with the utilisation and throughput curves approximately overlapping throughout. No additional overheads occur and the application is able to scale effectively to the full size of the system, running 262,144 simultaneous instances of LArSoft without any further decrease in throughput. Unlike the matrix multiplication benchmark, no slowdown is observed running at 4096 nodes since the job run-time is longer and hence the throughput is reduced to a level that does not cause a bottleneck.

Table 7.3 shows the total throughput of the LArTPC reconstruction benchmark runs at each scale in terms of number of jobs completed and number of LArTPC events reconstructed. It also shows the efficiency of each run relative to the 8 node case where any scaling related bottlenecks would be negligible. Unlike for the matrix multiplication benchmark, no significant decrease in efficiency is seen between the 8 node and 128 node runs despite the longer start-up time. Since the time taken to run each job is longer, the time saved as a result of the reduced start-up at 8 nodes is insufficient to complete an additional batch of jobs. Between 128 and 4096 nodes, although there is some small variation between individual runs ($< 1\%$), no significant decrease in efficiency is observed. The 4096 node run achieves 99.4% efficiency relative to the 8 node run, demonstrating near perfect scaling despite the substantially more complex application used. At 4096 nodes

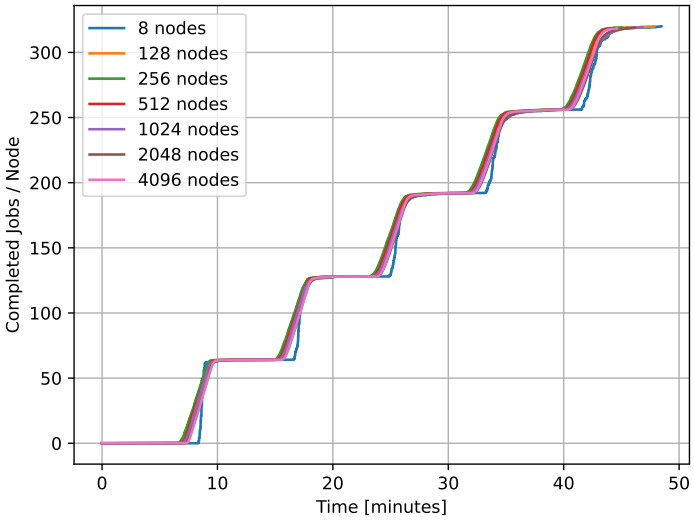

Figure 7.10: LArSoft reconstruction benchmark run throughput for runs at scales between 8 and 4096 nodes. The throughput is plotted relative to the start time of the first running job to allow easier comparison in the job completion rate at different scales.

Scale [nodes]	Throughput [jobs]	Throughput [events]	Relative efficiency
8	2560	12,800	-
128	40,951	204,755	99.9%
256	81,683	408,430	99.7%
512	163,175	815,875	99.6%
1024	327,015	1,635,075	99.8%
2048	652,877	3,264,385	99.6%
4096	1,303,383	6,516,915	99.4%

Table 7.3: LArTPC reconstruction benchmark throughput and relative efficiency for runs at scales between 8 and 4096 nodes. The efficiency is quoted relative to the throughput of the small scale run at 8 nodes where any scaling related bottlenecks would be negligible.

over 6.5 million LArTPC events are reconstructed in 1 hour. This far surpasses the equivalent achievable on existing grid-based resources, demonstrating the power of using systems like Theta even for predominantly single-threaded applications like LArSoft.

7.6 Realistic production workflow example

In the previous section, near 100% throughput efficiency was achieved running LArSoft at large scales on Theta. However, this was performed with a controlled benchmark workflow rather than a real application. In this section, the method developed

will be applied to a realistic production workflow generating the first large scale production sample with LArSoft on a HPC. This will include generating the sample, processing the output files, transferring them to FNAL, and archiving them to tape in an analogous manner to standard grid production samples.

7.6.1 Sample

For this test, a sample from an on-going SBND MC production campaign was duplicated using the same tagged and frozen version of LArSoft. The sample consisted of 250,000 cosmic-ray interaction events, now running the full simulation chain: generation with CORSIKA, particle propagation with Geant4, detector simulation, and then reconstruction. The various stages of this simulation and reconstruction are described in Chapter 5. This sample was chosen because the workflow is relatively simple yet it is typically one of the most computationally demanding samples to generate due to the large number of particles produced as a result of cosmic-ray interactions. The sample was split into single event jobs, each running the full simulation chain. As discussed in Section 7.4.3, this allowed the outputs from intermediate stages to be discarded only keeping the final reconstructed file, minimising the stress on the central file-system. The final single event files were then merged together into files of 50 events, approximately equivalent to the outputs of the grid-based production.

The run-time of the individual event jobs varied significantly depending on the number of particles from the simulated cosmic-ray showers that traverse the cryostat, spanning between around 15 minutes to over 2 hours. This test was performed with an older version of LArSoft than the benchmarking performed in Section 7.5 where simulation and reconstruction were significantly slower. This is both due to configuration differences, here a significantly larger drift window was simulated resulting in much larger events, and due to a significant amount of optimisation being performed and implemented into the newer version. The run times are therefore not representative of current SBND simulation, however the workflow is equivalent.

7.6.2 Generation on Theta

The bulk of the sample was generated in a single large scale job running at 1024 nodes for 3 hours. This approach was taken since larger jobs gain higher priority on Theta's queuing system. This first job completed just over 200,000 events, approximately 80% of the full sample. Next, several smaller make-up jobs were run at 256 and 128 nodes bringing the total number of completed events to 249,790, just shy of the target. Finally, a separate small scale job was run to merge the events together into the final files of 50 events each.

The large scale job ran with high initial throughput efficiency. However, this decreased during the final hour of running where a substantial fraction of jobs were being killed soon after starting. This is possibly as a result of a small subset of long-running jobs with very high memory usage restricting the resources available on some of the nodes. These are events with extremely high particle multiplicity that are challenging to simulate. Typically in grid production campaigns these high memory, high run-time, jobs would be terminated as they would exceed the maximum memory available on the grid node. It is possible that in the future a similar approach could be taken on Theta, although this risks biasing the sample. This issue is likely less prominent in more recent versions of the SBND simulation due the previously mentioned optimisation.

7.6.3 File transfer to FNAL

Once a sample has been generated on Theta, the resulting output files need to be transferred to FNAL where they can be accessed by analysers and archived to tape. The data transfer between ANL and FNAL was performed using the Globus file transfer system [321]. The total sample size was approximately 60 TB and the transfer to FNAL took approximately 14-16 hours at speeds of approximately 1.2 GB/s. The copying of the output files to FNAL is a significant bottleneck, taking substantially longer than the file generation. Recent optimisation efforts in SBND have dramatically reduced the output art-root file sizes. However, this remains a bottleneck and faster methods of data transfer need to be developed to make full scale production on systems like Theta viable.

For this sample, the transfer was performed after the sample generation had completed using a custom script that allowed both ANL and FNAL access permissions to be active simultaneously. In the longer term, a more robust method would be necessary. The most recent version of Balsam can be configured to perform these transfers as the jobs complete allowing the process to be automated and continuous, increasing efficiency. This feature was not available at the time the tests described here were performed, and hence the performance for this use case has not yet been assessed. Other potential solutions to this bottleneck, including storing the samples at ANL rather than FNAL, could also be explored.

7.6.4 Metadata, SAM and transfer to tape

Once the files have been copied to FNAL they need to be processed, added to the sample database used by the relevant experiment and transferred to long term storage. The database system used by most FNAL experiments is Serial Access to Metadata (SAM). In order for the files produced on Theta to be compatible

with this system, various file metadata must be added. The workflows used to simulate and reconstruct events on the grid are configured to perform this task automatically as part of the production. For Theta produced files, this is instead done in two stages. The majority of the file metadata can be added during the file generation, specifically at the merging phase when the final files are created. This can be configured to mimic the metadata of grid-produced files, although with an additional tag to identify the files as having originated from Theta. The final metadata entries require access to the FNAL network and SAM in order to generate a unique file address. This is therefore performed after the transfer to FNAL during the file processing.

The file processing was performed using the FNAL File Transfer Service (FTS) used in grid-based production. A custom FTS dropbox and metadata extraction script were created for Theta produced files. The files could then be copied directly into this dropbox where they would be automatically processed generating the final required metadata, declaring the files to the SAM database and transferring them to a staging area to be copied onto tape. For the 60 TB sample, it took approximately 4 days for the FTS to process the files and another 3-4 days to complete the transfer to tape. This was therefore by far the largest bottleneck in the production workflow. The FTS is also a significant bottleneck for grid-based production campaigns, and, since it is used by many different experiments, can often cause significant issues if multiple simultaneous production campaigns are occurring. Alternatives to using the FTS are being explored. In SBND these include running the metadata extraction scripts externally as part of separate jobs and copying files directly to persistent dCache storage. This serves as an intermediate area accessible to analysers immediately rather than requiring the FTS to process them first. The files can then be archived to tape as usual without analysers having to wait. Similar techniques could be implemented for the Theta produced files to avoid these kind of delays.

7.6.5 Summary

A large simulation sample was successfully generated on Theta and made available to analysers, demonstrating an end-to-end production workflow. This was the first time a large scale production sample for a LArTPC experiment has been generated on a HPC. The sample could be generated very quickly on Theta, demonstrating the power of the system. However, significant bottlenecks were identified in copying the files to FNAL and processing them once there. Solutions to these issues would need to be developed in order for Theta to be used more extensively in production campaigns. Many of these issues, especially with the FTS, also impact grid-based production and various solutions are being explored that could also be implemented

for Theta-based production.

7.7 Conclusions and future

High performance computers could provide a solution to the computing challenges faced by LArTPC neutrino experiments as they increase in scale and complexity, outstripping existing grid-based computing resources. Current homogeneous CPU-based systems like Theta can already provide substantial computing power and have the potential to be used in a manner complementary to grid-based production, especially for the most highly computationally demanding simulation samples. In this chapter, a method to run LArTPC simulation and data processing with LArSoft on high performance computers has been presented. The performance of this method was demonstrated using the Theta HPC at ANL, achieving efficient scaling to the full size of the system and job throughputs far surpassing the capabilities of traditional grid-based production. Additionally, an end-to-end production workflow was demonstrated including sample generation on Theta, transfer of the output files from ANL to FNAL, and processing and archiving the files to tape making them available for analysers to use. This was the first time a large scale LArTPC production sample had been generated on a HPC. A paper describing the work presented in this chapter is in preparation.

LArSoft in its current state, however, is poorly optimised to make use of these systems. Substantial efforts are on-going to allow event-level internal parallelisation and eventually full multi-core support. This would allow, for example, only a single instance of LArSoft to be run on each node that then internally generates separate events for each core, greatly reducing overheads. Significant bottlenecks were also identified in the file transfer to FNAL and subsequent processing that require further development to resolve. These issues also impact grid-based production and various solutions are being explored. If HPC-based production for LArTPC experiments becomes more widespread, an alternative approach could be to make use of the Eagle file-system at ANL. This is designed for high-bandwidth data-sharing between HPC centres and could be configured to allow read access from FNAL. Finally, as HPCs increasingly move towards heterogeneous architectures incorporating both CPUs and GPUs, further modification of LArSoft would be required to make use of the functionality provided by GPUs. If these systems are harnessed the computing power available could not only accommodate the increasing demands of the existing software, but allow more powerful techniques to be used. For example, rather than requiring simplified light simulation models such as that described in Chapter 6, full GPU-accelerated ray-tracing could be performed. The future exascale systems such as Aurora and Frontier have the potential to greatly enhance the US-based

LArTPC detector program through the orders-of-magnitude-greater computing resources available. A substantial amount of further development will be required to make this a reality, however.

Chapter 8

Searching for Heavy Neutral Leptons with ArgoNeuT

This chapter describes a search for Heavy Neutral Leptons with masses $\mathcal{O}(100)$ MeV performed using the ArgoNeuT experiment. The Heavy Neutral Leptons could be produced in the NuMI neutrino beam and then subsequently decay with the signature $N \to \nu\mu^+\mu^-$, which is possible to detect in ArgoNeuT and the MINOS near detector. The analysis described in this chapter has been published in Reference [5].

The search presented in this chapter was performed in collaboration with phenomenologists K. J. Kelly (FNAL, now at CERN) and A. de Gouvêa (Northwestern University) who were responsible for the theoretical parts of the analysis. The author was responsible for the experimental parts of the analysis.

8.1 Model and Simulation

The Heavy Neutral Lepton (HNL) model, production in the NuMI beam target and absorber, and subsequent decay in ArgoNeuT are described in detail in Sections 2.2.3 and 4.2. The generation of the simulated HNLs and their decays is performed using the method described in Section 5.1.3. The resulting decay products are then propagated through the standard ArgoNeuT and MINOS near detector (MINOS-ND) simulation and reconstruction chains, an overview of which is given in Chapter 5. Three scenarios are considered depending on the mixing of the HNL, N, with the standard model light neutrinos ν_α, $\alpha = e, \mu, \tau$: electron-coupled, where $|U_{eN}|^2 \neq 0$; muon-coupled, where $\left|U_{\mu N}\right|^2 \neq 0$; and tau-coupled, where $|U_{\tau N}|^2 \neq 0$. In each case the other two mixing angles are assumed to be zero. In all three scenarios it is assumed that the HNL is a Dirac particle. The decay channel $N \to \nu\mu^+\mu^-$ is selected as the signature due to ArgoNeuT's unique sensitivity to it when exploiting the downstream MINOS-ND. This signature is one of the dominant decay modes

Production mechanism	m_N [MeV]	Number of events
$\tau^{\pm}$ decay (tau-coupled)	400	50000
$K^{\pm}$ decay (electron/muon-coupled)	300	24750
$D^{\pm}_{(s)}$ decay (electron/muon-coupled)	1000	25000

Table 8.1: Details of the HNL MC simulation samples generated for the tau-coupled and electron/muon-coupled models. In each case, the HNLs were simulated with their expected position and energy distribution in the ArgoNeuT detector.

for tau-coupled HNLs, as discussed in Section 2.2.3. The tau-coupled scenario is therefore focused on throughout the development and optimisation of the analysis.

Two sets of HNL Monte Carlo (MC) simulation samples were created. The first consisted of HNL decays distributed throughout the ArgoNeuT detector with energies, E_N, sampled from the distributions predicted by the simulation of the HNL production in the NuMI beam target and absorber. This first set of samples were used to develop the experimental selection. Three samples were generated covering the different production modes: $\tau^{\pm}$, $K^{\pm}$ and $D^{\pm}_{(s)}$ decay. The samples are summarised in Table 8.1. The second set of samples consisted of events generated at fixed m_N and E_N and was used to study the selection efficiency in detail to allow the final sensitivity to be evaluated. These samples were not model-specific and could be used to evaluate the electron-coupled, muon-coupled and tau-coupled scenarios by weighting the energy distributions accordingly. Samples were generated with $m_N = 300 - 600\,\text{MeV}$ in steps of $25\,\text{MeV}$ and $E_N = 1 - 55\,\text{GeV}$ in steps of $1\,\text{GeV}$ up to $E_N = 10\,\text{GeV}$ then $E_N = 15, 25, 35, 45$ and $55\,\text{GeV}$. Additionally, the full set of mass-energy combinations was simulated with the decays occurring inside the detector and at different distances into the cavern between $z = -(12.5 - 62.5)\,\text{cm}$ in steps $12.5\,\text{cm}$, where $z = 0\,\text{cm}$ is the start of the instrumented volume in ArgoNeuT. For each of the 195 mass-energy combinations 2500 events were generated, for a total of 487,500 events at each of the 6 positions.

The ArgoNeuT simulation of the NuMI beam, described in Section 5.1.1, was also used during the development of the selection to model the expected backgrounds. The sample used consisted of 600,000 events corresponding to 9.1×10^{20} protons-on-target (POT), 7.28 times the total ArgoNeuT data POT. This sample will subsequently be referred to as the beam simulation sample.

8.2 Signature

8.2.1 Kinematics

The HNL energy, E_N, distributions are shown in Figure 8.1 for the $K^{\pm}$ decay electron/muon-coupled sample (top-left), the $D^{\pm}_{(s)}$ decay electron/muon-coupled sam-

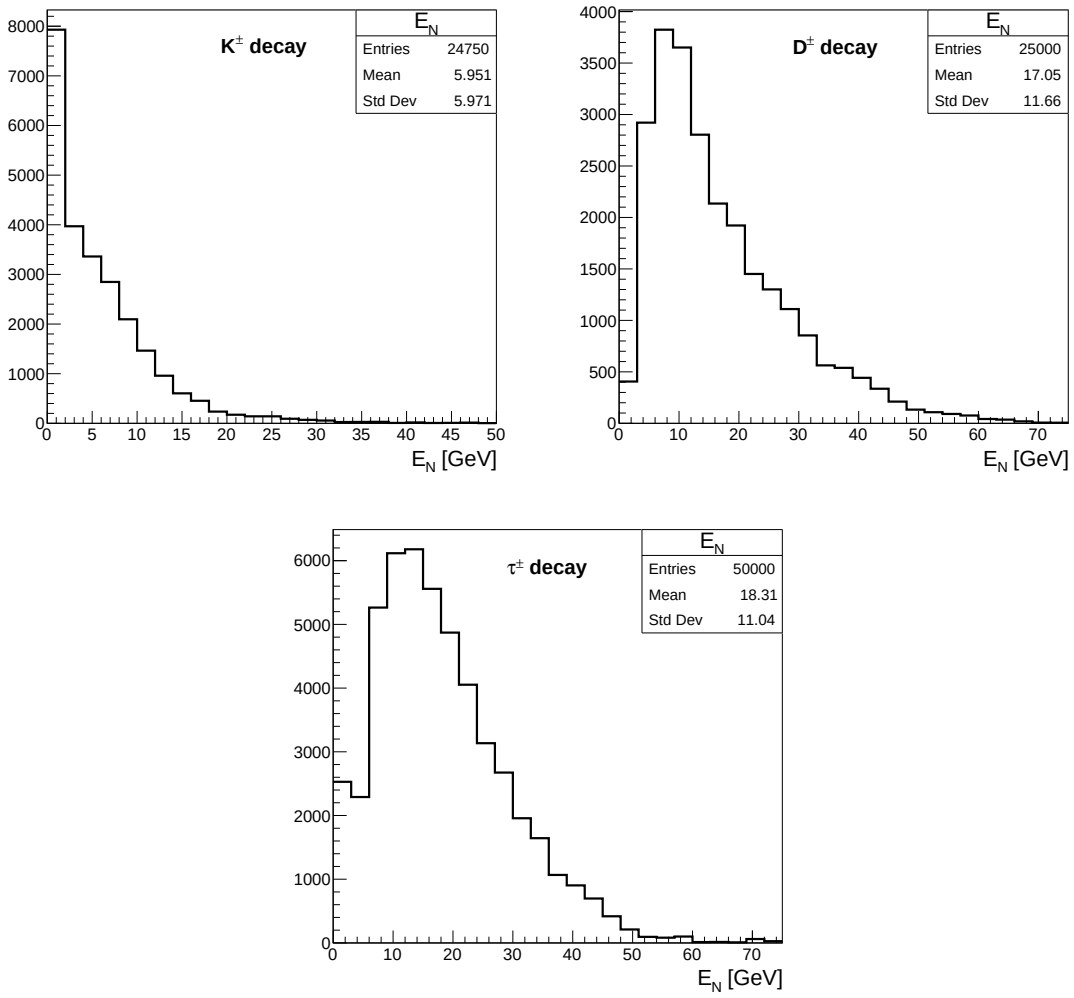

Figure 8.1: HNL energies, E_N, for the $K^{\pm}$ decay electron/muon-coupled sample (top-left), the $D^{\pm}_{(s)}$ decay electron/muon-coupled sample (top-right) and the $\tau^{\pm}$ decay tau-coupled sample (bottom). The average energies and their distributions vary substantially between the three different production mechanisms.

ple (top-right) and the $\tau^{\pm}$ decay tau-coupled sample (bottom)[1]. The average HNL energies are $\langle E_N \rangle \sim 6\,\text{GeV}$, $\langle E_N \rangle \sim 17\,\text{GeV}$ and $\langle E_N \rangle \sim 18\,\text{GeV}$ for the $K^{\pm}$ decay, $D^{\pm}_{(s)}$ decay and $\tau^{\pm}$ decay samples, respectively. The average energies and their distributions vary substantially between the three different production mechanisms.

The HNL decay signature $N \to \nu\mu^{-}\mu^{+}$ in ArgoNeuT consists of two highly-forward-going muons that exit the detector towards the MINOS-ND. Figure 8.2 shows the truth energies (top-left), angles with respect to the beam direction (top-right), and opening angles (bottom) of the muons from the tau-coupled $\tau^{\pm}$ decay

[1]In these samples the HNLs are assumed to be long-lived relative to the target-ArgoNeuT distance. However, it should be noted that a small subset of the HNL phase-space of interest in ArgoNeuT would result in HNL decay lengths of a similar scale to this distance. These HNLs would have a slightly different energy distribution, which will be accounted for appropriately in the final selection efficiency evaluation. The long-lived HNL samples are representative of the majority of the phase-space of interest and are therefore used for the development of the selection.

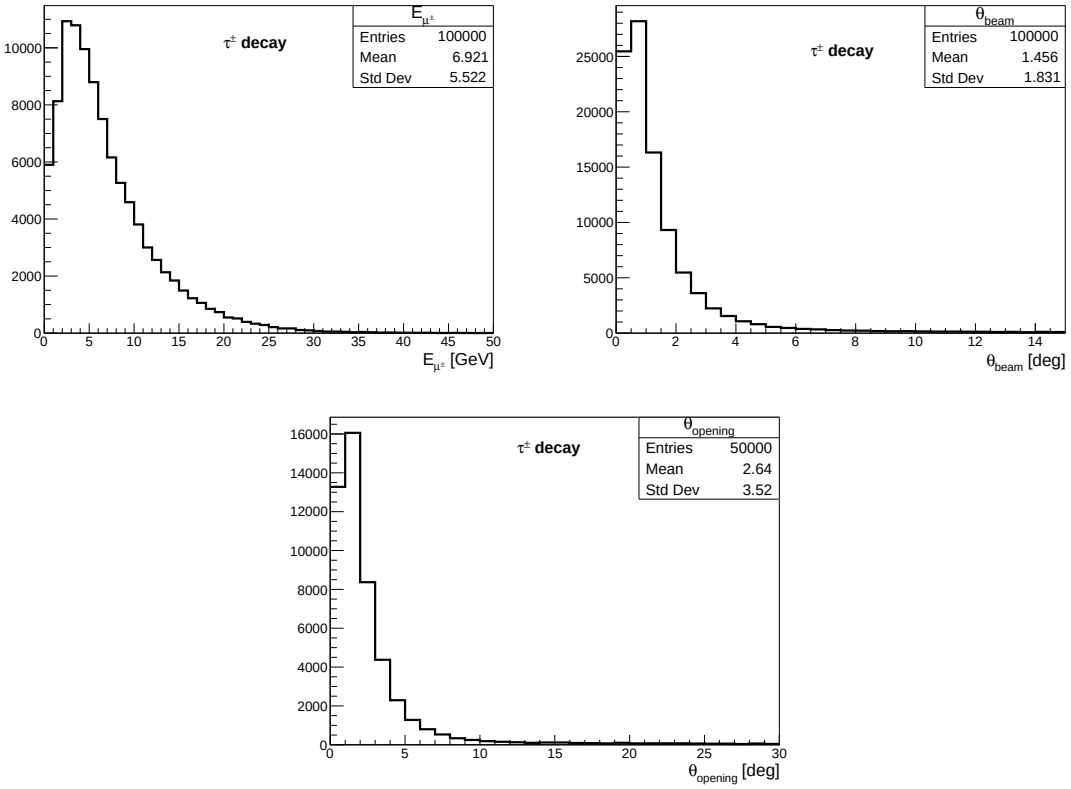

Figure 8.2: Truth energies (top-left), angles with respect to the beam direction (top-right), and opening angles (bottom) of the muons from the tau-coupled $\tau^{\pm}$ decay HNL sample with $m_N = 400\,\text{MeV}$. The resulting muons are energetic and highly forward-going.

HNL sample with $m_N = 400\,\text{MeV}$. The muons are energetic and highly forward going: with an average energy of $\langle E_{\mu^{\pm}} \rangle \sim 7\,\text{GeV}$; an average angle with respect to the beam direction of $\langle \theta_{beam} \rangle \sim 1.5^{\circ}$; and an average opening angle of $\langle \theta_{opening} \rangle \sim 2.6^{\circ}$.

Plots of the truth muon energies, angles with respect to the beam direction, and opening angles for the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples (electron/muon-coupled scenarios) can be found in Appendix A. In these cases $m_N = 300\,\text{MeV}$ and $m_N = 1000\,\text{MeV}$, respectively. For these samples the resulting muons have average energies of $\langle E_{\mu^{\pm}} \rangle \sim 2.5\,\text{GeV}$ and $\langle E_{\mu^{\pm}} \rangle \sim 5.7\,\text{GeV}$, respectively. They are also highly forward-going, although less so than for the $\tau^{\pm}$ decay mechanism, with average angles with respect to the beam direction of $\langle \theta_{beam} \rangle \sim 4.6^{\circ}$ and $\langle \theta_{beam} \rangle \sim 5.9^{\circ}$, respectively; and average opening angles of $\langle \theta_{opening} \rangle \sim 8.2^{\circ}$ and $\langle \theta_{opening} \rangle \sim 10.5^{\circ}$, respectively.

8.2.2 Decays inside of ArgoNeuT

ArgoNeuT's angular resolution is approximately 1-3°, depending on the track orientation [270]. This sets a limit to how small of an opening angle between two

tracks can be resolved. For the highly-boosted HNL decays, the resulting muons frequently overlap and may be reconstructed in ArgoNeuT as a single track for part or all of their length. Two in-ArgoNeuT decay signatures are therefore considered, both illustrated in Figure 8.3.

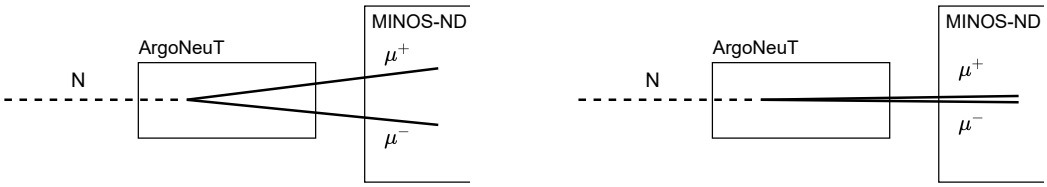

Figure 8.3: Diagrams of HNL decays occurring inside the ArgoNeuT detector with the two-track (left) and double-MIP (right) decay signatures. Diagrams not to scale.

In the first case, if the opening angle between the two muons is sufficiently large, the muons are reconstructed in ArgoNeuT as two distinct tracks originating from a common vertex. Each of these tracks can then be matched separately to oppositely charged tracks in the MINOS-ND. This topology is illustrated in Figure 8.3 (left) and will subsequently be referred to as a *two-track* type event. An example event display

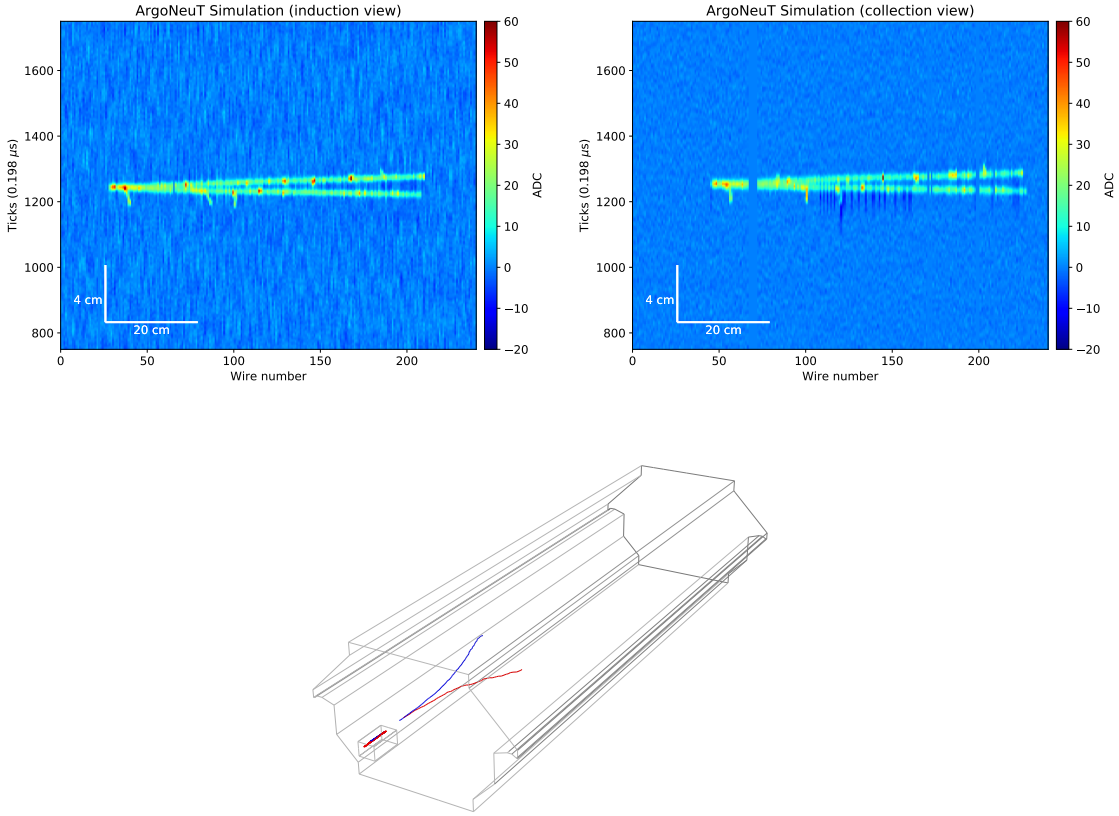

Figure 8.4: Simulated HNL decay with the two-track topology. The pair of muon tracks are initially overlapping ArgoNeuT (top) in both the induction (left) and collection (right) wire-plane views, but then separate allowing them to each be matched to tracks in the MINOS-ND (bottom).

in ArgoNeuT and the MINOS-ND with this signature can be seen in Figure 8.4.

In the second case, where the opening angle is smaller, the pair of muons may be reconstructed as a single track close to the vertex. Then, as the muons propagate and begin to separate further, they split into two separate tracks. Depending on how forward-going the muons are, this can occur either inside of ArgoNeuT or only once they reach the MINOS-ND where they are separated by the applied magnetic field. ArgoNeuT's reconstruction is not optimised to identify these splitting events leading to frequent mis-reconstruction, complicating the development of the selection. To work around this, the region close to the interaction vertex where the muons are most likely to overlap is focused on. In this region they may be reconstructed as having a dE/dx approximately double that of a single minimally ionising particle (MIP), allowing them to be identified. This single more-highly-ionising track in ArgoNeuT can then be matched to two tracks in the MINOS-ND. This topology is illustrated in Figure 8.3 (right) and will subsequently be referred to as a *double-MIP* type event. An example event display in ArgoNeuT and the MINOS-ND with this

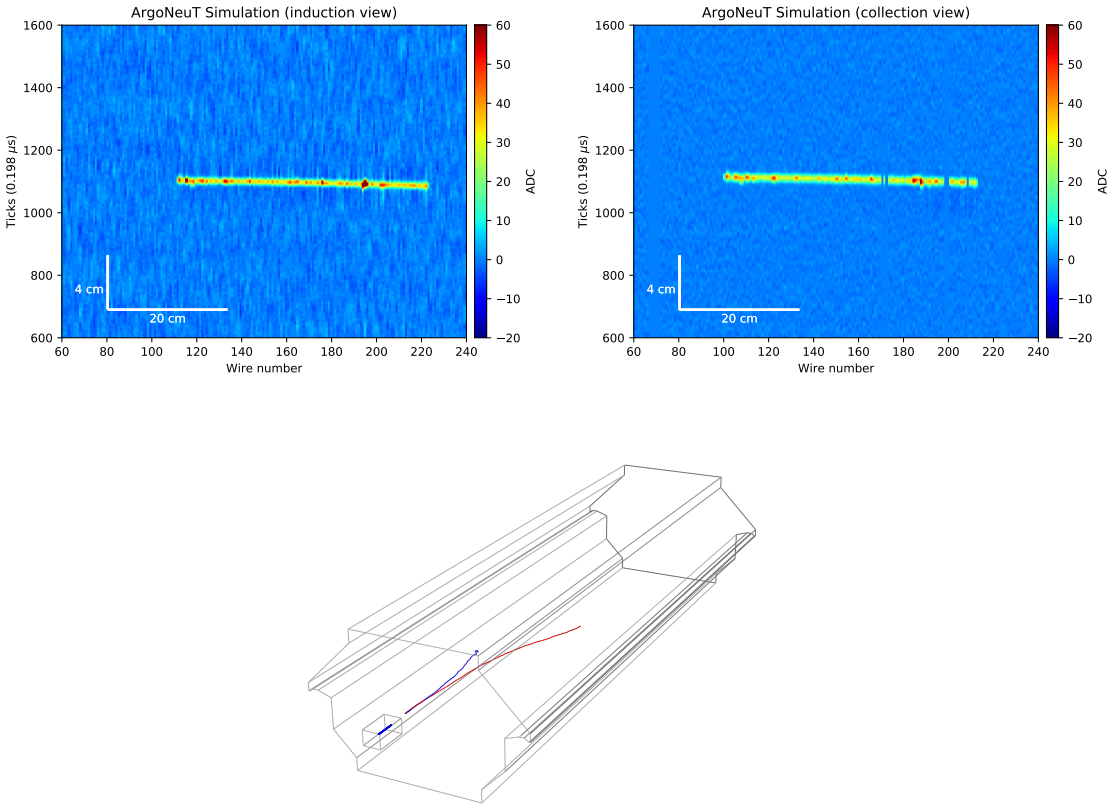

Figure 8.5: Simulated HNL decay with the double-MIP topology. The pair of muon tracks are fully overlapping in ArgoNeuT (top) in both the induction (left) and collection (right) wire-plane views, then split once reaching the MINOS near detector (bottom) due to the applied magnetic field. The single reconstructed track in ArgoNeuT can therefore be matched to a pair of tracks in the MINOS-ND.

signature can be seen in Figure 8.5, where the two muons fully overlap in ArgoNeuT then split once reaching the MINOS-ND.

8.2.3 Decays upstream of ArgoNeuT

In addition to decays occurring within the ArgoNeuT detector, decays occurring upstream of ArgoNeuT where the resulting muons then pass through the detector can be considered. This scenario is illustrated in Figure 8.6. During the ArgoNeuT physics run, the MINERvA detector [257] was under construction in the cavern upstream of ArgoNeuT. The downstream end of the MINERvA detector closest to ArgoNeuT consisted of steel plates and scintillator strips, similar to the MINOS-ND. These would substantially impact the propagation of muons through the detector resulting in them being deflected or, for lower energy muons, stopping. Since the MINERvA detector was under construction the amount of material present varied throughout ArgoNeuT's running and is therefore challenging to simulate. Instead, only decays that occur in the 63 cm region between the end of the MINERvA detector and the start of the ArgoNeuT instrumented volume are considered. This region includes the ArgoNeuT cryostat and non-instrumented argon along with the small air gap between the two detectors. The presence of the MINERvA detector does not impact the propagation of the HNLs prior to their decay.

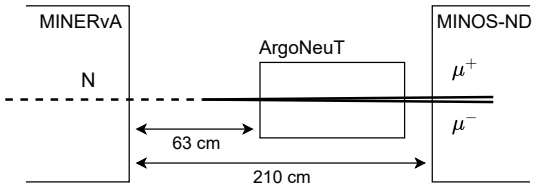

Figure 8.6: Diagram of a HNL decay occurring in the cavern between ArgoNeuT and the MINERvA detector with the double-MIP signature. The resulting muons pass through the ArgoNeuT detector and can be matched to the MINOS-ND. Diagram not to scale.

For decays upstream of ArgoNeuT only the double-MIP topology is considered, where the muons are sufficiently forward-going that they are still overlapping by the time they reach the ArgoNeuT detector. This choice is made because the double-MIP signature is unique and can be selected despite the absence of vertex information. Whereas in the two-track case the single muons may be difficult to distinguish from background neutrino-induced through-going muons originating from interactions occurring upstream of the detector in the dirt.

8.3 Pre-selection

A series of pre-selection cuts are first applied to remove poorly reconstructed events along with obvious non-HNL interactions. The pre-selection also identifies the *candidate tracks* in each event that will be considered in the subsequent selection.

8.3.1 Reconstruction quality

The first step of the selection chain is to remove events that have been poorly reconstructed since these either cannot be reliably selected in the case of signal events or cannot be reliably rejected in the case of backgrounds. This is achieved by assessing the fraction of reconstructed collection plane energy depositions (hits) in the events that are associated with reconstructed tracks, f_{hits}. This metric is shown in Figure 8.7 applied to the $\tau^{\pm}$ decay HNL sample (left) and to the beam simulation sample (right). Equivalent plots for the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples can be found in Appendix A. Note that events with zero reconstructed tracks in ArgoNeuT are not included in these plots. The majority of simulated HNL decays are well reconstructed, with more than 90% of hits associated with tracks. A cut is applied removing any events with $f_{hits} < 0.8$, illustrated by the dashed lines. This cut is also effective at removing background events that contain electromagnetic showers. The lack of a dedicated shower reconstruction being used in this analysis causes these events to have a large fraction of energy depositions that are not associated with reconstructed objects.

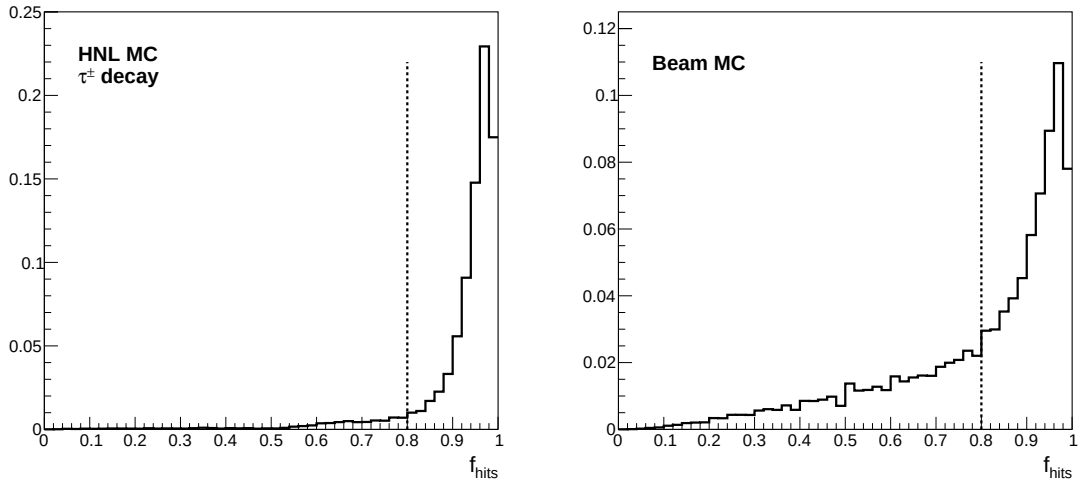

Figure 8.7: Fraction of hits associated with tracks, f_{hits}, for the $\tau^{\pm}$ decay HNL sample (left) and for the beam simulation sample (right). The dashed lines show the cut applied at $f_{hits} = 0.8$.

8.3.2 Fiducial volume

A minimal fiducial volume is used in ArgoNeuT, defined in Table 8.2. In this analysis, the fiducial volume is used when distinguishing whether a track originates from an interaction/decay occurring within ArgoNeuT or from the upstream cavern. It is also used to define whether a track exits the detector.

Coordinate	Lower bound [cm]	Upper bound [cm]
Drift direction, x	1	46
Vertical, y	-19	19
Beam direction, z	3	85

Table 8.2: Fiducial volume definition in this analysis.

8.3.3 Number and quality of reconstructed tracks

Track-based selection in ArgoNeuT

For each event the total number of reconstructed tracks in ArgoNeuT with $L \geq 5\,\text{cm}$ is counted. Shorter tracks are not assessed to avoid removing events that may have δ-rays originating from the muons. Events with zero or more than three reconstructed tracks are removed. A harsher requirement of a maximum of two tracks is not applied to compensate for a common reconstruction failure mode seen in simulated HNL decays. An additional vertex may be placed at the transition from a double-MIP-like track into two separate single-MIP tracks, resulting in incorrect splitting of the muon tracks and/or additional tracks being reconstructed. These events can often still be selected, however, provided that the double-MIP region around the true vertex is well reconstructed and for this reason are not removed outright.

Once events with too few or too many tracks are removed, candidate tracks that could have originated from HNL decays are identified. These tracks are then used in the subsequent selection. In ArgoNeuT, the candidate track requirements are:

- The track is forward-going with respect to the beam direction.
- The track has length $L \geq 5\,\text{cm}$.

Track-based selection in MINOS-ND

The number of tracks in the MINOS-ND that could have originated from a HNL decay in ArgoNeuT are counted. Events with fewer than two such tracks are removed. In the MINOS-ND, the candidate tracks requirements are:

- The track is forward-going with respect to the beam direction.

- The track starts within the fully-instrumented calorimeter region of the MINOS-ND. Tracks originating from ArgoNeuT start within this region unless they are very off-axis with respect to the beam direction, which is improbable for the very forward-going muons resulting from HNL decays.

- The track starts within 20 cm of the upstream (towards ArgoNeuT) end of the MINOS-ND. Tracks starting further inside the detector may have originated from, for example, neutrino interactions within the detector rather than propagating from ArgoNeuT. A harsher cut is not applied because the initial hits of tracks on the first few planes in the MINOS-ND are sometimes not reconstructed resulting in the start of the track being missed. Additionally, at the start of the detector, the pair of muons may still be overlapping having not yet been separated by the magnetic field.

- The track is at least 1 m long. Very short tracks are typically either originating from δ-rays, very low energy muons, or may be protons or pions originating from neutrino interactions in ArgoNeuT.

8.3.4 Vertex and opening angle requirements

Events with more than two tracks originating from a common vertex are removed. These events have additional particles present and hence are not consistent with candidate HNL decays. Two tracks are considered to have originated from a common vertex if their start positions are within $r \leq 4\,\text{cm}$, where r is the separation in three-dimensions.

For events where two tracks with a common vertex are identified, the opening angle between the tracks is assessed. The highly-boosted muons from HNL decays would have a small opening angle between them. Figure 8.8 (left) shows the reconstructed opening angle between pairs of muons from the $\tau^{\pm}$ decay HNL sample. Equivalent plots for the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples can be found in Appendix A. The pair of tracks are required to have an opening angle between them of $\theta_{opening} \leq 10^{\circ}$, illustrated by the dashed line. Figure 8.8 (right) shows this cut applied to a sample of charged current ν_{μ} interactions producing single pions from the beam simulation sample. These could form a significant background, producing two approximately minimally ionising tracks that could be mistaken for the two-track signature. However, the majority of these interactions have much larger opening angles between the resulting tracks and hence are removed.

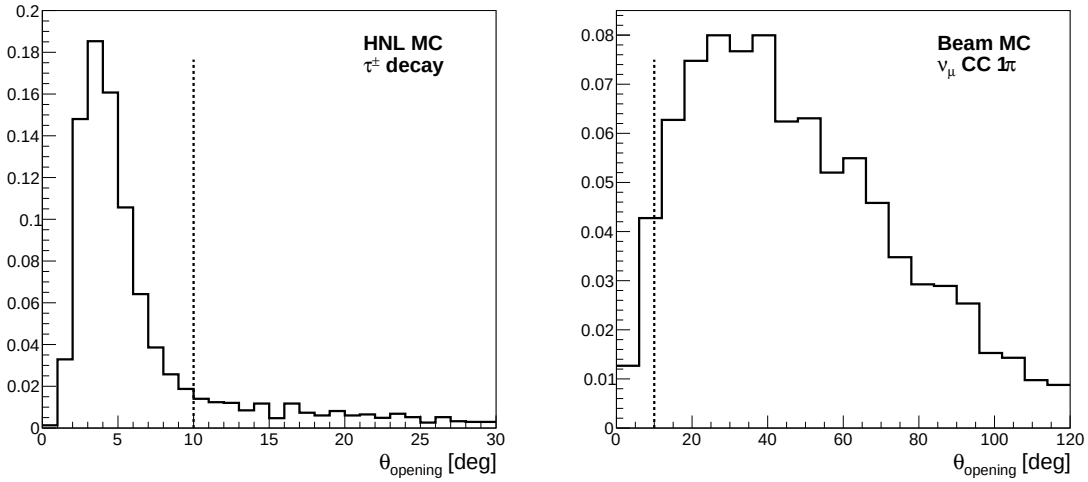

Figure 8.8: Reconstructed opening angle, $\theta_{opening}$, between muons from the $\tau^{\pm}$ decay HNL sample (left) and for charged current ν_{μ} interactions producing a single $\pi^{\pm}$ from the beam simulation sample (right, note different x-scale). The dashed lines show the cut applied at $\theta_{opening} = 10^{\circ}$.

8.4 Topological and calorimetric selection in ArgoNeuT

Events that pass the pre-selection are then assessed against the two-track and double-MIP selection criteria in ArgoNeuT. In cases where the event could be selected with either signature, it is preferentially selected as being of the two-track signature.

8.4.1 dE/dx

The strongest identifier of whether a pair of overlapping muons, or only a single muon, is present is provided by the average energy deposited per unit length, dE/dx, along the track. To identify pairs of overlapping muons, the average dE/dx is calculated over the first 10 collection plane hits ($\sim 5\,\text{cm}$) of each candidate track. Only the start of each track is used since the overlapping muons may begin to split further along. Any individual hits with $dE/dx > 10\,\text{MeV/cm}$ are discarded. This removes hits with anomalously high dE/dx, e.g. due to the presence of δ-rays, that can significantly skew the prediction. If hits are removed, additional hits are added along the track such that the mean dE/dx prediction is always made with 10 hits.

Figure 8.9 (left) shows the resulting average dE/dx of tracks from the $\tau^{\pm}$ decay HNL sample. Equivalent plots for the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples can be found in Appendix A. Two distinct peaks are visible. The first is at $dE/dx \sim 1.9\,\text{MeV/cm}$, approximately the dE/dx of a single MIP. These are events where the opening angle

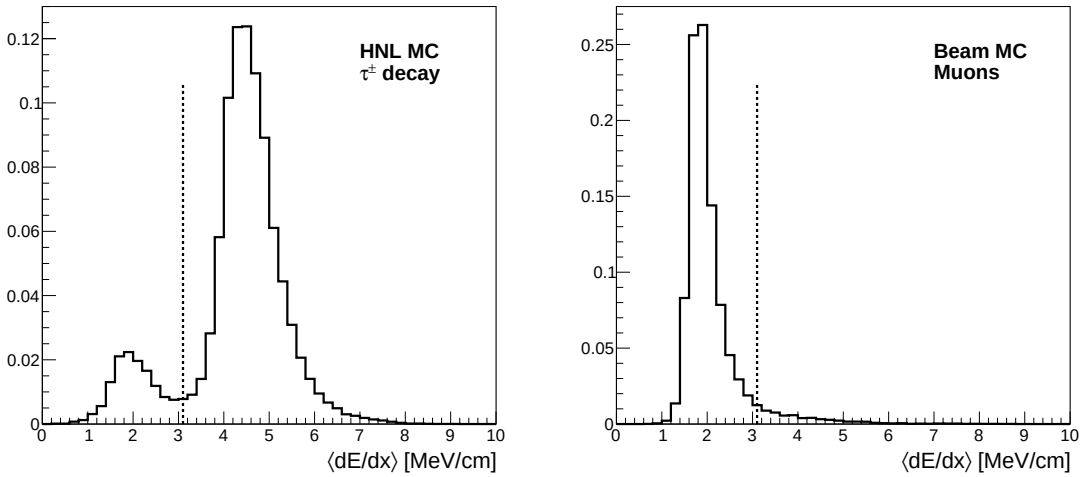

Figure 8.9: Average dE/dx for tracks from the $\tau^{\pm}$ decay HNL sample (left) and for through-going neutrino-induced background muons (right). The dashed lines represent the threshold applied at $dE/dx = 3.1\,\mathrm{MeV/cm}$.

between the muons is sufficiently large that it is possible to properly reconstruct them as two separate tracks. The second peak is at $dE/dx \sim 4.4\,\mathrm{MeV/cm}$, approximately double the single MIP dE/dx, indicating that there are two overlapping muons present[2]. To distinguish between single and double MIP tracks, a threshold is applied between these peaks at $dE/dx = 3.1\,\mathrm{MeV/cm}$, illustrated by the dashed line. Any track above this threshold is considered to be double-MIP-like and any track below is considered to be MIP-like. No upper threshold is applied to the dE/dx, therefore other forward-going highly-ionising particles, such as protons, can also be selected. However, these are unlikely to propagate to the MINOS-ND before stopping and hence would not pass the later stages of the selection.

Figure 8.9 (right) shows the same metric applied to single through-going neutrino-induced background muons from the beam simulation sample. The threshold applied at $dE/dx = 3.1\,\mathrm{MeV/cm}$, illustrated by the dashed line, is able to effectively distinguish these background single muons from pairs of overlapping muons. Approximately 94% of the background muons are rejected by this threshold.

8.4.2 Two-track signature

Each event is first assessed against the two-track topology selection criteria. Events are identified that have at least two candidate tracks, defined in Section 8.3.3, start-

[2]The second peak is at a slightly larger dE/dx value than double the first peak due to the way that the events are simulated versus detected and reconstructed. In Geant4 two separate muon tracks are simulated and recombination is accounted for treating them independently. However, in these highly-forward-going cases, they are then detected and reconstructed as a single track. Therefore, recombination is accounted for as a single energy deposition resulting in an over-correction due to its non-linear dependence as a function of dQ/dx, as discussed in Section 3.3.1.

ing within the fiducial volume described in Section 8.3.2. Tracks starting outside of the fiducial volume are removed. This allows the identification and removal of through-going muons from interactions upstream of ArgoNeuT that form an otherwise difficult to distinguish background. The two-track selection is therefore only sensitive to HNL decays occurring within the ArgoNeuT detector.

Each pair of tracks is assessed to see whether they originate from a common vertex. Two tracks are considered to have originated from a common vertex if the start positions are within $r \leq 4\,\text{cm}$, where r is the separation in three-dimensions. Since the muons from HNL decays are typically highly forward-going, close to the vertex only a single track may be reconstructed that then later splits into two distinct tracks. To identify these cases, the track starting furthest along the beam direction (at larger z) is projected back to the z position of the other track. They are then assessed to see whether they would intersect within the 4 cm tolerance. This allows intermediate cases where the muons are only partially overlapping to be selected with the two-track signature. However, in many cases the second track is poorly reconstructed in this region and instead the event can only be selected with the double-MIP signature.

Both tracks are required to exit towards the MINOS-ND. A track is defined as exiting if its endpoint is outside of the fiducial volume defined in Section 8.3.2. Additionally, the tracks are required to have an opening angle between them of $\theta_{opening} \leq 10^{\circ}$, as discussed in Section 8.3.4. Finally, for this signature the muons from the HNL decay are clearly separated and hence the average dE/dx of each track should be MIP-like. Both tracks are therefore required to have average $dE/dx < 3.1$ MeV/cm, as described in Section 8.4.1. Unlike in the double-MIP case, however, the average dE/dx is calculated over the full length of each track since the muons have already separated.

8.4.3 Double-MIP signature

Events that fail the two-track selection are then assessed to see if they pass the double-MIP selection. In this case, each candidate track is treated separately to see if it passes the threshold to be considered a double-MIP. Tracks starting outside of the fiducial volume are not removed since, for this topology, background through-going muons can be effectively identified and rejected based on their single-MIP-like dE/dx, as discussed in Section 8.4.1. This enables the double-MIP selection to be sensitive to highly-boosted HNL decays occurring in the cavern upstream of ArgoNeuT in addition to decays occurring within the detector volume.

Events that are reconstructed as double-MIPs are highly forward-going and therefore the resulting tracks have a small angle with respect to the beam direc-

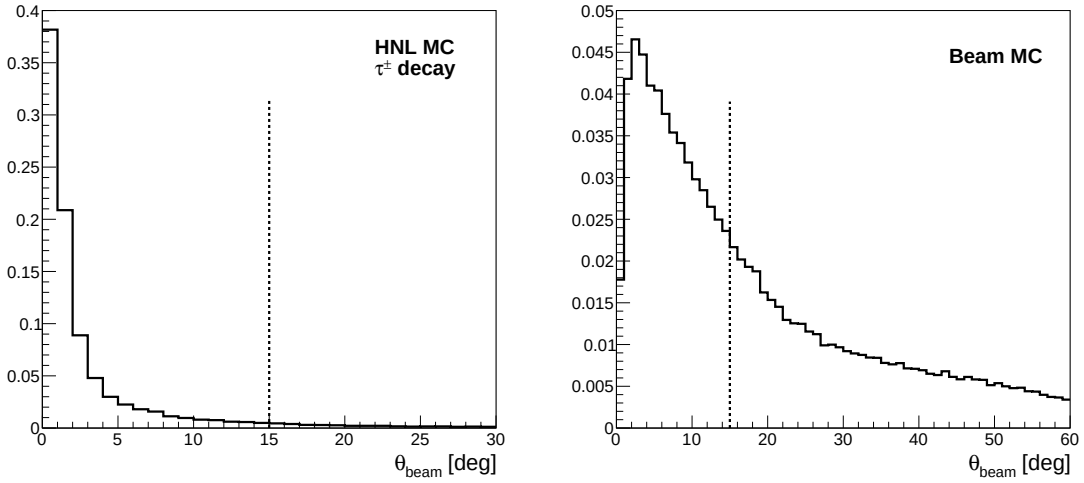

Figure 8.10: Reconstructed angle with respect to the beam direction, θ_{beam}, for double-MIP-like tracks from the $\tau^{\pm}$ decay HNL sample (left) and for background tracks from the beam simulation sample (right, note different x-scale). The dashed lines show the cut applied at $\theta_{beam} = 15^{\circ}$.

tion, θ_{beam}. The distribution of θ_{beam} is shown in Figure 8.10 for double-MIP-like tracks from the $\tau^{\pm}$ decay HNL sample (left) and for background tracks from the beam simulation sample (right). Equivalent plots for the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples can be found in Appendix A. The tracks are required to have an angle with respect to the beam direction of $\theta_{beam} \leq 15^{\circ}$, illustrated by the dashed lines. This cut is relatively loose, but helps to reject backgrounds that can have much larger angles.

Next, the average dE/dx at the start of each of the candidate tracks is calculated as described in Section 8.4.1. The candidate track is required to have an average $dE/dx > 3.1$ MeV/cm, consistent with a pair of overlapping muons.

8.5 Topological and calorimetric selection in the MINOS-ND

Once candidate events are identified with either the two-track or double-MIP signature in ArgoNeuT, MINOS-ND matching is performed followed by several further selection cuts in the MINOS-ND on the matched tracks. These cuts are the same for both the two-track and double-MIP scenarios.

8.5.1 MINOS-ND matching

As described in Section 5.5.2, each ArgoNeuT track is projected to the start of the MINOS-ND and the radial, r_{diff}, and angular, θ_{diff}, off-sets between the projected

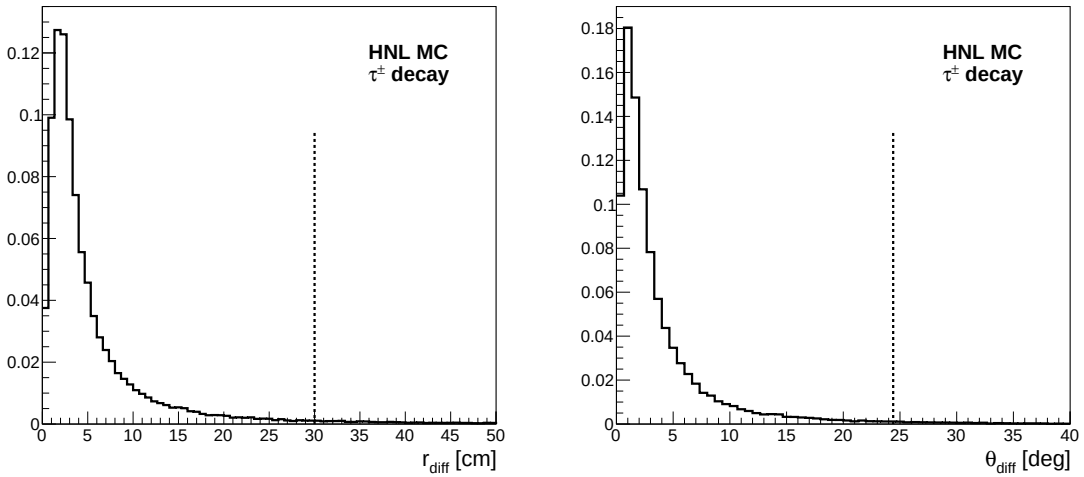

Figure 8.11: Double-MIP MINOS-ND matching r_{diff} (left) and θ_{diff} (right) for events with a double-MIP-like topology from the $\tau^{\pm}$ decay HNL sample. The dashed lines show the double-MIP matching tolerances of $r_{diff} = 30\,\text{cm}$ and $\theta_{diff} = 24.4^{\circ}$.

tracks and each reconstructed MINOS-ND track are compared. Only candidate tracks in the MINOS-ND, defined in Section 8.3.3, are considered.

In the two-track case the standard ArgoNeuT-MINOS-ND radial and angular matching tolerances are used: $r_{diff} = 12.0\,\text{cm}$ and $\theta_{diff} = 9.74^{\circ}$ [235]. Events are selected that have two MIP-like tracks passing the ArgoNeuT two-track selection that match with unique MINOS-ND tracks. In the case of more than two tracks matching within tolerance, for example if another nearby neutrino-induced background muon is present, the best matching tracks are considered subsequently.

In the double-MIP case, a single track in ArgoNeuT is matched to a pair of tracks in the MINOS-ND. The reconstructed trajectory of this track in ArgoNeuT is often less accurate: it may be reconstructed along the trajectory of one of the two truth tracks, or between them. Therefore, looser matching tolerances are used compared with the two-track scenario. For the double-MIP matching, tolerances of 2.5 times the standard matching are used: $r_{diff} = 30\,\text{cm}$ and $\theta_{diff} = 24.4^{\circ}$. Figure 8.11 shows the double-MIP MINOS-ND matching r_{diff} (left) and θ_{diff} (right) for events with a double-MIP-like topology from the $\tau^{\pm}$ decay HNL sample. Equivalent plots for the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples can be found in Appendix A. The dashed lines show the matching tolerances applied. It is feasible to keep these tolerances relatively loose and inclusive since most mismatched tracks can be removed by subsequent cuts. Events are selected that have a track passing the double-MIP selection that matches with two MINOS-ND tracks. As before, if more than two tracks match within tolerance the best matching tracks are considered.

8.5.2 Charge

The magnetic field in the MINOS-ND allows the charge of the muons to be reconstructed. The matched tracks are required to have opposite charges.

8.5.3 Timing

ArgoNeuT lacks timing information more sensitive than the NuMI beam spill window of 8 μs. This is due to the absence of a photon-detection system, unlike in more modern LArTPC detectors. As a result of this, it is difficult to identify whether a pair of muons in ArgoNeuT have originated from a HNL decay or are background muons that passed through the detector at different times in the readout window. The MINOS-ND, however, saves a timestamp of the start of each track, t_0, with a resolution of ± 5 ns [237]. Figure 8.12 shows the start time difference, Δt_0, for random pairs of tracks in the MINOS-ND data spanning the beam spill window. A cut is applied removing any pair of matched tracks with $|\Delta t_0| \geq 20$ ns. This slightly looser threshold is chosen to allow for the potential few nanosecond difference in arrival time at the MINOS-ND resulting from the different paths that the muons could take between ArgoNeuT and the MINOS-ND. The timing cut is highly effective at removing background muons originating from the cavern, removing more than 99% of random pairs of background tracks in the MINOS-ND.

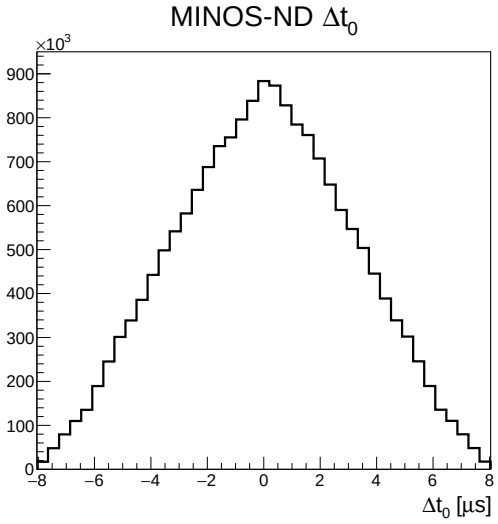

Figure 8.12: Time difference between pairs of tracks in the MINOS-ND data. Tracks with $|\Delta t_0| \geq 20$ ns are cut, removing more than 99% of background tracks.

8.5.4 Track length and dE/dx

The majority of background pions (∼93%) and protons (∼99%) produced in the interactions of neutrinos in ArgoNeuT stop before reaching the MINOS-ND. There-

fore, the MINOS-ND matching requirement is a powerful particle identification tool with the majority of matched tracks being muons. The small subset of background pions and protons that reach the MINOS-ND can be distinguished from muons based on the length of the tracks and their dE/dx.

Figure 8.13 shows the length of tracks in the MINOS-ND, $L_{MINOS-ND}$, for muons from the $\tau^{\pm}$ decay HNL sample (left) and for pions and protons from the beam simulation sample (right). Equivalent plots for the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples can be found in Appendix A. The large peak at $L_{MINOS-ND} \sim 16.6\,\text{m}$ is from tracks that exit the end of the MINOS-ND. The smaller peak at $L_{MINOS-ND} \sim 7\,\text{m}$ corresponds to the end of the calorimeter region of the detector and is a binning effect: beyond this point, only every 5th plane is instrumented resulting in larger jumps in the track lengths. Pions and protons that reach the MINOS-ND typically form much shorter tracks than the muons originating from HNL decays. Any tracks with $L_{MINOS-ND} < 1\,\text{m}$ are rejected in the matching process, illustrated by the dashed lines. A harsher cut was found to be unnecessary due to the very small number of background pions and protons passing other parts of the selection.

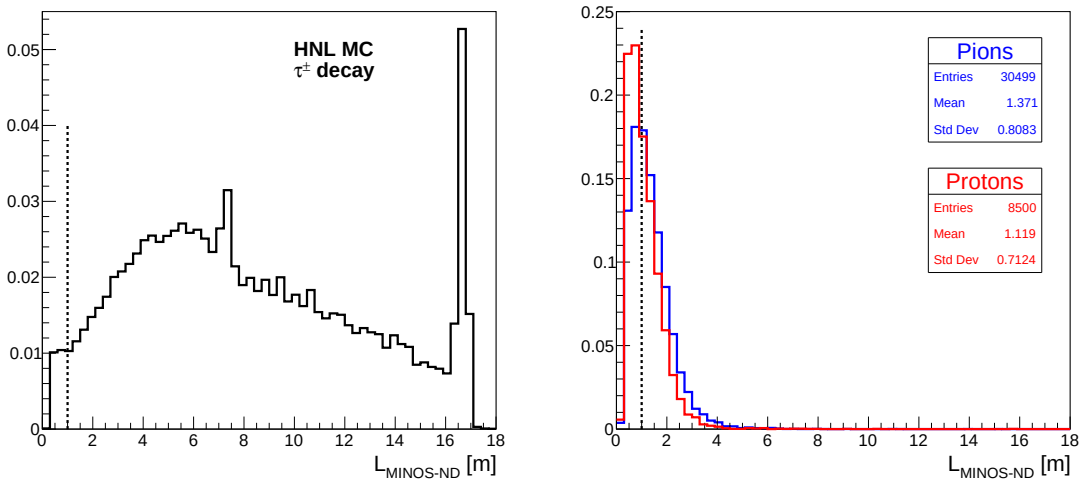

Figure 8.13: MINOS-ND track length, $L_{MINOS-ND}$, for muons from the $\tau^{\pm}$ decay HNL sample (left) and for pions and protons from the beam simulation sample (right). The dashed lines illustrate the cut applied at $L_{MINOS-ND} = 1\,\text{m}$.

Tracks originating from ArgoNeuT start within the fully-instrumented calorimeter region of the MINOS-ND. The dE/dx of these tracks can therefore also be used to separate the more highly ionising stopping protons and pions from muons. Figure 8.14 shows the average dE/dx of muons, protons and pions from the beam simulation sample that are contained within the calorimeter region of the MINOS-ND. The muon tracks have $dE/dx \sim 10\,\text{MeV/cm}$, approximately corresponding to a minimally ionising particle in steel once the steel plate spacing and thickness is considered. The protons and pions tracks have much broader dE/dx distributions with

higher mean values. A cut on the MINOS-ND dE/dx is applied on tracks that are contained within the calorimeter region of the detector, $4 \leq dE/dx \leq 18\,\mathrm{MeV/cm}$, removing any tracks that are inconsistent with being a muon. This cut also removes poorly reconstructed tracks in the MINOS-ND, which commonly have anomalously low dE/dx. Tracks that exit the calorimeter region have to propagate several metres and are hence highly unlikely to be pions or protons since these typically stop rapidly, as seen in Figure 8.13. Therefore, no dE/dx requirement is applied to these tracks.

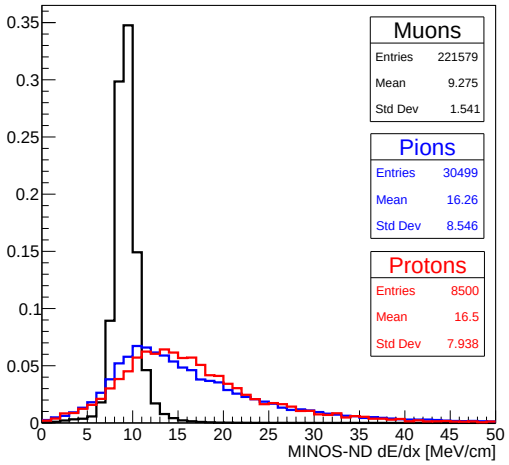

Figure 8.14: MINOS-ND dE/dx for muons, protons and pions from the beam simulation sample contained within the calorimeter region of the detector.

8.5.5 Reconstruction failure correction

A failure was identified that affects a subset of MINOS-ND reconstructed tracks. The ordering of the hits saved for these tracks is jumbled, such that there are apparent non-physical jumps in the track position within the detector. This affects the ArgoNeuT-MINOS-ND matching since the first stored hit may not actually be the start of the track. In addition, the momentum and charge of the tracks are not correctly reconstructed. Figure 8.15 shows two event displays of a simulated interaction with a MINOS-ND track impacted by this failure. In the left event display a line is drawn that connects each of the hits in the order that they are saved in the hit vector. Since these hits are jumbled, this line jumps erratically back and forth. The right event display shows the same event, but without the line drawn. It can be seen that the reconstructed hits form a normal looking track, suggesting that the issue is simply that they are stored out of order. Events affected by this issue were found in both simulation and data.

The affected tracks were found to be those that pass either close to or though

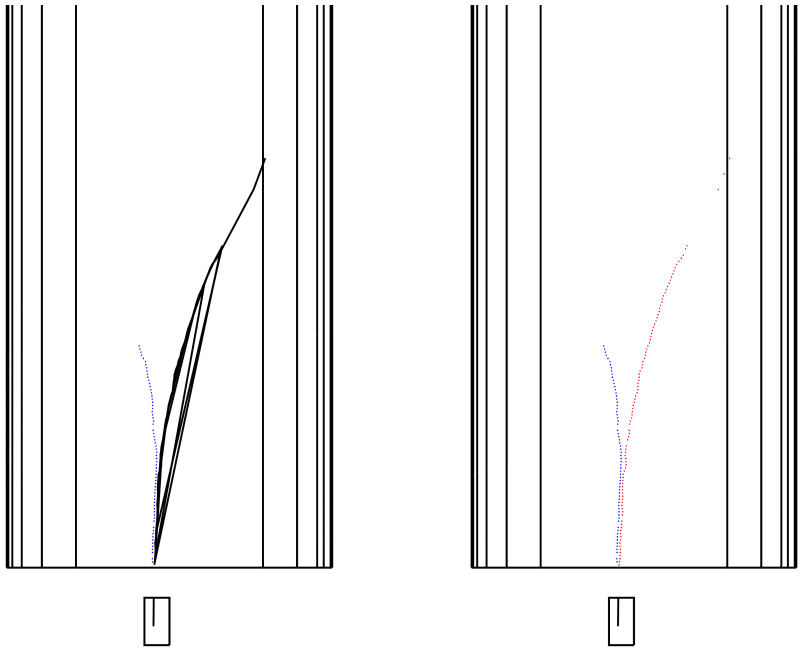

Figure 8.15: Simulated MINOS-ND event containing a track affected by the reconstruction failure. The two plots show the same event. The left plot has a line drawn connecting each of the hits of the mis-reconstructed track in the order that they are saved in. Since the hits are jumbled, the resulting line jumps back and forth repeatedly. However, the hits themselves appear to form a normal looking track as can be seen in the right plot without the line drawn.

the MINOS-ND magnetic coil. The issue results from a bug in the version of the MINOS-ND reconstruction that was made available for ArgoNeuT to use (Dogwood-5) and hence could not be fixed by ArgoNeuT. However, the affected events could instead be recovered during the selection. The hit ordering is jumbled, but the hits themselves are reconstructed correctly and form a sensible track. They could therefore be sorted, based on the hit coordinates in the beam-direction (z). This allowed the ArgoNeuT-MINOS-ND matching to then be performed as normal. The standard charge and momentum reconstruction is not recoverable since this requires use of the MINOS-ND magnetic field mapping. Instead, for forward-going muons starting within the calorimeter region of the MINOS-ND (all muons resulting from HNL decays) only positively charged particles curve towards the coil whereas negatively charged particles curve away. The charge can be therefore be effectively predicted based on geometry for this type of event. In the simulated HNL samples, approximately 10% of events were found to have a track with this issue, always affecting the μ^+. Using this recovery approach, the position and charge was able to be correctly identified for 100% of these tracks comparing with truth information. This reconstruction failure has an especially large impact in this analysis due to the topology of the events. It would have been present in previous ArgoNeuT analyses,

however the impact was likely minimal.

8.6 Summary of the selection

A summary of the cuts applied and their impact on the selection efficiency is presented in Tables 8.3, 8.4 and 8.5. The selection is applied to the three samples described in Section 8.1 corresponding to HNL production from $K^{\pm}$, $D^{\pm}_{(s)}$ and $\tau^{\pm}$ decays. In each case the HNLs are produced in the beam target and the decays occur within the ArgoNeuT detector[3]. The selection is also applied to the beam simulation sample modelling the expected backgrounds. These will be discussed further in Section 8.8.

Table 8.3 shows the pre-selection cuts in the order that they are applied. Following the pre-selection, 43.8%, 51.7% and 67.0% of the $K^{\pm}$, $D^{\pm}_{(s)}$ and $\tau^{\pm}$ decay HNL samples remain, respectively. The lower efficiency in the $K^{\pm}$ case is as a result of the much lower average E_N of these events, as shown in Figure 8.1, resulting in a larger fraction of events with low energy muons that do not reach the MINOS-ND and hence cannot be selected. The pre-selection cuts remove 85.9% of the ArgoNeuT beam simulation events.

	HNL MC $K^{\pm}$ decay	HNL MC $D^{\pm}_{(s)}$ decay	HNL MC $\tau^{\pm}$ decay	Beam MC
Initial events	24750	25000	50000	600000
Zero reconstructed activity	24111	24387	48457	322024
Hit reconstruction fraction	22241	22662	44696	172951
Number tracks ArgoNeuT	20774	20949	41348	136475
Number tracks MINOS-ND	11519	15717	35494	96251
Number tracks common vertex	11513	15689	35454	94743
Opening angle at vertex	10852	12933	33502	84492

Table 8.3: Pre-selection cut flow table.

Table 8.4 shows the two-track selection cuts in the order that they are applied. The two-track selection is applied to all of the events that pass the pre-selection. In total, 5.6%, 13.9% and 6.1% of the $K^{\pm}$, $D^{\pm}_{(s)}$ and $\tau^{\pm}$ decay HNL samples are selected with the two-track signature, respectively. From the beam simulation, 1 event passes the selection with this signature.

Table 8.5 shows the double-MIP selection cuts in the order that they are applied. The double-MIP selection is applied to all of the events that pass the pre-selection,

[3]As noted in Section 8.2.1, these samples assume that the HNLs are long-lived relative to the target-ArgoNeuT distance. This assumption does not hold true for all of the HNL phase-space of interest in ArgoNeuT, and for this reason they cannot be used for the final sensitivity evaluation. However, they are illustrative of the approximate selection efficiency for each production mechanism for the majority of the HNL phase-space.

	HNL MC $K^{\pm}$ decay	HNL MC $D^{\pm}_{(s)}$ decay	HNL MC $\tau^{\pm}$ decay	Beam MC
Pass pre-selection	10852	12933	33502	84492
Two tracks, start in fiducial volume and exit	3668	7015	9552	4404
ArgoNeuT MIP dE/dx	2212	5514	4981	2536
MINOS-ND matching	1703	4760	3881	90
MINOS-ND MIP dE/dx	1703	4760	3880	89
ArgoNeuT vertex (MINOS matched tracks)	1384	3483	3031	1
MINOS-ND charge	1384	3483	3031	1
MINOS-ND timing	1384	3483	3031	1

Table 8.4: Two-track selection cut flow table.

	HNL MC $K^{\pm}$ decay	HNL MC $D^{\pm}_{(s)}$ decay	HNL MC $\tau^{\pm}$ decay	Beam MC
Pass pre-selection and fail two-track selection	9468	9450	30471	84491
Angle with respect to beam	8498	8734	27434	52636
ArgoNeuT double-MIP dE/dx	8149	7952	26543	4871
MINOS-ND matching	8108	7494	26440	695
MINOS-ND MIP dE/dx	8107	7492	26435	678
MINOS-ND charge	7929	7400	26227	353
MINOS-ND timing	7929	7400	26227	2

Table 8.5: Double-MIP selection cut flow table.

but do not pass the two-track selection. In total, 32.0%, 29.6% and 52.5% of the $K^{\pm}$, $D^{\pm}_{(s)}$ and $\tau^{\pm}$ decay HNL samples are selected with the double-MIP signature, respectively. From the beam simulation, 2 events pass the selection with this signature.

The total selection efficiencies are summarised in Table 8.6, in each case showing the contributions from the two-track and double-MIP signatures. In all three cases, the majority of events are selected with the double-MIP signature. This is especially true for the $\tau^{\pm}$ decay produced HNLs that are the most forward-going. The total selection efficiency is 37.6% for HNLs produced from $K^{\pm}$ decays, 43.5% for HNLs produced from $D^{\pm}_{(s)}$ decays and 58.5% for HNLs produced from $\tau^{\pm}$ decays. The selection is also highly effective at removing backgrounds: from the beam simulation sample, 1 event passes with the two-track signature and 2 events pass with the double-MIP signature.

	HNL MC $K^{\pm}$ decay	HNL MC $D^{\pm}_{(s)}$ decay	HNL MC $\tau^{\pm}$ decay
Two-Track	5.6%	13.9%	6.1%
Double-MIP	32.0%	29.6%	52.5%
Total	37.6%	43.5%	58.6%

Table 8.6: Summary of the selection efficiencies for HNLs produced from decays of $K^{\pm}$, $D^{\pm}_{(s)}$ and $\tau^{\pm}$. In each case, the contributions from the two-track and double-MIP signatures are shown.

8.7 Selection efficiency

8.7.1 Decays inside of ArgoNeuT

To evaluate the selection efficiency for HNL decays occurring inside the ArgoNeuT detector, decays distributed throughout the instrumented volume were simulated at fixed m_N and E_N covering the phase-space of interest, as described in Section 8.1. This allows the selection efficiencies to be applied to each scenario (electron-coupled, muon-coupled or tau-coupled), weighting by the respective energy distributions. The efficiency is defined as the fraction of the events that pass the full selection.

Figure 8.16 shows the selection efficiencies as a function of E_N for m_N = 300, 400, 500 and 600 MeV. The total efficiency is around $60 - 65\%$ and relatively flat above $E_N \sim 10\,\text{GeV}$. However, it drops significantly at lower energies. In each case the contributions from two-track and double-MIP events are shown. At low energies there are approximately equal numbers of two-track and double-MIP type events, whereas at higher energies the double-MIP type events dominate as the muons are increasingly forward-going. The fraction of events passing with the two-track signature also increases for higher m_N. In these cases, the HNLs are less boosted at lower energies due to their larger mass ($\gamma_N = E_N/m_N$) resulting in slightly larger opening angles between the muons.

Losses at lower energies

As seen in Figure 8.16, at low HNL energies the selection efficiency sharply decreases. The dominant cause of this is one or both muons being too low energy to propagate to the MINOS-ND and create a long enough track to have an accurately reconstructed charge. Figure 8.17 (left) shows the fraction of simulated events where both muon tracks are reconstructed in the MINOS-ND as a function of the HNL energy. At higher energies, above $E_N \sim 10\,\text{GeV}$, nearly 100% of events have both tracks reconstructed in MINOS-ND. Below $E_N \sim 10\,\text{GeV}$, this starts to significantly decrease and at very low energies, $E_N < 4\,\text{GeV}$, the majority of events are not reconstructed. Since a significant fraction of the HNLs have low energy, as seen in

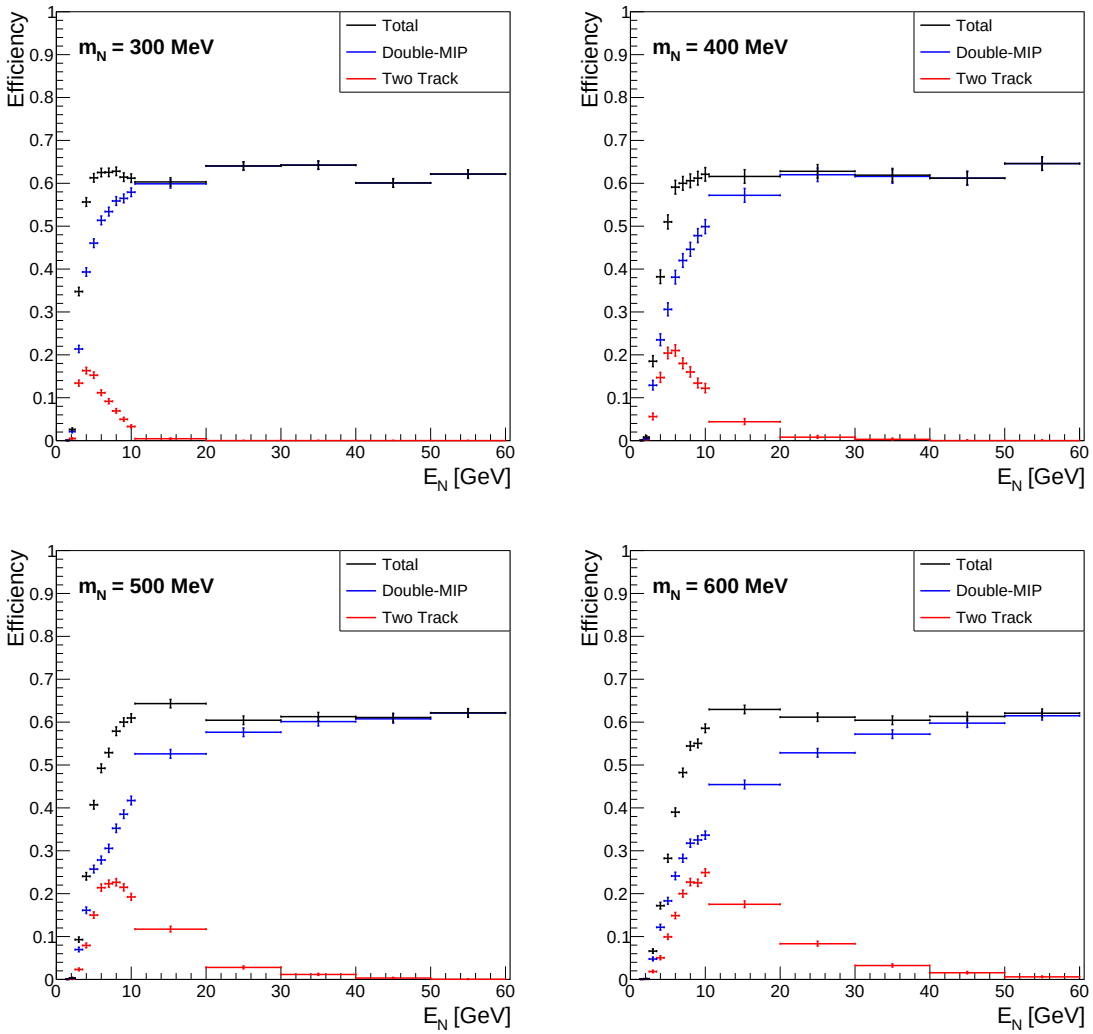

Figure 8.16: Selection efficiencies for HNL decays occurring inside the ArgoNeuT detector as a function of E_N for m_N = 300, 400, 500 and 600 MeV. In each case the contributions from two-track and double-MIP events are shown. Note that for m_N = 400 MeV the statistical uncertainties shown are larger since there are only 1000 events per mass-energy combination rather than 2500.

Figure 8.1, this significantly impacts the overall selection efficiency. This has an especially large impact on $K^{\pm}$ decay-produced HNLs that have the lowest average energy, where $\sim 30\%$ of events have one or both tracks stopping before reaching the MINOS-ND.

Furthermore, for lower energy muons that reach the MINOS-ND, the resulting tracks may be too short to effectively determine their charge. Figure 8.17 (right) shows the charge reconstruction efficiency for μ^- and μ^+ resulting from HNL decays as a function of the track length in the MINOS-ND, $L_{MINOS-ND}$. MINOS-ND tracks affected by the reconstruction issue are included, with charges reconstructed as described in Section 8.5.5. For tracks longer than $L_{MINOS-ND} \sim 2$ m, the charge reconstruction is near 100% efficient. However, for shorter tracks this rapidly de-

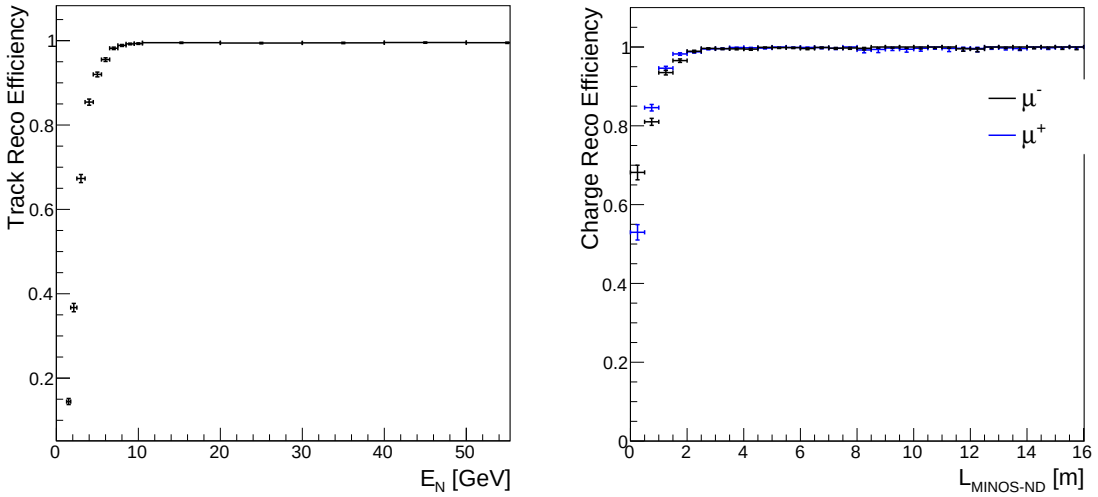

Figure 8.17: Left: fraction of events where both muon tracks are reconstructed in the MINOS-ND as a function of the HNL energy, E_N. For lower energy HNLs, one or both of the muons may stop before reaching the MINOS-ND. Right: charge reconstruction efficiency for tracks that reach the MINOS-ND as a function of the track length, $L_{MINOS-ND}$. The charge reconstruction performs worse for tracks shorter than $\sim 2\,\text{m}$, originating from lower energy HNLs.

creases. It should be noted that since particles are always reconstructed as either positive or negative an efficiency of near 50%, as is seen for very short tracks, is close to random chance. The average dE/dx of minimally ionising muon tracks in the MINOS-ND is $dE/dx \sim 10\,\text{MeV/cm}$. Therefore, to reconstruct the muon charges with high efficiency both muons must reach the MINOS-ND with $E_{\mu^\pm} \sim 2\,\text{GeV}$.

In the full selection, the efficiency is further decreased at lower energies due to a combination of reconstruction failures in either ArgoNeuT or MINOS-ND along with additional selection constraints such as the requirement for the MINOS-ND tracks to be at least 1 m long. This most strongly impacts HNLs produced from $K^\pm$ decays where approximately 45% of events are rejected during the pre-selection as a result of too few tracks being reconstructed in the MINOS-ND that pass the requirements, as can be seen in Table 8.3. These efficiency losses at low energies are a fundamental limitation of matching between the two detectors. To recover these events would require selecting the lower energy HNL decays using ArgoNeuT only. This, however, would be very challenging due to the small size of the ArgoNeuT detector and the difficulty distinguishing between muons and backgrounds such as pions in the absence of the MINOS-ND.

Losses at higher energies

At medium to high energy, the selection efficiency flattens off at around $\sim 60-65\%$. The dominant loss of efficiency in these energy ranges is events being insufficiently

well reconstructed in ArgoNeuT to be identified as HNL candidates. The highly forward-going muons from HNL decays can be challenging to reconstruct correctly in LArTPC detectors. This is because the ionisation tracks are close to parallel to the readout wire planes and hence the drifted ionisation charge arrives on the wires at approximately the same time. This can introduce ambiguities in the matching between wire planes during three-dimensional reconstruction, resulting in inaccurate track trajectories. This is partially compensated for in the selection by allowing looser matching tolerances between ArgoNeuT and the MINOS-ND, in particular for the double-MIP signature. However, in many cases, the reconstruction failures are not recoverable and prevent matching between the two detectors leading to the events being rejected.

A second common reconstruction failure occurs in cases where the muon-pair may begin to separate part way along the track while still within the ArgoNeuT detector. The region around where the two tracks split is often poorly reconstructed resulting in incorrect track trajectories, multiple split tracks being reconstructed with a vertex placed near the separation, or even the separating tracks failing to be reconstructed at all. This failure mode is partially compensated for in the selection in two ways: by allowing up to three tracks to be present to avoid rejecting events with split tracks; and by focusing on the start of the tracks in the double-MIP selection and not applying any constraints on whether the tracks begin to separate.

It is likely that improvements could be made to the efficiency at higher energies by addressing these reconstruction failures. In particular, by designing a reconstruction algorithm that specifically checks for this topology – a single track from two particles that then splits into a pair of tracks. This optimisation was not attempted for this analysis, however, as it would be challenging to do due to the age of ArgoNeuT's code-base. If future searches for this signature are performed in other LArTPC detectors, such as the SBN program detectors or the DUNE near detector, this could be revisited. In larger LArTPCs, it would be more likely for the initially overlapping muons to separate before exiting the detector and therefore correctly reconstructing this region could play an important role in identifying between this signature and background.

8.7.2 Decays upstream of ArgoNeuT

To evaluate the selection efficiency for HNL decays occurring upstream of the ArgoNeuT detector in the cavern, decays were simulated at fixed offset positions in the beam direction, z, as described in Section 8.1. Events are then identified where both of the resulting muons would intersect with the start of the ArgoNeuT instrumented volume at truth level. The selection efficiency is then defined as the fraction of these

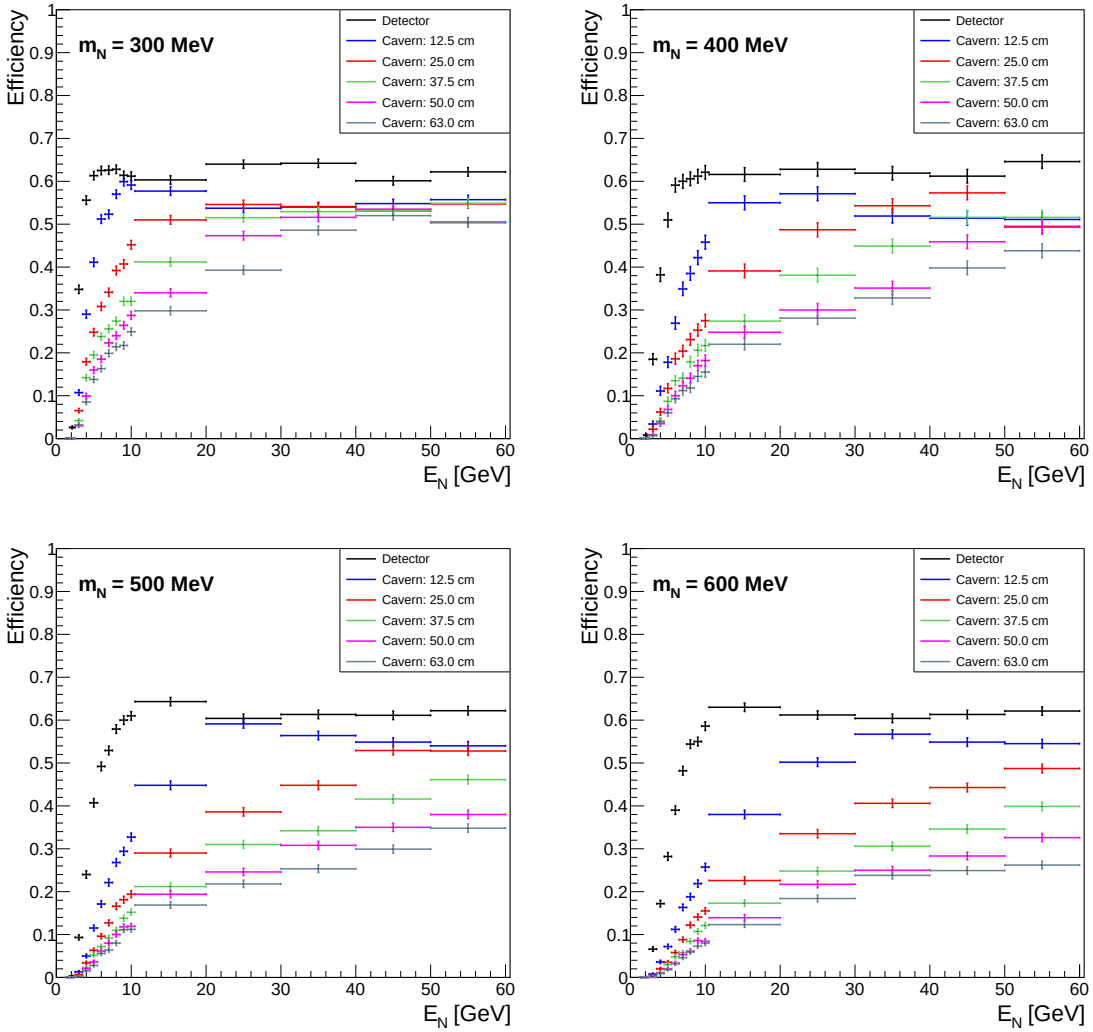

Figure 8.18: Selection efficiencies for HNL decays upstream of the ArgoNeuT detector in the cavern as a function of E_N for m_N = 300, 400, 500 and 600 MeV. In each case the efficiencies are shown for decays inside the detector and decays 12.5, 25, 37.5, 50 and 63 cm upstream of the detector. Note that for m_N = 400 MeV the statistical uncertainties shown are larger since there are only 1000 events per mass-energy combination rather than 2500.

events that pass the full selection. As described in Section 8.2, only events that pass the double-MIP selection are considered. Figure 8.18 shows the selection efficiencies as a function of E_N for m_N = 300, 400, 500 and 600 MeV. In each case decays occurring at five different offset positions are shown: z = -12.5, -25, -37.5, -50 and -63 cm, where z = 0 cm is the start of the ArgoNeuT instrumented volume. The efficiencies for decays inside the detector are also shown for comparison, these are the same as the total efficiencies shown in Figure 8.16. The efficiencies for the decays upstream of ArgoNeuT are lower than for decays within the detector primarily because only double-MIP type events are selected. This means that only the most forward-going of the muon pairs can be selected. This has the largest impact at lower energies

and higher masses where the resulting muons are less boosted and hence more likely to have split into two tracks before reaching ArgoNeuT. As the distance into the cavern increases, the efficiency further declines since the muons are less likely to still be overlapping by the time they reach the detector.

8.8 Background estimation

The primary backgrounds in this search originate from mis-reconstructed muon neutrino interactions occurring within ArgoNeuT, and from neutrino-induced through-going muons arising from interactions upstream of the detector. The backgrounds are evaluated using ArgoNeuT's simulation of the NuMI beam, described in Section 5.1.1. A sample of 600,000 events corresponding to 9.1×10^{20} POT was used. This is approximately 7.28 times the ArgoNeuT data POT.

8.8.1 Two-track signature

In the two-track scenario, the dominant form of observed background events are charged current ν_μ interactions that produce a single charged pion along with the outgoing muon. These are most likely as a result of either coherent or resonant pion production [235]. Charged pions are challenging to distinguish from muons in LArTPC detectors, especially those as small as ArgoNeuT, since they produce tracks with similar dE/dx as discussed in Section 5.3.3. The outgoing pion therefore may be mistaken for a second muon in the ArgoNeuT detector, and may reach the MINOS-ND enabling a successful matching. An example of this type of event in ArgoNeuT can be seen in Figure 8.19. However, this type of event can typically be removed based on the topology of the interaction in either ArgoNeuT or the MINOS-ND. In ArgoNeuT they can be removed either based on the opening angle between the two tracks – HNL decays are typically much more forward going, as shown in Figure 8.8 – or, in the case of resonant pion production, the presence of additional proton tracks originating from the vertex. In the example shown, there are multiple low energy protons present, however, these may be too short to be reconstructed correctly. Instead this event would be rejected in ArgoNeuT based on the large opening angle between the tracks. In the MINOS-ND this type of event would also likely be rejected. The majority of pion tracks from beam neutrino interactions either do not reach the MINOS-ND or do not form a track at least 1 m long and hence would be rejected during the matching. The small fraction that do may then be rejected based on the track dE/dx in the MINOS-ND, as described in Section 8.5.4. Finally, it is also possible that the pion track could be mismatched to a second nearby neutrino-induced background muon in the MINOS-ND. In this

case, the event would most likely be rejected in the MINOS-ND based on the timing information as described in Section 8.5.3.

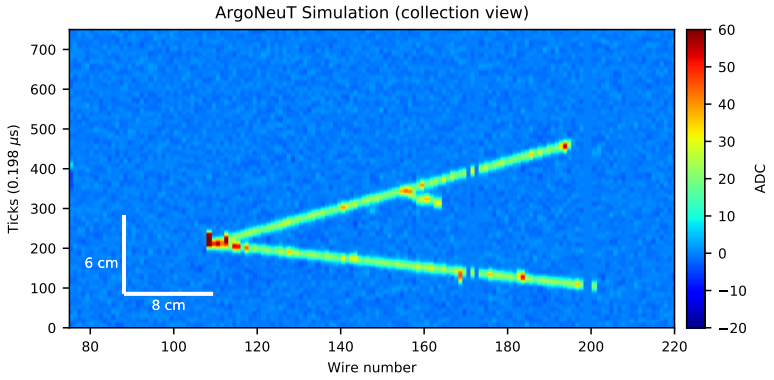

Figure 8.19: Simulated example background event in ArgoNeuT. A ν_μ interaction produces a muon and a single charged pion that could be mistaken for the two-track signature.

In the beam simulation sample 1 event passes the selection with the two-track signature, as shown in Table 8.4. This corresponds to a background expectation of 0.1 ± 0.1 events of this topology in the data, once accounting for the POT scaling. The uncertainty on the background expectation is overwhelmingly dominated by statistical uncertainties due to the small number of background events remaining despite the large initial statistics.

8.8.2 Double-MIP signature

In the double-MIP scenario, the dominant type of background events are single reconstructed muons that have either low energy δ-rays or low energy protons near the track vertex causing them to have a double-MIP-like dE/dx. These can then be incorrectly matched to a pair of muons in the MINOS-ND if a second beam-induced background muon is passing near the ArgoNeuT detector at approximately the same time. An example of this type of event in the MINOS-ND can be seen in Figure 8.20. In this example, a single track in ArgoNeuT is matched to a pair of muon tracks in the MINOS-ND. The second matched track, however, is a background neutrino-induced through-going muon that passes close to the ArgoNeuT detector and is only seen in the MINOS-ND. A third, shorter, track can also be seen in the MINOS-ND, however this is not matched. The vast majority of this type of event can be rejected based on the timing information in the MINOS-ND, as described in Section 8.5.3. If the pair of muons have originated from a single decay or interaction they should arrive in the MINOS-ND at approximately the same time ($\Delta t < 20\,\text{ns}$),

whereas neutrino-induced background muons are spread approximately uniformly throughout the 8 μs beam spill window.

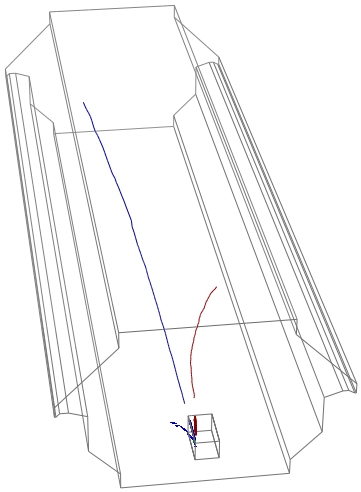

Figure 8.20: Simulated example background event in the MINOS-ND. A single track in ArgoNeuT is incorrectly matched to a pair of muon tracks in the MINOS-ND due to the presence of a second nearby beam-induced background muon. This type of event could be mistaken for the double-MIP signature.

In the beam simulation, the MINOS-ND timing information, despite being simulated in the MINOS-ND software, is not accessible in the ArgoNeuT software and had not been used in any previous ArgoNeuT analyses. Instead, to estimate the expected background the selection is performed without the timing cut in place, and then the result scaled down based on the time window cut applied relative to the full beam spill window size. The timing information is present in the data and was validated for both randomly distributed background muons – found to be uniformly distributed across the beam spill window as expected – and for neutrino interactions with two matched tracks – found to have $\Delta t << 20$ ns as expected.

In the beam simulation sample 2 events pass the selection with the double-MIP signature after the MINOS-ND timing scaling has been applied, as shown in Table 8.5. This corresponds to a background expectation of 0.3 ± 0.2 events of this type in the data, once accounting for the POT scaling. As with the two-track case, the uncertainty on the background expectation is statistically dominated.

8.8.3 Summary of the expected background rates

Table 8.7 summarises the expected background rates for each signature. The total expected background is 0.4 ± 0.2 events. The combination of the two detector

technologies – the high precision LArTPC of ArgoNeuT and the large magnetised muon-spectrometer of the MINOS-ND – allows the vast majority of background events to be rejected. This enables a very clean, near-background-free, search to be performed for the $N \to \nu\mu^+\mu^-$ signature.

Signature	Predicted events
Two-track	0.1 ± 0.1
Double-MIP	0.3 ± 0.2
Total	0.4 ± 0.2

Table 8.7: Predicted number of background events for each signature.

8.9 Systematic uncertainties

8.9.1 HNL production

$D^{\pm}_{(s)}$ production

The production cross-sections of $D^{\pm}$ and $D^{\pm}_{(s)}$ mesons in the NuMI environment have never been measured. However, fixed-target measurements at similar energies were performed by the NA32 [322] and E769 [323] experiments. These measurements are summarised and re-evaluated with updated branching ratios in Reference [324] and found to be consistent with the predictions from Pythia. Based on these measurements, a 20% systematic uncertainty on the $D^{\pm}_{(s)}$ production is applied. This leads to a 20% uncertainty on the flux of HNLs produced from $D^{\pm}_{(s)}$ decays.

$K^{\pm}$ production

In the electron-coupled and muon-coupled scenarios, HNLs are also produced from $K^{\pm}$ decays. The production cross-section of $K^{\pm}$ mesons has been measured at 120 GeV with a carbon target with an uncertainty of $^{+1.9}_{-5.8}$% [325].

$D^{\pm}_{(s)} \to \tau^{\pm} + \nu_{\tau}$ branching ratios

The branching ratios $D^{\pm} \to \tau^{\pm} + \nu_{\tau}$ and $D^{\pm}_s \to \tau^{\pm} + \nu_{\tau}$ have a 22.5% and 4.2% uncertainty, respectively [27]. Using the ratio of the contributions of $D^{\pm}$ and $D^{\pm}_s$ to the $\tau^{\pm}$ flux [85], this corresponds to an additional 5.7% uncertainty in the tau-coupled HNL flux.

Summary of the systematic uncertainties on production

For electron-coupled and muon-coupled HNLs the total production systematic uncertainty varies depending on the production mechanism. For HNLs produced from

$K^{\pm}$ decays the production systematic uncertainty is $^{+1.9}_{-5.8}$%. For HNLs produced from $D^{\pm}_{(s)}$ the production systematic uncertainty is 20%.

For tau-coupled HNLs the production systematic uncertainty is 20.8%, combining the impact of the uncertainty in the $D^{\pm}_{(s)}$ meson production and their subsequent decay to $\tau^{\pm}$ leptons.

8.9.2 Reconstruction effects

To determine the impact of each reconstruction related uncertainty on the final result, the analysis is repeated varying each parameter according to its uncertainty individually. This is performed for the tau-coupled scenario, as that is the primary focus of this analysis. The impact for the electron-coupled and muon-coupled scenarios would be approximately the same.

ArgoNeut calorimetry

There is a 3% uncertainty on the tuning of the calorimetry in ArgoNeuT [235]. This affects the calculated dE/dx used in the selection for both signatures. However, the impact of this on the tau-coupled HNL selection efficiencies is small: $^{+0.07}_{-0.09}$ %.

ArgoNeuT angle reconstruction

There is a 1-3% uncertainty on the track angular reconstruction in ArgoNeuT [270]. To be conservative, a 3% systematic uncertainty on the resulting reconstructed track angle is considered. The impact on the tau-coupled HNL selection efficiencies from this uncertainty is small: $^{+0.11}_{-0.06}$ %.

MINOS-ND energy reconstruction

The uncertainty on the reconstructed energy of tracks in the MINOS-ND is 2% for stopping muons, 4% for exiting muons and 5.6% for hadronic particles [239]. This reconstructed energy is used to determine the average dE/dx over the full track length for contained tracks. To be conservative, a 5.6% systematic uncertainty on the reconstructed track energy in the MINOS-ND is considered. The impact of this, however, on the selection efficiencies for the tau-coupled HNL sample is negligible: ± 0.002%.

MINOS-ND magnetic field modelling

Uncertainty on the modelling of the magnetic field in the MINOS-ND could impact the efficiency of the charge reconstruction, especially in the region close to the coil where the field is most complex. To assess the impact of mismodelling in different

regions of the MINOS-ND detector, the MINOS collaboration performed a run with the magnetic field polarity reversed. This results in charged particles being focused in the opposite direction, for example a positively charged particle that previously would have been curved towards the coil now curves away. This data-set was then used in an inclusive cross-section analysis, finding a difference on the level of 1% in the final cross-section [239] compared with the nominal magnetic field polarity. It should be noted that this discrepancy also incorporates the impact on the momentum reconstruction, rather than solely due to charge mis-identification. However, to be conservative, a 1% systematic uncertainty on the charge reconstruction efficiency in the MINOS-ND due to uncertainty in the magnetic field modelling is considered. The impact of this on the selection efficiencies for the tau-coupled HNL sample is $\pm 0.47\%$.

Summary of the systematic uncertainties on reconstruction

Combining the impact of the performed variations in quadrature leads to a total systematic uncertainty of 0.5% due to reconstruction effects. This uncertainty is relatively small since the majority of the selection cuts are fairly loose, and hence small variations in the reconstructed variables do not significantly impact the selection efficiency.

8.9.3 Other sources of systematic uncertainty

Impact of through-going muons

Neutrino-induced through-going muons were not simulated in the HNL samples used to determine the selection efficiencies. If backgrounds such as an additional through-going muon were present, these could lead to a reduction of efficiencies due to events being discarded for having additional tracks present. In ArgoNeuT a rate of approximately 60 through-going muons per hour is observed, corresponding to approximately 3.3% of triggers [159, 172]. The fraction of events with other interactions, such as from neutrinos, is sub-dominant. If all of these events are discarded due to having additional tracks present, this could lead to a 3.3% reduction in the expected efficiencies. This is accounted for by added a 3.3% systematic uncertainty on the selection efficiency.

Detector volume

There is an uncertainty on the size of the ArgoNeuT instrumented volume. In the y and z dimensions this arises from the 1 mm uncertainty in the locations where the wires intersect. In the x dimension this arises from the determination of the electron

drift velocity, resulting in an uncertainty of 1 cm. This results in a 2.2% uncertainty on the size of the ArgoNeuT instrumented volume [270].

POT counting

There is a 1% uncertainty assigned to the number of POT collected by ArgoNeuT [232].

Background modelling

The estimation of the expected number of background events is highly statistically limited despite the large initial size of the beam simulation sample used to evaluate them. Therefore, the impact of systematic uncertainties on the background estimation, such as on the modelling of neutrino interactions in GENIE, would be negligible and as a result are not considered.

8.9.4 Summary of the systematic uncertainties

The systematic uncertainties are summarised in Table 8.8 for the tau-coupled scenario. For tau-coupled HNLs originating from the decays of secondary $\tau^{\pm}$ leptons, the total systematic uncertainty is $\pm 21.2\%$.

Systematic Uncertainty	Impact (%)
HNL flux	20.8
Reconstruction effects	0.5
Selection efficiency	3.3
Instrumented volume	2.2
POT counting	1.0
Total	21.2

Table 8.8: The systematic uncertainty impact on the sensitivity for the tau-coupled scenario.

The systematic uncertainty in the electron-coupled and muon-coupled scenarios is lower due to the lower uncertainty on the $K^{\pm}$ cross-section. For electron-coupled and muon-coupled HNLs originating from $K^{\pm}$ decays, the total systematic uncertainty is $^{+4.5}_{-7.1}\%$. For electron-coupled and muon-coupled HNLs originating from $D^{\pm}_{(s)}$ decays, the total systematic uncertainty is $\pm 20.4\%$.

8.10 Results

8.10.1 Data

The selection has been applied to ArgoNeuT's full anti-neutrino mode data-set corresponding to an exposure of 1.25×10^{20} POT. In total, 0 events pass the selection with either the two-track or the double-MIP signature. In each case, this is consistent with the expected background rate predicted by the beam simulation.

8.10.2 Exclusion limits

The absence of an observed signal allows an exclusion limit to be set on the HNL phase-space. The limits are evaluated using a Bayesian approach with a uniform prior [326]. The 90% confidence level (CL) requirement when 0.4 ± 0.2 background events are expected and 0 are observed is 2.30 events.

Electron-coupled model

Figure 8.21 shows the constraint on the parameter space at 90% CL from 1.25×10^{20} POT at ArgoNeuT (area enclosed by the black line) for the electron-coupled model, assuming production from $K^{\pm}$ and $D^{\pm}_{(s)}$ decays. The $\pm 1\sigma$ uncertainty on the expected constraint includes both the uncertainty on the background expectation and the systematic uncertainty on the signal production, combined conservatively.

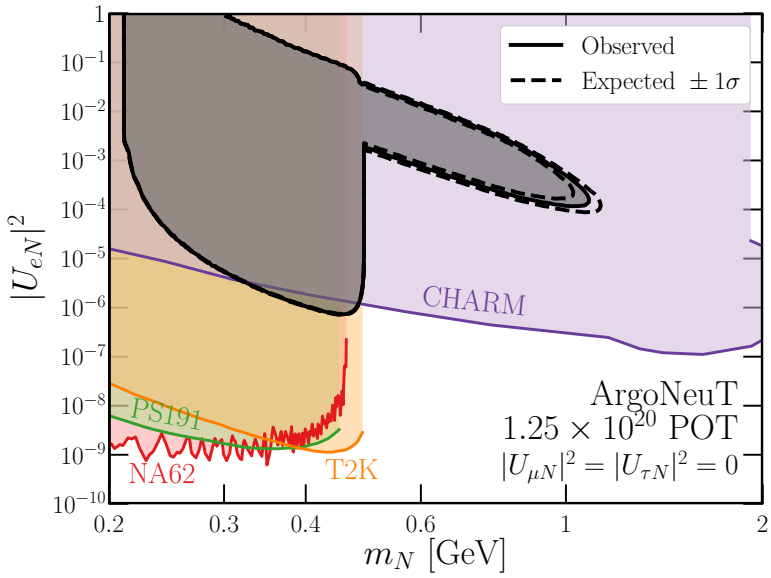

Figure 8.21: Constraint on the parameter space at 90% CL from 1.25×10^{20} POT at ArgoNeuT (black) for the electron-coupled model assuming production from $K^{\pm}$ and $D^{\pm}_{(s)}$ decays. The existing limits are shown from CHARM (purple) [88], PS191 (green) [89, 90], T2K (orange) [87] and NA62 (red) [91].

The strongest existing limits are shown from CHARM (purple) [88], PS191 (green) [89, 90], T2K (orange) [87] and NA62 (red) [91]. At lower masses production from both $K^{\pm}$ and $D^{\pm}_{(s)}$ is possible, whereas at higher masses only production from $D^{\pm}_{(s)}$ is kinematically allowed. This results in a sharp decrease in the sensitivity for HNL masses $m_N > m_{K^{\pm}} - m_{e^{\pm}}$. The extracted limit from ArgoNeuT is found to be considerably weaker than other existing limits across the full mass range. This is due to a combination of the significantly greater exposure of the other experiments along with access to final-states with larger branching ratios, such as $N \to e^{-}\pi^{+}$.

Muon-coupled model

Figure 8.22 shows the constraint on the parameter space at 90% CL from 1.25×10^{20} POT at ArgoNeuT (area enclosed by the black line) for the muon-coupled model, assuming production from $K^{\pm}$ and $D^{\pm}_{(s)}$ decays. The $\pm 1\sigma$ uncertainty on the expected constraint includes both the uncertainty on the background expectation and the systematic uncertainty on the signal production, combined conservatively. The strongest existing limits are shown from NuTeV (blue) [94], E949 (purple) [92] and NA62 (red) [93]. Similar to the electron-coupled scenario, there is a sharp decrease in the sensitivity for $m_N > m_{K^{\pm}} - m_{\mu^{\pm}}$ beyond which production from $K^{\pm}$ decay is not possible. As with the electron-coupled scenario, the ArgoNeuT constraint is found to be significantly weaker than existing limits. This is due to a combination of the greater exposure of the other experiments along with access to

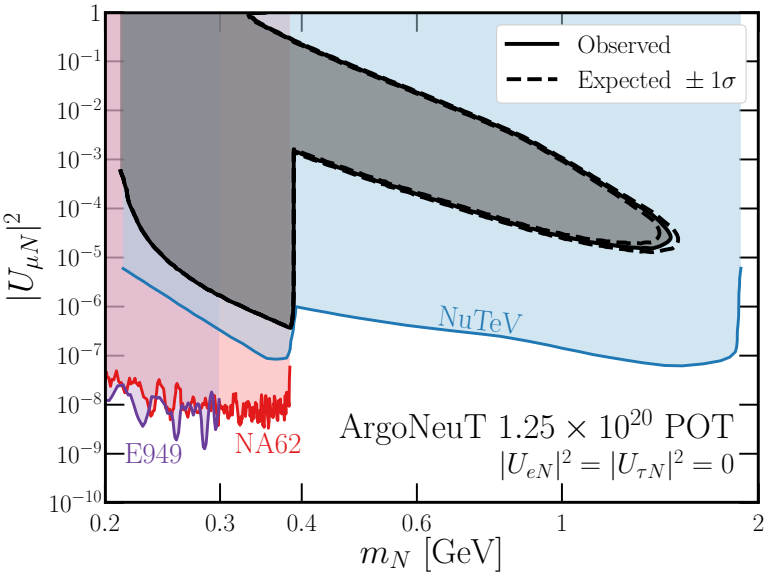

Figure 8.22: Constraint on the parameter space at 90% CL from 1.25×10^{20} POT at ArgoNeuT (black) for the muon-coupled model assuming production from $K^{\pm}$ and $D^{\pm}_{(s)}$ decays. The existing limits are shown from NuTeV (blue) [94], E949 (purple) [92] and NA62 (red) [93].

final states with larger branching ratios, such as $N \to \mu^- \pi^+$.

Tau-coupled model

Figure 8.23 shows the constraint on the parameter space at 90% CL from 1.25×10^{20} POT at ArgoNeuT (area enclosed by the black line) for the tau-coupled model, assuming production from $\tau^\pm$ decays. The $\pm 1\sigma$ uncertainty on the expected constraint includes both the uncertainty on the background expectation and the 21.2% systematic uncertainty on the signal production, combined conservatively. The existing limits are also shown from CHARM [96] and DELPHI [97] in purple and blue, respectively. Unlike in the electron-coupled and muon-coupled scenarios, ArgoNeuT is sensitive to regions of phase-space that have not previously been excluded. The results of this search lead to a significant increase in the exclusion region on the mixing angle $|U_{\tau N}|^2$ of tau-coupled Dirac HNLs with masses $m_N = 280$ - $970\,\text{MeV}$, assuming $|U_{eN}|^2 = |U_{\mu N}|^2 = 0$. For part of the phase-space, the excluded region is a closed contour rather than an upper-limit. This is due to the $\sim$1033 (318) m that the HNLs must propagate between the beam target (absorber) and ArgoNeuT before decaying in order to be detected. At larger mixing angles, particularly for larger masses, the HNLs would decay before reaching ArgoNeuT and hence this region of phase-space cannot be excluded. The slight increase in sensitivity at $m_N \sim 1000\,\text{MeV}$ seen in the $+1\sigma$ contour is as a result of several additional HNL decay channels becoming kinematically allowed and having a significant contribution, for example $N \to \nu\rho^0$,

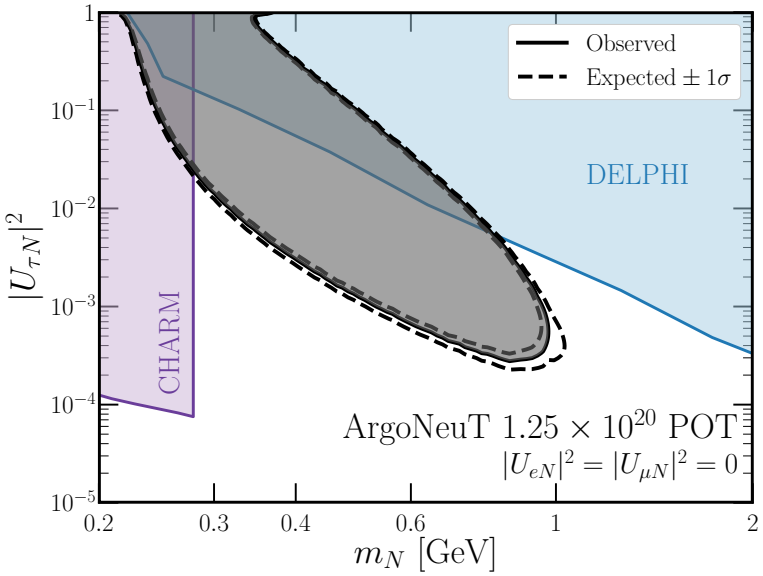

Figure 8.23: Constraint on the parameter space at 90% CL from 1.25×10^{20} POT at ArgoNeuT (black) for the tau-coupled model assuming production from $\tau^\pm$ decays. The existing limits from CHARM [96] and DELPHI [97] are shown in purple and blue, respectively.

altering the total rate [85].

8.11 Conclusions and future

In this chapter a search for Heavy Neutral Leptons with masses $\mathcal{O}(100)$ MeV performed using the ArgoNeuT experiment has been presented. The HNLs are produced in the NuMI neutrino beam from high-energy proton–fixed-target collisions in the NuMI target and absorber. They can then propagate to the ArgoNeuT detector, or the cavern immediately upstream of ArgoNeuT, before decaying. The decay signature $N \to \nu\mu^+\mu^-$ was searched for. Three scenarios were considered: electron-coupled, $|U_{eN}|^2 \neq 0$; muon-coupled, $\left|U_{\mu N}\right|^2 \neq 0$; and tau-coupled, $|U_{\tau N}|^2 \neq 0$. In each case assuming that the mixing with the other two flavours was zero. A novel technique was developed to identify the pairs of overlapping highly-forward-going oppositely-charged muons from these decays using the dE/dx and interaction topology in ArgoNeuT combined with matching to the magnetised MINOS near detector. The combination of the two detector technologies – ArgoNeuT as a high resolution LArTPC detector and the MINOS near detector as a large magnetised muon-spectrometer – enabled a near-background-free search to be performed for this signature. In the data, corresponding to an exposure to 1.25×10^{20} POT, zero passing events were observed consistent with the expected background of 0.4 ± 0.2 events. The results of this search lead to a significant increase in the exclusion region on the mixing angle $|U_{\tau N}|^2$ of tau-coupled Dirac HNLs with masses $m_N =$ 280 - 970 MeV, assuming $|U_{eN}|^2 = \left|U_{\mu N}\right|^2 = 0$. New constraints are also applied on electron-coupled and muon-coupled HNLs, however these are weaker than existing limits. This analysis is the first search for Heavy Neutral Leptons decaying with the signature $N \to \nu\mu^+\mu^-$ in a LArTPC detector. It is also the first search for tau-coupled HNLs in a LArTPC detector. The techniques developed in this analysis could be applied in future HNL searches performed in larger mass LArTPC experiments. The analysis described in this chapter has been published in Reference [5].

Future searches in other LArTPC detectors are expected to have significant sensitivity to these models. In particular, the DUNE near detector (DUNE-ND) [253] will be able to constrain a substantial new region of parameter space [84, 85]. This will surpass the new constraint on tau-coupled HNLs from this analysis as a result of the much larger anticipated exposure due to both the larger detector size and larger incident POT. The DUNE-ND will have a multi-detector design consisting of a LArTPC along with multiple magnetised muon-spectrometers, as described in Section 3.6.3. The techniques developed in this analysis could therefore be applied in the DUNE-ND to search for the $N \to \nu\mu^+\mu^-$ signature. The DUNE-ND is also expected to be sensitive to other complimentary signatures, such as $N \to \nu e^+ e^-$.

The MicroBooNE [158] and ICARUS [3] experiments could also be sensitive to these models in the same mass range using the NuMI beam located off-axis to these detectors. The SBN program detectors [3] would also be sensitive to lower mass HNLs produced in the 8 GeV Booster Neutrino Beam [242]. Finally, since the publication of this analysis, a re-cast extending the CHARM constraint on tau-coupled HNLs has been performed [327]. This surpasses the constraint set by ArgoNeuT. However, this analysis was performed without access to the CHARM data or simulation. It therefore does not model experimental efficiencies or the corresponding uncertainties realistically.

Chapter 9

Searching for Heavy QCD Axions with ArgoNeuT

This chapter describes a search for Heavy QCD Axions with masses $\mathcal{O}(100)$ MeV performed using the ArgoNeuT experiment. The Heavy QCD Axions could be produced in the NuMI neutrino beam and then subsequently decay with the signature $a \to \mu^+\mu^-$, which is possible to detect in ArgoNeuT and the MINOS near detector. A paper describing this search is available as a preprint [6] and has been submitted to a journal.

The search presented in this chapter was performed in collaboration with phenomenologists Z. Liu, R. Co, K. Lyu (University of Minnesota) and S. Kumar (University of California, Berkeley) who were responsible for the theoretical parts of the analysis. The author was responsible for the experimental parts of the analysis.

9.1 Model and simulation

The Heavy QCD Axion model, production in the NuMI beam target and absorber via meson-mixing, and subsequent decay in ArgoNeuT are described in detail in Sections 2.3.3 and 4.3. The generation of the simulated Heavy QCD axions and their decays is performed using the method described in Section 5.1.3. The resulting decay products are then propagated through the standard ArgoNeuT and MINOS near detector (MINOS-ND) simulation and reconstruction chains, an overview of which is given in Chapter 5. The Heavy QCD Axions will subsequently be referred to as axions to be concise.

Axion simulation samples were generated for production occurring at the NuMI beam target and at the hadron absorber. In both cases samples were generated for axions masses m_a = 220, 360, 500, 640, 780 and 920 MeV, covering range of axion masses that could be produced in the NuMI beam. For production at the target

separate samples were generated for production via π^0, η and η' meson mixing with 50,000, 50,000 and 25,000 events per mass point, respectively. For production at the absorber 100,000 events were generated at each mass point including production from all three types of meson-mixing weighted according to their expected contribution. In each case, the axion decays were simulated between $-63 < z < 90\,\text{cm}$ where $z = 0\,\text{cm}$ is the start of the ArgoNeuT instrumented volume. The position distribution of the decays within this volume and the kinematics of the muons were determined from the simulation of the axion kinematics and decay probability discussed in Section 5.1.3. In total 1,350,000 events were simulated.

The ArgoNeuT simulation of the NuMI beam, described in Section 5.1.1, was also used during the development of the selection to model the expected backgrounds. A sample of 600,000 events was used, corresponding to 9.1×10^{20} POT, 7.28 times the total ArgoNeuT data POT. This is the same sample that was used in the Heavy Neutral Lepton search described in Chapter 8. It will subsequently be referred to as the beam simulation sample.

9.2 Signature

The axion decay signature $a \to \mu^+\mu^-$ is similar to the Heavy Neutral Lepton (HNL) decay signature $N \to \nu\mu^+\mu^-$ in the search presented in Chapter 8. However, the kinematics of the axions and the resulting muons are different. The production mechanism via meson-mixing (as opposed to the meson decays for the HNLs) enables significantly more energetic axions to be generated in the NuMI beam. Additionally, the axion decay is two-body, rather than three-body, resulting in different muon kinematics.

9.2.1 Kinematics

Figure 9.1 shows the energy distributions, E_a, of $m_a = 500\,\text{MeV}$ axions produced from π^0 mixing (top-left), η mixing (top-right) and η' mixing (bottom) for production occurring at the NuMI beam target. The axions from all three production mechanisms are significantly more energetic than the HNLs described in Section 8.2.1, with energies up to nearly the beam energy of 120 GeV possible. The η and η' mixing mechanisms result in, on average, more energetic axions. However, the production rates from these mechanisms are subdominant compared with the π^0 mixing mechanism, as discussed in Section 5.1.3. The subsequent analysis development and optimisation will therefore focus on the kinematics of the π^0 case. The energy distributions of axions produced at the NuMI hadron absorber are approximately the same. The energy distributions also do not vary significantly as a function of the

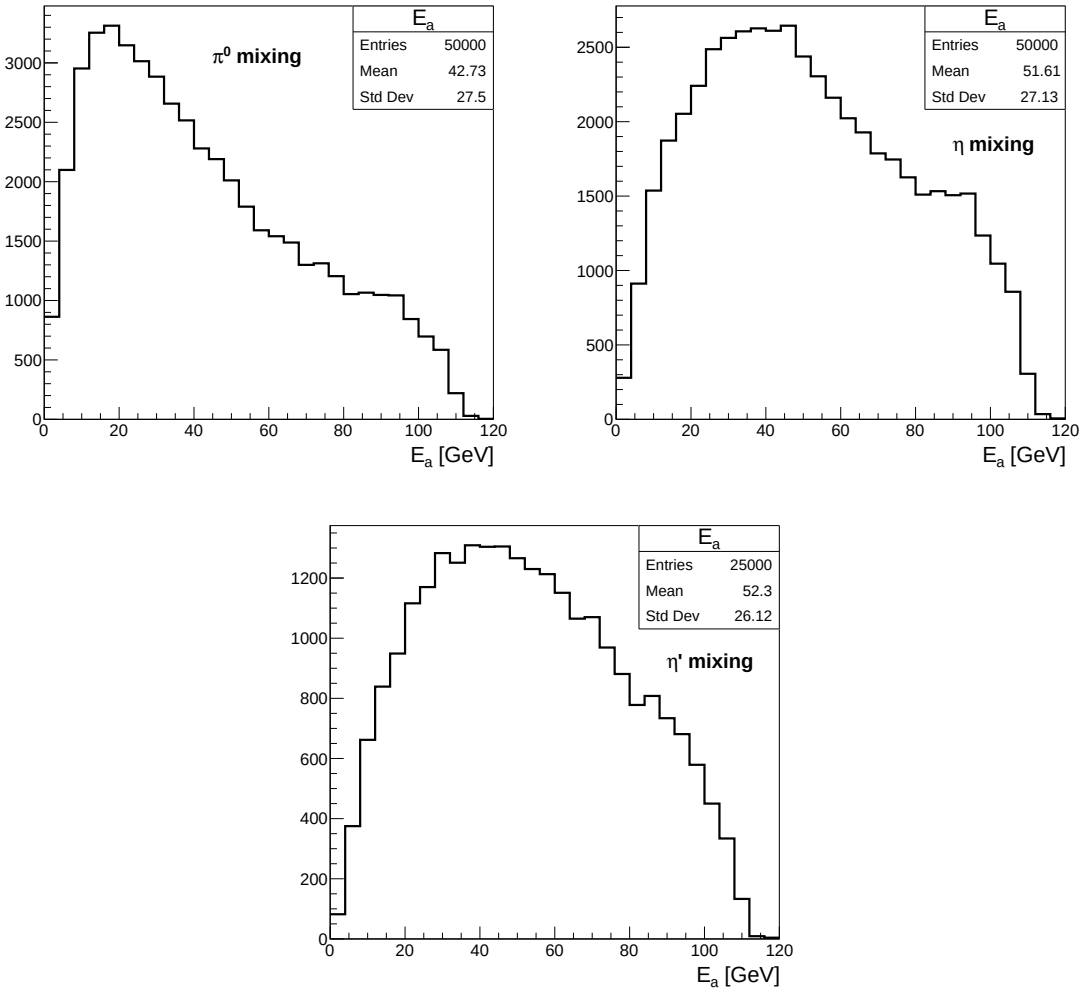

Figure 9.1: Axion energies, E_a, for production occurring from π^0 mixing (top-left), η mixing (top-right) and η' mixing (bottom). In each case $m_a = 500\,\text{MeV}$ and production is occurring at the beam target.

axion mass.

The resulting muons from the $a \to \mu^+\mu^-$ decay are highly energetic and forward-going across all axion masses. Figure 9.2 shows the truth energies (top-left), angles with respect to the beam direction (top-right), and opening angles (bottom) of the muons resulting from decays of $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue) axions. In each case, the axions are produced via π^0 mixing at the beam target. Equivalent plots for production via η mixing and η' mixing can be found in Appendix B. The kinematics of the muons vary depending on m_a since this impacts how boosted ($\gamma_a = E_a/m_a$) the parent axion was. The average energy of the muons for both masses is $\langle E_{\mu^\pm} \rangle \sim 20\,\text{GeV}$. The angle with respect to the beam direction, θ_{beam}, and the opening angle, $\theta_{opening}$, vary depending on the mass: with $\langle \theta_{beam} \rangle \sim 0.75^\circ$ and $\langle \theta_{opening} \rangle \sim 1.5^\circ$ at $m_a = 360\,\text{MeV}$; and $\langle \theta_{beam} \rangle \sim 2.5^\circ$ and $\langle \theta_{opening} \rangle \sim 5^\circ$ at $m_a = 780\,\text{MeV}$.

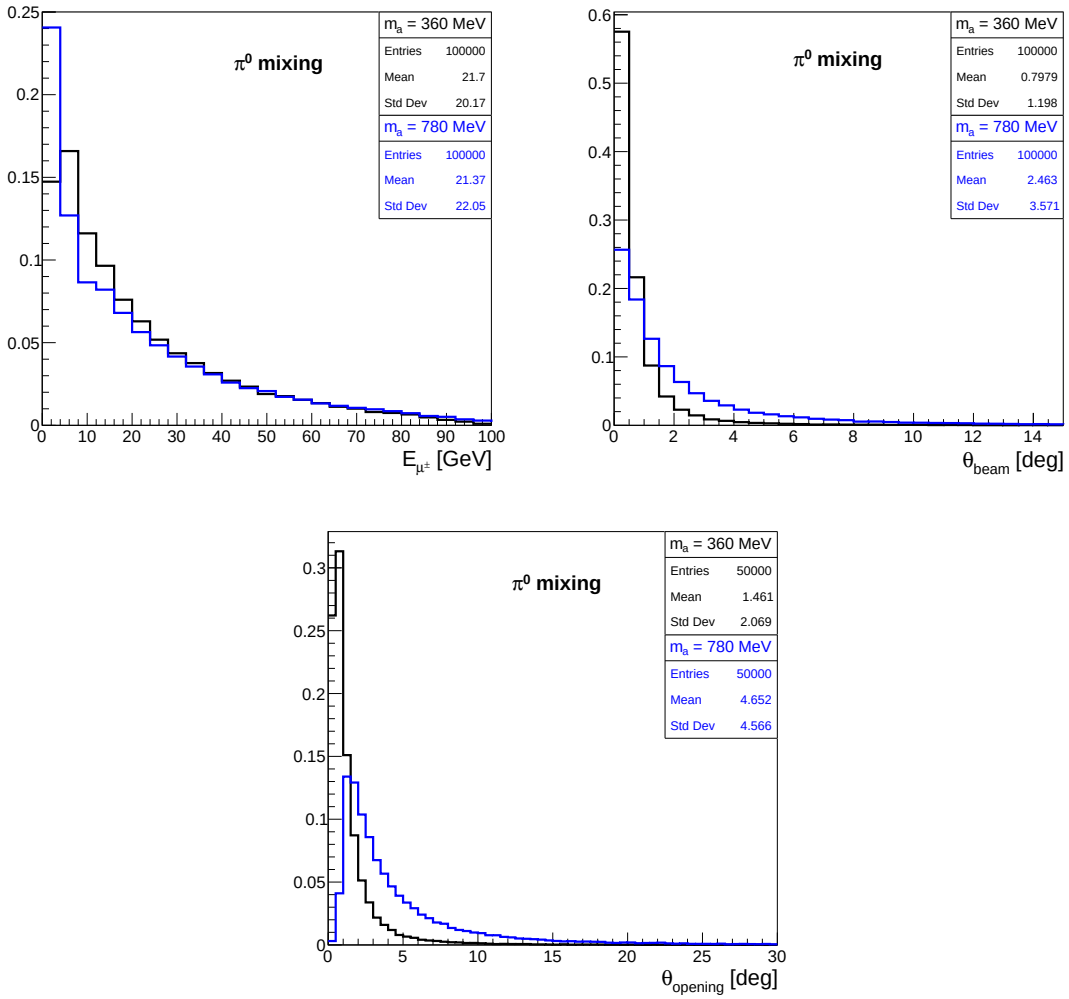

Figure 9.2: Truth energies (top-left), angles with respect to the beam direction (top-right), and opening angles (bottom) of the muons resulting from axion decays. In each case the axions are produced from π^0 mixing at the beam target and two different masses are shown, $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue). The resulting muons are energetic and highly forward-going, with their kinematics varying depending on m_a.

9.2.2 Signatures in ArgoNeuT

The axion decay in ArgoNeuT and the MINOS-ND is similar to that of the HNL decay and the same set of are signatures are therefore considered. These are described in detail in Sections 8.2.2 and 8.2.3 for the HNL search and are summarised here. Two in-ArgoNeuT decay signatures are considered, each of which are illustrated in Figure 9.3. If the opening angle between the muons is sufficiently large, two distinct minimally ionising tracks are reconstructed in ArgoNeuT originating from a common vertex. This topology will subsequently be referred to as the *two-track* signature and is illustrated in Figure 9.3 (left). In cases where the opening angle is smaller, the pair of muons may be reconstructed as a single track for part, or all, of their length

in ArgoNeuT. This track would have double the ionisation of the minimally ionising single muon tracks. This topology will be referred to as the *double-MIP* signature and is illustrated in Figure 9.3 (right). In both cases the track(s) in ArgoNeuT can be matched to a pair of muon tracks in the MINOS-ND where the muons become clearly separated due to their opposite charges in the magnetic field.

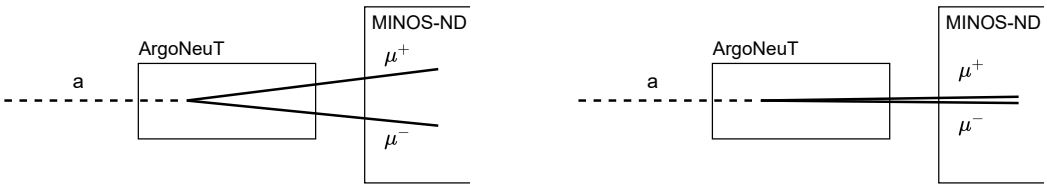

Figure 9.3: Diagrams of axion decays occurring inside the ArgoNeuT detector with the two-track (left) and double-MIP (right) decay signatures. Diagrams not to scale.

As with the HNL search, axion decays occurring in the upstream cavern are also considered. In this scenario the resulting muons then pass through the ArgoNeuT detector and can be matched to the MINOS-ND, as illustrated in Figure 9.4. As with the HNL search, only the double-MIP signature is considered in this scenario as a result of the easier background mitigation compared with the two-track signature. Additionally, only decays occurring in the 63 cm between the end of the MINERvA detector and the start of the ArgoNeuT instrumented volume are considered.

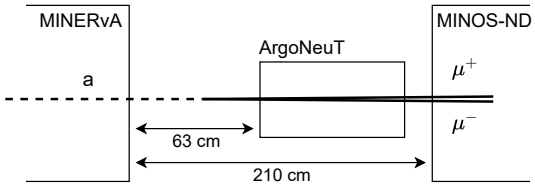

Figure 9.4: Diagram of axion decays occurring in the cavern between ArgoNeuT and the MINERvA detector with the double-MIP signature. The resulting muons pass through the ArgoNeuT detector and can be matched to the MINOS-ND. Diagram not to scale.

9.3 Selection

Since the axion decay signature is similar to the HNL decay signature an analogous selection strategy can be followed. The selection has, however, been modified and re-optimised to account for the differing kinematics of the axions and the impact of the two-body rather than three-body nature of the decay. This section will provide an overview of the selection, in particular highlighting the changes from the previous HNL selection. A detailed discussion of the HNL selection can be found in Sections 8.3, 8.4 and 8.5.

9.3.1 Pre-selection

A pre-selection is first applied to remove poorly reconstructed interactions and obvious non-axion backgrounds. Analogous to the HNL search, at least 80% of reconstructed energy depositions in ArgoNeuT are required to be associated with reconstructed tracks. This removes events that cannot be reliably identified as a result of incomplete reconstruction. It is also effective at removing background events containing electromagnetic showers since the lack of a dedicated shower reconstruction used in this analysis causes these to have large regions that are not reconstructed.

Next, events with more than two reconstructed tracks in ArgoNeuT are removed as the additional reconstructed particles are incompatible with axion decays. This is harsher than in the HNL analysis, where up to three tracks were allowed to compensate for common reconstruction failures as the pair of muons begin to separate. Since the muons originating from axion decays are on average more highly boosted, this occurs less frequently. The requirement to have only two tracks additionally allows unambiguous reconstruction of the axion invariant mass, which will be discussed further in Section 9.4. As with the HNL search, tracks shorter than $L = 5\,\text{cm}$ are not considered to avoid removing events containing δ-rays originating from the muons. In cases where two tracks are present, a further constraint is applied on the opening angle between them if they originate from a common vertex. Two tracks are considered to have originated from a common vertex if the start positions are within $r \leq 4\,\text{cm}$, where r is the separation in three-dimensions. Figure 9.5 (right) shows the reconstructed opening angles of muons originating from the decays of $m_a = 500\,\text{MeV}$ axions produced via π^0 mixing. Any events with a pair of tracks with an opening angle $\theta_{opening} > 15^\circ$ are removed, illustrated by the dashed line. These tracks are much more likely to have originated from a background neutrino interaction than from an axion decay. This cut is slightly looser than in the HNL search to allow better identification of higher mass axions where a subset of the muons may have a larger opening angle.

Finally, the number of *candidate tracks* in ArgoNeuT and the MINOS-ND are also identified, following the criteria described in Section 8.3. These are then considered in the subsequent selections. In addition to the constraint on the number of tracks in ArgoNeuT, at least two candidate tracks must be present in the MINOS-ND.

9.3.2 Topological and calorimetric selection in ArgoNeuT

Events that pass the pre-selection are then assessed against the two-track and double-MIP selection criteria. As with the HNL search, in cases where the event could be selected with either signature it is preferentially selected with the two-track signature.

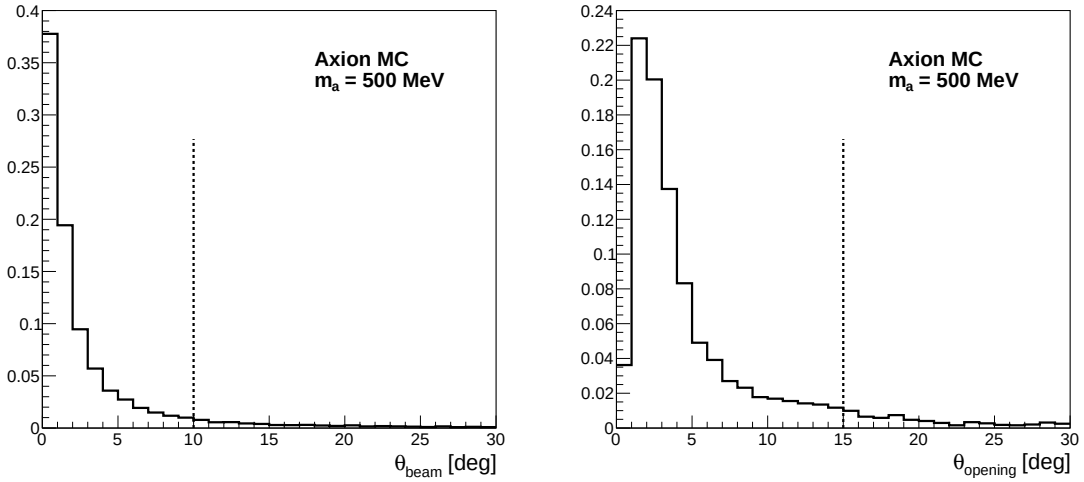

Figure 9.5: Reconstructed angles with respect to the beam direction (left) and opening angles (right) of muons originating from $m_a = 500\,\text{MeV}$ axions produced via π^0 mixing. The dashed lines show the cuts applied at $\theta_{beam} = 10^\circ$ and $\theta_{opening} = 15^\circ$, respectively.

In the two-track scenario a fiducial volume cut is applied to remove backgrounds originating from the cavern. The fiducial volume is defined as: $1 \leq x \leq 46\,\text{cm}$ (drift), $-19 \leq y \leq 19\,\text{cm}$ (vertical) and $z \geq 3\,\text{cm}$ (beam direction). Events with two tracks that originate from, or can be projected back to, a common vertex within the fiducial volume are then identified. The tracks are required to be forward-going with respect to the beam direction, have length $L \geq 5\,\text{cm}$, exit ArgoNeuT towards the MINOS-ND, have an average dE/dx over their full length consistent with being minimally ionising ($dE/dx < 3.1\,\text{MeV/cm}$), and have an opening angle between them of $\theta_{\text{opening}} \leq 15^\circ$. This selection is unchanged from the HNL search, except for the opening angle requirement previously mentioned.

In the double-MIP scenario, axion decays occurring both inside of ArgoNeuT and in the upstream cavern are considered. As discussed in Section 8.4.1, the strongest identifier of whether a pair of overlapping muons is present is provided by the dE/dx towards the beginning of the track prior to the muon pair potentially separating. The average dE/dx is calculated over the first 10 hits ($\sim 5\,\text{cm}$) of each track, where any anomalously large hits ($dE/dx > 10\,\text{MeV/cm}$) are discarded. This metric applied to simulated $m_a = 500\,\text{MeV}$ axions produced via π^0 mixing is shown in Figure 9.6. Similar to the HNL decays shown in Figure 8.9, two peaks are visible corresponding to single minimally-ionising muon tracks and to tracks with approximately double the minimally-ionising dE/dx where the pair of muons are overlapping. The threshold applied at $dE/dx = 3.1\,\text{MeV/cm}$ is shown by the dashed line. Candidate tracks are required to have an angle with respect to the beam direction $\theta_{\text{beam}} \leq 10^\circ$ and an average $dE/dx > 3.1\,\text{MeV/cm}$. Figure 9.5 (left) shows the reconstructed

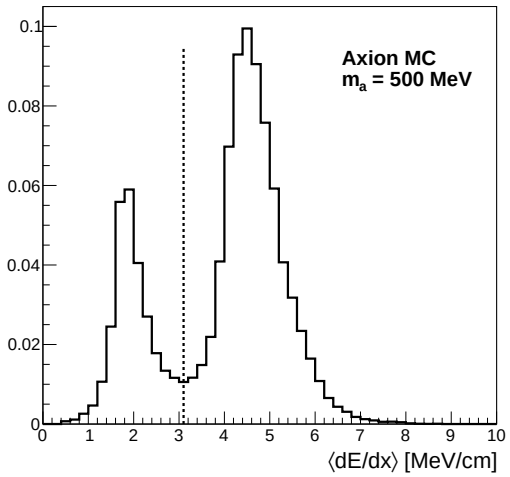

Figure 9.6: Average track dE/dx for decays of $m_a = 500\,\text{MeV}$ axions produced via π^0 mixing. The dashed line illustrates the threshold applied at $dE/dx = 3.1\,\text{MeV/cm}$.

angles with respect to the beam direction of muons originating from the decays of $m_a = 500\,\text{MeV}$ axions produced via π^0 mixing. The cut applied at $\theta_{\text{beam}} \leq 10^\circ$ is shown by the dashed line. The θ_{beam} requirement is tightened relative the HNL search since the muons from axion decays are on average more highly forward-going. The small subset of muons with larger angles with respect to the beam direction would instead be selected with the two-track signature. This selection is otherwise unchanged from the HNL search.

9.3.3 Topological and calorimetric selection in the MINOS-ND

Once candidate events are identified with either the two-track or double-MIP signatures in ArgoNeuT, MINOS-ND matching is performed. As described in Section 5.5.2, tracks exiting ArgoNeuT are projected to their expected start position in the MINOS-ND and compared with each reconstructed MINOS-ND track. In the two-track case the standard ArgoNeuT-MINOS-ND radial and angular matching tolerances are used: $r_{diff} = 12.0\,\text{cm}$ and $\theta_{diff} = 9.74^\circ$ [235]. In the double-MIP case the tolerances are doubled: $r_{diff} = 24\,\text{cm}$ and $\theta_{diff} = 19.48^\circ$. This accounts for the single track in ArgoNeuT being matched to two tracks in the MINOS-ND. This requirement is slightly tighter than in the HNL search, where a tolerance of 2.5 times the standard was used. Using a tighter matching requirement allows better rejection of other background, neutrino-induced, through-going muons that may be incorrectly matched, with negligible impact on the selection efficiency. As before, the matched tracks are required to be forward-going with respect to the beam direction, start in the calorimeter region of the MINOS-ND and within the first 20 cm

of its upstream face, and be at least 1 m long. This removes tracks that are unlikely to have originated from ArgoNeuT or short tracks that may have originated from a proton or pion rather than a muon.

Finally, several selection cuts are applied in the MINOS-ND. These cuts are the same for both the two-track and double-MIP signatures. The matched tracks are required to be reconstructed with opposite charges and to have start times, t_0, consistent with having originated from the same interaction or decay: $|\Delta t_0| \leq 20\,\text{ns}$. Pairs of tracks with larger Δt_0 could not have originated from a single axion decay and instead are likely random coincidences with neutrino-induced background muons. The average dE/dx of the tracks is also required to be consistent with a muon: $6 \leq dE/dx \leq 12\,\text{MeV/cm}$. This is significantly tighter than the previous very loose cut applied in the HNL search, allowing further rejection of potential backgrounds from protons and pions while having minimal impact on the selection efficiency.

9.3.4 Summary of the selection

A summary of the cuts applied and their impact on the selection efficiency is presented in Tables 9.1, 9.2 and 9.3 for two different axion masses, $m_a = 360\,\text{MeV}$ and $m_a = 780\,\text{MeV}$. In each case the axions are produced via mixing with π^0 mesons at the beam target and decay within the ArgoNeuT detector. The selection is also applied to the beam simulation sample modelling the expected backgrounds.

Table 9.1 shows the pre-selection cuts in the order that they are applied. Following the pre-selection, 65.5% of the $m_a = 360\,\text{MeV}$ and 54.5% of the $m_a = 780\,\text{MeV}$ samples remain. The pre-selection cuts remove 87.2% of the ArgoNeuT beam simulation events.

	Axion MC $m_a = 360\,\text{MeV}$	Axion MC $m_a = 780\,\text{MeV}$	Beam MC
Initial events	25000	25000	600000
Zero reconstructed activity	23977	24234	322024
Hit reconstruction fraction	22323	22005	172951
Number tracks ArgoNeuT	19130	17167	117468
Number tracks MINOS-ND	16410	13900	82715
Opening angle at vertex	16385	13623	77042

Table 9.1: Pre-selection cut flow table.

Table 9.2 shows the two-track selection cuts in the order that they are applied. The two-track selection is applied to all of the events that pass the pre-selection. In total, 2.4% of the $m_a = 360\,\text{MeV}$ and 10.8% of the $m_a = 780\,\text{MeV}$ samples are selected with the two-track signature. From the beam simulation, 0 events pass the

selection with this signature.

	Axion MC $m_a = 360\,\text{MeV}$	Axion MC $m_a = 780\,\text{MeV}$	Beam MC
Pass pre-selection	16385	13623	77042
Two tracks, start in fiducial volume and exit	2700	5611	2869
ArgoNeuT MIP dE/dx	985	3733	1638
MINOS-ND matching	738	3211	58
MINOS-ND MIP dE/dx	730	3192	57
ArgoNeuT vertex (MINOS matched tracks)	606	2708	0
MINOS-ND charge	606	2708	0
MINOS-ND timing	606	2708	0

Table 9.2: Two-track selection cut flow table.

Table 9.3 shows the double-MIP selection cuts in the order that they are applied. The double-MIP selection is applied to all of the events that pass the pre-selection, but do not pass the two-track selection. In total, 51.5% of the $m_a = 360\,\text{MeV}$ and 32.8% of the $m_a = 780\,\text{MeV}$ samples are selected with the double-MIP signature. From the beam simulation, 1 event passes the selection with this signature.

The total selection efficiencies are summarised in Table 9.4 for the two different example axion masses, in each case showing the contributions from the two-track and double-MIP signatures. The majority of the events are selected with the double-MIP signature in both cases. The two-track signature does, however, contribute significantly for the higher mass, less-boosted ($\gamma_a = E_a/m_a$), case, where the resulting muons have a larger average opening angle. The total selection efficiencies for the $m_a = 360\,\text{MeV}$ and $m_a = 780\,\text{MeV}$ samples are 53.9% and 43.6%, respectively. The selection is also highly effective at removing backgrounds: from the beam simulation sample 0 events pass with the two-track signature and 1 event passes with the double-MIP signature.

	Axion MC $m_a = 360\,\text{MeV}$	Axion MC $m_a = 780\,\text{MeV}$	Beam MC
Pass pre-selection and fail two-track selection	15779	10915	77042
Angle with respect to beam	13290	9135	36457
ArgoNeuT double-MIP dE/dx	12991	8513	2562
MINOS-ND matching	12957	8332	229
MINOS-ND MIP dE/dx	12919	8288	217
MINOS-ND charge	12866	8209	109
MINOS-ND timing	12866	8209	1

Table 9.3: Double-MIP selection cut flow table.

	Axion MC $m_a = 360\,\text{MeV}$	Axion MC $m_a = 780\,\text{MeV}$
Two-Track	2.4%	10.8%
Double-MIP	51.5%	32.8%
Total	53.9%	43.6%

Table 9.4: Summary of the selection efficiencies for the two example Heavy QCD Axion samples with masses $m_a = 360\,\text{MeV}$ and $m_a = 780\,\text{MeV}$. In each case, the contributions from the two-track and double-MIP signatures are shown.

9.4 Axion invariant mass

Since the axion decay is two-body, $a \to \mu^+\mu^-$, the invariant mass of the parent axion can be reconstructed. This was not possible in the HNL search due to the presence of an additional out-going neutrino. A constraint is not applied on the invariant mass in the selection since the search is performed across a significant axion mass range. However, if a signal were to be observed in the data, the invariant mass could be a powerful tool to further characterise and understand it.

The invariant mass is calculated using the track trajectories reconstructed in ArgoNeuT along with the track momentum reconstructed in the MINOS-ND. Depending on the signature this is calculated slightly differently. In the two-track case both of the muon trajectories are reconstructed in ArgoNeuT. In the double-MIP case, however, only a single track is reconstructed in ArgoNeuT. The muon trajectories are therefore identified by projecting between vertex of double-MIP track in ArgoNeuT and the start of each matched track in the MINOS-ND. The momentum is then reconstructed in the MINOS-ND using a range-based approach for contained tracks and curvature-based for exiting tracks, as described in Section 5.5.1.

The invariant mass resolution depends on the signature and whether the tracks are contained within the MINOS-ND. Figure 9.7 shows the performance of the invariant mass reconstruction for the two-track (left) and double-MIP (right) signatures where in both cases the muon tracks are contained within the MINOS-ND. In each case, the invariant mass is reconstructed effectively with a resolution of approximately 100 MeV. However, there is a slight bias to higher masses resulting from a long tail from cases where the track trajectories are poorly reconstructed. This is slightly worse for the double-MIP case. Figure 9.8 shows the performance of the invariant mass reconstruction where one or both muon tracks exit the MINOS-ND. In this case only the double-MIP signature is shown since a negligible number of two-track events contain sufficiently high-energy muons to exit the MINOS-ND. For exiting tracks the performance degrades significantly as a result of poorer reconstruction of the muon momentum. A resolution of approximately 200 MeV is obtained, with a slightly larger bias towards higher masses. Since the muons originating from

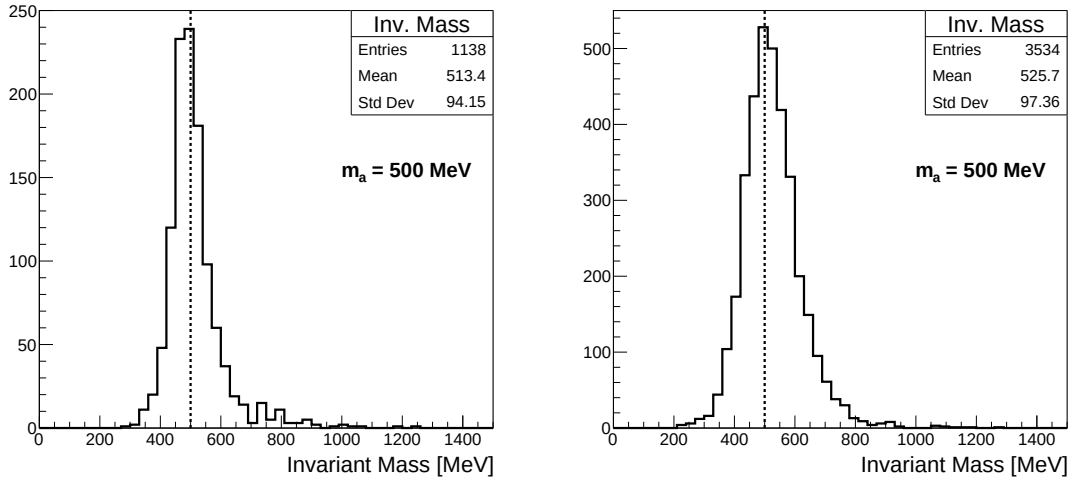

Figure 9.7: Reconstructed invariant mass of simulated $m_a = 500\,\text{MeV}$ axions for the two-track signature (left) and double-MIP signature (right). In both cases the muons are contained within the MINOS-ND.

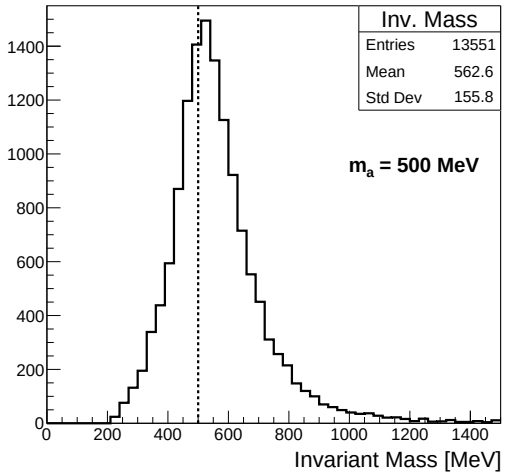

Figure 9.8: Reconstructed invariant mass of simulated $m_a = 500\,\text{MeV}$ axions with the double-MIP signature where one or both of the muons exit the MINOS-ND.

axion decays are highly energetic, this category of event is dominant. However, the majority of background neutrino-induced muons originating from the beam would be lower energy and likely contained within the MINOS-ND. Therefore most potential background signals would correspond to the scenario where both matched tracks are contained, where the axion invariant mass reconstruction performs better.

9.5 Selection efficiency

9.5.1 Decays inside ArgoNeuT

Figure 9.9 shows the selection efficiencies for axion decays occurring inside the ArgoNeuT detector as a function of E_a for m_a = 360, 500, 640 and 780 MeV. In each case, the contributions from two-track and double-MIP events are shown. The total efficiency is around 45 − 55% and relatively flat at higher energies, decreasing slightly at larger m_a. However, it drops significantly at lower energies. At low energies there are approximately equal numbers of two-track and double-MIP type events, whereas at higher energies the double-MIP type events dominate as the muons are increasingly forward-going. The fraction of events passing with the two-track signature also increases for higher m_a. In these cases, the axions are less boosted at lower energies

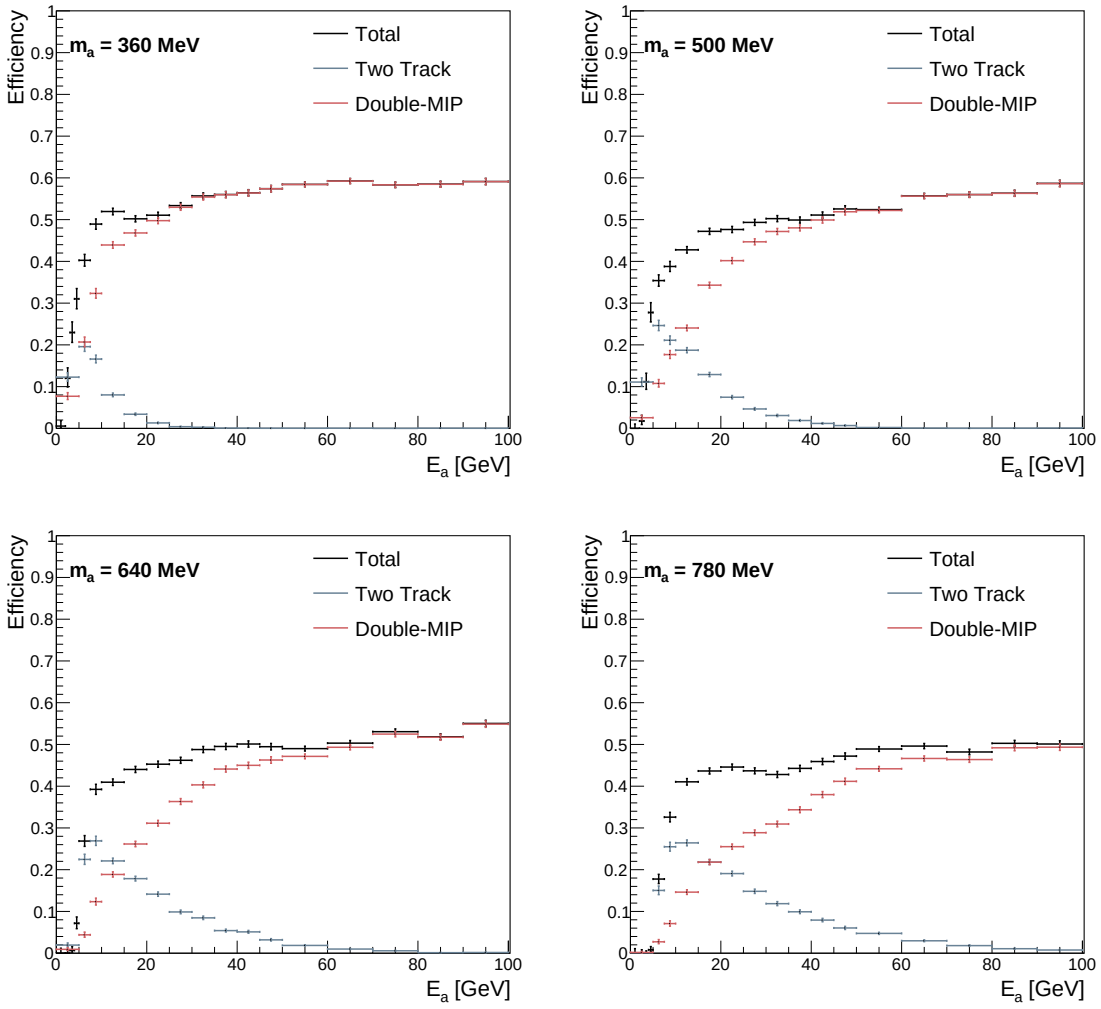

Figure 9.9: Selection efficiencies for axion decays occurring inside the ArgoNeuT detector as a function of E_a for m_a = 360, 500, 640 and 780 MeV. In each case the contribution from two-track and double-MIP events is shown.

due to their larger mass ($\gamma_a = E_a/m_a$) resulting in larger opening angles between the muons. The primary losses in efficiency are the same as those in the HNL search: at low axion energies one or both of the muons may stop before reaching the MINOS-ND; and at high energies reconstruction failures may occur resulting from the muons being approximately parallel to the wire planes. These are discussed in detail in Section 8.7.

Across all masses, the axion selection efficiencies are slightly lower than in the HNL analysis. As the axions are on average more boosted, and hence the muons are more forward-going, the reconstruction failures arising from the muons being close to parallel to the wire-planes occur more frequently reducing the efficiency. Additionally, the requirement to only have two tracks in ArgoNeuT reduces the efficiency at moderate energies. At these energies, in a subset of events, the muons may separate before exiting the ArgoNeuT detector, often resulting in reconstruction failures leading to extra split tracks. In the HNL selection this scenario is compensated for by allowing up to three tracks to be present, whereas in the axion selection these events are rejected. However, due to the higher average energies of the axions these events make up a smaller fraction of the total events.

9.5.2 Decays upstream of ArgoNeuT

Figure 9.10 shows the selection efficiencies for axion decays upstream of the ArgoNeuT detector in the cavern as a function of E_a for m_a = 360, 500, 640 and 780 MeV. In each case the efficiencies are shown for decays inside the detector and for decays in cavern occurring 10-20 cm, 30-40 cm and 50-60 cm upstream of the start of the instrumented volume. The efficiencies for decays upstream of ArgoNeuT are lower than for decays within the detector. This is primarily because only events with the double-MIP signature are selected. Therefore, only the most forward-going of the muon pairs can be selected. This has the largest impact at lower energies and higher masses where the resulting muons are less boosted and hence more likely to have split into two tracks. As the distance into the cavern increases, the efficiency further declines since the muons are less likely to still be overlapping by the time they reach the detector.

9.6 Backgrounds and systematic uncertainties

9.6.1 Backgrounds

The primary backgrounds in this search are the same as those for the HNL search, and are discussed in detail in Section 8.8. They predominantly arise from misreconstructed muon neutrino interactions occurring within the ArgoNeuT detector

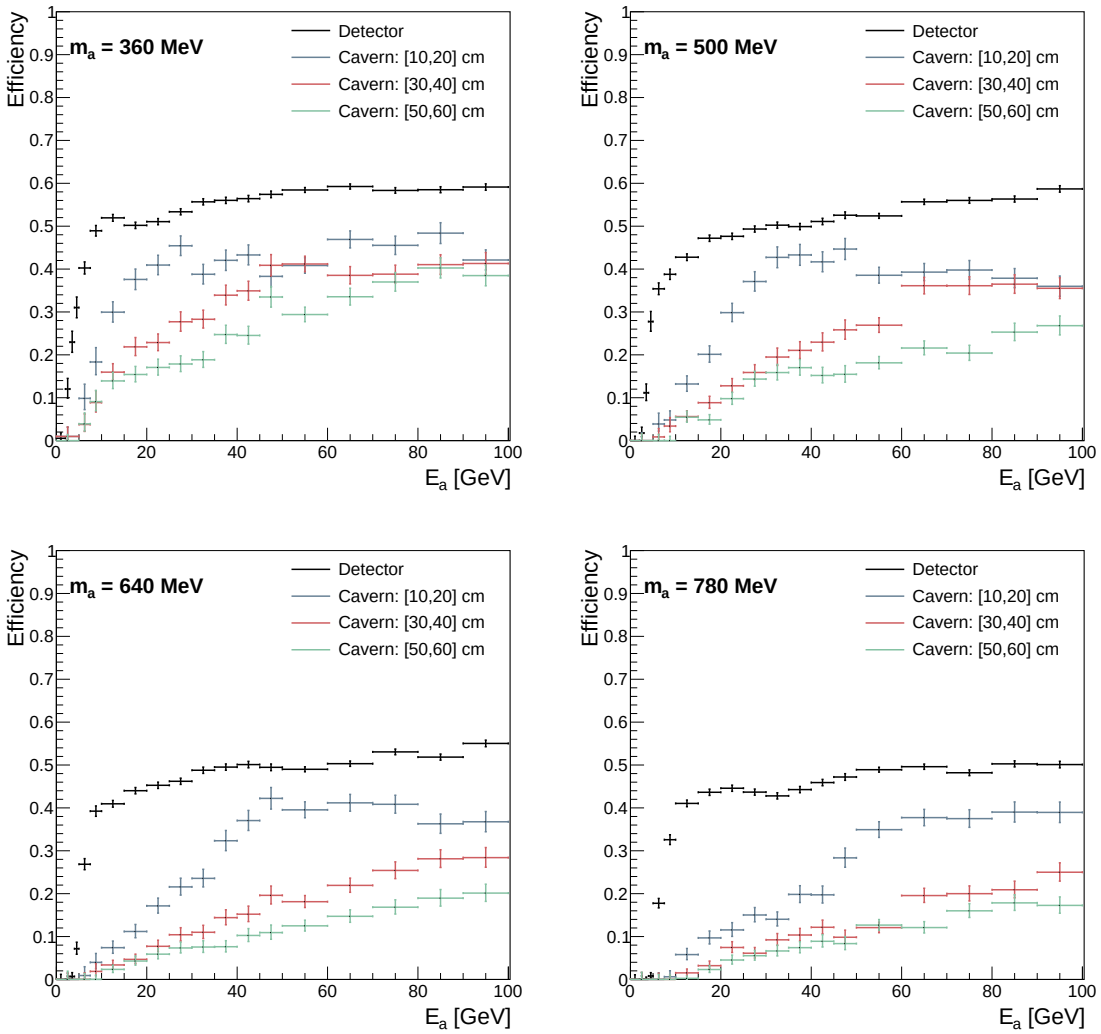

Figure 9.10: Selection efficiencies for axion decays upstream of the ArgoNeuT detector in the cavern as a function of E_a for m_a = 360, 500, 640 and 780 MeV. In each case the efficiencies are shown for decays inside the detector and for decays in the cavern occurring 10-20 cm, 30-40 cm and 50-60 cm upstream of the detector.

and from incorrectly matched, neutrino-induced, through-going muons arising from interactions upstream of the detector. The vast majority of these, however, can be removed based on the topological and calorimetric information in either ArgoNeuT or the MINOS-ND along with the timing information in the MINOS-ND. The backgrounds are evaluated using the beam simulation sample described in Section 9.1. This sample consists of 600,000 events corresponding to 9.1×10^{20} POT, 7.28 times the data POT. As shown in Tables 9.2 and 9.3, 0 events pass the selection with the two-track signature and 1 event passes with the double-MIP signature. Accounting for the POT scaling, this corresponds to a total background expectation of 0.1 ± 0.1 events in the ArgoNeuT data. This is slightly lower than in the HNL selection as a result of the stricter requirements on the number of tracks in ArgoNeuT and the

tighter MINOS-ND matching tolerances. The uncertainty on the background estimation is overwhelmingly statistically dominated due to the small number of events passing the selection despite the high initial statistics.

9.6.2 Systematic uncertainties

Axion production

The uncertainty on the axion production is dominated by theoretical uncertainties in the model. These are conservatively estimated to be approximately 30%, which is a typical size for QCD uncertainties [153]. The impact of these uncertainties on the expected event rate, and hence the sensitivity, varies depending on the model parameters chosen and the region of phase space probed. The effects of these are therefore shown as an error band in the final constraint in Section 9.7. Future refinements in the axion theoretical modelling would help to improve the precision of the derived limit.

Reconstruction effects

The impact of reconstruction related uncertainties are assessed by repeating the analysis with each parameter varied according to its uncertainty individually. The same sources of uncertainties are considered as those in the HNL search, described in Section 8.9. Uncertainties of 3% are applied on the tuning of the calorimetry [235], 3% on the track angular reconstruction [270], 6% on the energy reconstruction in the MINOS-ND [239] and 1% on the charge reconstruction due to the modelling of the magnetic field [239]. Combining the impact of the performed variations leads to a 0.5% systematic uncertainty due to reconstruction effects.

Additional sources of uncertainty

Several additional sources of uncertainty are accounted for. These are the same as those described in Section 8.9. A 3.3% systematic uncertainty is assigned to the selection efficiency to account for the potential impact of neutrino-induced through-going muons [159, 172]. There is a 2.2% uncertainty in the size of the ArgoNeuT instrumented volume originating from uncertainty in the electron drift velocity [270]. Finally, there is a 1% uncertainty in the number of collected POT [232]. Systematic uncertainties on the background estimation are negligible relative to the statistical uncertainties and are therefore not evaluated, similar to the HNL analysis.

9.7 Results

9.7.1 Data

The selection has been applied to ArgoNeuT's full anti-neutrino mode data-set, corresponding to an exposure of 1.25×10^{20} POT. In the data 0 events pass the selection, consistent with the expected background rate of 0.1 ± 0.1 events.

9.7.2 Exclusion limits

The absence of an observed signal allows an exclusion limit to be set on the axion phase-space. The limits are evaluated using a Bayesian approach with a uniform prior [326]. The 95% confidence level (CL) requirement when 0.1 ± 0.1 background events are expected and 0 are observed is 3.00 events. Figure 9.11 shows the exclusion constraint on the axion parameter space at 95% confidence level from 1.25×10^{20} POT at ArgoNeuT. The limits are evaluated for the two theory benchmarks dis-

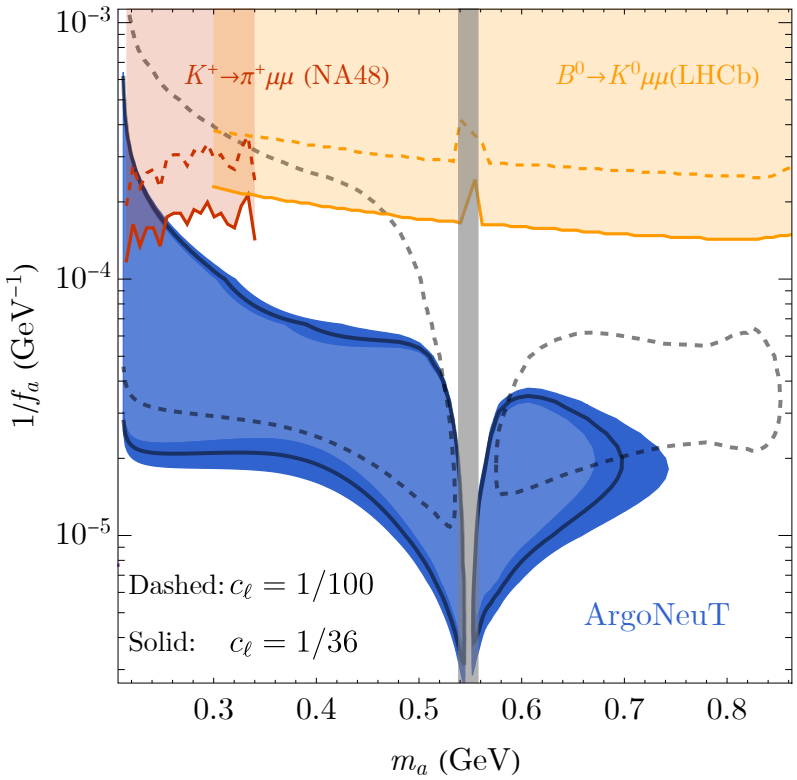

Figure 9.11: Constraints on the axion model parameter space at 95% CL from 1.25×10^{20} POT at ArgoNeuT (blue shaded region and black contours). The derived limits for $c_\ell = 1/36$ and $c_\ell = 1/100$ are shown by the solid and dashed contours. The impact of the theoretical uncertainties on the axion flux along with the experimental uncertainties on the expected constraint is shown by the dark blue band. The strongest existing limits are shown in red and orange. The grey-shaded band indicates a region with increased theoretical uncertainty around the η mass.

cussed in Section 2.3.3: $c_\ell = 1/36$ (blue shaded region and solid black contour) and $c_\ell = 1/100$ (dashed black contour). The impact of the theoretical and experimental uncertainties on the expected constraint is shown for the $c_\ell = 1/36$ case at $\pm 1\sigma$ by the dark blue band[1]. The strongest re-cast existing constraints from the LHCb experiment (orange) [148] and the NA48/2 experiment (red) [141] are shown. In each case, the bounds for the two different benchmark models are shown in the solid and dashed lines. The results of this search lead to a significant increase in the exclusion region of Heavy QCD Axions with masses above the di-muon threshold and below 1 GeV. Above 1 GeV, hadronic decays that ArgoNeuT is not sensitive to dominate. For the benchmark model with $c_\ell = 1/36$ the coverage of the axion decay constant is around $f_a \sim 50$ TeV in the mass range up to 0.65 GeV. For the benchmark model with $c_\ell = 1/100$ the coverage of the axion decay constant is around $f_a \sim 20$ TeV in the mass range up to 0.84 GeV. In both cases, ArgoNeuT provides a significant improvement to the constraints on the relevant parameter space compared with existing searches.

9.8 Conclusions and future

In this chapter a search for Heavy QCD Axions with masses 0.2-0.9 GeV performed using the ArgoNeuT experiment has been presented. The axions are produced in the NuMI neutrino beam from high-energy proton–fixed-target collisions in the NuMI target and absorber. They can then propagate to the ArgoNeuT detector, or the cavern upstream of ArgoNeuT, before decaying with the signature $a \to \mu^+\mu^-$. Dedicated simulations were performed to compute the axion production, along with the detector response and reconstruction. The selection developed for ArgoNeuT's Heavy Neutral Lepton search, presented in Chapter 8, was re-optimised for the differing kinematics of the muons resulting from the axion decays. As with the Heavy Neutral Lepton search, the combination of ArgoNeuT and the MINOS-ND enabled a near-background-free search to be performed for the axion decay signature. In the data, corresponding to an exposure to 1.25×10^{20} POT, zero passing events were observed consistent with the expected background of 0.1 ± 0.1 events. The results of this search exclude a significant previously unexplored region of parameter space for axion masses above the di-muon threshold. This is evaluated for two benchmark models. For the benchmark model with $c_\ell = 1/36$ the coverage of the axion decay constant is around $f_a \sim 50$ TeV in the mass range up to 0.65 GeV. For the benchmark model with $c_\ell = 1/100$ the coverage of the axion decay constant is around $f_a \sim 20$ TeV in the mass range up to 0.84 GeV. This analysis is the first

[1] The size of the uncertainty band for the $c_\ell = 1/100$ scenario would be similar, but is not shown to aid the clarity of the figure.

search for Heavy QCD Axions in a LArTPC detector. The search could be extended to various new Heavy QCD Axion models and paves the way for Heavy QCD Axion searches at future neutrino facilities. A paper describing this search is available as a preprint [6] and has been submitted to a journal.

Future searches in other LArTPC detectors are expected to have significant sensitivity to this and similar Heavy QCD Axion models. A detailed evaluation of the expected sensitivity of the DUNE near detector (DUNE-ND) was performed in Reference [108] for a similar model in the absence of the couplings to standard model leptons probed in this analysis. It found that the DUNE-ND will have significant sensitivity to Heavy QCD Axions with masses between 20 MeV and 2 GeV. This could be expanded to incorporate lepton couplings, and would be expected to significantly improve upon the constraints set by ArgoNeuT. The techniques developed to select the di-muon signature in this analysis, and in the previous Heavy Neutral Lepton search, could also be applied in the DUNE-ND. As described in Section 3.6.3, the DUNE-ND will consist of a LArTPC along with multiple magnetised muon spectrometers allowing a similar matching between sub-detectors to be performed. Searches for Heavy QCD Axions could also be performed in the SBN program detectors [3] and MicroBooNE [158] using either the on-axis 8 GeV Booster Neutrino Beam to probe lower mass axions, or the off-axis NuMI beam. A search for an alternative leptophilic axion-like particle model was also performed in the context of ArgoNeuT in Reference [328], recasting the Heavy Neutral Lepton search described in Chapter 8. It searched for a different, more niche, axion-like model that does not have the gluon-couplings required to resolve the strong CP problem. It was able to set significant new constraints on the model, but was performed without access to the ArgoNeuT simulation or data and was therefore less rigorous than a dedicated experimental search.

Chapter 10

Conclusions and outlook

In addition to their primary neutrino program, Liquid Argon Time Projection Chamber (LArTPC) neutrino detectors are powerful tools to probe beyond the standard model physics. In this thesis searches for two well-motivated dark-sector models using the ArgoNeuT experiment have been presented: Heavy Neutral Leptons and Heavy QCD Axions. These searches are both the first of their kind in LArTPC neutrino detectors and set constraints on significant regions of previously unexplored parameter space. Alongside these searches, this thesis has presented two models and methods that provide solutions to the simulation and computational challenges faced as LArTPC detectors increase in scale and complexity. The first of these was a new model to perform light simulation and the second a novel method to make use of high performance computing resources.

A novel semi-analytical model to predict the transport of scintillation-light photons in large scale LArTPC detectors was presented in Chapter 6. This model uses geometry along with parameterisations to predict both the number of arriving photons and their arrival time distribution. It is able to model the propagation of direct light and light reflected off wavelength-shifting reflector foils for argon, xenon and xenon-doped argon scintillation light. It is able to achieve a resolution of 10-15% or better, while being dramatically faster (∼30-70 times) than the equivalent Geant4 simulation. It can be applied in any large scale liquid argon detector, as well as in liquid xenon or xenon-doped argon detectors, with a simple tuning of the model parameters. The model presented in this chapter has been published in European Physical Journal C [4]. It is currently the standard fast optical simulation approach used in the SBND, DUNE horizontal drift and DUNE vertical drift detectors. The photon transport time parameterisation is also in use in the ICARUS detector, in conjunction with an optical look-up library. Additionally, the model is currently in development for the MicroBooNE detector.

A new method to run LArTPC simulation and data processing with LArSoft on high performance computers (HPCs) was presented in Chapter 7. The approach

was demonstrated using the Theta HPC at the Argonne Leadership Computing Facility (ALCF) at Argonne National Laboratory (ANL). Efficient scaling to the full size of the system was achieved, running over 250,000 simultaneous instances of LArSoft. An end-to-end production workflow was also demonstrated including sample generation on Theta, transfer of the output files from ANL to FNAL, and finally processing and archiving the files to tape. This was the first time a large scale LArTPC production sample had been generated on a HPC. The use of HPCs in this way could provide a potential solution to the computing challenges faced by LArTPC detectors as they increase in scale and complexity, in particular moving towards the multi-kiloton-scale DUNE far detectors. This work is the first step towards harnessing the upcoming exascale HPC systems such as Aurora and Frontier that can provide orders of magnitude greater computing resources than the existing grid-based methods. A paper describing this work is in preparation.

A search for Heavy Neutral Leptons (HNLs) with masses $\mathcal{O}(100)$ MeV performed using the ArgoNeuT experiment was presented in Chapter 8. The HNLs can be produced in the NuMI beam from high-energy proton–fixed-target collisions and can then decay with the signature $N \to \nu\mu^+\mu^-$. Three different mixing scenarios were considered: electron-coupled, $|U_{eN}|^2 \neq 0$; muon-coupled, $\left|U_{\mu N}\right|^2 \neq 0$; and tau-coupled, $|U_{\tau N}|^2 \neq 0$. In each case the mixing with the other two flavours was assumed to be zero. A novel technique was developed to identify the pairs of overlapping highly-forward-going oppositely-charged muons from these decays using the dE/dx and interaction topology in ArgoNeuT combined with matching to the magnetised MINOS near detector. The combination of the high-resolution LArTPC and the matching to a large magnetised muon-spectrometer enabled a near-background-free search to be performed. In the data, corresponding to an exposure to 1.25×10^{20} POT, zero passing events were observed consistent with the expected background of 0.4 ± 0.2 events. The results of this search lead to a significant increase in the exclusion region on the mixing angle $|U_{\tau N}|^2$ of tau-coupled Dirac HNLs with masses $m_N = 280$ - $970\,\text{MeV}$, assuming $|U_{eN}|^2 = \left|U_{\mu N}\right|^2 = 0$. New constraints were also evaluated for electron-coupled and muon-coupled HNLs, however these are weaker than existing limits. This analysis is the first search for Heavy Neutral Leptons decaying with the signature $N \to \nu\mu^+\mu^-$ in a LArTPC detector. It is also the first search for tau-coupled HNLs in a LArTPC detector. The techniques developed in this analysis could be applied in future HNL searches performed in larger mass LArTPC experiments, such as the DUNE near detector. The search presented in this chapter has been published in Physical Review Letters [5].

Finally, a search for Heavy QCD Axions with masses 0.2-0.9 GeV performed using the ArgoNeuT experiment was presented in Chapter 9. The axions can be produced in the NuMI beam from high-energy proton–fixed-target collisions and subsequently

decay with the signature $a \to \mu^+\mu^-$. The selection developed for ArgoNeuT's Heavy Neutral Lepton search, presented in Chapter 8, was re-optimised for the differing kinematics of the muons resulting from the axion decays. Using a combination of the high resolution calorimetric and topological reconstruction in ArgoNeuT and matching to the magnetised MINOS-ND muon spectrometer, a near-background-free search was also possible for this signature. In the data, corresponding to an exposure to 1.25×10^{20} POT, zero passing events were observed consistent with the expected background of 0.1 ± 0.1 events. The results of this search exclude a significant, previously unexplored, region of parameter space for axion masses above the di-muon threshold. This was evaluated for two benchmark models. For the benchmark model with axion-lepton coupling parameter $c_\ell = 1/36$ the coverage of the axion decay constant is around $f_a \sim 50$ TeV in the mass range up to 0.65 GeV. For the benchmark model with $c_\ell = 1/100$ the coverage of the axion decay constant is around $f_a \sim 20$ TeV in the mass range up to 0.84 GeV. This analysis is the first search for Heavy QCD Axions in a LArTPC detector. The search can be extended to various new Heavy QCD Axion models and paves the way for Heavy QCD Axion searches at future neutrino facilities. A paper describing this search is available as a preprint [6] and has been submitted to a journal.

The studies presented in this thesis provide significant steps both towards resolving the simulation and computational challenges faced as LArTPC detectors increase in scale and complexity and towards demonstrating the strength of LArTPC detectors in performing dark-sector searches. The light simulation model developed as part of this work provides the current state-of-the-art fast simulation of the propagation of scintillation light in multiple LArTPC detectors. The development of an approach to run production-grade simulation and reconstruction on a high performance computer provides the first step towards exploiting the exascale computing resources that will soon be available. The dark-sector searches for Heavy Neutral Leptons and Heavy QCD Axions demonstrate the strength of LArTPC detectors setting world-leading constraints despite being performed with the decade-old ArgoNeuT detector. These analyses are both the first of their kind in LArTPC neutrino detectors and pave the way for future dark-sector searches in upcoming detectors, including in the DUNE near detector.

Appendix A

Electron-coupled and muon-coupled Heavy Neutral Leptons

This appendix contains plots of the Heavy Neutral Lepton (HNL) kinematics and selection for the electron-coupled and muon-coupled scenarios. The plots are made using the $K^{\pm}$ and $D^{\pm}_{(s)}$ decay HNL samples, described in Section 8.1.

A.1 Kinematics

Figure A.1 shows the truth energies of muons originating from the decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The muons have average energies of $\langle E_{\mu^{\pm}} \rangle \sim 2.5\,\text{GeV}$ and $\langle E_{\mu^{\pm}} \rangle \sim 5.7\,\text{GeV}$, respectively.

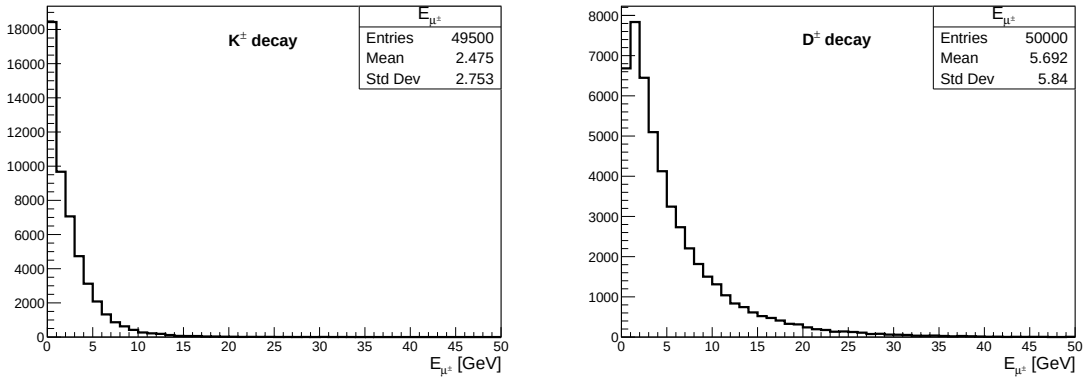

Figure A.1: Truth energies of muons originating from the decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right).

Figure A.2 shows the truth angles with respect to the beam direction of muons originating from the decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay

(left) and $D^{\pm}_{(s)}$ decay (right). The muons have average angles with respect to the beam direction of $\langle\theta_{beam}\rangle \sim 4.6^{\circ}$ and $\langle\theta_{beam}\rangle \sim 5.9^{\circ}$, respectively.

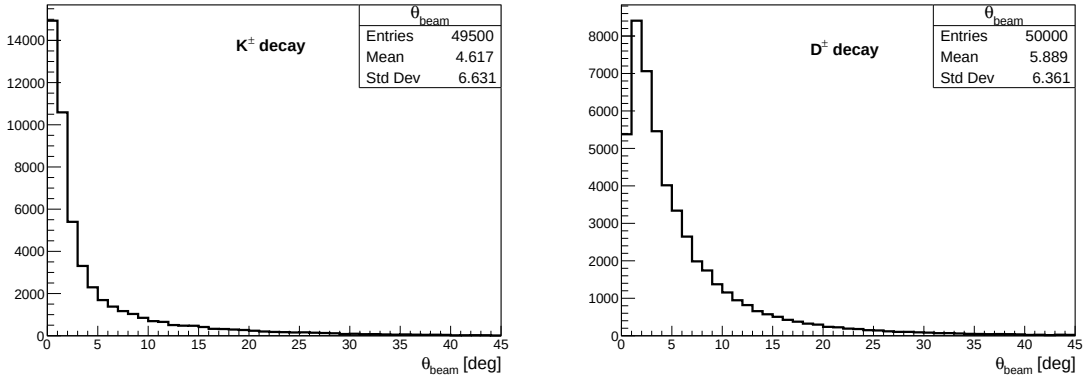

Figure A.2: Truth angles with respect to the beam direction of muons originating from the decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right).

Figure A.3 shows the truth opening angles between muons originating from the decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The muons have average opening angles of $\langle\theta_{opening}\rangle \sim 8.2^{\circ}$ and $\langle\theta_{opening}\rangle \sim 10.5^{\circ}$, respectively.

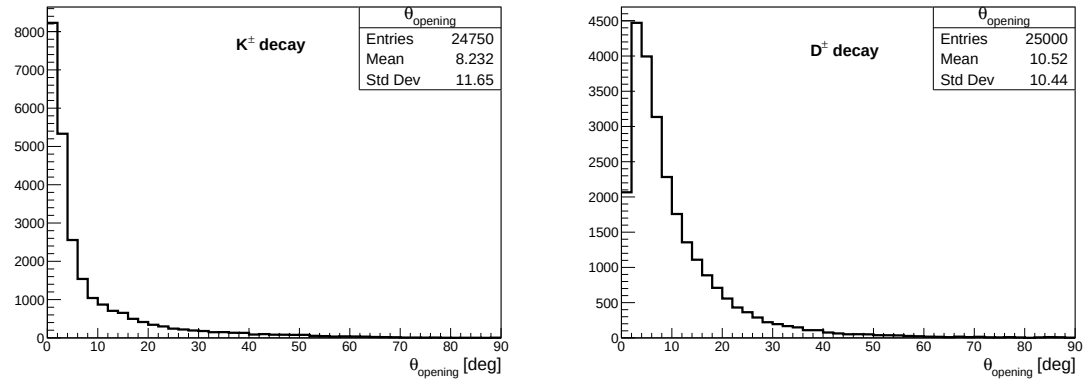

Figure A.3: Truth opening angles between muons originating from the decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right).

A.2 Selection

Figure A.4 shows the fraction of hits associated with tracks, f_{hits}, for decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $f_{hits} = 0.8$.

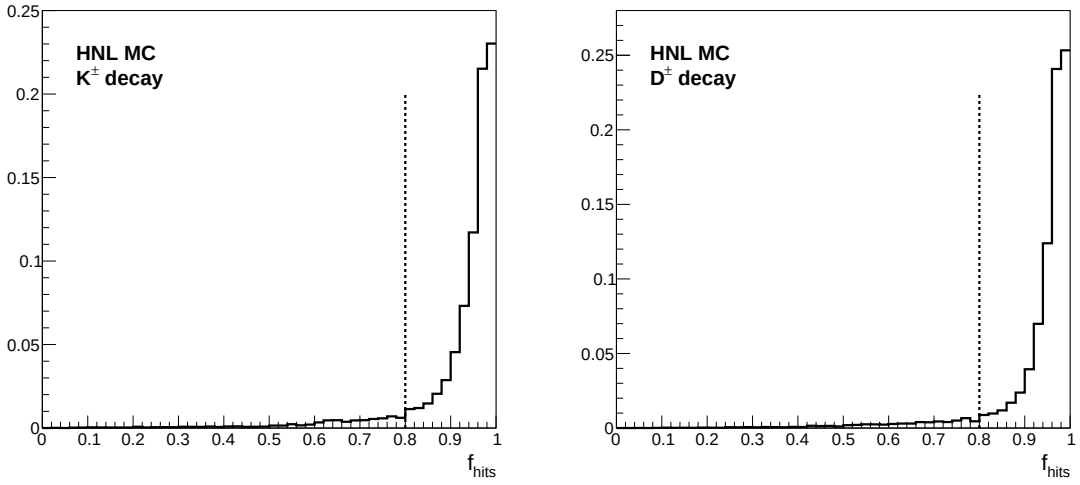

Figure A.4: Fraction of hits associated with tracks, f_{hits}, for electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $f_{hits} = 0.8$.

Figure A.5 shows the reconstructed opening angle, $\theta_{opening}$, between pairs of muons originating from decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $\theta_{opening} = 10^{\circ}$.

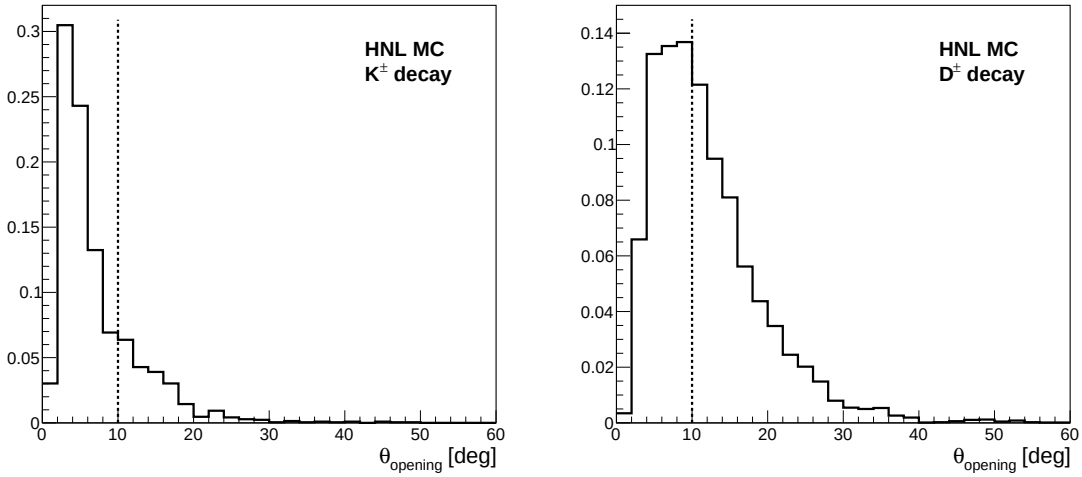

Figure A.5: Reconstructed opening angle, $\theta_{opening}$, between pairs of muons originating from decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $\theta_{opening} = 10^{\circ}$.

Figure A.6 shows the average track dE/dx for decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the threshold applied at $dE/dx = 3.1\,\mathrm{MeV/cm}$.

Figure A.7 shows the reconstructed angle with respect to the beam direction,

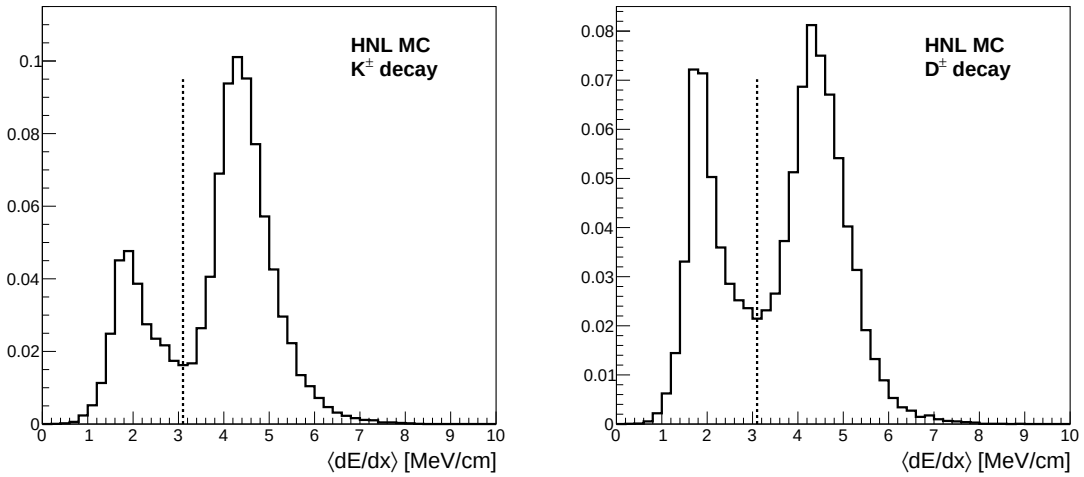

Figure A.6: Average track dE/dx for decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the threshold applied at $dE/dx = 3.1\,\mathrm{MeV/cm}$.

θ_{beam}, of muons originating from decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $\theta_{beam} = 15^{\circ}$.

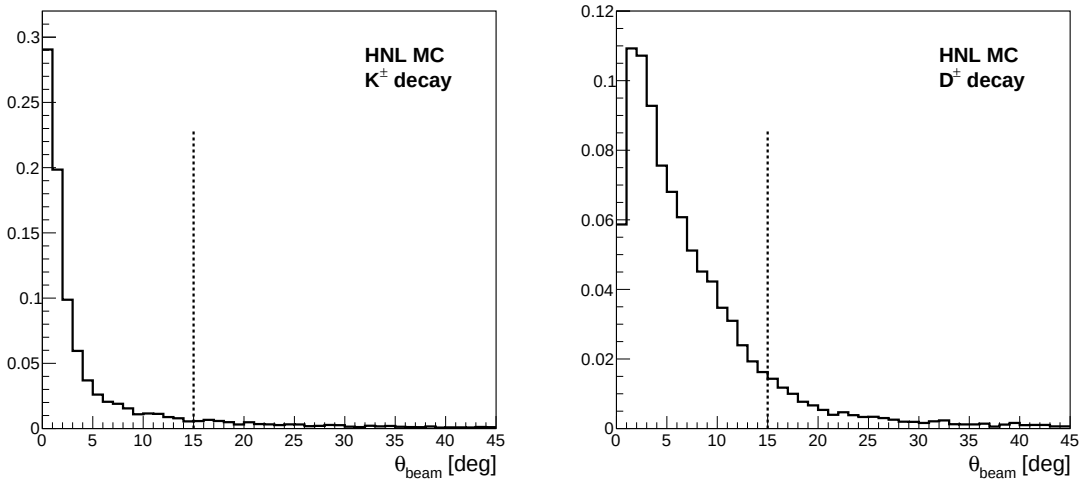

Figure A.7: Reconstructed angle with respect to the beam direction, θ_{beam}, of muons originating from decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $\theta_{beam} = 15^{\circ}$.

Figure A.8 shows the double-MIP MINOS-ND matching radial offset, r_{diff}, for decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right) reconstructed with double-MIP-like topologies. The dashed lines show the double-MIP radial matching tolerance of $r_{diff} = 30\,\mathrm{cm}$.

Figure A.9 shows the double-MIP MINOS-ND matching angular offset, θ_{diff},

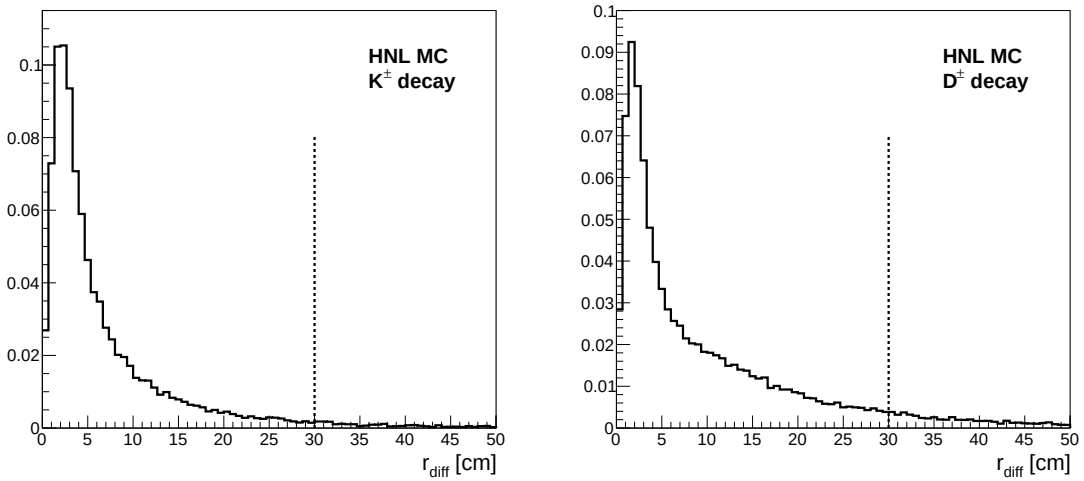

Figure A.8: Double-MIP MINOS-ND matching radial offset, r_{diff}, for decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right) reconstructed with double-MIP-like topologies. The dashed lines show the double-MIP radial matching tolerance of $r_{diff} = 30\,\text{cm}$.

for decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right) reconstructed with double-MIP-like topologies. The dashed lines show the double-MIP angular matching tolerance of $\theta_{diff} = 24.4^{\circ}$.

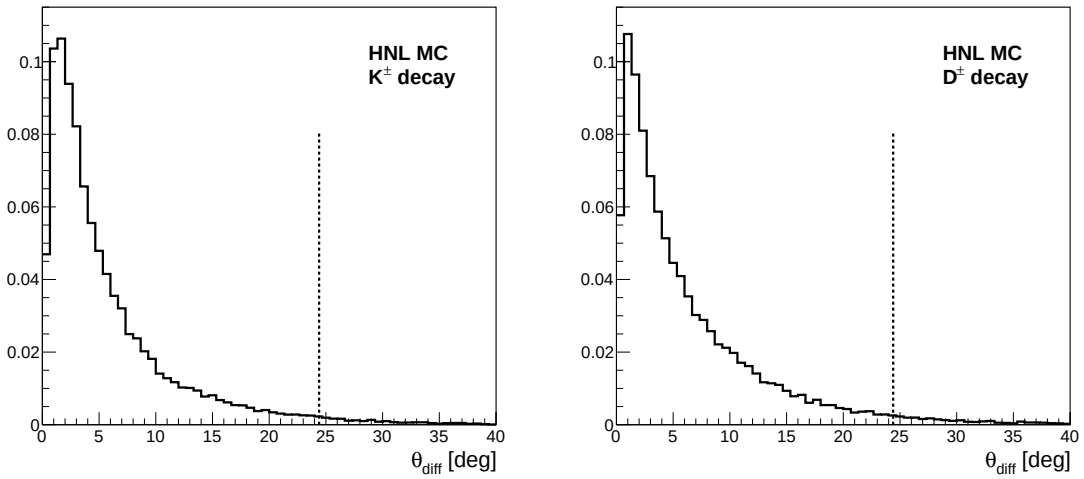

Figure A.9: Double-MIP MINOS-ND matching angular offset, θ_{diff}, for decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right) reconstructed with double-MIP-like topologies. The dashed lines show the double-MIP angular matching tolerance of $\theta_{diff} = 24.4^{\circ}$.

Figure A.10 shows the MINOS-ND track length, $L_{MINOS-ND}$, of muons originating from decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $L_{MINOS-ND} = 1\,\mathrm{m}$.

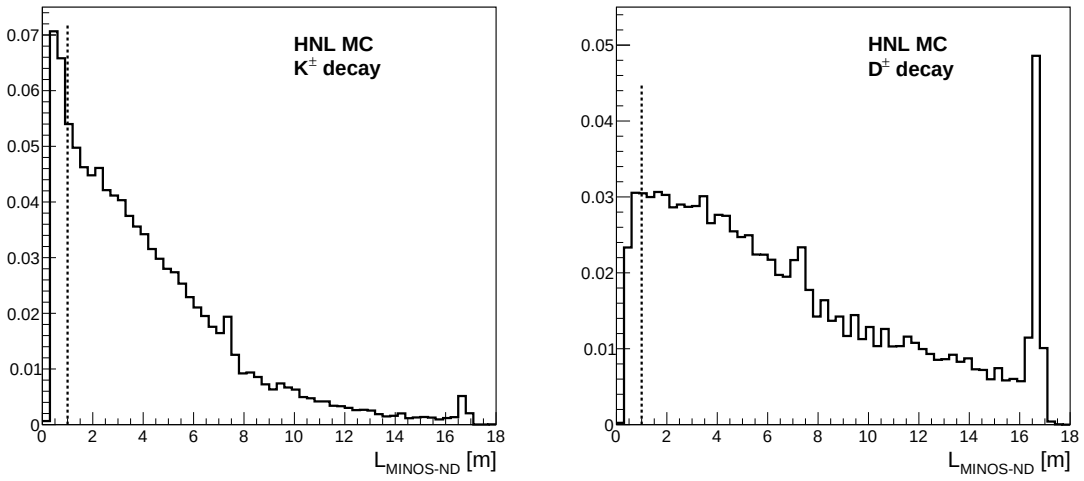

Figure A.10: MINOS-ND track length, $L_{MINOS-ND}$, of muons originating from decays of electron/muon-coupled HNLs produced via $K^{\pm}$ decay (left) and $D^{\pm}_{(s)}$ decay (right). The dashed lines show the cut applied at $L_{MINOS-ND} = 1\,\mathrm{m}$.

Appendix B

Additional Heavy QCD Axion kinematic distributions

This appendix contains additional plots of the Heavy QCD Axion kinematics. The plots are made using the Heavy QCD Axion samples corresponding to production via η and η' mixing, described in Section 9.1.

Figure B.1 shows the truth energies of the muons originating from decays of axions produced via η mixing (left) and η' mixing (right). In each case the axions are produced at the beam target and two different masses are shown, $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue). For both production mechanisms and masses, the resulting muons are highly energetic with average energies $\langle E_{\mu^\pm} \rangle \sim 26\,\text{GeV}$.

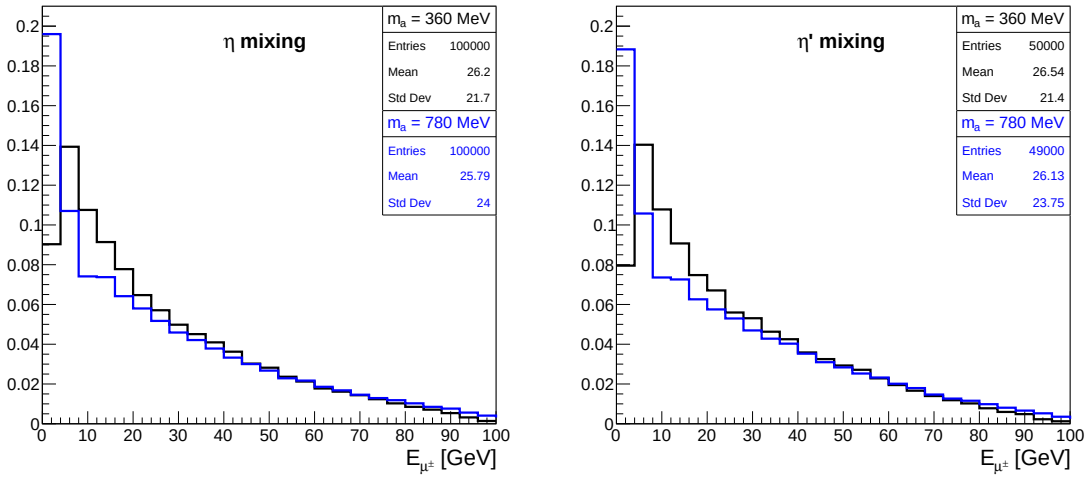

Figure B.1: Truth energies of the muons originating from decays of axions produced via η mixing (left) and η' mixing (right). In each case the axions are produced at the beam target and two different masses are shown, $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue).

Figure B.2 shows the truth angles with respect to the beam direction of the muons originating from decays of axions produced via η mixing (left) and η' mixing (right). In each case the axions are produced at the beam target and two different masses are shown, $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue). In the η mixing case, the muons have average angles with respect to the beam direction of $\langle\theta_{beam}\rangle \sim 0.6^\circ$ for $m_a = 360\,\text{MeV}$ and $\langle\theta_{beam}\rangle \sim 1.5^\circ$ for $m_a = 780\,\text{MeV}$. In the η' mixing case, the muons have average angles with respect to the beam direction of $\langle\theta_{beam}\rangle \sim 0.5^\circ$ for $m_a = 360\,\text{MeV}$ and $\langle\theta_{beam}\rangle \sim 1.7^\circ$ for $m_a = 780\,\text{MeV}$.

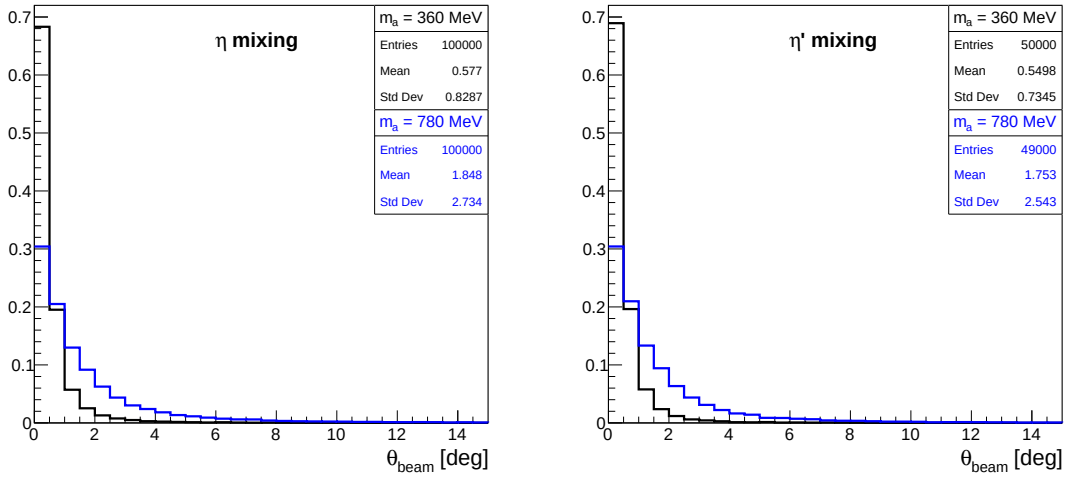

Figure B.2: Truth angles with respect to the beam direction of the muons originating from decays of axions produced via η mixing (left) and η' mixing (right). In each case the axions are produced at the beam target and two different masses are shown, $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue).

Figure B.3 shows the truth opening angles between the muons originating from decays of axions produced via η mixing (left) and η' mixing (right). In each case the axions are produced at the beam target and two different masses are shown, $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue). In the η mixing case, the muons have average opening angles of $\langle\theta_{opening}\rangle \sim 1.0^\circ$ for $m_a = 360\,\text{MeV}$ and $\langle\theta_{opening}\rangle \sim 3.5^\circ$ for $m_a = 780\,\text{MeV}$. In the η' mixing case, the muons have average opening angles of $\langle\theta_{opening}\rangle \sim 0.9^\circ$ for $m_a = 360\,\text{MeV}$ and $\langle\theta_{opening}\rangle \sim 3.4^\circ$ for $m_a = 780\,\text{MeV}$.

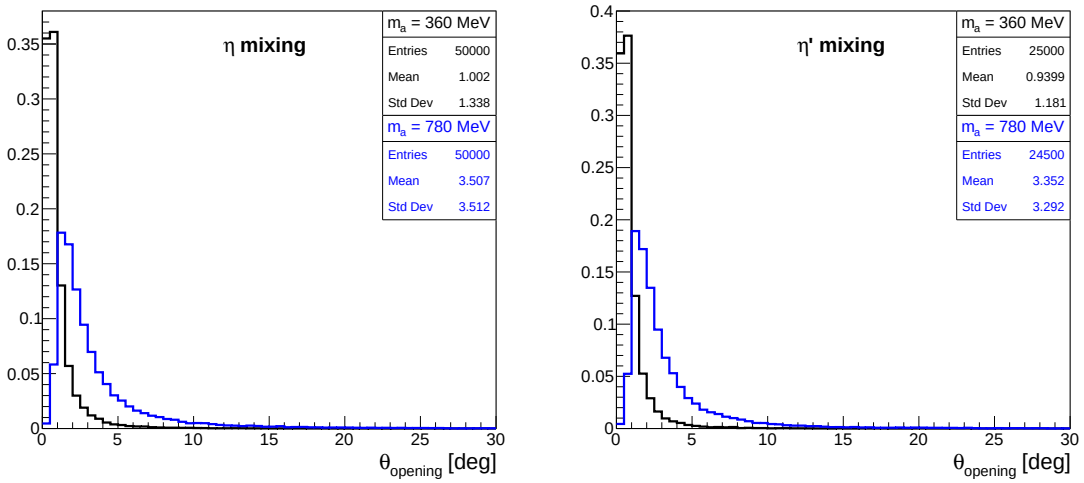

Figure B.3: Truth opening angles between the muons originating axion decays produced via η mixing (left) and η' mixing (right). In each case the axions are produced at the beam target and two different masses are shown, $m_a = 360\,\text{MeV}$ (black) and $m_a = 780\,\text{MeV}$ (blue).

Bibliography

[1] B. Abi et al. (DUNE). "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume I Introduction to DUNE". *JINST* 15.08 (2020), T08008. arXiv: 2002.02967 [physics.ins-det].

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

[3] P. A. N. Machado, O. Palamara, and D. W. Schmitz. "The Short-Baseline Neutrino Program at Fermilab". *Ann. Rev. Nucl. Part. Sci.* 69 (2019), pp. 363–387. arXiv: 1903.04608 [hep-ex].

[4] D. Garcia-Gamez, P. Green, and A. M. Szelc. "Predicting Transport Effects of Scintillation Light Signals in Large-Scale Liquid Argon Detectors". *Eur. Phys. J. C* 81.4 (2021), p. 349. arXiv: 2010.00324 [physics.ins-det].

[5] R. Acciarri et al. (ArgoNeuT). "New Constraints on Tau-Coupled Heavy Neutral Leptons with Masses mN=280–970 MeV". *Phys. Rev. Lett.* 127.12 (2021), p. 121801. arXiv: 2106.13684 [hep-ex].

[6] R. Acciarri et al. (ArgoNeuT). "First constraints on Heavy QCD Axions with a Liquid Argon Time Projection Chamber using the ArgoNeuT Experiment" (2022). arXiv: 2207.08448 [hep-ex].

[7] W. Pauli. "Pauli letter collection: letter to Lise Meitner". Typed copy.

[8] E. Fermi. "Tentativo di una teoria dell'emissione dei raggi beta". *Ric. Sci.* 4 (1933), pp. 491–495.

[9] H. Bethe and R. Peierls. "The 'neutrino'". *Nature* 133 (1934), p. 532.

[10] F. Reines and C. L. Cowan. "Detection of the free neutrino". *Phys. Rev.* 92 (1953), pp. 830–831.

[11] C. L. Cowan et al. "Detection of the free neutrino: A Confirmation". *Science* 124 (1956), pp. 103–104.

[12] J. C. Street and E. C. Stevenson. "New Evidence for the Existence of a Particle of Mass Intermediate Between the Proton and Electron". *Phys. Rev.* 52 (1937), pp. 1003–1004.

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

[14] M. L. Perl et al. "Evidence for Anomalous Lepton Production in e+ - e- Annihilation". *Phys. Rev. Lett.* 35 (1975), pp. 1489–1492.

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

[16] S. Schael et al. (ALEPH, DELPHI, L3, OPAL, SLD, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group). "Precision electroweak measurements on the Z resonance". *Phys. Rept.* 427 (2006), pp. 257–454. arXiv: hep-ex/0509008.

[17] B. Pontecorvo. "Inverse beta processes and nonconservation of lepton charge". *Zh. Eksp. Teor. Fiz.* 34 (1957), p. 247.

[18] M. Gell-Mann and A. Pais. "Behavior of neutral particles under charge conjugation". *Phys. Rev.* 97 (1955), pp. 1387–1389.

[19] Z. Maki, M. Nakagawa, and S. Sakata. "Remarks on the unified model of elementary particles". *Prog. Theor. Phys.* 28 (1962), pp. 870–880.

[20] B. Pontecorvo. "Neutrino Experiments and the Problem of Conservation of Leptonic Charge". *Zh. Eksp. Teor. Fiz.* 53 (1967), pp. 1717–1725.

[21] R. Davis. "Solar neutrinos. II: Experimental". *Phys. Rev. Lett.* 12 (1964), pp. 303–305.

[22] J. N. Bahcall. "Solar neutrinos. I: Theoretical". *Phys. Rev. Lett.* 12 (1964), pp. 300–302.

[23] V. N. Gribov and B. Pontecorvo. "Neutrino astronomy and lepton charge". *Phys. Lett. B* 28 (1969), p. 493.

[24] Q. R. Ahmad et al. (SNO). "Measurement of the rate of $\nu_e + d \to p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory". *Phys. Rev. Lett.* 87 (2001), p. 071301. arXiv: nucl-ex/0106015.

[25] Q. R. Ahmad et al. (SNO). "Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory". *Phys. Rev. Lett.* 89 (2002), p. 011301. arXiv: nucl-ex/0204008.

[26] Y. Fukuda et al. (Super-Kamiokande). "Evidence for oscillation of atmospheric neutrinos". *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. arXiv: hep-ex/9807003.

[27] P. A. Zyla et al. (Particle Data Group). "Review of Particle Physics". *Prog. Theor. Exp. Phys.* 2020.8 (2020), p. 083C01.

[28] C. Athanassopoulos et al. (LSND). "Evidence for $\overline{\nu}_\mu \to \overline{\nu}_e$ Oscillations from the LSND Experiment at the Los Alamos Meson Physics Facility". *Phys. Rev. Lett.* 77 (15 1996), pp. 3082–3085.

[29] A. A. Aguilar-Arevalo et al. (MiniBooNE). "Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations". *Phys. Rev. Lett.* 105 (2010), p. 181801. arXiv: 1007.1150 [hep-ex].

[30] A. A. Aguilar-Arevalo et al. (MiniBooNE). "Improved Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment". *Phys. Rev. Lett.* 110 (2013), p. 161801. arXiv: 1303.2588 [hep-ex].

[31] A. A. Aguilar-Arevalo et al. (MiniBooNE). "Significant Excess of ElectronLike Events in the MiniBooNE Short-Baseline Neutrino Experiment". *Phys. Rev. Lett.* 121.22 (2018), p. 221801. arXiv: 1805.12028 [hep-ex].

[32] M. A. Acero et al. (NOvA). "An Improved Measurement of Neutrino Oscillation Parameters by the NOvA Experiment" (2021). arXiv: 2108.08219 [hep-ex].

[33] K. Abe et al. (T2K). "The T2K Experiment". *Nucl. Instrum. Meth. A* 659 (2011), pp. 106–135. arXiv: 1106.1238 [physics.ins-det].

[34] K. J. Kelly. "The MiniBooNE Anomaly". Presented at the Mini SBN-Theory workshop. 2021.

[35] M. Aker et al. (KATRIN). "Direct neutrino-mass measurement with sub-electronvolt sensitivity". *Nature Phys.* 18.2 (2022), pp. 160–166. arXiv: 2105.08533 [hep-ex].

[36] A. Loureiro et al. "On The Upper Bound of Neutrino Masses from Combined Cosmological Observations and Particle Physics Experiments". *Phys. Rev. Lett.* 123.8 (2019), p. 081301. arXiv: 1811.02578 [astro-ph.CO].

[37] Y. Abe et al. (Double Chooz). "Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment". *Phys. Rev. Lett.* 108 (2012), p. 131801. arXiv: 1112.6353 [hep-ex].

[38] F. P. An et al. (Daya Bay). "Observation of electron-antineutrino disappearance at Daya Bay". *Phys. Rev. Lett.* 108 (2012), p. 171803. arXiv: 1203.1669 [hep-ex].

[39] J. K. Ahn et al. (RENO). "Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment". *Phys. Rev. Lett.* 108 (2012), p. 191802. arXiv: 1204.0626 [hep-ex].

[40] L. Wolfenstein. "Neutrino Oscillations in Matter". *Phys. Rev. D* 17 (1978), pp. 2369–2374.

[41] S. P. Mikheyev and A. Yu. Smirnov. "Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos". *Sov. J. Nucl. Phys.* 42 (1985), pp. 913–917.

[42] S. N. Ahmed et al. (SNO). "Measurement of the total active B-8 solar neutrino flux at the Sudbury Neutrino Observatory with enhanced neutral current sensitivity". *Phys. Rev. Lett.* 92 (2004), p. 181301. arXiv: nucl-ex/0309004.

[43] K. Eguchi et al. (KamLAND). "First results from KamLAND: Evidence for reactor anti-neutrino disappearance". *Phys. Rev. Lett.* 90 (2003), p. 021802. arXiv: hep-ex/0212021.

[44] S. F. King and C. Luhn. "Neutrino Mass and Mixing with Discrete Symmetry". *Rept. Prog. Phys.* 76 (2013), p. 056201. arXiv: 1301.1340 [hep-ph].

[45] K. Abe et al. (T2K). "Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations". *Nature* 580.7803 (2020). [Erratum: Nature 583, E16 (2020)], pp. 339–344. arXiv: 1910.03887 [hep-ex].

[46] K. Abe et al. (Hyper-Kamiokande). "Hyper-Kamiokande Design Report" (2018). arXiv: 1805.04163 [physics.ins-det].

[47] P. A. M. Dirac. "The quantum theory of the electron". *Proc. Roy. Soc. Lond. A* 117 (1928), pp. 610–624.

[48] E. Majorana. "Teoria simmetrica dell'elettrone e del positrone". *Nuovo Cim.* 14 (1937), pp. 171–184.

[49] M. Goldhaber, L. Grodzins, and A. W. Sunyar. "Helicity of Neutrinos". *Phys. Rev.* 109 (1958), pp. 1015–1017.

[50] D. Gorbunov and M. Shaposhnikov. "How to find neutral leptons of the νMSM?" *JHEP* 10 (2007). [Erratum: JHEP 11, 101 (2013)], p. 015. arXiv: 0705.1729 [hep-ph].

[51] J. Schechter and J. W. F. Valle. "Neutrinoless Double beta Decay in SU(2) x U(1) Theories". *Phys. Rev. D* 25 (1982), p. 2951.

[52] M. Agostini et al. (GERDA). "Final Results of GERDA on the Search for Neutrinoless Double-β Decay". *Phys. Rev. Lett.* 125.25 (2020), p. 252502. arXiv: 2009.06079 [nucl-ex].

[53] S. Abe et al. (KamLAND-Zen). "First Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen" (2022). arXiv: 2203.02139 [hep-ex].

[54] D. Q. Adams et al. (CUORE). "New direct limit on neutrinoless double beta decay half-life of ^{128}Te with CUORE" (2022). arXiv: 2205.03132 [nucl-ex].

[55] K. N. Abazajian et al. "Light Sterile Neutrinos: A White Paper" (2012). arXiv: 1204.5379 [hep-ph].

[56] F. Bezrukov, H. Hettmansperger, and M. Lindner. "keV sterile neutrino Dark Matter in gauge extensions of the Standard Model". *Phys. Rev. D* 81 (2010), p. 085032. arXiv: 0912.4415 [hep-ph].

[57] P. Minkowski. "$\mu \to e\gamma$ at a Rate of One Out of 10^9 Muon Decays?" *Phys. Lett. B* 67 (1977), pp. 421–428.

[58] M. Gell-Mann, P. Ramond, and R. Slansky. "Complex Spinors and Unified Theories". *Conf. Proc. C* 790927 (1979), pp. 315–321. arXiv: 1306.4669 [hep-th].

[59] T. Yanagida. "Horizontal gauge symmetry and masses of neutrinos". *Conf. Proc. C* 7902131 (1979). Ed. by O. Sawada and A. Sugamoto, pp. 95–99.

[60] S. L. Glashow. "The Future of Elementary Particle Physics". *NATO Sci. Ser. B* 61 (1980), p. 687.

[61] R. N. Mohapatra and G. Senjanovic. "Neutrino Mass and Spontaneous Parity Nonconservation". *Phys. Rev. Lett.* 44 (1980), p. 912.

[62] J. Schechter and J. W. F. Valle. "Neutrino Masses in SU(2) x U(1) Theories". *Phys. Rev. D* 22 (1980), p. 2227.

[63] R. Foot et al. "Seesaw Neutrino Masses Induced by a Triplet of Leptons". *Z. Phys. C* 44 (1989), p. 441.

[64] P. Abratenko et al. (MicroBooNE). "Search for an Excess of Electron Neutrino Interactions in MicroBooNE Using Multiple Final-State Topologies". *Phys. Rev. Lett.* 128.24 (2022), p. 241801. arXiv: 2110.14054 [hep-ex].

[65] P. Abratenko et al. (MicroBooNE). "Search for Neutrino-Induced Neutral-Current Δ Radiative Decay in MicroBooNE and a First Test of the MiniBooNE Low Energy Excess under a Single-Photon Hypothesis". *Phys. Rev. Lett.* 128.11 (2022), p. 111801. arXiv: 2110.00409 [hep-ex].

[66] M. Drewes et al. "A White Paper on keV Sterile Neutrino Dark Matter". *JCAP* 01 (2017), p. 025. arXiv: 1602.04816 [hep-ph].

[67] A. Boyarsky et al. "Sterile neutrino Dark Matter". *Prog. Part. Nucl. Phys.* 104 (2019), pp. 1–45. arXiv: 1807.07938 [hep-ph].

[68] A. Boyarsky et al. "Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster". *Phys. Rev. Lett.* 113 (2014), p. 251301. arXiv: 1402.4119 [astro-ph.CO].

[69] S. Dodelson and L. M. Widrow. "Sterile-neutrinos as dark matter". *Phys. Rev. Lett.* 72 (1994), pp. 17–20. arXiv: hep-ph/9303287.

[70] K. Abazajian, G. M. Fuller, and M. Patel. "Sterile neutrino hot, warm, and cold dark matter". *Phys. Rev. D* 64 (2001), p. 023501. arXiv: astro-ph/0101524.

[71] T. Venumadhav et al. "Sterile neutrino dark matter: Weak interactions in the strong coupling epoch". *Phys. Rev. D* 94.4 (2016), p. 043515. arXiv: 1507.06655 [astro-ph.CO].

[72] J. Ghiglieri and M. Laine. "Improved determination of sterile neutrino dark matter spectrum". *JHEP* 11 (2015), p. 171. arXiv: 1506.06752 [hep-ph].

[73] J. Ghiglieri and M. Laine. "Sterile neutrino dark matter via GeV-scale leptogenesis?" *JHEP* 07 (2019), p. 078. arXiv: 1905.08814 [hep-ph].

[74] A. Boyarsky et al. "Constraints on sterile neutrino as a dark matter candidate from the diffuse x-ray background". *Mon. Not. Roy. Astron. Soc.* 370 (2006), pp. 213–218. arXiv: astro-ph/0512509.

[75] M. Viel et al. "How cold is cold dark matter? Small scales constraints from the flux power spectrum of the high-redshift Lyman-alpha forest". *Phys. Rev. Lett.* 100 (2008), p. 041304. arXiv: 0709.0131 [astro-ph].

[76] M. Fukugita and T. Yanagida. "Baryogenesis Without Grand Unification". *Phys. Lett. B* 174 (1986), pp. 45–47.

[77] E. K. Akhmedov, V. A. Rubakov, and A. Yu. Smirnov. "Baryogenesis via neutrino oscillations". *Phys. Rev. Lett.* 81 (1998), pp. 1359–1362. arXiv: hep-ph/9803255.

[78] S. Davidson, E. Nardi, and Y. Nir. "Leptogenesis". *Phys. Rept.* 466 (2008), pp. 105–177. arXiv: 0802.2962 [hep-ph].

[79] E. K. Akhmedov and A. Y. Smirnov. "Paradoxes of neutrino oscillations". *Phys. Atom. Nucl.* 72 (2009), pp. 1363–1381. arXiv: 0905.1903 [hep-ph].

[80] T. Asaka, S. Blanchet, and M. Shaposhnikov. "The nuMSM, dark matter and neutrino masses". *Phys. Lett. B* 631 (2005), pp. 151–156. arXiv: hep-ph/0503065.

[81] T. Asaka and M. Shaposhnikov. "The νMSM, dark matter and baryon asymmetry of the universe". *Phys. Lett. B* 620 (2005), pp. 17–26. arXiv: hep-ph/0505013.

[82] J. Tastet, O. Ruchayskiy, and I. Timiryasov. "Why interpretation matters for BSM searches: a case study with Heavy Neutral Leptons at ATLAS". *PoS* EPS-HEP2021 (2022), p. 703. arXiv: 2110.11907 [hep-ph].

[83] P. D. Bolton, F. F. Deppisch, and P. S. Bhupal D. "Neutrinoless double beta decay versus other probes of heavy sterile neutrinos". *JHEP* 03 (2020), p. 170. arXiv: 1912.03058 [hep-ph].

[84] J. M. Berryman et al. "Searches for Decays of New Particles in the DUNE Multi-Purpose Near Detector". *JHEP* 02 (2020), p. 174. arXiv: 1912.07622 [hep-ph].

[85] P. Coloma et al. "GeV-scale neutrinos: interactions with mesons and DUNE sensitivity". *Eur. Phys. J. C* 81.1 (2021), p. 78. arXiv: 2007.03701 [hep-ph].

[86] P. Abratenko et al. (MicroBooNE). "Search for Heavy Neutral Leptons Decaying into Muon-Pion Pairs in the MicroBooNE Detector". *Phys. Rev. D* 101.5 (2020), p. 052001. arXiv: 1911.10545 [hep-ex].

[87] K. Abe et al. (T2K). "Search for heavy neutrinos with the T2K near detector ND280". *Phys. Rev. D* 100.5 (2019), p. 052006. arXiv: 1902.07598 [hep-ex].

[88] F. Bergsma et al. (CHARM). "A Search for Decays of Heavy Neutrinos in the Mass Range 0.5-{GeV} to 2.8-{GeV}". *Phys. Lett. B* 166 (1986), pp. 473–478.

[89] G. Bernardi et al. "Search for Neutrino Decay". *Phys. Lett. B* 166 (1986), pp. 479–483.

[90] G. Bernardi et al. "FURTHER LIMITS ON HEAVY NEUTRINO COUPLINGS". *Phys. Lett. B* 203 (1988), pp. 332–334.

[91] E. Cortina Gil et al. (NA62). "Search for heavy neutral lepton production in K^+ decays to positrons". *Phys. Lett. B* 807 (2020), p. 135599. arXiv: 2005.09575 [hep-ex].

[92] A. V. Artamonov et al. (E949). "Study of the decay $K^+ \to \pi^+ \nu\bar{\nu}$ in the momentum region $140 < P_\pi < 199$ MeV/c". *Phys. Rev. D* 79 (2009), p. 092004. arXiv: 0903.0030 [hep-ex].

[93] E. Cortina Gil et al. (NA62). "Search for K+ decays to a muon and invisible particles". *Phys. Lett. B* 816 (2021), p. 136259. arXiv: 2101.12304 [hep-ex].

[94] A. Vaitaitis et al. (NuTeV, E815). "Search for neutral heavy leptons in a high-energy neutrino beam". *Phys. Rev. Lett.* 83 (1999), pp. 4943–4946. arXiv: hep-ex/9908011.

[95] K. J. Kelly and P. A. N. Machado. "MicroBooNE experiment, NuMI absorber, and heavy neutral leptons". *Phys. Rev. D* 104.5 (2021), p. 055015. arXiv: 2106.06548 [hep-ph].

[96] J. Orloff, A. N. Rozanov, and C. Santoni. "Limits on the mixing of tau neutrino to heavy neutrinos". *Phys. Lett. B* 550 (2002), pp. 8–15. arXiv: hep-ph/0208075.

[97] P. Abreu et al. (DELPHI). "Search for neutral heavy leptons produced in Z decays". *Z. Phys. C* 74 (1997). [Erratum: Z.Phys.C 75, 580 (1997)], pp. 57–71.

[98] R. D. Peccei and H. R. Quinn. "CP Conservation in the Presence of Instantons". *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443.

[99] R. D. Peccei and H. R. Quinn. "Constraints Imposed by CP Conservation in the Presence of Instantons". *Phys. Rev. D* 16 (1977), pp. 1791–1797.

[100] G. 't Hooft. "Symmetry Breaking Through Bell-Jackiw Anomalies". *Phys. Rev. Lett.* 37 (1976). Ed. by M. A. Shifman, pp. 8–11.

[101] C. A. Baker et al. "An Improved experimental limit on the electric dipole moment of the neutron". *Phys. Rev. Lett.* 97 (2006), p. 131801. arXiv: hep-ex/0602020 [hep-ex].

[102] J. M. Pendlebury et al. "Revised experimental upper limit on the electric dipole moment of the neutron". *Phys. Rev. D* 92.9 (2015), p. 092003. arXiv: 1509.04411 [hep-ex].

[103] B. Graner et al. "Reduced Limit on the Permanent Electric Dipole Moment of Hg199". *Phys. Rev. Lett.* 116.16 (2016). [Erratum: Phys.Rev.Lett. 119, 119901 (2017)], p. 161601. arXiv: 1601.04339 [physics.atom-ph].

[104] C. Abel et al. (nEDM). "Measurement of the permanent electric dipole moment of the neutron". *Phys. Rev. Lett.* 124.8 (2020), p. 081803. arXiv: 2001.11966 [hep-ex].

[105] S. Weinberg. "A New Light Boson?" *Phys. Rev. Lett.* 40 (1978), pp. 223–226.

[106] F. Wilczek. "Problem of Strong P and T Invariance in the Presence of Instantons". *Phys. Rev. Lett.* 40 (1978), pp. 279–282.

[107] G. Grilli di Cortona et al. "The QCD axion, precisely". *JHEP* 01 (2016), p. 034. arXiv: 1511.02867 [hep-ph].

[108] K. J. Kelly, S. Kumar, and Z. Liu. "Heavy axion opportunities at the DUNE near detector". *Phys. Rev. D* 103.9 (2021), p. 095002. arXiv: 2011.05995 [hep-ph].

[109] J. R. Ellis and K. A. Olive. "Constraints on Light Particles From Supernova Sn1987a". *Phys. Lett.* B193 (1987), p. 525.

[110] G. Raffelt and D. Seckel. "Bounds on Exotic Particle Interactions from SN 1987a". *Phys. Rev. Lett.* 60 (1988), p. 1793.

[111] M. S. Turner. "Axions from SN 1987a". *Phys. Rev. Lett.* 60 (1988), p. 1797.

[112] R. Mayle et al. "Constraints on Axions from SN 1987a". *Phys. Lett.* B203 (1988), pp. 188–196.

[113] G. Raffelt. "Astrophysical axion bounds". *Lect. Notes Phys.* 741 (2008), pp. 51–71. arXiv: hep-ph/0611350 [hep-ph].

[114] J. H. Chang, R. Essig, and S. D. McDermott. "Supernova 1987A Constraints on Sub-GeV Dark Sectors, Millicharged Particles, the QCD Axion, and an Axion-like Particle". *JHEP* 09 (2018), p. 051. arXiv: 1803.00993 [hep-ph].

[115] P. Carenza et al. "Improved axion emissivity from a supernova via nucleon-nucleon bremsstrahlung". *JCAP* 10.10 (2019). [Erratum: JCAP 05, E01 (2020)], p. 016. arXiv: 1906.11844 [hep-ph].

[116] M. Buschmann et al. "Upper Limit on the QCD Axion Mass from Isolated Neutron Star Cooling". *Phys. Rev. Lett.* 128.9 (2022), p. 091102. arXiv: 2111.09892 [hep-ph].

[117] J. Preskill, M. B. Wise, and F. Wilczek. "Cosmology of the Invisible Axion". *Phys. Lett. B* 120 (1983). Ed. by M. A. Srednicki, pp. 127–132.

[118] L. F. Abbott and P. Sikivie. "A Cosmological Bound on the Invisible Axion". *Phys. Lett. B* 120 (1983). Ed. by M. A. Srednicki, pp. 133–136.

[119] M. Dine and W. Fischler. "The Not So Harmless Axion". *Phys. Lett. B* 120 (1983). Ed. by M. A. Srednicki, pp. 137–141.

[120] L. Di Luzio et al. "The landscape of QCD axion models". *Phys. Rept.* 870 (2020), pp. 1–117. arXiv: 2003.01100 [hep-ph].

[121] S. B. Giddings and A. Strominger. "Loss of Incoherence and Determination of Coupling Constants in Quantum Gravity". *Nucl. Phys.* B307 (1988), pp. 854–866.

[122] S. R. Coleman. "Why There Is Nothing Rather Than Something: A Theory of the Cosmological Constant". *Nucl. Phys.* B310 (1988), pp. 643–668.

[123] G. Gilbert. "WORMHOLE INDUCED PROTON DECAY". *Nucl. Phys.* B328 (1989), pp. 159–170.

[124] J. Alvey and M. Escudero. "The axion quality problem: global symmetry breaking and wormholes". *JHEP* 01 (2021), p. 032. arXiv: 2009.03917 [hep-ph].

[125] V. A. Rubakov. "Grand unification and heavy axion". *JETP Lett.* 65 (1997), pp. 621–624. arXiv: hep-ph/9703409.

[126] Z. Berezhiani, L. Gianfagna, and M. Giannotti. "Strong CP problem and mirror world: The Weinberg-Wilczek axion revisited". *Phys. Lett. B* 500 (2001), pp. 286–296. arXiv: hep-ph/0009290.

[127] A. Hook. "Anomalous solutions to the strong CP problem". *Phys. Rev. Lett.* 114.14 (2015), p. 141801. arXiv: 1411.3325 [hep-ph].

[128] H. Fukuda et al. "Model of visible QCD axion". *Phys. Rev. D* 92.1 (2015), p. 015021. arXiv: 1504.06084 [hep-ph].

[129] T. Gherghetta, N. Nagata, and M. Shifman. "A Visible QCD Axion from an Enlarged Color Group". *Phys. Rev. D* 93.11 (2016), p. 115010. arXiv: 1604.01127 [hep-ph].

[130] A. Hook et al. "High Quality QCD Axion and the LHC". *Phys. Rev. Lett.* 124.22 (2020), p. 221801. arXiv: 1911.12364 [hep-ph].

[131] Y. Nomura and J. Thaler. "Dark Matter through the Axion Portal". *Phys. Rev. D* 79 (2009), p. 075008. arXiv: 0810.5397 [hep-ph].

[132] R. T. Co, T. Gherghetta, and K. Harigaya. "Axiogenesis with a Heavy QCD Axion" (2022). arXiv: 2206.00678 [hep-ph].

[133] F. Takahashi and W. Yin. "Challenges for heavy QCD axion inflation". *JCAP* 10 (2021), p. 057. arXiv: 2105.10493 [hep-ph].

[134] K. Van Bibber et al. "Proposed experiment to produce and detect light pseudoscalars". *Phys. Rev. Lett.* 59 (1987), pp. 759–762.

[135] R. Ballou et al. (OSQAR). "New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall". *Phys. Rev. D* 92.9 (2015), p. 092002. arXiv: 1506.08082 [hep-ex].

[136] K. Ehret et al. "New ALPS Results on Hidden-Sector Lightweights". *Phys. Lett. B* 689 (2010), pp. 149–155. arXiv: 1004.1313 [hep-ex].

[137] K. Isleif (ALPS). "The Any Light Particle Search experiment at DESY" (2022). arXiv: 2202.07306 [hep-ex].

[138] E. Goudzovski et al. "New Physics Searches at Kaon and Hyperon Factories" (2022). arXiv: 2201.07805 [hep-ph].

[139] A. V. Artamonov et al. (E949). "Search for the decay K+ to pi+ gamma gamma in the pi+ momentum region P > 213 MeV/c". *Phys. Lett. B* 623 (2005), pp. 192–199. arXiv: hep-ex/0505069.

[140] J. R. Batley et al. (NA48/2). "A new measurement of the $K^{\pm} \to \pi^{\pm}\gamma\gamma$ decay at the NA48/2 experiment". *Phys. Lett. B* 730 (2014), pp. 141–148. arXiv: 1310.5499 [hep-ex].

[141] J. R. Batley et al. (NA48/2). "Searches for lepton number violation and resonances in $K^{\pm} \to \pi\mu\mu$ decays". *Phys. Lett. B* 769 (2017), pp. 67–76. arXiv: 1612.04723 [hep-ex].

[142] C. Lazzeroni et al. (NA62). "Study of the $K^{\pm} \to \pi^{\pm}\gamma\gamma$ decay by the NA62 experiment". *Phys. Lett. B* 732 (2014), pp. 65–74. arXiv: 1402.4334 [hep-ex].

[143] E. Cortina Gil et al. (NA62). "Measurement of the very rare $K^+ \to \pi^+\nu\overline{\nu}$ decay". *JHEP* 06 (2021), p. 093. arXiv: 2103.15389 [hep-ex].

[144] A. Alavi-Harati et al. (KTEV). "Search for the Decay $K_L \to \pi^0\mu^+\mu^-$". *Phys. Rev. Lett.* 84 (2000), pp. 5279–5282. arXiv: hep-ex/0001006.

[145] A. Alavi-Harati et al. (KTeV). "Search for the rare decay K(L) —> pi0 e+ e-". *Phys. Rev. Lett.* 93 (2004), p. 021805. arXiv: hep-ex/0309072.

[146] B. Aubert et al. (BaBar). "Study of B Meson Decays with Excited eta and eta-prime Mesons". *Phys. Rev. Lett.* 101 (2008), p. 091801. arXiv: 0804.0411 [hep-ex].

[147] J. P. Lees et al. (BaBar). "Search for an Axionlike Particle in B Meson Decays". *Phys. Rev. Lett.* 128.13 (2022), p. 131802. arXiv: 2111.01800 [hep-ex].

[148] R. Aaij et al. (LHCb). "Search for hidden-sector bosons in $B^0 \to K^{*0}\mu^+\mu^-$ decays". *Phys. Rev. Lett.* 115.16 (2015), p. 161802. arXiv: 1508.04094 [hep-ex].

[149] M. Dine, W. Fischler, and M. Srednicki. "A Simple Solution to the Strong CP Problem with a Harmless Axion". *Phys. Lett.* 104B (1981), pp. 199–202.

[150] A. R. Zhitnitsky. "On Possible Suppression of the Axion Hadron Interactions. (In Russian)". *Sov. J. Nucl. Phys.* 31 (1980). [Yad. Fiz.31,497(1980)], p. 260.

[151] M. A. Buen-Abad et al. "Challenges for an axion explanation of the muon $g-2$ measurement". *JHEP* 09 (2021), p. 101. arXiv: 2104.03267 [hep-ph].

[152] D. Aloni, Y. Soreq, and M. Williams. "Coupling QCD-Scale Axionlike Particles to Gluons". *Phys. Rev. Lett.* 123.3 (2019), p. 031803. arXiv: 1811.03474 [hep-ph].

[153] R. Co, S. Kumar, and Z. Liu. *Private communication.* 2022.

[154] C. Rubbia. "The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors" (1977).

[155] J. N. Marx and D. R. Nygren. "The Time Projection Chamber". *Phys. Today* 31N10 (1978), pp. 46–53.

[156] D. R. Nygren. "Origin and development of the TPC idea". *Nucl. Instrum. Meth. A* 907 (2018), pp. 22–30.

[157] W. J. Willis and V. Radeka. "Liquid Argon Ionization Chambers as Total Absorption Detectors". *Nucl. Instrum. Meth.* 120 (1974), pp. 221–236.

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

[159] C. Anderson et al. (ArgoNeuT). "Analysis of a Large Sample of Neutrino-Induced Muons with the ArgoNeuT Detector". *JINST* 7 (2012), P10020. arXiv: 1205.6702 [physics.ins-det].

[160] L. Landau. "On the energy loss of fast particles by ionization". *J. Phys. (USSR)* 8 (1944), pp. 201–205.

[161] E. Aprile, W.H. Ku, and J. Park. "Delta electron production and the ultimate energy resolution of liquid argon ionization detectors". *IEEE Trans. Nucl. Sci.* 35.1 (1988), pp. 37–41.

[162] K. V. J. Mistry. "First Measurement of the Flux-Averaged Differential Charged-Current Electron-Neutrino and Antineutrino Cross Section on Argon with the MicroBooNE Detector". PhD thesis. The University of Manchester, 2021.

[163] R. Acciarri et al. (ArgoNeuT). "First measurement of electron neutrino scattering cross section on argon". *Phys. Rev. D* 102.1 (2020), p. 011101. arXiv: 2004.01956 [hep-ex].

[164] M. Miyajima et al. "Average energy expended per ion pair in liquid argon". *Phys. Rev. A* 9 (1974), pp. 1438–1443.

[165] E. Shibamura et al. "Drift velocities of electrons, saturation characteristics of ionization and W-values for conversion electrons in liquid argon, liquid argon-gas mixtures and liquid xenon". *Nucl. Instrum. Meth.* 131 (1975), pp. 249–258.

[166] P. Cennini et al. "Performance of a 3-ton liquid argon time projection chamber". *Nucl. Instrum. Meth. A* 345 (1994), pp. 230–243.

[167] S. Kubota et al. "Dynamical behavior of free electrons in the recombination process in liquid argon, krypton, and xenon". *Phys. Rev. B* 20.8 (1979), p. 3486.

[168] S. Amoruso et al. (ICARUS). "Study of electron recombination in liquid argon with the ICARUS TPC". *Nucl. Instrum. Meth. A* 523 (2004), pp. 275–286.

[169] R. Acciarri et al. (ArgoNeuT). "A Study of Electron Recombination Using Highly Ionizing Particles in the ArgoNeuT Liquid Argon TPC". *JINST* 8 (2013), P08005. arXiv: 1306.1712 [physics.ins-det].

[170] J. B. Birks. "Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations". *Proc. Phys. Soc. A* 64 (1951), pp. 874–877.

[171] J. Thomas and D. A. Imel. "Recombination of electron-ion pairs in liquid argon and liquid xenon". *Phys. Rev. A* 36 (1987), pp. 614–616.

[172] C. Anderson et al. (ArgoNeuT). "The ArgoNeuT Detector in the NuMI Low-Energy beam line at Fermilab". *JINST* 7 (2012), P10019. arXiv: 1205.6747 [physics.ins-det].

[173] W. Walkowiak. "Drift velocity of free electrons in liquid argon". *Nucl. Instrum. Meth. A* 449 (2000), pp. 288–294.

[174] R. Acciarri et al. (WArP). "Oxygen contamination in liquid Argon: Combined effects on ionization electron charge and scintillation light". *JINST* 5 (2010), P05003. arXiv: 0804.1222 [nucl-ex].

[175] A. Bettini et al. "A Study of the factors affecting the electron lifetime in ultrapure liquid argon". *Nucl. Instrum. Meth. A* 305 (1991), pp. 177–186.

[176] B. Abi et al. (DUNE). "First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform". *JINST* 15.12 (2020), P12004. arXiv: 2007.06722 [physics.ins-det].

[177] P. Abratenko et al. (MicroBooNE). "Measurement of the longitudinal diffusion of ionization electrons in the MicroBooNE detector". *JINST* 16.09 (2021), P09025. arXiv: 2104.06551 [physics.ins-det].

[178] A. Lister and M. Stancari. "Investigations on a Fuzzy Process: Effect of Diffusion on Calibration and Particle Identification in Liquid Argon Time Projection Chambers" (2022). arXiv: 2201.09773 [physics.ins-det].

[179] N. Gee et al. "Ion and electron mobilities in cryogenic liquids: argon, nitrogen, methane, and ethane". *J. Appl. Phys.* 57.4 (1985), pp. 1097–1101.

[180] P. Abratenko et al. (MicroBooNE). "Measurement of space charge effects in the MicroBooNE LArTPC using cosmic muons". *JINST* 15.12 (2020), P12037. arXiv: 2008.09765 [physics.ins-det].

[181] R. Acciarri et al. (LArIAT). "The Liquid Argon In A Testbeam (LArIAT) Experiment". *JINST* 15.04 (2020), P04026. arXiv: 1911.10379 [physics.ins-det].

[182] B. Abi et al. (DUNE). "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV: Far Detector Single-phase Technology". *JINST* 15.08 (2020), T08010. arXiv: 2002.03010 [physics.ins-det].

[183] S. Amerio et al. (ICARUS). "Design, construction and tests of the ICARUS T600 detector". *Nucl. Instrum. Meth. A* 527 (2004), pp. 329–410.

[184] R. Acciarri et al. (SBND). "Construction of precision wire readout planes for the Short-Baseline Near Detector (SBND)". *JINST* 15.06 (2020), P06033. arXiv: 2002.08424 [physics.ins-det].

[185] T. Doke, K. Masuda, and E. Shibamura. "Estimation of absolute photon yields in liquid argon and xenon for relativistic (1 MeV) electrons". *Nucl. Instrum. Meth. A* 291 (1990), pp. 617–620.

[186] P. Benetti et al. "First results from a Dark Matter search with liquid Argon at 87 K in the Gran Sasso Underground Laboratory". *Astropart. Phys.* 28 (2008), pp. 495–507. arXiv: astro-ph/0701286.

[187] P. Agnes et al. (DarkSide). "First Results from the DarkSide-50 Dark Matter Experiment at Laboratori Nazionali del Gran Sasso". *Phys. Lett. B* 743 (2015), pp. 456–466. arXiv: 1410.0653 [astro-ph.CO].

[188] R. Ajaj et al. (DEAP Collaboration). "Search for dark matter with a 231-day exposure of liquid argon using DEAP-3600 at SNOLAB". *Phys. Rev. D* 100 (2 2019), p. 022004.

[189] M. Sorel. "Expected performance of an ideal liquid argon neutrino detector with enhanced sensitivity to scintillation light". *JINST* 9 (2014), P10002. arXiv: 1405.0848 [physics.ins-det].

[190] A.M. Szelc (SBND). "Developing LAr Scintillation Light Collection Ideas in the Short Baseline Neutrino Detector". *JINST* 11.02 (2016), p. C02018.

[191] T. Doke. "Fundamental Properties of Liquid Argon, Krypton and Xenon as Radiation Detector Media". *Portugal. Phys.* 12 (1981), pp. 9–48.

[192] E. Morikawa et al. "Argon, krypton, and xenon excimer luminescence: From the dilute gas to the condensed phase". *J. Chem. Phys. 91, 1469* (1989).

[193] A. Hitachi et al. "Effect of ionization density on the time dependence of luminescence from liquid argon and xenon". *Phys. Rev. B* 27 (1983), pp. 5279–5285.

[194] E. Segreto. "Evidence of delayed light emission of TetraPhenyl Butadiene excited by liquid Argon scintillation light". *Phys. Rev. C* 91.3 (2015), p. 035503. arXiv: 1411.4524 [physics.ins-det].

[195] T. Doke et al. "Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles". *Jap. J. Appl. Phys.* 41 (2002), pp. 1538–1545.

[196] H. Cao et al. (SCENE). "Measurement of Scintillation and Ionization Yield and Scintillation Pulse Shape from Nuclear Recoils in Liquid Argon". *Phys. Rev. D* 91 (2015), p. 092007. arXiv: 1406.4825 [physics.ins-det].

[197] P. Agnes et al. "Measurement of the liquid argon energy response to nuclear and electronic recoils". *Phys. Rev. D* 97.11 (2018), p. 112005. arXiv: 1801.06653 [physics.ins-det].

[198] M.G. Boulay and A. Hime. "Technique for direct detection of weakly interacting massive particles using scintillation time discrimination in liquid argon". *Astropart. Phys.* 25 (2006), pp. 179–182.

[199] R. Acciarri et al. (WArP). "Effects of Nitrogen contamination in liquid Argon". *JINST* 5 (2010), P06003. arXiv: 0804.1217 [nucl-ex].

[200] M. Antonello et al. "Detection of Cherenkov light emission in liquid argon". *Nucl. Instrum. Meth. A* 516 (2004), pp. 348–363.

[201] V. Basque. "Understanding and Enhancing The Performance of LArTPC Neutrino Detectors for Neutrino Measurements and Beyond the Standard Model Searches". PhD thesis. The University of Manchester, 2021.

[202] E. Grace et al. "Index of refraction, Rayleigh scattering length, and Sellmeier coefficients in solid and liquid argon and xenon". *Nucl. Instrum. Meth. A* 867 (2017), pp. 204–208. arXiv: 1502.04213 [physics.ins-det].

[203] J. Calvo et al. (ArDM). "Measurement of the attenuation length of argon scintillation light in the ArDM LAr TPC". *Astropart. Phys.* 97 (2018), pp. 186–196. arXiv: 1611.02481 [astro-ph.IM].

[204] P. Agnes et al. (DarkSide). "Simulation of argon response and light detection in the DarkSide-50 dual phase TPC". *JINST* 12.10 (2017), P10015. arXiv: 1707.05630 [physics.ins-det].

[205] A. Neumeier et al. "Attenuation measurements of vacuum ultraviolet light in liquid argon revisited". *Nucl. Instrum. Meth. A* 800 (2015), pp. 70–81. arXiv: 1511.07726 [physics.ins-det].

[206] M. Babicz et al. "A measurement of the group velocity of scintillation light in liquid argon". *JINST* 15.09 (2020), P09009. arXiv: 2002.09346 [physics.ins-det].

[207] D. Garcia-Gamez. *Private communication*. 2022.

[208] B. J. P. Jones et al. "A Measurement of the Absorption of Liquid Argon Scintillation Light by Dissolved Nitrogen at the Part-Per-Million Level". *JINST* 8 (2013). [Erratum: JINST 8, E09001 (2013)], P07011. arXiv: 1306.4605 [physics.ins-det].

[209] B. J. P. Jones et al. "The Effects of Dissolved Methane upon Liquid Argon Scintillation Light". *JINST* 8 (2013), P12015. arXiv: 1308.3658 [physics.ins-det].

[210] A. Ankowski et al. "Characterization of ETL 9357FLA photomultiplier tubes for cryogenic temperature applications". *Nucl. Instrum. Meth. A* 556 (2006), pp. 146–157.

[211] J. A. Nikkel, W. H. Lippincott, and D. N. McKinsey. "Demonstration of photomultiplier tube operation at 29 K". *JINST* 2 (2007), P11004. arXiv: astro-ph/0702202.

[212] R. Acciarri et al. "Demonstration and Comparison of Operation of Photomultiplier Tubes at Liquid Argon Temperature". *JINST* 7 (2012), P01016. arXiv: 1108.5584 [physics.ins-det].

[213] A. Gola et al. "NUV-Sensitive Silicon Photomultiplier Technologies Developed at Fondazione Bruno Kessler". *Sensors* 19.2 (2019), p. 308.

[214] B. Howard et al. "A Novel Use of Light Guides and Wavelength Shifting Plates for the Detection of Scintillation Photons in Large Liquid Argon Detectors". *Nucl. Instrum. Meth. A* 907 (2018), pp. 9–21. arXiv: 1710.11233 [physics.ins-det].

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

[216] E. Segreto et al. "First liquid argon test of the X-ARAPUCA". *JINST* 15.05 (2020), p. C05045.

[217] M. Kuźniak and A. M. Szelc. "Wavelength Shifters for Applications in Liquid Argon Detectors". *Instruments* 5.1 (2020), p. 4. arXiv: 2012.15626 [physics.ins-det].

[218] M. Kuźniak et al. "Polyethylene naphthalate film as a wavelength shifter in liquid argon detectors". *Eur. Phys. J. C* 79.4 (2019), p. 291. arXiv: 1806.04020 [physics.ins-det].

[219] B. Ali-Mohammadzadeh et al. (ICARUS). "Design and implementation of the new scintillation light detection system of ICARUS T600" (2020). arXiv: 2006.05261 [physics.ins-det].

[220] C. Wahl et al. "Pulse-shape discrimination and energy resolution of a liquid-argon scintillator with xenon doping". *JINST* 9 (2014), P06013. arXiv: 1403.0525 [physics.ins-det].

[221] N. Gallice (DUNE). "Xenon doping of liquid argon in ProtoDUNE single phase". *JINST* 17.01 (2022), p. C01034. arXiv: 2111.00347 [physics.ins-det].

[222] S. Kubota et al. "Evidence of the existence of exciton states in liquid argon and exciton-enhanced ionization from xenon doping". *Phys. Rev. B* 13.4 (1976), p. 1649.

[223] E. Aprile and T. Doke. "Liquid Xenon Detectors for Particle Physics and Astrophysics". *Rev. Mod. Phys.* 82 (2010), pp. 2053–2097. arXiv: 0910.4956 [physics.ins-det].

[224] K. Lung et al. "Characterization of the Hamamatsu R11410-10 3-Inch Photomultiplier Tube for Liquid Xenon Dark Matter Direct Detection Experiments". *Nucl. Instrum. Meth. A* 696 (2012), pp. 32–39. arXiv: 1202.2628 [physics.ins-det].

[225] L. Baudis et al. "Characterisation of Silicon Photomultipliers for Liquid Xenon Detectors". *JINST* 13.10 (2018), P10022. arXiv: 1808.06827 [physics.ins-det].

[226] L. Paulucci (DUNE). "The DUNE vertical drift photon detection system". *JINST* 17.01 (2022), p. C01067. arXiv: 2112.09520 [physics.ins-det].

[227] A. Abed Abud et al. (DUNE). "Scintillation light detection in the 6-m drift-length ProtoDUNE Dual Phase liquid argon TPC" (2022). arXiv: 2203.16134 [physics.ins-det].

[228] V. M. Gehman et al. "Fluorescence Efficiency and Visible Re-emission Spectrum of Tetraphenyl Butadiene Films at Extreme Ultraviolet Wavelengths". *Nucl. Instrum. Meth. A* 654 (2011), pp. 116–121. arXiv: 1104.3259 [astro-ph.IM].

[229] M. G. Boulay et al. "Direct comparison of PEN and TPB wavelength shifters in a liquid argon detector". *Eur. Phys. J. C* 81.12 (2021), p. 1099. arXiv: 2106.15506 [physics.ins-det].

[230] P. Adamson et al. "The NuMI Neutrino Beam". *Nucl. Instrum. Meth. A* 806 (2016), pp. 279–306.

[231] C. Anderson et al. (ArgoNeuT). "First Measurements of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon". *Phys. Rev. Lett.* 108 (2012), p. 161802. arXiv: 1111.0103 [hep-ex].

[232] R. Acciarri et al. (ArgoNeuT). "Measurements of Inclusive Muon Neutrino and Antineutrino Charged Current Differential Cross Sections on Argon in the NuMI Antineutrino Beam". *Phys. Rev. D* 89.11 (2014), p. 112003. arXiv: 1404.4809 [hep-ex].

[233] R. Acciarri et al. (ArgoNeuT). "First Measurement of Neutrino and Antineutrino Coherent Charged Pion Production on Argon". *Phys. Rev. Lett.* 113.26 (2014). [Erratum: Phys.Rev.Lett. 114, 039901 (2015)], p. 261801. arXiv: 1408.0598 [hep-ex].

[234] R. Acciarri et al. (ArgoNeuT). "Measurement of ν_μ and $\bar{\nu}_\mu$ neutral current $\pi^0 \to \gamma\gamma$ production in the ArgoNeuT detector". *Phys. Rev. D* 96.1 (2017), p. 012006. arXiv: 1511.00941 [hep-ex].

[235] R. Acciarri et al. (ArgoNeuT). "First measurement of the cross section for ν_μ and $\bar{\nu}_\mu$ induced single charged pion production on argon using ArgoNeuT". *Phys. Rev. D* 98.5 (2018), p. 052002. arXiv: 1804.10294 [hep-ex].

[236] R. Acciarri et al. (ArgoNeuT). "Improved Limits on Millicharged Particles Using the ArgoNeuT Experiment at Fermilab". *Phys. Rev. Lett.* 124.13 (2020), p. 131801. arXiv: 1911.07996 [hep-ex].

[237] D. G. Michael et al. (MINOS). "The Magnetized steel and scintillator calorimeters of the MINOS experiment". *Nucl. Instrum. Meth. A* 596 (2008), pp. 190–228. arXiv: 0805.3170 [physics.ins-det].

[238] J. Evans (MINOS). "The MINOS Experiment: Results and Prospects". *Adv. High Energy Phys.* 2013 (2013), p. 182537. arXiv: 1307.0721 [hep-ex].

[239] P. Adamson et al. (MINOS). "Neutrino and Antineutrino Inclusive Charged-current Cross Section Measurements with the MINOS Near Detector". *Phys. Rev. D* 81 (2010), p. 072002. arXiv: 0910.2201 [hep-ex].

[240] I. Stancu et al. "Technical Design Report for the 8 GeV Beam" (2001).

[241] M. Antonello et al. (MicroBooNE, LAr1-ND, ICARUS-WA104). "A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam" (2015). arXiv: 1503.01520 [physics.ins-det].

[242] P. Ballett, S. Pascoli, and M. Ross-Lonergan. "MeV-scale sterile neutrino decays at the Fermilab Short-Baseline Neutrino program". *JHEP* 04 (2017), p. 102. arXiv: 1610.08512 [hep-ph].

[243] M. Nunes. *Private communication*. 2022.

[244] V. Pandey. *Private communication*. 2022.

[245] M. Auger et al. "A Novel Cosmic Ray Tagger System for Liquid Argon TPC Neutrino Detectors". *Instruments* 1.1 (2017), p. 2. arXiv: 1612.04614 [physics.ins-det].

[246] M. Ball et al. "The PIP-II Conceptual Design Report" (2017).

[247] B. Abi et al. (DUNE). "Long-baseline neutrino oscillation physics potential of the DUNE experiment". *Eur. Phys. J. C* 80.10 (2020), p. 978. arXiv: 2006.16043 [hep-ex].

[248] F. Capozzi et al. "DUNE as the Next-Generation Solar Neutrino Experiment". *Phys. Rev. Lett.* 123.13 (2019), p. 131803. arXiv: 1808.08232 [hep-ph].

[249] B. Abi et al. (DUNE). "Supernova neutrino burst detection with the Deep Underground Neutrino Experiment". *Eur. Phys. J. C* 81.5 (2021), p. 423. arXiv: 2008.06647 [hep-ex].

[250] B. Abi et al. (DUNE). "Prospects for beyond the Standard Model physics searches at the Deep Underground Neutrino Experiment". *Eur. Phys. J. C* 81.4 (2021), p. 322. arXiv: 2008.12769 [hep-ex].

[251] B. Abi et al. (DUNE). "The DUNE Far Detector Conceptual Design Report, Single-Phase Vertical Drift Technology" (2021). eprint: https://edms.cern.ch/ui/file/2619631/1/Vertical_Single_Phase_FD_Technology_CDR.pdf.

[252] B. Abi et al. (DUNE). "The DUNE Far Detector Interim Design Report, Volume 3: Dual-Phase Module" (2018). arXiv: 1807.10340 [physics.ins-det].

[253] A. Abed Abud et al. (DUNE). "Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report". *Instruments* 5.4 (2021), p. 31. arXiv: 2103.13910 [physics.ins-det].

[254] F. Sauli. "The gas electron multiplier (GEM): Operating principles and applications". *Nucl. Instrum. Meth. A* 805 (2016), pp. 2–24.

[255] B. Baibussinov et al. "Operation of a LAr-TPC equipped with a multilayer LEM charge readout". *JINST* 13.03 (2018), T03001. arXiv: 1711.06781 [physics.ins-det].

[256] C. Vázquez et al. "Power over Fiber in Radio over Fiber Systems in 5G Scenarios". *2019 21st International Conference on Transparent Optical Networks (ICTON)*. 2019, pp. 1–4.

[257] L. Aliaga et al. (MINERvA). "Design, Calibration, and Performance of the MINERvA Detector". *Nucl. Instrum. Meth. A* 743 (2014), pp. 130–159. arXiv: 1305.5199 [physics.ins-det].

[258] G. Tzanankos et al. (MINOS+). "MINOS+: a Proposal to FNAL to run MINOS with the medium energy NuMI beam" (2011).

[259] D. S. Ayres et al. (NOvA). "NOvA: Proposal to Build a 30 Kiloton Off-Axis Detector to Study $\nu_\mu \to \nu_e$ Oscillations in the NuMI Beamline" (2004). arXiv: hep-ex/0503053.

[260] Google Earth. 2018.

[261] L. Aliaga et al. (MINERvA). "Neutrino Flux Predictions for the NuMI Beam". *Phys. Rev. D* 94.9 (2016), p. 092005. arXiv: 1607.00704 [hep-ex].

[262] M. Bauer, M. Neubert, and A. Thamm. "Collider Probes of Axion-Like Particles". *JHEP* 12 (2017), p. 044. arXiv: 1708.00443 [hep-ph].

[263] F. Ertas and F. Kahlhoefer. "On the interplay between astrophysical and laboratory probes of MeV-scale axion-like particles". *JHEP* 07 (2020), p. 050. arXiv: 2004.01193 [hep-ph].

[264] E.L. Snider and G. Petrillo. "LArSoft: Toolkit for Simulation, Reconstruction and Analysis of Liquid Argon TPC Neutrino Detectors". *J. Phys. Conf. Ser.* 898.4 (2017). Ed. by R. Mount and C. Tull, p. 042057.

[265] C. Green et al. "The Art Framework". *J. Phys. Conf. Ser.* 396 (2012). Ed. by M. Ernst et al., p. 022020.

[266] R. Brun and F. Rademakers. "ROOT: An object oriented data analysis framework". *Nucl. Instrum. Meth. A* 389 (1997). Ed. by M. Werlen and D. Perret-Gallix, pp. 81–86.

[267] C. Andreopoulos et al. "The GENIE Neutrino Monte Carlo Generator". *Nucl. Instrum. Meth. A* 614 (2010), pp. 87–104. arXiv: 0905.2517 [hep-ph].

[268] D. Heck et al. "CORSIKA: A Monte Carlo code to simulate extensive air showers" (1998).

[269] S. Agostinelli et al. (GEANT4). "GEANT4–a simulation toolkit". *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303.

[270] J. Spitz. "Measuring Muon-Neutrino Charged-Current Differential Cross Sections with a Liquid Argon Time Projection Chamber". PhD thesis. Yale University, 2011.

[271] R. Acciarri et al. (SBND). "Cosmic Ray Background Removal With Deep Neural Networks in SBND". *Front. Artif. Intell.* 4 (2021), p. 649917. arXiv: 2012.01301 [physics.data-an].

[272] P. Abratenko et al. (MicroBooNE). "Measurement of the atmospheric muon rate with the MicroBooNE Liquid Argon TPC". *JINST* 16.04 (2021), P04004. arXiv: 2012.14324 [physics.ins-det].

[273] P. Abratenko et al. (MicroBooNE). "Novel Approach for Evaluating Detector-Related Uncertainties in a LArTPC Using MicroBooNE Data" (2021). arXiv: 2111.03556 [hep-ex].

[274] P. deNiverville et al. "Light dark matter in neutrino beams: production modelling and scattering signatures at MiniBooNE, T2K and SHiP". *Phys. Rev. D* 95.3 (2017), p. 035006. arXiv: 1609.01770 [hep-ph].

[275] J. Isaacson et al. "A novel event generator for the automated simulation of neutrino scattering" (2021). arXiv: 2110.15319 [hep-ph].

[276] C. Bierlich et al. "A comprehensive guide to the physics and usage of PYTHIA 8.3" (2022). arXiv: 2203.11601 [hep-ph].

[277] T. Sjöstrand et al. "An introduction to PYTHIA 8.2". *Comput. Phys. Commun.* 191 (2015), pp. 159–177. arXiv: 1410.3012 [hep-ex].

[278] J. Spitz. *Private communication*. 2021.

[279] G. Altarelli, R. Kleiss, and C. Verzegnassi. "Workshop on Z Physics at LEP1" (1989), pp. 327–330.

[280] B. Baller. "Liquid argon TPC signal formation, signal processing and reconstruction techniques". *JINST* 12.07 (2017), P07010. arXiv: 1703.04024 [physics.ins-det].

[281] R. Veenhof. "GARFIELD, recent developments". *Nucl. Instrum. Meth. A* 419 (1998). Ed. by M. Krammer et al., pp. 726–730.

[282] B. Jones. "Sterile Neutrinos in Cold Climates". PhD thesis. MIT, 2015.

[283] W. Mu, A. I. Himmel, and B. Ramson. "Photon detection probability prediction using one-dimensional generative neural network". *Mach. Learn. Sci. Tech.* 3.1 (2022), p. 015033. arXiv: 2109.07277 [physics.ins-det].

[284] R. Acciarri et al. (ArgoNeuT). "A deep-learning based raw waveform region-of-interest finder for the liquid argon time projection chamber". *JINST* 17.01 (2022), P01018. arXiv: 2103.06391 [physics.ins-det].

[285] M. Antonello et al. "Precise 3D track reconstruction algorithm for the ICARUS T600 liquid argon time projection chamber detector". *Adv. High Energy Phys.* 2013 (2013), p. 260820. arXiv: 1210.5089 [physics.ins-det].

[286] R. Acciarri et al. (LArIAT). "The Liquid Argon In A Testbeam (LArIAT) Experiment". *JINST* 15.04 (2020), P04026. arXiv: 1911.10379 [physics.ins-det].

[287] W. Foreman et al. (LArIAT). "Calorimetry for low-energy electrons using charge and light in liquid argon". *Phys. Rev. D* 101.1 (2020), p. 012010. arXiv: 1909.07920 [physics.ins-det].

[288] R. Acciarri et al. (ArgoNeuT). "First Observation of Low Energy Electron Neutrinos in a Liquid Argon Time Projection Chamber". *Phys. Rev. D* 95.7 (2017), p. 072005. arXiv: 1610.04102 [hep-ex].

[289] J. S. Marshall and M. A. Thomson. "The Pandora Software Development Kit for Pattern Recognition". *Eur. Phys. J. C* 75.9 (2015), p. 439. arXiv: 1506.05348 [physics.data-an].

[290] R. Acciarri et al. (MicroBooNE). "The Pandora multi-algorithm approach to automated pattern recognition of cosmic-ray muon and neutrino events in the MicroBooNE detector". *Eur. Phys. J. C* 78.1 (2018), p. 82. arXiv: 1708.03135 [hep-ex].

[291] X. Qian et al. "Three-dimensional Imaging for Large LArTPCs". *JINST* 13.05 (2018), P05032. arXiv: 1803.04850 [physics.ins-det].

[292] P. Abratenko et al. (MicroBooNE). "Neutrino event selection in the MicroBooNE liquid argon time projection chamber using Wire-Cell 3D imaging, clustering, and charge-light matching". *JINST* 16.06 (2021), P06043. arXiv: 2011.01375 [physics.ins-det].

[293] P. Abratenko et al. (MicroBooNE). "Search for an anomalous excess of inclusive charged-current ν_e interactions in the MicroBooNE experiment using Wire-Cell reconstruction". *Phys. Rev. D* 105.11 (2022), p. 112005. arXiv: 2110.13978 [hep-ex].

[294] P. Abratenko et al. (MicroBooNE). "First Measurement of Energy-Dependent Inclusive Muon Neutrino Charged-Current Cross Sections on Argon with the MicroBooNE Detector". *Phys. Rev. Lett.* 128.15 (2022), p. 151801. arXiv: 2110.14023 [hep-ex].

[295] R. Acciarri et al. (MicroBooNE). "Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber". *JINST* 12.03 (2017), P03011. arXiv: 1611.05531 [physics.ins-det].

[296] C. Adams et al. (MicroBooNE). "Deep neural network for pixel-level electromagnetic particle identification in the MicroBooNE liquid argon time projection chamber". *Phys. Rev. D* 99.9 (2019), p. 092001. arXiv: 1808.07269 [hep-ex].

[297] P. Abratenko et al. (MicroBooNE). "Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE". *Phys. Rev. D* 103.5 (2021), p. 052012. arXiv: 2012.08513 [physics.ins-det].

[298] P. Abratenko et al. (MicroBooNE). "Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber". *Phys. Rev. D* 103.9 (2021), p. 092003. arXiv: 2010.08653 [hep-ex].

[299] P. Abratenko et al. (MicroBooNE). "Search for an anomalous excess of charged-current quasielastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction". *Phys. Rev. D* 105.11 (2022), p. 112003. arXiv: 2110.14080 [hep-ex].

[300] R. Chytracek et al. "Geometry Description Markup Language for Physics Simulation and Analysis Applications". *IEEE Trans. Nucl. Sci.* 53.5 (2006), pp. 2892–2896.

[301] R. Francini et al. "Tetraphenyl-butadiene films: VUV-Vis optical characterization from room to liquid argon temperature". *JINST* 8 (2013), p. C09010.

[302] J. Zwinkels, M. Noël, and C. Dodd. "Procedures and standards for accurate spectrophotometric measurements of specular reflectance". *Appl. Opt.* 33.34 (1994), pp. 7933–7944.

[303] P. Adamson et al. (MINOS). "A Study of Muon Neutrino Disappearance Using the Fermilab Main Injector Neutrino Beam". *Phys. Rev. D* 77 (2008), p. 072002. arXiv: 0711.0769 [hep-ex].

[304] R. Brun et al. "GEANT Detector Description and Simulation Tool" (1994).

[305] P. V. C. Hough. "Machine Analysis of Bubble Chamber Pictures". *Conf. Proc. C* 590914 (1959), pp. 554–558.

[306] R. E. Kalman. "A New Approach to Linear Filtering and Prediction Problems". *Trans. ASME, J. Basic Eng.* 82.Series D (1960), pp. 35–45.

[307] F. Paxton. "Solid Angle Calculation for a Circular Disk". *Rev. Sci. Instrum.* 30.04 (1959).

[308] A. Khadjavi. "Calculation of Solid Angle Subtended by Rectangular Apertures". *J. Opt. Soc. Am.* 58.10 (1968), pp. 1417–1418.

[309] T. K. Gaisser and A. M. Hillas. "Reliability of the Method of Constant Intensity Cuts for Reconstructing the Average Development of Vertical Showers". *International Cosmic Ray Conference*. Vol. 8. 1977, p. 353.

[310] L. Michel. "Interaction between four half spin particles and the decay of the μ meson". *Proc. Phys. Soc. A* 63 (1950), pp. 514–531.

[311] D. Garcia-Gamez, P. Green, and A. M. Szelc. "Predicting Transport Effects of Scintillation Light Signals in Large-Scale Liquid Argon Detectors". Presented at LIDINE. 2021.

[312] J. Li. "Measuring LAr Light Yield using Isolated Protons". Poster presented at LIDINE. 2021.

[313] S. Timm et al. "Grids, virtualization, and clouds at Fermilab". *J. Phys. Conf. Ser.* 513.3 (2014), p. 032037.

[314] R. Pordes et al. "The Open Science Grid". *J. Phys. Conf. Ser.* 78 (2007), p. 012057.

[315] M. Fahey et al. "Contemporary High Performance Computing From Petascale toward Exascale, Volume 3". 2019. Chap. Theta and Mira at Argonne National Laboratory, pp. 31–61.

[316] G. M. Kurtzer, V. Sochat, and M. W. Bauer. "Singularity: Scientific containers for mobility of compute". *PLoS ONE* 12.5 (2017).

[317] P. Buncic et al. "CernVM: A virtual software appliance for LHC applications". *J. Phys. Conf. Ser.* 219 (2010). Ed. by Jan Gruntorad and Milos Lokajicek, p. 042003.

[318] D. Weitzel et al. "StashCache: A Distributed Caching Federation for the Open Science Grid". *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning)*. PEARC '19. Association for Computing Machinery, 2019.

[319] M Salim et al. "Balsam: Automated Scheduling and Execution of Dynamic, Data-Intensive HPC Workflows" (2019). arXiv: 1909.08704 [cs.DC].

[320] C. R. Harris et al. "Array programming with NumPy". *Nature* 585.7825 (2020), pp. 357–362.

[321] I. Foster and C. Kesselman. "Globus: a Metacomputing Infrastructure Toolkit". *Int. J. High Perform. Comput. Appl.* 11.2 (1997), pp. 115–128.

[322] S. Barlag et al. (ACCMOR). "Production of D, D* and D(s) Mesons in 200-GeV/c pi-, K- and p Si Interactions". *Z. Phys. C* 39 (1988), p. 451.

[323] G. A. Alves et al. (E769). "Forward cross-sections for production of D+, D0, D(s), D*+ and Lambda(c) in 250-GeV pi+-, K+-, and p - nucleon interactions". *Phys. Rev. Lett.* 77 (1996). [Erratum: Phys. Rev. Lett. 81, 1537 (1998)], pp. 2388–2391.

[324] C. Lourenco and H. K. Wohri. "Heavy flavour hadro-production from fixed-target to collider energies". *Phys. Rept.* 433 (2006), pp. 127–180. arXiv: hep-ph/0609101.

[325] A. Aduszkiewicz et al. (NA61/SHINE). "Measurements of production and inelastic cross sections for p+C , p+Be , and p+Al at 60 GeV/c and p+C and p+Be at 120 GeV/c". *Phys. Rev. D* 100.11 (2019), p. 112001. arXiv: 1909.03351 [hep-ex].

[326] O. Helene. "Upper Limit of Peak Area". *Nucl. Instrum. Meth.* 212 (1983), p. 319.

[327] I. Boiarska et al. "Constraints from the CHARM experiment on heavy neutral leptons with tau mixing". *Phys. Rev. D* 104.9 (2021), p. 095019. arXiv: 2107.14685 [hep-ph].

[328] E. Bertuzzo et al. "New limits on leptophilic ALPs and Majorons from ArgoNeuT" (2022). arXiv: 2202.12317 [hep-ph].