Using a two-phase xenon time projection chamber for improved background rejection in searches for neutrinoless double-beta decay

by

Callan Jessiman

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in partial fulfilment of the requirements for the degree of

Master of Science

in

Physics

Department of Physics
Carleton University
Ottawa, Ontario
January 2020

Abstract

The nature of the neutrino masses is an important open question in particle physics; neutrinos were conceived and implemented into the Standard Model as massless, and their masses, now known to be nonzero, represent an area of new physics. Essential to the investigation of this area are mechanisms by which one might measure these masses, including the hypothetical neutrinoless double-beta decay, which has a lifetime related to the neutrino masses. Experiments searching for this rare decay, such as EXO, are impacted heavily by sources of background radiation. Here, a prototype two-phase time projection chamber is described, having superior temporal resolution to the existing EXO architecture. A machine-learning analysis of data from this prototype is used for pulse-shape discrimination, which has the potential to significantly increase EXO's sensitivity; the preliminary efforts described here are able to reduce backgrounds by 94%, while rejecting only 21% of the signal.

Acknowledgements

First, and foremost, I express my gratitude towards my supervisor, David Sinclair, and my colleague Braeden Veenstra. The entire project was of David's devising, as was the leadership that saw it through. Meanwhile, getting the thing to actually *work*, and extracting results from it, was Braeden's task as much as it was mine.

I would also like to thank Warren Cree and Caio Licciardi for providing patient instruction to a newcomer to EXO, Jeff Mason and Matt Bowcock for lending their technical expertise and hard work to the building of our detector, and Philippe Gravelle for trusting me to return the tools I borrowed from his machine shop.

I spent September of 2018 in Carlsbad, New Mexico, and underground at WIPP, working on the calibration and maintenance of the EXO-200 detector. I found the entire experience wonderfully interesting and enjoyable, and would like to thank my hosts and overseers, Brian Mong and Jon Davis, for their friendly hospitality.

Finally, I'd like to thank my best friend, for reminding me not to get too serious.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

$\beta\beta0\nu$ neutrinoless double-beta decay (a hypothetical rare nuclear decay that is the subject of several experimental searches; described in section 1.2.3)

BSM beyond the Standard Model (the forefront of fundamental science is the search for "new" physics, beyond the Standard Model, such as $\beta\beta0\nu$; described with a focus on neutrino physics in section 1.2)

CDF cumulative distribution function (the integral of a probability distribution from $-\infty$ to x, equal to the probability of obtaining a value less than x)

CERN European Organization for Nuclear Research (a leading organization for particle physics research which, among other things, develops software for detector simulation and data analysis)

DEAP Dark matter Experiment using Argon Pulse-shape discrimination (DEAP-3600 is an experiment attempting the direct detection of dark matter that may be methodologically relevant to EXO for reasons outlined in section 3.2)

EL electroluminescence (the emission of light by a medium due to the passage of energetic electrons; described in section 3.1)

EXO Enriched Xenon Observatory (a series of experiments, consisting of EXO-200 and nEXO, that search for $\beta\beta0\nu$ using ^{136}Xe time projection chambers; described in chapter 2)

LXe liquid xenon (a scintillator used in several particle physics experiments, including EXO)

S1 TPC signal 1 (the prompt scintillation signal in a TPC, consisting of light emitted due to the passage of ionizing ratiation; described in section 2.2)

S2 TPC signal 2 (the delayed ionization signal in a TPC, consisting of free electrons liberated by the passage of ionizing radiation and conveyed to a readout by an applied electric field; described in sections 2.2, 3.1, and 4.2.1)

SATP standard ambient temperature and pressure (a standard defined as a temperature of 25°C and pressure of 100 kPa, typical for the interior of a building)

SiPM silicon photomultiplier (a type of photodetector; used in the prototype described in section 3.3)

SM Standard Model (the Standard Model of particle physics explains most observed phenomena through a single quantum field theory, and, along with general relativity, comprises our current understanding of fundamental physics; introduced in section 1.1)

SNOLAB Sudbury Neutrino Observatory laboratory (an underground physics facility, originally the site of the SNO experiment, located in a nickel mine in Ontario; SNOLAB is now the site of several physics experiments, and may be the future site of nEXO)

TPC time projection chamber (a type of ionization detector, used by EXO to search for $\beta\beta0\nu$, and by this work; described in chapters 2 and 3)

WIPP Waste Isolation Pilot Plant (WIPP is a salt mine in New Mexico that is used as a deep geological repository for nuclear waste, and was the site of the EXO-200 experiment)

Introduction

Despite being one of the first components of the Standard Model of particle physics to be discovered, neutrinos remain one of the least well-understood, due to the rarity of their interactions. Originally conceived as massless, these particles are now known to have very small but nonzero masses, a fact which could have significant implications for the future of fundamental physics. Central to the study of neutrinos are their masses themselves (which have been constrained below upper limits, but not yet directly measured) and the physical processes through which measurement of the masses may be possible. One such process is a hypothetical rare nuclear decay known as neutrinoless double-beta decay, or $\beta\beta0\nu$.

The Enriched Xenon Observatory is an experimental effort to measure $\beta\beta0\nu$ in xenon-136. A first detector, EXO-200, completed its experimental run this past winter, setting upper limits on the rate of $\beta\beta0\nu$ but not observing the process. R&D for a larger and more advanced detector is ongoing; this nEXO detector is projected to be 200 times more sensitive to $\beta\beta0\nu$ than EXO-200, and advanced background rejection techniques with the potential to increase sensitivity further by a factor of up to 4.3 are being investigated. One such technique, pulse-shape discrimination, is the topic of this thesis.

Chapter 1 provides a background in neutrino physics, covering their discovery in 1933, the realization between 1968 and 2002 that they do, in fact, have mass, and the implications of these masses on the future of theoretical and experimental particle physics, including their relationship with $\beta\beta0\nu$.

Chapter 2 explains how $\beta\beta0\nu$ could be observed, and describes the current state of the experimental search for it, with a focus on the EXO-200 experiment. It then goes on to describe this experiment's successor, nEXO, as well as some of the factors that limit its sensitivity to $\beta\beta0\nu$ and potential strategies to overcome these limits.

Chapter 3 describes a prototype EXO-like detector constructed at Carleton University as part of the nEXO R&D effort. It explains how this detector differs from EXO-200 and the planned nEXO detector and how these differences may be of use to the search for $\beta\beta0\nu$, and lays out a method for evaluating this usefulness.

Chapter 4 covers the analysis of data taken using the prototype detector, and of the results of a detector simulation. Machine learning is used to attempt pulse-shape discrimination on a partial dataset, with some success, and simulation provides a proof-of-concept for further analysis of a full dataset. However, due to challenges in detector operation, a full dataset has not yet been obtained; these difficulties are described, along with planned efforts to address them.

Author's contributions

The first two chapters of this thesis contain no original material; they merely relate relevant theoretical and experimental background material, which is arranged here to provide context for the work described in the later chapters.

Chapter 3 describes the design and operation of the prototype detector; many of the ideas presented are those of my supervisor, Dr. David Sinclair. My own contributions to this include the assembly and mounting of the detectors themselves (the SiPM and NaI detectors) and their readouts, wiring and calibration of both methods of level sensing, design and implementation of the data acquisition hardware and software, and simulation of gamma-ray transport. Additionally, the final assembly of the vessels, assembly of the high-voltage circuit, and operation of the full system were shared between myself and my colleague Braeden Veenstra.

The fourth chapter describes analyses of data from the prototype detector. The overall strategies of analysis and simulation were conceived by my supervisor, Dr. David Sinclair; the details and implementations are my own.

Chapter 1

Neutrinoless Double-Beta Decay

1.1 Neutrinos in the Standard Model

The SM of particle physics is the culmination of over a century of research in the fields of quantum physics and relativity. This mathematical framework, finalized in the late 20th century, unites the menagerie of particles and interactions discovered in the preceding decades - that is, all of fundamental physics other than gravity and a few other loose ends[1] - into a single quantum field theory of a few matter particles that make up everything and a few mediator particles responsible for their interactions [1]. With the experimental discovery of Higgs boson - the last of these particles, responsible for the others' mass - in 2012, the SM has been thoroughly probed and found to be an excellent description of fundamental physics up to very high energies [2]. However, one of the particles in this model has proven rather elusive, and has recently been found not to behave quite as predicted.

1.1.1 Beta decay

Study of this particle (the neutrino), while very difficult, actually predates knowledge of most of the SM. In the early 1900s, before the advent of particle accelerators, the best source of high-energy particles was the natural decay of radioactive material. Early study of the phenomenon of decay identified three types: alpha, beta, and gamma decay (see figure 1.1), characterized by the emission of α, β, and γ particles, which were later identified as ^{4}He nuclei, electrons, and photons respectively.

Two laws govern these decays. First, each type of decay causes a nucleus to transition from a specific "parent" nuclide into a specific "daughter" nuclide; for example, the α decay of ^{238}U always results in an atom of ^{234}Th. Second, energy is conserved. The mass of an α particle plus that of a ^{234}Th nucleus is less than that

[1] See section 1.2

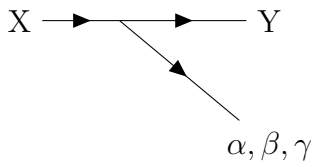

Figure 1.1: Nuclear decay, in which an atom of X decays to an atom of Y by emitting an α, β, or γ particle. Y is lighter than X, and their difference in energy and charge is carried off by the emitted particle.

of the initial ^{238}U nucleus, and the difference is precisely balanced by the kinetic energy of the α (4.27 MeV, in this case); if the difference was less than zero, the decay could not occur. This conservation holds for all α decays and all γ decays (and, for that matter, all phenomena that have ever been observed).

However, β decays did *not* appear to conserve energy, and by 1927 it had been shown that the electrons emitted by this process had an energy anywhere between the mass difference and zero [3]. Wolfgang Pauli suggested that a new, undetectable particle he called a "neutron" was responsible, but the name was stolen by Chadwick for his new nucleon before Pauli's idea had any real weight behind it [1]. Chadwick's neutron, incidentally, is also relevant to β decay; atoms of X contain some characteristic number of protons and neutrons, and the decay turns one of the neutrons *into* a proton, changing X into a different element.

Finally, in 1933, Fermi proposed a theory reproducing with the observed electron spectrum: that the missing energy was taken by a new particle he called a *neutrino* (meaning "little neutral one", as the particle he predicted had no mass or charge) that was emitted in the decay alongside the electron [4, 5]. This brings us to the almost-modern understanding of β decay shown in figure 1.2b.

We need two more concepts from the SM; the first is flavour conservation. All SM interactions conserve a quantity called *electron number*, which is the number of electrons (e^-) and neutrinos (ν) minus the number of antimatter electrons (antielectrons, also called positrons, e^+) and antineutrinos ($\overline{\nu}$). The neutrino

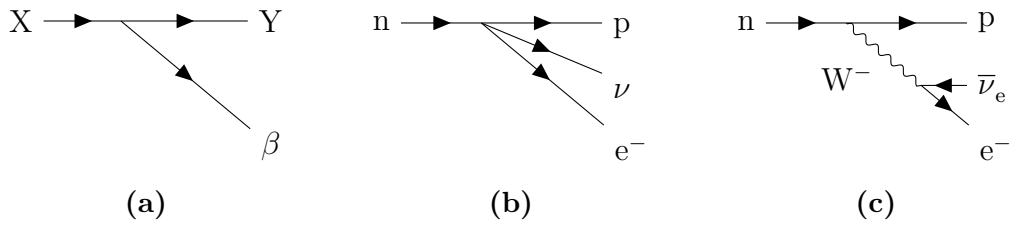

Figure 1.2: Models of beta decay: (a) shows the simple decay of figure 1.1. (b) is Fermi's interaction, adding ν and identifying n, p, and e. Finally, (c) shows the SM interaction, adding W^- and $\overline{\nu}_e$ (see below).

in figure 1.2c must therefore really be an *anti*neutrino (this is why its arrow is pointing backwards). Furthermore, we now know that there exist two particles μ and τ, which are identical to the electron, except heavier. They each have their own flavour conservation law, and their own massless neutrino counterpart, meaning that the neutrino in figure 1.2 is specifically an *electron* antineutrino $\overline{\nu}_{\mathrm{e}}$. The particles e, μ, τ, ν_{e}, ν_{μ}, and ν_{τ} are known collectively as *leptons*.

Finally, note the new particle that appeared in 1.2c: the W boson. Fermi's interaction, which ignores this mediator, is approximately correct at low energies, but the W is essential for the accurate treatment of the SM. This particle, along with the neutral Z boson, mediates the aptly-named *weak interaction*, which is typically weaker than the other two SM interactions (strong and electromagnetic), though much stronger than the last of the four fundamental interactions, gravity. Neutrinos are only known to interact via the W or Z, and because of this they interact so rarely with other particles that most of those produced by fusion in the sun pass unnoticed through our entire planet [6]. However, as discussed in 1.2.1, neutrinos will, though rarely, scatter off common matter through processes similar to beta decay, as was first observed - confirming the existence of the neutrino - by Cowan and Reines in 1959 [7].

1.2 Neutrinos beyond the Standard Model

The history of our understanding of beta decay holds an important lesson: anomalies are fantastic for physics. The discrepancy between observed decays and 1920s theory led to the discovery of the neutrino, only the third SM particle known when Fermi proposed it. The two great twentieth-century revolutions in physics came about in the same way; the first piece of quantum theory described thermal radiation where classical theory had failed, and relativity was proposed after experiments intended to probe the luminiferous aether found that it does not, in fact, exist. Similarly, the most promising leads in the search for physics BSM are the holes *in* the Standard Model, the observations that directly contradict the theory. There are also conceptual issues with the SM; some parameters seem to have oddly specific values, and the framework of the model simply does not admit a description of gravity. However, despite substantial theoretical work, there is as yet no good evidence on how to resolve these issues; it seems that at least the shortest path forwards to new physics lies in the anomalies. The most dramatic of these is the existence of dark matter; the particles in the SM only appear to account for one sixth of the matter in the universe, while the remainder is some new particle, which has only ever been **seen** to interact via gravitation. This, while interesting, is by its nature very difficult to study. Another anomaly, per-

haps the easiest to probe, is that the neutrinos in fact *do* have mass. Additionally, it appears that these masses allow neutrinos to violate flavour conservation, and possibly even more fundamental aspects of the Standard Model.

1.2.1 Neutrino oscillations

When Fermi first developed his theory of beta decay, he noted that the shape of the electron energy spectrum provided information about the mass of the neutrino [4]. Specifically, the high end of the spectrum (where the electron carries most of the energy, and little is "missing" via the neutrino) contains two terms: one proportional to $(Q-E)^2$ (where Q is the total decay energy) and one proportional to $m_\nu\sqrt{Q-E}$. The square-root term gives the spectrum infinite slope at $E=Q$, unless $m_\nu=0$, in which case the slope is zero (see figure 1.3). Contemporary experiments were consistent with a slope of zero, so Fermi concluded that the neutrino mass was zero, or at least too small to be resolved. Indeed, no experiment to date has made any direct measurement of a nonzero neutrino mass, through the beta decay spectrum or otherwise. However, several experiments have made *indirect* observations which imply that the neutrino masses, while small, are nonzero.

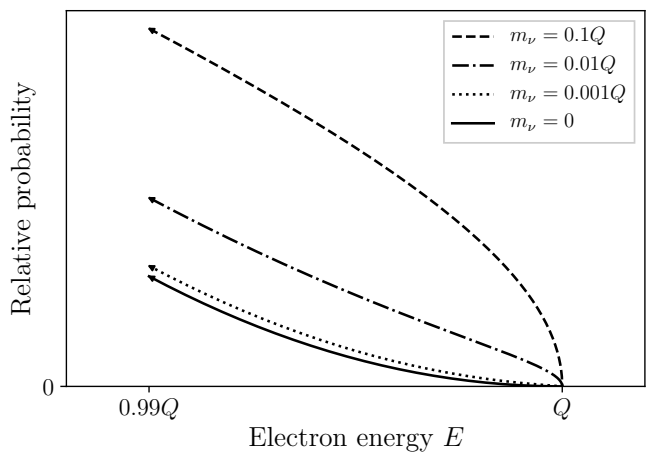

Figure 1.3: The tail of Fermi's beta-decay spectrum, for various values of the neutrino mass. With the modern limit of $\frac{m_\nu}{Q} \lesssim 10^{-7}$, the $m_\nu = 0$ curve is accurate within a few parts per billion. Adapted from [4].

These experiments begin with the Solar Neutrino Paradox, another interesting anomaly in particle physics [1]. By the 1960s, physicists had a fairly good idea of how the sun works: hydrogen in its core fuses, through a variety of steps, into helium. This requires the conversion of protons into neutrons, which occurs via a W boson in a process akin to beta decay in reverse, shown in figure 1.4.

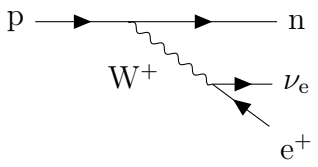

Figure 1.4: Positron emission, as in solar fusion; a proton becomes a neutron, emitting a positron and an electron neutrino via a W boson.

The energy released by this process is substantial, and so the sun is heated and incandesces, providing the Earth with light. However, just like in Earth-based beta decay, some of the energy is carried off by the neutrino; even the sun is not large enough to substantially attenuate the neutrino flux. These solar neutrinos escape the sun's core, and about eight minutes later a small fraction of them (about 10^{29} per second) reach the Earth.

In 1968, with the famous Homestake experiment, the solar neutrino flux was measured for the first time. Located underground to attenuate the cosmic-ray background, this experiment observed the interaction of solar neutrinos with several hundred tonnes of C_2Cl_4 [8]. Occasionally, the neutrinos were expected to cause chlorine atoms to transmute to argon ($\nu_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e^-$), in a process identical to beta decay but with the outgoing $\overline{\nu}_e$ replaced by an incoming ν_e (figure 1.5a). After a few months, the experiment had produced a few dozen argon atoms (the process occurs via the *weak* interaction, remember), but this was only about a third the number predicted by contemporary models of solar fusion [9]. While this was widely thought to be the result of some experimental or theoretical error, rather than a legitimate discrepancy, it did hold up under scrutiny [1], and - almost immediately - a solution was proposed: Pontecorvo's *neutrino oscillations*.

Pontecorvo's solution was simple: he suggested that the mass eigenstates of the neutrino are different from the weak eigenstates, as was already known to be the case for quarks [10]. The mass eigenstates are the sorts of neutrino that are unchanged under propagation (for example, if ν_1 is a mass eigenstate, then a ν_1 in free space will always be a ν_1). The weak eigenstates, on the other hand, are unchanged in the weak interaction (if a ν_e emits a Z boson, it stays as a ν_e and does not become a ν_μ, because ν_e and ν_μ are eigenstates of the weak interaction). Specifically, Pontecorvo proposed that

$$\begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} \tag{1.1}$$

where θ is a mixing angle, relating the mass eigenstates ν_1 and ν_2 to the flavour

eigenstates ν_e and ν_μ (ν_τ was not yet known).[2] If θ is zero, the two sets of states are equivalent, and there is no oscillation. Otherwise, however, the states mix, and it's perfectly possible that a ν_1 originated as a ν_e in solar fusion, but interacted with the Homestake detector as a ν_μ (that is, did not interact at all).

Calculation of the oscillation probability is straightforward; applying equation 1.1 and then the Schrödinger equation $i\dot{\psi} = H\psi$ to a state that starts as a ν_e gives the time evolution of the ν_1 and ν_2 components, and the probability of observing it as a ν_μ at a later time t:

$$P_{\nu_\mu}(t) = |\nu_\mu(t)| = \sin^2(2\theta)\sin^2\left(\frac{E_2 - E_1}{2}t\right) \tag{1.2}$$

From this, it can be seen that what starts as an electron neutrino will oscillate back and forth between ν_e and ν_μ, with an amplitude and frequency determined respectively by θ and $E_2 - E_1$, the energy difference between ν_2 and ν_1. This can easily be seen to be outside the Standard Model; flavour conservation is violated when the electron neutrino ν_e turns into a muon neutrino ν_μ.

The Standard Model also fails to account for the dependence of the oscillation frequency on $E_2 - E_1$. This difference must be nonzero for oscillations to occur, but for highly relativistic particles like solar neutrinos, travelling close to the speed of light, it works out to

$$E_2 - E_1 \approx \frac{m_2^2 - m_1^2}{2E} \tag{1.3}$$

The dependence on the mass-squared difference $\Delta m^2 = m_2^2 - m_1^2$ implies that, for oscillations to occur, the neutrino masses must be different and, therefore, that not all of them can be zero.

The Homestake experiment specifically observed the scattering of an *electron* neutrino (because an electron is involved in the scattering, and flavour is conserved in the SM), which made sense at the time; after all, the sun only produces electron neutrinos. More recently, taking the possibility of oscillations into account, experiments have sought to measure the *total* solar neutrino flux, in addition to its ν_e and non-ν_e components, via the scattering processes shown in figure 1.5.

The total flux was measured indirectly by Super-Kamiokande in 1998, and directly by the Sudbury Neutrino Observatory (SNO) in 2002. Their results were consistent with the theory of solar fusion, with the addendum that about two-thirds of the solar ν_e flux oscillated into ν_μ and ν_τ before reaching Earth. This

[2]Despite the omission of ν_τ, the two-flavour approximation of oscillations is quite accurate, because ν_e turns out to be mostly ν_1, and is much simpler than the full treatment.

showed that neutrinos do indeed have mass, and the experiments were jointly awarded the Nobel Prize in 2015.

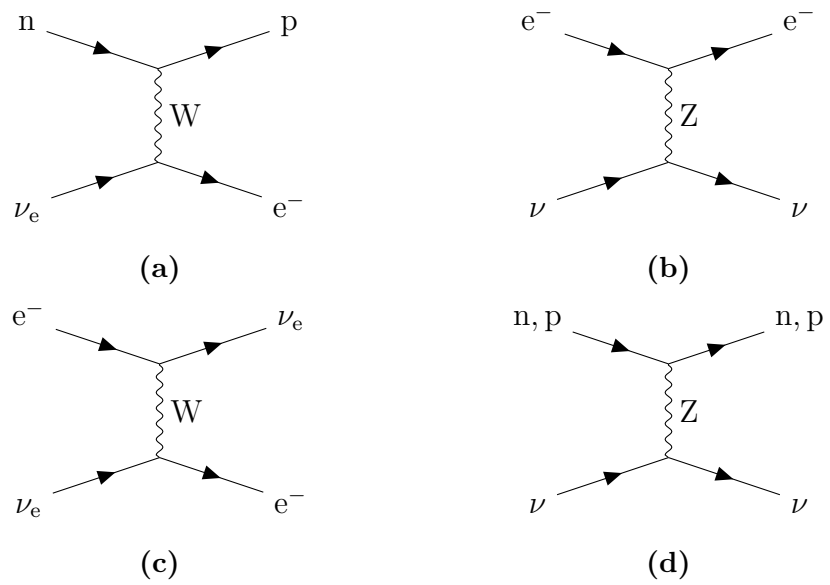

Figure 1.5: Interactions of neutrinos with matter. Process (a) is the reaction studied at Homestake, only sensitive to ν_e. Process (b), observed at Super-Kamiokande, is sensitive to all neutrinos, but is indistinguishable from the more common process (c), which is not. Finally, SNO observed all four processes, adding (d), which is sensitive to the total neutrino flux and can be distinguished from processes which are not.

1.2.2 Mass terms

Section 1.2.1 described oscillations between two species of neutrinos, ν_e and ν_μ. While this turns out to be a good approximation for solar neutrino oscillation, a full treatment requires the addition of the ν_τ and a third mass eigenstate ν_3, resulting in three mixing angles and mass-squared differences; $\Delta m_{21}^2 = m_2^2 - m_1^2$ and the corresponding θ_{12} roughly correspond to the Δm^2 and θ of solar oscillations, while the others can be measured in the short-range propagation of neutrinos originating in cosmic-ray showers, nuclear reactors, and particle accelerators. However, measuring Δm cannot give a measurement of any individual mass. The two masses involved in solar oscillation are known to be much closer to each other than they are to the third, but this provides no information on whether the third is the heaviest or the lightest,[3] leaving two possible scenarios: the *normal hierarchy*, in which $m_1 < m_2 \ll m_3$, and the *inverted hierarchy*, with $m_3 \ll m_1 < m_2$. The parameters as they are currently known are shown in table 1.1.

[3]by convention, m_2 is larger of the close pair and m_1 is the smaller

Parameter		Value
	Δm^2_{21}	$73.7^{+5.9}_{-4.4}$ meV^2
	$\sin^2\theta_{12}$	$0.297^{+0.057}_{-0.047}$
	Δm^2_{31}	2560^{+130}_{-110} meV^2
Normal hierarchy:	$\sin^2\theta_{23}$	$0.425^{+0.190}_{-0.044}$
	$\sin^2\theta_{13}$	0.0215 ± 0.0025
	Δm^2_{23}	2540 ± 120 meV^2
Inverted hierarchy:	$\sin^2\theta_{23}$	$0.589^{+0.047}_{-0.205}$
	$\sin^2\theta_{13}$	0.0216 ± 0.0026

Table 1.1: Neutrino mass-squared differences and mixing angles from oscillation experiments, as compiled in [11]. Uncertainties are to 3σ.

The refinement of these measurements is under way, as is the effort to measure the neutrino masses themselves; these masses and their ordering are outside the Standard Model, and as such are of immediate interest to the search for BSM physics. Furthermore, the fundamental nature of the masses themselves remains an open question, with substantial implications for the future of particle physics.

The core object of a quantum field theory - such as the SM, its predecessors, and most of its potential successors - is the Lagrangian density, $\mathcal{L}$. The contents of $\mathcal{L}$ encode the mechanics of the theory; for example, in some theory describing three scalar particles (A, B, and C), $\mathcal{L}$ might include ($\supset$) the following terms:

$$\mathcal{L} \supset g_1 ABC + g_2 A^2 B^2 - \frac{1}{2} m^2 C^2 \tag{1.4}$$

The first term describes an interaction with strength g_1 between one A, one B, and one C, and the second describes an interaction between two As and two Bs, with strength g_2. The third, however, quadratic in C and independent of A and B, is a special case; it identifies the mass of C, as m.

In the case of real-world fermions, the mass term becomes more complicated; there are *two* sorts of possible Lagrangian terms that describe mass:

$$\mathcal{L} \supset -m_D \bar{\phi}\phi + m_M \bar{\phi}\phi^c \tag{1.5}$$

The first of these terms describes a *Dirac* mass, like that in equation 1.4, while the second describes a *Majorana* mass [1,12]. The Dirac mass term is straightforward: it describes a process where a ϕ comes in ($\bar{\phi}$) and a ϕ goes out (ϕ); that is, the

propagation of a ϕ, governed by ϕ's mass, m_D. On the other hand, the Majorana mass term has ϕ^c, the charge conjugate - or antiparticle - of ϕ. This term describes the same thing as the Dirac term, except that the propagation includes a mixing between particle and antiparticle, like the mixing we've seen between ν_e and ν_μ. While there is no *a priori* reason to favour one sort of mass over the other, it can easily be seen that we cannot have Majorana masses in the Standard Model: the model's massive fermions (quarks and charged leptons) have charges opposite to those of their antiparticles and so cannot mix with them without violating the conservation of charge, so all SM masses are Dirac. However, if we venture a little outside the SM, we notice that the neutrinos in fact *do* have mass, and - since they are neutral fermions - their mass could potentially be Majorana.

The nature of the neutrino mass is interesting for two reasons. First, it is generally considered "unnatural" in particle physics for free parameters of a model to take on oddly specific values. For example, the *hierarchy problem* considers the mass of the Higgs boson. This particle gains extra mass-like terms from its interactions with other particles, and its observable mass is the sum of its "bare" mass and these extra terms. However, the extra terms are vastly larger than the observed mass, implying that the bare mass (which is, problematically, a free parameter) is "fine-tuned" to cancel them out almost exactly. A similar problem exists concerning the masses of the neutrinos. It is strange that these masses are so small; the twelve nonzero particle masses of the Standard Model are more or less evenly distributed between 10^6 and 10^{11} eV, while the neutrinos have masses well below 1 eV. This oddity suggests that the source of the SM mass terms - the Higgs field - is *not* the source of the neutrino masses, and that some other mechanism is at play. One potential solution to this is the *seesaw mechanism*, which would add new, massive neutrinos to the SM. Through interactions with the SM neutrinos and Higgs field, these new neutrinos would quite naturally acquire very large masses, while the old neutrinos would gain a very small mass - specifically, a Majorana mass [11].

Additionally, the seesaw mechanism in the early universe could have introduced an asymmetry between matter and antimatter, causing the universe to become dominated by one sort over the other. This feature is of great interest, because the actual universe displays such an asymmetry: it appears to consist almost entirely of ordinary matter (hence "ordinary"), while antimatter only exists as a short-lived product of high-energy processes.[4] This asymmetry is not accounted for in the SM, making it - along with those mentioned earlier in this chapter - an important hole in our current understanding of the universe.

[4] with the exception of the weakly-interacting antineutrinos

1.2.3 Double beta decay

As we have seen, it would be quite interesting to know if the neutrino is, contrary to the SM, a Majorana particle, with a strange mass term that mixes it with its antiparticle and permits the seesaw mechanism. However, both neutrino physics and physics beyond the Standard Model are notoriously difficult to investigate. In fact, only one signature of Majorana neutrinos is thought to be within current experimental reach: a process known as *neutrinoless double-beta decay* ($\beta\beta0\nu$) [11].

This chapter started with an introduction to beta decay. An atomic nucleus undergoing this decay transmutes into one of a different element, with one less neutron, one more proton, and less energy. The energy lost by the nucleus is released in the form of an electron and a neutrino. This sort of decay is only allowed in specific nuclides, with sufficient energy left over after the transmutation to produce the electron and neutrino.

An interesting consequence of this is illustrated by the decay of ^{136}Xe [13]. One might suspect that this nuclide could beta-decay into ^{136}Cs, but it in fact cannot; the energy difference between the two is only 0.4 MeV [11], not enough to create an electron (which has a mass of 0.5 MeV). However, this does not mean that ^{136}Xe is stable; as it happens, ^{136}Ba, the next nuclide over from ^{136}Cs, is 3 MeV less massive, so the decay of ^{136}Xe to ^{136}Ba leaves more that enough energy to create the two necessary electrons. This process is called *double* beta decay (figure 1.6a). Because two simultaneous beta decays require twice as many simultaneous weak interactions, double beta decay is extremely rare; many other isotopes of xenon single-beta decay with half-lives of seconds to hours, but the half-life of ^{136}Xe is vastly longer, at $2 \cdot 10^{21}$ years [14].

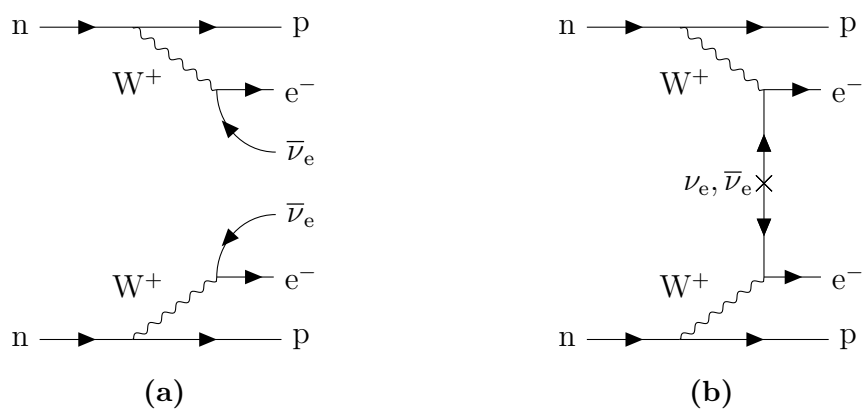

Figure 1.6: Double beta decay: (a) shows the ordinary mode, emitting two electrons and two neutrinos, while (b) shows the conjectured neutrinoless mode. The Majorana mixing of $\overline{\nu}_e$ with ν_e is denoted by $\times$.

Double beta decay is of note to us because the two electrons are accompanied by two neutrinos. Standard Model neutrinos would simply escape the detector, carrying off missing energy. However, if neutrinos have Majorana mass, there is a small but non-negligible chance that one of the neutrinos (which, remember, are specifically $\overline{\nu}_e$) would become a ν_e through mixing (the $\times$ in figure 1.6b), and would annihilate with the second $\overline{\nu}_e$. The energy spectrum of this process is starkly different from that of ordinary double-beta decay (see section 2.1), providing a clear signature of Majorana neutrinos. Additionally, the ratio of $\beta\beta0\nu$ to ordinary double-beta decay ($\beta\beta2\nu$) depends on the Majorana mass of the neutrino [15], and measurement of the half-life of this process is potentially the best way to obtain a direct measurement of that mass. This relationship between mass and half-life is shown in figure 1.7, along with the ranges in which the masses might be. Experiments attempting to do this are under way, but the task is very difficult; $\beta\beta0\nu$ is expected to be orders of magnitude rarer than any known process, with half-lives approaching 10^{30} years.

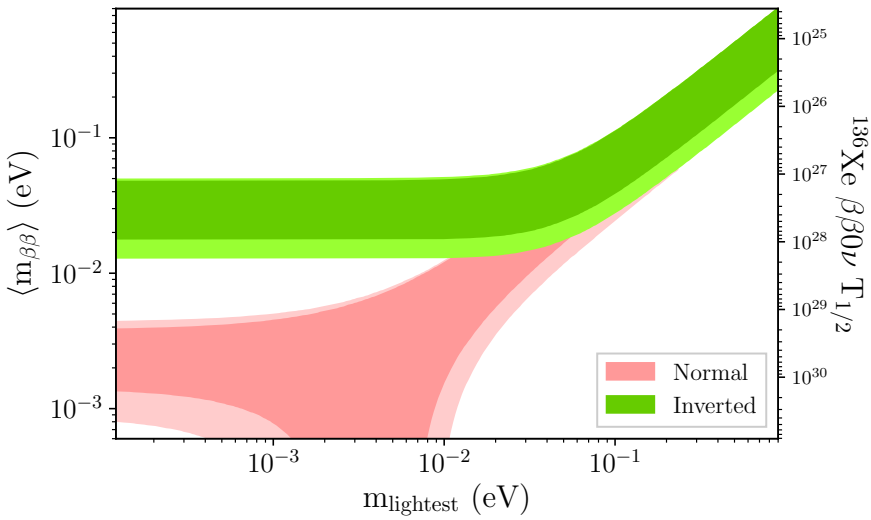

Figure 1.7: The neutrino mass parameter-space, projected to the lightest mass and the effective Majorana mass. The normal and inverted hierarchies as constrained by oscillation experiments are shown in red and green, with lighter areas indicating 3σ uncertainties. The secondary vertical axis shows conservative calculations of the $\beta\beta0\nu$ half-life of ^{136}Xe corresponding to the effective Majorana mass. Adapted from [16].

Chapter 2

The Enriched Xenon Observatory

2.1 Searches for $\beta\beta0\nu$

As with the investigations of single-beta decay that led up to the discovery of the neutrino, the investigation into neutrinoless double-beta decay revolves around the missing energy carried off by neutrinos, which is much easier to detect than the neutrinos themselves. If the two neutrinos were to annihilate, the full decay energy would be in the electrons; none would be missing. This would cause an excess of decays with electron energies that sum to the Q-value of the decay, as shown in figure 2.1, though this excess would be very small and difficult to detect.

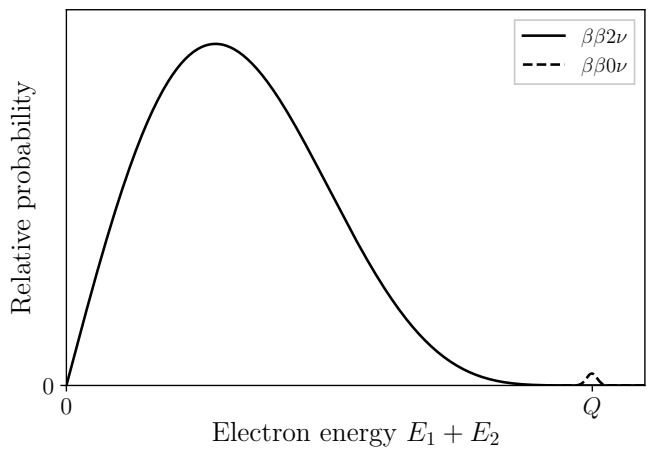

Figure 2.1: The spectra of the two modes of double beta decay, including a 1% detector resolution. The $\beta\beta0\nu$ peak is exaggerated for visibility; the experimental upper limit on the rate of neutrinoless decays is smaller than shown by a factor of 100.

There are a number of different approaches to the search for $\beta\beta0\nu$, largely due to the wide variety of double-beta-decaying nuclides. Different elements and isotopes, with different Q-values, availabilities, and chemical properties, can be

exploited in different ways to improve experimental sensitivity [13]. Some experiments have used the traditional layout of a radioactive source surrounded by detectors for tracking and calorimetry, which has good background rejection but poor energy resolution, efficiency, and scalability. However, certain candidate nuclides permit the use of the $\beta\beta0\nu$ source *as* the detection medium, which can allow for large, efficient, high-resolution searches, though further effort is required to give this approach competitive background rejection capabilities.

2.2 EXO

The EXO is an experimental effort of the second sort mentioned above, using liquid ^{136}Xe as both source and detector. A 100kg-scale experiment, EXO-200, completed its experimental run in late 2018, and its tonne-scale successor, nEXO, is currently under R&D.

2.2.1 The EXO-200 detector

EXO-200, the first iteration of EXO, monitored an isotopically enriched LXe TPC for potential $\beta\beta0\nu$ events. Xenon was chosen for a variety of reasons; in addition to the scalability offered by its dual role as source and detector, and its lack of background radioisotopes, xenon's fluid nature makes it rather unique among $\beta\beta0\nu$ candidates. It is naturally a gas, and can be condensed into a liquid at reasonable temperature and pressure, allowing for the continuous radiopurification required for a rare-event search, and potentially for reuse of the xenon in a later, more advanced experiment. Additionally, its gaseous nature at SATP lends itself to relatively easy isotopic enrichment via centrifugation. The fraction of ^{136}Xe in the EXO-200 LXe was enriched to 80.6% from a natural abundance of 8.9%, to benefit detector size and performance [15].

EXO-200 was located at the Waste Isolation Pilot Plant (WIPP) in New Mexico, an experimental deep geological repository for transuranic waste from the manufacture of nuclear weapons. Located 660m below the surface in a purpose-built salt mine, WIPP also offers a dedicated area for unrelated underground science. Rare event searches like EXO are impossible on the Earth's surface, due to the substantial background produced by cosmic rays; WIPP's overburden of rock provides 1585 m.w.e (metres of water equivalent) of shielding from cosmic rays, reducing the cosmic-ray background by a factor of $\sim10^8$ [11]. The relatively radio-pure salt body also provides a measure of shielding against the radionuclides present in the surrounding rock [17].

Xenon's electronic structure, high density, and high atomic number make it an effective material for a scintillation detector. High-energy particles incident on the xenon volume quickly deposit their energy into the detector, producing, on average, an electron-ion pair for every 15.6eV deposited. Excited xenon dimers Xe_2^* are also produced (both directly and through recombination of the electron-ion pairs), which then relax to the ground state through emission of a photon. This scintillation light is emitted with a wavelength of 178nm, plus or minus a few percent. On average, it takes at most 13.8eV of deposited energy to produce one scintillation photon; higher ionization densities increase the probability of recombination, increasing the scintillation light yield by a factor of up to 3 [18]. The EXO-200 TPC exploits these two processes, ionization and scintillation, to generate a 3D reconstruction of decays.

The time projection chamber consists of a cylindrical volume of liquid xenon with a strong axial electric field applied. Ionizing particles generated in the decay of ^{136}Xe and background radionuclides deposit their energy along short tracks ($\sim$1cm) in the LXe. The resulting scintillation light is collected by an array of avalanche photodiodes on either endcap of the TPC, while the ionization electrons are extracted to one endcap by the electric field. Once there, a grid of wires collect the charge, providing a two-dimensional reconstruction of the event location. The third dimension is provided by the delay between the scintillation and ionization signals (which arecalled S1 and S2, respectively) based on the known speed of drifting electrons; this is the "time projection". In EXO, this drift velocity is on the order of millimetres per microsecond, giving drift times on the order of milliseconds (see table 2.1).

For each event, the amount of scintillation and ionization detected is used to reconstruct the total energy deposited into the xenon. The signals are also used for background rejection; alpha decays leave short, dense tracks of ionization in the detector, and can be efficiently identified by their high ratio of S1 to S2. Two additional lines of background rejection are performed based on the position reconstruction. Particles coming from outside the TPC are attenuated by the LXe, so events near the edge of the volume are less likely to be genuine double-beta decays than events closer to the center. Also, γs have a higher tendency than βs to produce multi-site events, with separate, resolvable regions of ionization, so these events are discarded for their lower signal fraction.

EXO searches for $\beta\beta0\nu$ by looking for an excess of decays around the Q-value of ^{136}Xe, but a useful energy spectrum requires a substantial amount of exposure. Even in EXO-200's $\sim10^{27}$ atoms of ^{136}Xe, a double-beta decay is expected every couple of minutes, and no more than one $\beta\beta0\nu$ every few months [17].

2.2.2 EXO-200 results

EXO-200 was operational from 2011 to 2018; and an analysis of the full dataset was published in 2019 (see figure 2.2). The first two months of data-taking gave the world's-first observation of the decay of ^{136}Xe [19] and, after further operation, the half-life of this process was refined to [20]

$$T_{1/2} = 2.165 \pm 0.016\,(\text{stat}) \pm 0.059\,(\text{sys}) \cdot 10^{21}\text{y}$$

So far, neutrinoless double-beta decay has not been observed by EXO, nor any other experiment. Analysis of the full dataset taken by EXO-200 limits the half-life of the process to above $3.5 \cdot 10^{25}$y at a 90% CL, while the experimental sensitivity (the experiment's median *expected* half-life limit, given the observed backgrounds and assuming no $\beta\beta0\nu$ signal) is $5.0 \cdot 10^{25}$y [21].

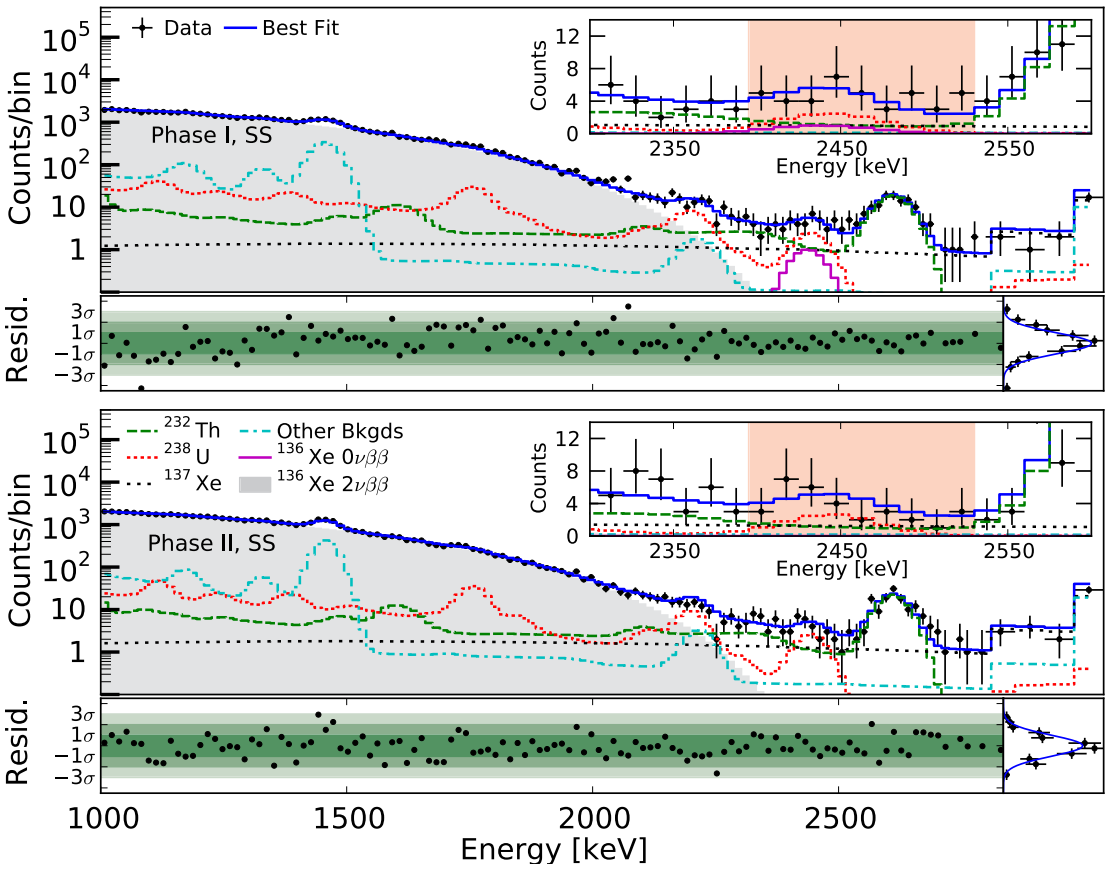

Figure 2.2: Results of the full EXO-200 $\beta\beta0\nu$ analysis, showing the measured energy spectrum and fits for signal and background sources. Insets show the $\beta\beta0\nu$ region of interest, and residual discrepancy between fits and data are also shown. Reprinted from [21].

2.2.3 Other $\beta\beta0\nu$ results

The highest half-life sensitivity to date in the search for neutrinoless double-beta decay has been obtained by the KamLAND-Zen experiment, at $5.6 \cdot 10^{25}$y [16], constraining the Majorana neutrino mass to below 61 to 165 meV, as shown in figure 2.3. The uncertainty on the constraint arises from the operators describing the nuclear transition of the decay. This constraint is approaching the region between 20 and 50 meV, in which the mass is expected to be in the inverted hierarchy (see section 1.2.2). The next generation of $\beta\beta0\nu$ experiments are expected to reach this region and allow probing or exclusion of the inverted hierarchy, making them of substantial importance to theoretical particle physics and cosmology.

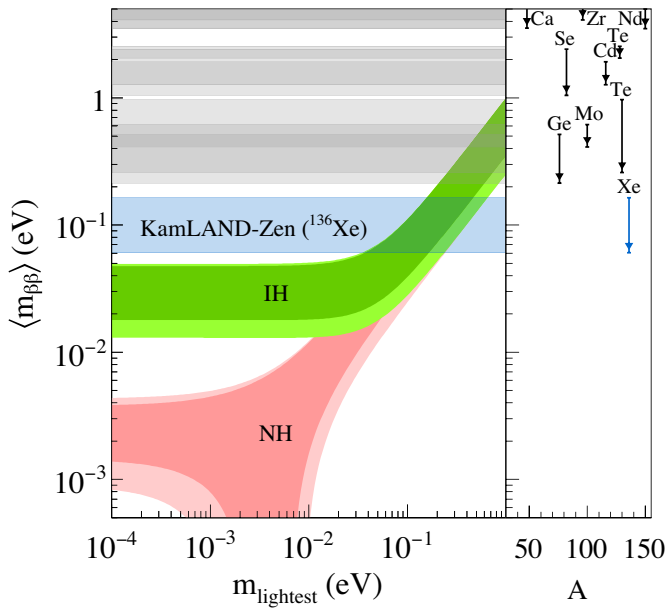

Figure 2.3: The neutrino mass parameter-space, projected to the lightest mass and the effective $\beta\beta0\nu$ mass. The red and green regions show the normal and inverted hierarchies as constrained by oscillation experiments, with lighter areas indicating 3σ uncertainties. Horizontal bars show 90% CL upper limits from $\beta\beta0\nu$ experiments, with nuclides shown on the right. Reprinted from [16].

2.3 nEXO

The next-generation EXO (nEXO) experiment, currently undergoing R&D, will be many times larger and more sensitive than current $\beta\beta0\nu$ experiments, nominally observing a five-tonne volume of enriched LXe over ten years. This will give it a baseline $\beta\beta0\nu$ half-life sensitivity of 10^{28} years and Majorana mass sensitivity between 6 and 18 meV, resulting in a very strong potential to shed light on the mass ordering, if not the masses themselves.

2.3.1 The nEXO detector

The experiment will follow the same architecture as EXO-200, with a LXe TPC in a low-background environment precisely measuring the decay spectrum of ^{136}Xe. Physical parameters of the two experiments are shown in table 2.1. The key advantage of nEXO is the 33-fold increase in the amount of ^{136}Xe and thereby the exposure to double beta decay, which will be many times larger than that of any existing experiment. Additionally, the larger TPC will have superior self-shielding, whereby background radiation from external sources is attenuated in the outer LXe. The attenuation length is at most 10 cm [14], which is more than a quarter the length and diameter of EXO-200 but less than 10% those of nEXO, increasing the fraction of fiducial[1] xenon in the larger experiment.

Parameter	Value (EXO-200)	Value (nEXO)
LXe mass (total)	175 kg	5109 kg
LXe mass (fiducial)	110 kg	4038 kg
^{136}Xe fraction	80.6%	90%
TPC Length	38.4 cm	125 cm
TPC Diameter	36.6 cm	116 cm
Readout pitch	9.0 mm	3.0 mm
Electric field	380 V/cm	400 V/cm
Electron drift velocity	1.71 mm/μs	similar
Rock overburden	660 m	2070 m
Water-equivalent overburden	1585 m	6010 m
Lifetime ^{136}Xe exposure	234.1 kg · y	∼8000 kg · y

Table 2.1: Physical parameters of the EXO-200 detector compared with design values for nEXO. Overburden values for nEXO are for its nominal location, the Cryopit at SNOLAB. Values from [14, 17, 22].

Optimizations over the EXO-200 architecture and potentially entirely new technologies will also be present in nEXO. A four-fold increase in overburden will reduce cosmic-ray backgrounds, a finer charge readout and more mature xenon purification techniques will improve energy resolution, and the experience gained in constructing a clean and radio-pure EXO-200 will further reduce nEXO's backgrounds. New, aggressive techniques for background rejection are also under development, which could be implemented in an upgrade to the initial nEXO detector, a larger experiment succeeding nEXO, or both.

[1] the *fiducial* xenon is that which is actually studied, after events which are too close to the edge of the volume are cut from the analysis.

2.3.2 Background rejection

Despite significant shielding and background discrimination, EXO-200 suffers from substantial backgrounds near the ^{136}Xe Q-value. As figure 2.2 (on page 17) shows, the decays of ^{232}Th , ^{238}U, their daughter nuclides, and ^{137}Xe are significant to the EXO-200 analysis.

A detailed study of nEXO's sensitivity to $\beta\beta0\nu$ has been published [23], and includes analysis of the backgrounds that will be present in the experiment. The small fraction of $\beta\beta2\nu$ events with electron energies near the Q-value cannot be separated from the neutrinoless signal by any physical or analytical means, and are taken as a best-case background. Taking backgrounds smaller than $\beta\beta2\nu$ as negligible, there are three sources with significance to the analysis:

1. **Beta decay of cosmogenic ^{137}Xe.** Incident cosmic rays can transmute ^{136}Xe into ^{137}Xe, which then decays with a half-life of 3.8 min and a Q-value of 4.2 MeV. A significant fraction of these decays will have electron energies near the ^{136}Xe Q-value, and these events are difficult to tag during analysis for several reasons. Single and double-beta events look much the same in the EXO-200 detector, and the cosmogenic isotope is uniformly distributed in the LXe, precluding positional and S1/S2-based rejection. Conventional cosmic-ray vetoing is also ineffective, as the ^{137}Xe lifetime is larger than the timescale of double beta decay.

2. **Trace ^{222}Rn present in the LXe volume.** Radon is part of the ^{238}U decay chain, and as such emanates slowly from natural materials, notably the inner surfaces of the xenon circulation system, from where it is transported into the TPC. A daughter of ^{222}Rn, ^{214}Bi, is especially problematic; it decays via a γ of 2.448 MeV, only 10 keV lower than the Q-value of ^{136}Xe. Unlike alpha particles, gamma rays are difficult to distinguish from betas in analysis.

3. **Long-lived radionuclides in detector materials.** The same nuclides that emanate ^{222}Rn into circulating LXe are present in the detector itself, most importantly ^{232}Th and ^{238}U in the copper TPC vessel. The decay chains of these two nuclides introduce a variety of β and γ backgrounds, with the most notable again being the ^{214}Bi line.

While the last of these sources can be partially mitigated using site information, they are otherwise quite stubborn and have a dramatic effect on the sensitivity of the experiment, as illustrated by figure 2.4.

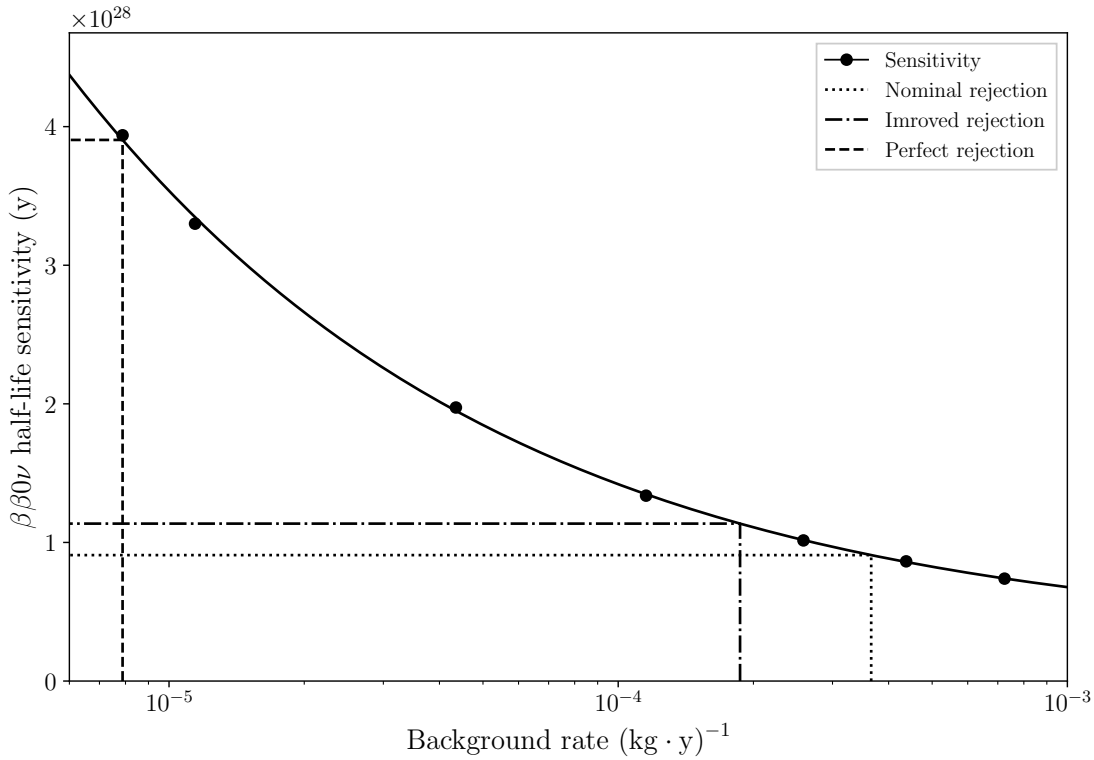

Figure 2.4: $\beta\beta0\nu$ half-life sensitivity of nEXO as a function of background rate. Simulated sensitivities are plotted, with a trendline. Broken lines show particular values of the background rate and resulting sensitivity, corresponding to nominal and improved rejection in nEXO, and to hypothetical "perfect" rejection in which the only remaining background is $\beta\beta2\nu$. The shown backgrounds are in the inner 2000 kg of LXe, with energy in the $\beta\beta0\nu$ region of interest. Adapted from [23].

Evidently, while conventional techniques can meaningfully improve background rejection and give a 25% increase over nEXO's baseline sensitivity, it is conceivable that backgrounds could be reduced much further, increasing sensitivity by a factor of more than four. This would double the experiment's Majorana mass sensitivity, placing it very solidly below the inverted hierarchy and overlapping with the top of the normal hierarchy (see figure 2.3). This makes "perfect" background rejection a very appealing concept, and multiple methods to realize it are currently being researched.

One such technique, which has been in development for nearly three decades, is *barium tagging* [24]. While the electrons produced in the decay of ^{136}Xe behave much the same as a background single-β or γ, the daughter nuclide, ^{136}Ba can in principle provide a unique tag to identify $\beta\beta$ events. One recent R&D effort [25] demonstrated the ability to identify single atoms of barium trapped in a matrix of solid xenon. The xenon, containing some atoms of barium, was frozen onto a sapphire window. A laser and photodetector located behind the window were used to induce and measure fluorescence in the barium, accurately detecting and

counting the atoms. The matrix was then evaporated to make way for the next sample. However, before this technique can be implemented in a $\beta\beta0\nu$ search, the apparatus would have to be mounted in a probe capable of quickly conveying it through the TPC to the location of a candidate decay, and would have to be able to efficiently capture the single Ba^+ ion, requiring significant further development.

A perhaps simpler approach to achieving "perfect" background rejection is via directly addressing EXO-200's inability to distinguish $\beta\beta$, β, and γ events by their signals in the TPC. The research presented in this thesis concerns the development of such an approach; specifically, the characterization of a novel *two-phase* xenon TPC with superior resolution to the existing single-phase EXO TPC, and its potential to allow for pulse-shape discrimination in the search for $\beta\beta0\nu$.

Chapter 3

A Two-phase EXO Prototype

3.1 Two-phase time projection chambers

A two-phase TPC is an extension of the standard TPC architecture, substantially improving energy resolution at some cost to detector complexity. While ordinary noble-liquid TPCs, such as EXO-200, have proven useful for rare event searches, they are limited by their energy resolution; the charge readouts generally have a resolution no better than 10keV, and attempts to introduce charge amplification in the xenon have not been particularly successful [26]. The two-phase architecture bypasses this difficulty by converting the charge signal to an amplified light signal while still in the detector volume, and then reading out this light rather than the charge. The gain and resulting resolution allowed by this technique have led to its adoption by experiments such as XENON1T and LUX, which search for dark-matter nuclear recoils at energies on the order of 1 keV, low enough to render conventional charge readouts useless [18].

The unique ability of a two-phase TPC to amplify the ionization signal arises from a layer of gas that is allowed to form above the main liquid volume. When electrons reach the phase boundary, they are extracted by the electric field into the gas.[1] An additional field-shaping electrode slightly below the phase boundary maintains an electric field in the gas which is significantly higher than that in the rest of the TPC, accelerating the electrons through the gas until they reach the anode. Collisions of these accelerated electrons induce EL in the gas, with achievable yields of thousands of photons per electron [26]. The array of photodetectors which collect the scintillation signal then also collect this electroluminescence signal, which is of the same wavelength.

Two-phase xenon TPCs have seen notable success in rare event searches. XENON1T and LUX are leading experiments in the effort to directly detect dark

[1] extraction efficiency is close to 100%, for electric fields above 1 kV/mm in xenon [18]

matter [27, 28], and achieve excellent background rejection by exploiting the two-phase architecture. The ratio of S2 (which cannot be measured at such low energies in a single-phase detector) to S1 provides them with an efficient tag for discriminating between signal-like nuclear recoils and background-like electron recoils, just as it allows EXO-200 to reject background α events. However, as we saw in section 2.3.2, this ratio does not provide an useful tag to distinguish EXO's signal-like $\beta\beta$ events from the background of βs and γs.

3.2 The potential of a two-phase EXO

While a simple S2/S1 cut is of limited usefulness in $\beta\beta0\nu$ searches, a two-phase detector could offer a different advantage to EXO. In the current architecture, the relatively small number of counts in the ionization signal necessitates slow shaping circuits to obtain a useful signal; EXO-200 uses shapers with time constants on the order of 1-10 μs, and digitizes the signals at 1 MHz [17]. The large gain offered by electroluminescence circumvents this difficulty, allowing for faster readouts and for a detector to reach the intrinsic temporal resolution limit (due to electron diffusion) of a LXe TPC. In particular, the superior timing resolution of a two-phase EXO could provide the detector with two useful abilities:

1. **More accurate reconstruction of ionization density.** Recombination is a good signature of ionization density; more electron-ion pairs close together are more likely to recombine into atoms. However, measuring recombination via the resulting loss of ionization (as in EXO-200's α rejection) is relatively insensitive; in the high electric field of a LXe TPC, only around 10% of the ionization recombines [29], making it easy to lose the recombination signal in the larger ionization signal. It would be preferable to be able to measure the recombination directly, independent of the total ionization, which may be possible by examining S1 with good time resolution. When ionizing radiation deposits energy in xenon, two sorts of excited states are produced, which decay to the ground state via scintillation with lifetimes of 4 and 22 ns. The shorter-lived state is preferentially produced via recombination in dense ionization [30], so measurement of the shape of S1 could allow superior reconstruction of ionization density via the ratio of slow to fast scintillation. This method is precisely how the DEAP experiment achieves its titular pulse-shape discrimination [31].

2. **Resolution of detail within the ionization cluster.** A typical TPC reconstructs the axial position of an event based on the time of S2. By the same principle, the axial *shape* of the ionization cluster is encoded in

the temporal *substructure* of S2. Reconstructing this shape is impossible at EXO-200, which has a time resolution similar to the time-projected width of an ionization cluster (1μs). With a fast detector with two-phase amplification, however, one is able to reach diffusion-limited resolutions of tens of nanoseconds (equivalent to an axial resolution on the order of 10μm), giving a detailed one-dimensional profile of the ionization cloud.[2]

Both of these potential abilities are of interest due to the phenomenon of *Bragg peaks*. When high-energy charged particles pass through matter, they leave a long, narrow ionization track that curves due to interactions with the medium.2 The track is more or less uniform for most of its length, but ends in a much more densely-ionized region (called a Bragg peak) where the last of the particle's energy is deposited, due to an increase in the scattering cross-section as the particle loses energy [30, 32]. The most difficult background rejection for EXO is distinguishing between β and $\beta\beta$ events of the same energy, as they leave similar ionization trails and produce similar signals. However, single electrons produce a single Bragg peak, while pairs of electrons produce two. The size of the Bragg peak is reasonably independent of the initial energy, because it is the deposition of the last of the electron's energy; higher energy electrons leave longer trails, but end in similar peaks. This implies that $\beta\beta$ events have a larger fraction of their total ionization in the denser Bragg peaks, which may be detectable in a two-phase TPC as an increase in the short-lived component of the scintillation. This line of analysis is investigated in section 4.1. One might also be able to "image" the two Bragg peaks directly, as temporal structures visible in a high-resolution trace of the ionization signal; this possibility is explored in section 4.2.

3.3 The prototype

A small prototype two-phase EXO-like detector was constructed to directly evaluate the ability of the two-phase architecture to allow pulse-shape discrimination between β and $\beta\beta$ events in LXe. It is a minimalistic but complete two-phase xenon TPC, with S1 and S2 read out by a small four-pixel photodetector mounted above the center of the electroluminescent region. A γ source induces β-like and $\beta\beta$-like events in the LXe, and external detectors tag each event as one of the two types. An independent determination of the event type is then attempted, using only the scintillation and ionization signals from the TPC.

[2] see section 4.2.1 for more detail

3.3.1 Detector design

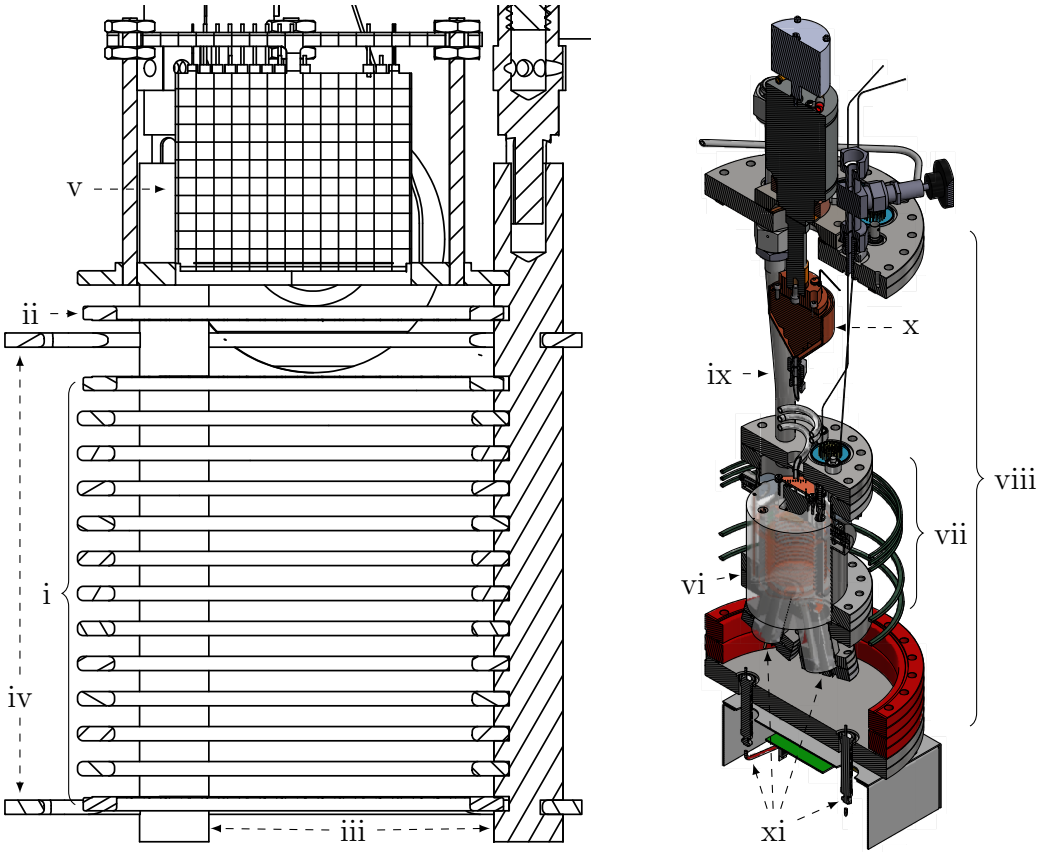

(a) **(b)**

Figure 3.1: Drawings of the two-phase TPC prototype. Details of the TPC are shown in (a), while (b) shows the TPC (vi) within the inner xenon vessel (vii) and outer vacuum vessel (viii). Individual components are described below.

TPC

The TPC itself is shown in figure 3.1a, and consists of the high-voltage electrodes that induce the electric field, the photodetector, and supporting structures. The drift field, which pushes electrons towards the readout, is applied by a set of 12 copper rings, labelled (i), which are connected in series through resistors. With a negative voltage applied to the bottommost of these rings relative to the topmost, the potential is evenly distributed through the resistors connecting the rings, applying a uniform electric field pointing downwards along the axis of the cylinder formed by the rings. Field uniformity at the ends of the cylinder is ensured by fine hexagonal copper grids bonded to the top faces of the top and bottom (i) rings (the *middle* and *bottom grids* respectively). Dimensions and electrical details of the field rings are provided in table 3.1. An identical 13th ring (ii), separated

from the top (i) ring by 8.0 mm and with a copper grid (the *top grid*) bonded to its bottom face, is held above the potential of the (i) rings to induce the high field for electroluminescence. The liquid level in the xenon is nominally halfway between this ring and the topmost field ring. The 13 rings are held in slots cut into the inside faces of three Macor ceramic support posts (two shown as iii), which are in turn held together by two copper retaining rings (iv) in slots on the posts' outside faces. The photodetector (v), 3.5mm above the top grid, is a Hamamatsu S13371-8554 four-pixel SiPM sensitive to the 178nm wavelength of scintillation in xenon. The anode of each pixel is connected in parallel to a 56 V bias (through a 1 kΩ resistor) and to ground (through a 100 nF capacitor) for low-frequency noise filtering, as advised by the manufacturer.

Also present in the xenon vessel is an ultraviolet LED for SiPM testing, as well as a capacitive liquid-level probe. The level probe consists of two long, narrow copper plates held apart by five disk-shaped Teflon spacers, all enclosed in a glass tube, which is mounted vertically to the outside of the retaining rings and extends from below the bottom grid to above the top one. Dimensions are given in table 3.2. As the TPC fills, LXe displaces the gaseous xenon between the plates, increasing their capacitance due to the liquid's higher permittivity. Calibration is performed using the spacers. As the probe fills, its capacitance increases proportionally to the fill level; however, because the spacers block the gas, the increase in capacitance pauses as the level passes a spacer. This probe is used to roughly measure the level as the TPC is filled with liquid; when the level is between the top

Parameter	Value
Inner diameter	61.6 mm
Outer diameter	51.1 mm
Thickness	2.03 mm
Pitch	5.00 mm
Inter-ring resistance	4 MΩ

Table 3.1: Dimensions and resistive properties of the field-shaping rings of the prototype two-phase TPC. The rings are shown in figure 3.1a (i).

Part	Dimension	Value
Copper plates	length	100 mm
	width	3.0 mm
	thickness	1.5 mm
Spacers	diameter	3.0 mm
	thickness	1.5 mm
	pitch	19.4 mm

Table 3.2: Dimensions of the level probe. The probe's length is along the axial direction of the TPC.

and middle grids, it is measured more accurately via the capacitance between the grids, which similarly increases as the gap between them fills with liquid. Traces of the two capacitances during filling are shown in figures 3.2 and 3.3, on page 31.

Vessels and support systems

The TPC is housed in a pair of coaxial, cylindrical steel vessels, shown in figure 3.1b. The dimensions of the vessels are given in table 3.3. The inner xenon vessel, shown as (vii), contains the xenon and the TPC (vi), while the outer vacuum vessel (endcaps shown as viii; curved face omitted for visibility) is evacuated to 10^{-10} atm to provide convective and conductive insulation. A sheet of Mylar between the two vessels provides radiative insulation.

The inner vessel is held within the outer by a pair of steel pipes (one shown as ix), which also passively extracts gaseous xenon from the top of the vessel. These pipes join outside the outer vessel and feed the xenon to a cryocooler (x), which condenses the xenon and allows it to drip back into the vessel, via a spout that extends into the liquid to minimize splashing. This loop can also be opened to an external system of pipes, which provides the functions necessary to fill the xenon vessel, evacuate the vacuum vessel, purify the xenon, and recover the LXe to a 520L storage tank if necessary. By manipulating valves in the pipes, one can:

- open a high-pressure xenon bottle to the TPC and the storage tank, for initial filling of the system with external xenon
- open the storage tank to the TPC via a cold getter (purifier), to fill the inner vessel with purified xenon
- open the storage tank to the TPC via a cold getter, a diaphragm pump, another cold getter, and a hot getter; pumping to raise the pressure in the TPC becomes useful to increase filling speed and prevent freezing as the pressure in the storage tanks drops below 1.3 atm
- circulate gas passively extracted from the TPC loop through a diaphragm pump, a cold getter, and a hot getter for purification, then back to the TPC
- allow gas to escape from the TPC back into the storage tank to lower the liquid level; the path can be opened manually, or through a pressure-relief check valve in case of overpressure above 2 atm
- evacuate the outer vessel, or any other segment of the gas system, using a turbopump

Finally, the cryocooler is complemented by a 32W resistive heater wrapped around the waist of the inner vessel, which is used to control the pressure in the TPC. The various modules used in the gas system are detailed in table 3.4. Sensors used to monitor the gas system and the TPC are listed in table 3.5.

Components inside the vessels are supported by a series of feedthroughs welded into the vessels. A pair of multipin connectors, one in the top of each vessel, allows access to the level sensor and the LED. An additional such connector in the top of the outer vessel allows access to the heater and the three temperature sensors. The four channels of the SiPM signal exit the vessels through SMA connectors, welded in the top of the inner vessel and gasketed to the cylindrical face of the outer. The top and bottom drift-field (i) rings are accessed via pairs of high-voltage ports in the bottoms of each vessel (shown as xi in figure 3.1b). Finally, the EL-field (ii) ring is connected to a BNC connector in the top of the outer vessel via the multipin connector in the top of the inner. This feedthrough can be shorted to the vessel to ground the grid for field-shaping, or connected in series with the middle grid to the capacitance meter for level measurement.

Vessel	Dimension	Value
Inner	diameter	97.38 mm
	length	152.40 mm
Outer	diameter	203.71 mm
	length	557.28 mm

Table 3.3: Interior dimensions of the two pressure vessels.

Function	Part	Model
Purification	Cold getter	SAES MicroTorr MC200-902FV
	Hot getter	SAES PS4-MT3-R-1
Circulation Pressurization	Diaphragm pump	KNF UNO35 STP
Evacuation	Turbopump	Edwards 750XT
Condensation	Cryocooler	SunPower Cryotel GT with AVC

Table 3.4: List of modules used in the gas system.

Parameter	Location	Sensor model
Gas pressure	TPC	Brooks CMCA73114BR
	Storage tank	Brooks GFD01A4BSM
Gas flow rate	Diaphragm pump	Omega FMA-A2308-55
Temperature	Cryocooler Inner vessel (top) Inner vessel (bottom)	LakeShore PT-103-AM
Capacitance	External	IET 1920

Table 3.5: List of sensors used in the TPC and gas system.

Slow control and detector operation

Four Omega Platinum CN81PT-220 PID controllers are used to collect data from the various sensors and relay commands to the heater. During normal operations, these controllers read from the pressure sensors in the TPC and storage tank and the temperature sensors on the inner vessel, and the TPC pressure drives a PID loop that controls the heater. The PID is set to maintain the pressure at a setpoint between 1.0 and 1.4 atm, and does so to within a few thousandths of an atmosphere. The cryocooler has an application-specific controller, which modulates cooling power based on input from the temperature sensor at the head of the cooler, where it interfaces with the xenon. The cooler is set to maintain this temperature at 165K, in the middle of the range where xenon exists as a liquid at the desired pressures, and does so to within a few hundredths of a Kelvin.

The filling of the TPC with liquid begins with the inner vessel open to the storage tank, the system containing roughly 2 atm of gaseous xenon at room temperature, and both the heater and cooler turned off. First, the cooler is turned on, and begins to condense the xenon, which evaporatively cools the warm piping and inner vessel. After about three hours, the vessel cools down sufficiently to begin filling with liquid. After two hours of filling, the level has risen to about half the nominal value, as shown in figure 3.2. At this point, a significant amount of the gas has condensed, causing the pressure to drop below 1 atm. This reduces the heat load on the cooler to below its minimum cooling power of 70W, which causes the xenon to freeze, blocking the pipe through the cooler. With the supply of cold LXe cut off, the liquid in the inner vessel will entirely evaporate over the next day or so. To prevent this, the diaphragm pump is used to periodically increase the pressure, which also heats the xenon as it is forced through pipes exposed to ambient conditions, stabilizing the cooler and allowing filling to continue. At this point, the capacitance meter is disconnected from the level probe and connected to the top and middle grids, and, about an hour later, the level reaches the grids and their capacitance begins to increase. The level is allowed to reach the top grid for calibration, and then the valve between the inner vessel and the storage tank is cracked open, causing the TPC pressure to slowly drop and the liquid to start evaporating. Once the capacitance reaches the midpoint between its minimum and maximum values, meaning that the level is halfway between the middle and top grids, the valve is closed, and the level is allowed to stabilize, with possible adjustments made by pumping more xenon into the vessel or venting it back into the storage tank. Traces of both capacitance are shown in figures 3.2 and 3.3.

Once the vessel has been filled to the appropriate level, the grids are disconnected from the capacitance meter to allow high voltage to be applied. The level

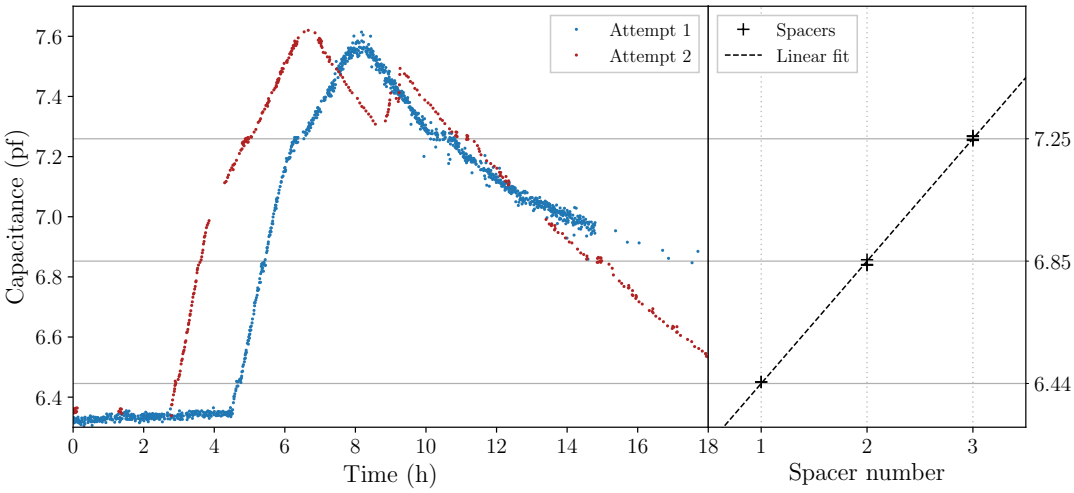

Figure 3.2: Traces of level probe capacitance for the first two attempts at filling the inner vessel. The traces show the capacitance rising with the level for several hours, until unforeseen icing around t = 7h causes the liquid to evaporate, reducing the capacitance. Both the rise and fall of the capacitance can be seen to slow periodically as the level passes each spacer. The bottom of the probe is just above the bottom of the TPC, while the top and middle grids are between the fourth and fifth spacers.

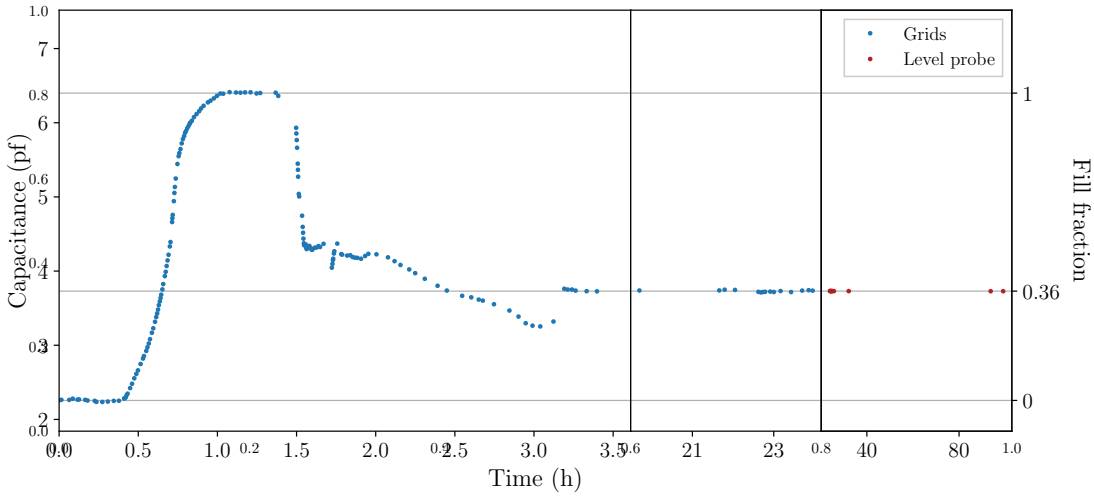

Figure 3.3: A trace of the grid capacitance, showing the TPC vessel being filled with LXe up to roughly the nominal level. From several hours before t = 0h until t = 0.5h, the capacitance is constant, as the liquid level rises below the grids. Between t = 0.5h and t = 1.0h, the level rises from the middle grid to the top grid, and the capacitance rises sharply. From t = 1.0h to t = 1.5h, the level is above both grids and the capacitance is again constant. At t = 1.5h, the level drops sharply as pressure is let out of the vessel, then wanders back and forth due to further evaporation and additional pumping. Finally, it stabilizes roughly between the grids, and remains there for the next four days.

probe is reconnected, to allow for continued monitoring. The level probe measurements can be calibrated to those of the grids using an offset and a scaling factor; the offset is determined by measuring each capacitance when switching from the grids to the probe after filling, while the scale can be found by comparing the change in capacitance as the level moves from one spacer to the next to the change as the grids are immersed. The heater is then set to maintain the desired pressure (between 1 and 1.4 atm), and high voltage is applied to the field-shaping rings.

Voltage is applied to the rings by a Bertan 380N -10kV supply. The supply is connected to the bottom of the drift-field rings, which, with the inter-ring resistors, have a series resistance of 44MΩ. The top ring is connected to ground through a 175MΩ resistor located outside the vessels. This serves to apply 20% of the potential difference across the drift-field rings and the remaining 80% between the top drift-field ring and the grounded top ring. This arrangement gives a drift field of up to 330V/cm (similar to that in EXO-200 [17]) and an electroluminescence field 30 times that, up to 10kV/cm (sufficient for efficient electron extraction and light amplification [18,26]). The ratio of the two fields can be adjusted by changing the resistance between the rings and ground.

3.3.2 Inducing events in the LXe

A gamma-ray source placed outside of the vessels is used to induce β-like and $\beta\beta$-like events in the xenon, which can be externally tagged as either type and used to simulate signal and background events in EXO. The source is a small disk of ^{228}Th, which has a γ spectrum dominated by a $Q = 2.614$ MeV line produced in the decay of its daughter nuclide ^{208}Tl. It currently has an activity of roughly 2 kBq. The dominant interactions of these gamma rays with xenon are Compton scattering, pair production, and photoelectric absorption, with cross-sections of 3.1, 0.58, and 0.081 mm^2/g, respectively [33].

Compton scattering is the elastic scattering of a photon off of a valence electron. The electron is ejected from its atom, with some energy imparted to it by the photon, and interacts with the medium until it comes to rest. The photon is deflected at an angle θ_{C}, with a lower energy E' and a higher wavelength λ' [34]:

$$\lambda' = \lambda + \frac{2h}{m_e c}\left(1 - \cos\theta_{\mathrm{C}}\right) \tag{3.1}$$

Pair production is the absorption of a high-energy photon by the electric field of an atom, resulting in the creation of an electron-positron pair sharing the full energy of the photon. The pair deposits their energy in the medium on a timescale

of 1 ps [32], with the positron behaving similarly to the electron [11]. Eventually, on a timescale of 100 ps [35], the positron, having deposited its kinetic energy into the medium, annihilates with an atomic electron, producing a pair of photons with opposite momenta and each with energy equal to $m_e c^2 = 0.511$keV, the mass of an electron or positron. These photons are emitted nearly isotropically.

Photoelectric absorption is the absorption of a photon by an electron, which is ejected from its atom with the full energy of the photon and then deposits this energy in the medium. No other particles are emitted.

Neither Compton scatters nor pair productions result in the net creation or annihilation of massive particles; Compton scattering is an elastic process, and pair production both creates and annihilates an electron and a positron. The net process for Compton scattering, then, is that a photon enters a medium and deposits some energy in the form of ionization due to a high-energy electron, and exits the medium with the remainder of the energy. As far as the medium is concerned, this looks just like a β decay. In pair production, the photon deposits energy in the form of ionization due to a high-energy electron *and* positron, and *two* photons exit the medium with the remainder, looking just like a $\beta\beta$ decay in which the two electrons have the same energy. These two processes are used to emulate background-like and signal-like events in the LXe.

It is ideal to select events of the two types with equal amounts of energy deposited as ionization, both to remove this quantity as a confounding variable and because the most important backgrounds at EXO are naturally those with energy similar to the $\beta\beta0\nu$ signal. This can be done fairly straightforwardly; in pair production, this energy is always that of the incoming photon minus the $2m_e c^2$ carried off by the outgoing photon pair, while in Compton scattering it is the energy of the incoming photon minus that out the outgoing one. The outgoing photon will sometimes have an energy near $2m_e c^2$; equation 3.1 tells us that this occurs when the scattering angle θ_{C} is near 46°.

Photon detection and tagging of the two sorts of events is accomplished using a pair of NaI scintillation detectors. These detectors, along with the ^{228}Th source, are placed just outside the outer vessel (∼20 cm from the TPC axis), at the same height as the center of the drift volume. An enclosure of 5×10×20 cm rubber-coated lead bricks is used to shield the detectors and TPC from external radiation, and to provide some degree of collimation to the source.

One of these detectors, henceforth referred to as the "rear" NaI detector, is placed with an angle of $\theta_{\mathrm{C}} = 46°$ separating it from the source. The other ("forward") NaI detector is placed directly across the vessel from the rear one. When photons are emitted by the source and Compton scatter with the desired energy in

the xenon, they will be scattered into a cone with angle θ_C; this cone intersects the forward detector. The photons produced by annihilation following pair production, on the other hand, are emitted in all directions, but always back-to-back, so that if one is emitted in the direction of the forward detector, the other will travel in the direction of the rear. The placement of the two detectors and topology of the gamma rays is visualized in figure 3.4. Events with both detectors registering a photon with $E \approx m_e c^2$ can therefore be tagged as pair productions, while events with the forward detector registering a photon with $E \approx 2m_e c^2$ can be tagged as Compton scatters of the appropriate energy.

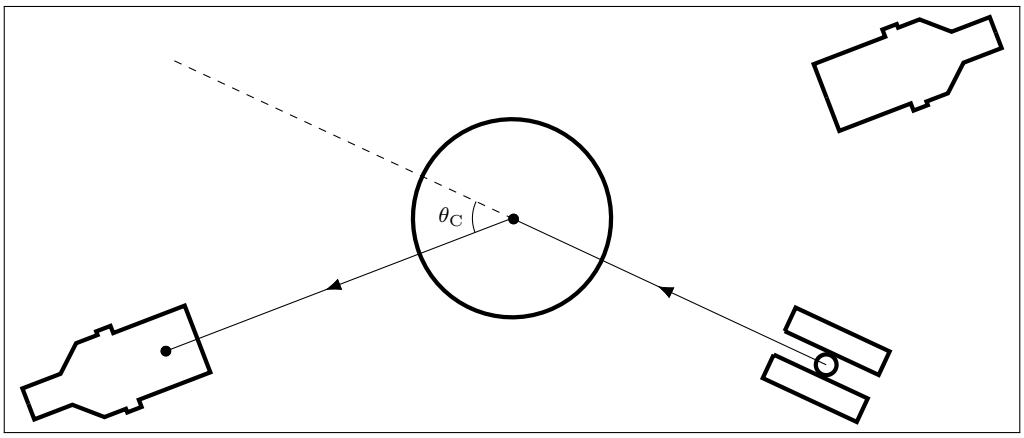

(a)

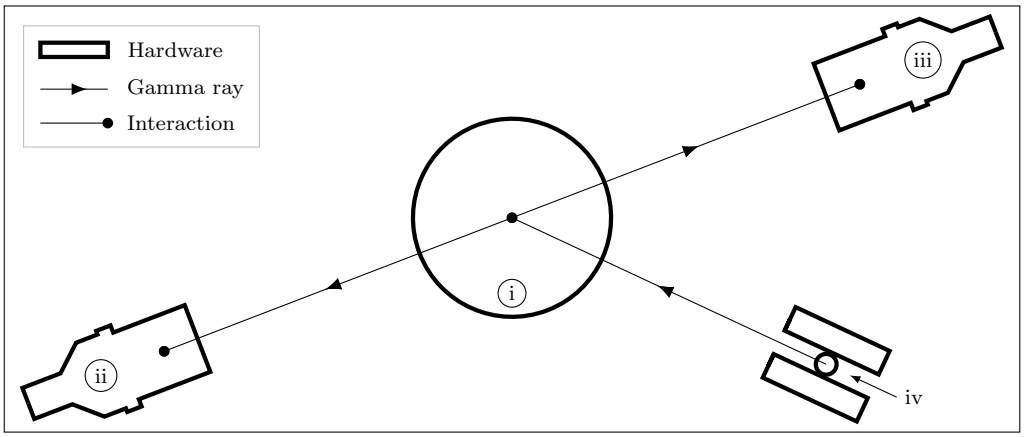

(b)

Figure 3.4: A digram of the apparatus described in section 3.3.2, showing the placement of the γ source and detectors around the TPC vessel (i). The forward and rear NaI detectors are labelled (ii) and (iii) respectively, and the source in its collimator is (iv). Compton-scattered photons of the desired energy fall in a cone with angle $\theta_C \approx 46°$; (a) shows one of these photons reaching the forward detector, which lies on that cone. Photon pairs due to pair production are emitted back to back along a random axis, which may intersect both NaI detectors, as shown in (b). The size of the TPC relative to the other components is exaggerated by a factor of 4 for clarity.

External backgrounds are judged to be acceptably low; with the data acquisition trigger described in section 3.3.3, the event rate with no source present was roughly two orders of magnitude lower than the rate with the ^{228}Th source in place. However, another background arises from within the detector itself: there is no guarantee that events will involve only a single interaction. For example, a photon may Compton scatter a second time in the xenon, producing the same final photon energy and amount of ionization as a single scatter, but leaving a very different pattern of ionization. It may be possible to identify these multi-site events using data from the TPC (as was done in the EXO-200 analysis [15]); however, additionally, a Monte-Carlo simulation was run to determine the magnitude of this effect in our prototype. The simulation modelled the transport and interactions of gamma-rays in xenon using cross-section data provided by [33] and the Compton-scatter sampling algorithm described in [36]. Interactions with the LXe and NaI volumes are simulated, producing energy deposits and outgoing photons, until all photons have escaped or been absorbed.

Events are tagged, based on the simulated signals in the NaI detectors, as pair productions, Compton scatters of the appropriate energy, or neither; the rate of multi-site events tagged as either type is shown in figure 3.5. They show that all but 26% of events tagged as Compton scatters are single-site Compton scatters and all but 22% of pair productions are single-site pair productions. The remaining multi-site events do likely introduce some systematic error into any attempt to identify event type based on TPC signals, but this does not compromise this experiment as a proof-of-concept; as long as significantly more than 26% of events can be classified correctly, the analysis is achieving some meaningful degree of background rejection.

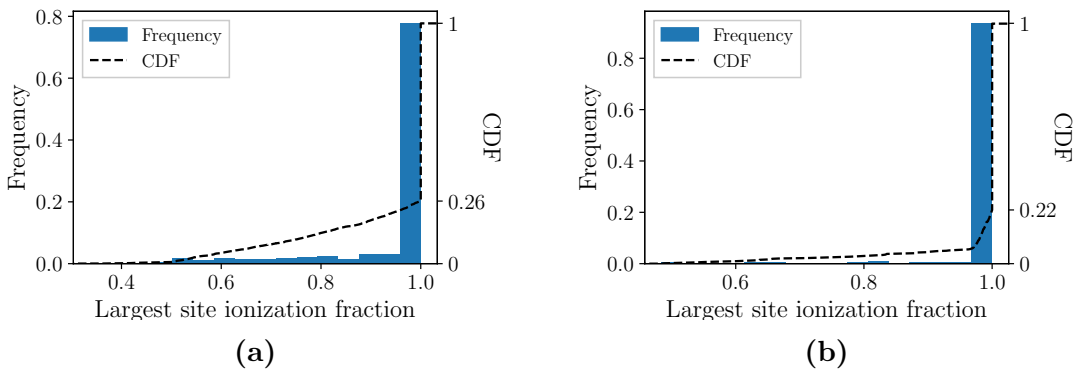

Figure 3.5: Histograms showing the extent to which events are multi-site, quantified as the fraction of the total energy which is deposited at the largest single ionization site. For example, this quantity will be 1 for a single-site event (all of the ionization is at the largest site), and 0.33 for an event with 3 equal sites. Compton scatters are shown in (a), pair productions in (b).

3.3.3 Data acquisition

The data acquisition system consists of six data channels (four channels for the four SiPM pixels, and two for the two NaI detectors) and an acquisition trigger. It functions to record a 65.5μs trace of all channels at 500MHz whenever a signal is seen simultaneously in the SiPM and the forward NaI detector. This trace is long and detailed enough to capture both S1 and S2 with high resolution, and the trigger records all events that are useful for the analysis described in section 4.1. With the ^{228}Th source in place, the trigger rate is roughly one count per second. A schematic of the system is shown in figure 3.6, and the modules used and their parameters are detailed in tables 3.6 and 3.7.

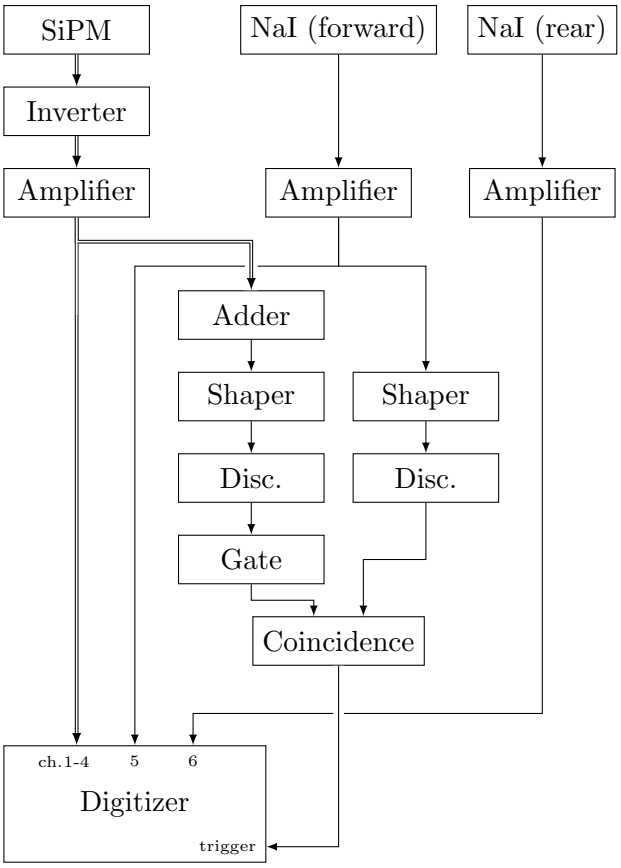

Figure 3.6: A schematic of the data acquisition system, showing the six data channels and the acquisition trigger. The four SiPM channels are processed identically, and so are shown grouped together as a double line for clarity. First, the SiPM channels are inverted to negative polarity, as required by the amplifiers, and then the six channels are amplified and digitized. The SiPM channels are summed by an adder, and this sum, as well as the forward NaI channel, are shaped to reduce noise. A pair of discriminators (shown as "Disc.") generate logic pulses when the shaped signals rise above a threshold. The SiPM pulse is widened by a gate generator to account for delays in the system, and the coincidence between the two pulses is used to trigger acquisition.

Function	Model	Parameters
Amplifier	LeCroy 612AM	gain: 10x
Inverter Adder	LeCroy 428F	inverting non-inverting
Shaper (SiPM) Shaper (NaI)	Ortec 474	integration time: 50 ns integration time: 20 ns
Discriminator	Phillips Scientific 710	threshold: 20 mV width: 160 ns
Gate generator	LeCroy 222	width: 204 ns
Coincidence	Ortec C314/NL	-
Digitizer	CAEN DT5730B	rate: 500 MHz samples: 32768 post-trigger: 95% trigger: external
Oscilloscope (for testing)	Tektronix TDS 3012	various
Pulser (for testing)	Berkeley Nucleonics PB-4	various

Table 3.6: Models and parameters of the components used for data processing, data acquisition, and testing. Note that both shapers used a proportional gain of 1x and zero differential gain, and that the digitizer's "post-trigger" setting refers to the fraction of the recorded trace taking place after the trigger.

Function	Model	Parameters
SiPM	Hamamatsu S13371-8554	-
SiPM bias	Keithley 6517B	55.9 V
NaI detector	Bicron 5H6/5L-HR	-
NaI bias	HP Harrison 6516A	900 V

Table 3.7: Models and parameters of the detector modules used in the TPC.

Chapter 4

Characterizing the Prototype

As described in section 3.2, the two-phase TPC prototype presents two separate avenues of pulse-shape discrimination to be investigated. One involves S1, the prompt, fast scintillation signal, and would make use of slight changes in the timing of the signal due to differing ionization density to attempt to distinguish β-like events from $\beta\beta$-like events. The other would use the late, prolonged ionization signal, S2, to reconstruct a profile of the initial ionization, and attempt to directly resolve the particles involved in the event by their signatures in this profile. Due to the more explicit detail and the amplification present in the ionization signal, analysis of S2 is the primary goal of this effort, and analysis of S1 is a secondary goal, to provide an additional tag and perhaps augment the discriminatory capabilities of the S2 analysis.

However, experimental science rarely proceeds as smoothly as one might hope; in 2014, for example, a barrel of waste stored at WIPP ruptured due to incorrect packing, releasing a cloud of plutonium dust and forcing EXO-200 to suspend operations for two years. Less dramatically, when the commissioning of the two-phase prototype was complete, the full system was switched on and no S2 was observed. No immediately obvious causes could be identified; S1 had been observed coincident with signals in the forward NaI detector since before the high-voltage system was operational, and applying high voltage did reduce the magnitude of S1, indicating that a drift field was indeed present (the electric field separates ionization electrons from their ions, reducing recombination and therefore scintillation [29]). Nevertheless, nothing clearly identifiable as S2 was seen, either in the SiPM via electroluminescence or as a direct charge signal on the grids. The working hypothesis is that the build-up of impurities in the xenon is larger than the purification system is capable of correcting, due to a flaw in the placement of the diaphragm pump and purifier, and that the ionization electrons are absorbed by electronegative impurities before reaching the readout. This is potentially supported by the presence of occasional odd pulses seen in the SiPM signal; they are similar in ap-

pearance to what is expected of S2, but always begin simultaneously with an event in the forward NaI detector, rather than after the expected ∼10μs delay that is characteristic of S2. It is possible that these signals in fact are S2; that they are the rare Compton scatters that occur close enough to the top of the liquid for the ionization to reach the gas before being absorbed.

A set of changes to the TPC and gas system to improve xenon purity are planned, including re-routing of the purification loop and stricter exclusion of outgassing materials from the TPC. In the meantime, this issue does not interfere with the production of S1, so this secondary objective can still be advanced using a small dataset taken for validation before the high-voltage system was completed; this data is analysed in section 4.1. Additionally, simulations of the ionization signal are presented in section 4.2, and used to evaluate the accuracy of reconstruction of the ionization profile from S2.

4.1 Analysis of scintillation signals

A small dataset originally taken for general validation purposes is instead used to evaluate the usefulness of the S1 pulse-shape for discrimination between β-like Compton scatter events and $\beta\beta$-like pair production events. The dataset consists of five subsets; four single-channel sets of events observed in the NaI detectors, for the calibration of these detectors, and one six-channel set recording the full output of the SiPM and NaI detectors on the usual coincident trigger with the ^{228}Th source in place. The calibration sets are each several minutes of exposure, and each contains roughly 8000 events. Two are from the forward detector: one, with no sources present, is used to evaluate backgrounds, while the other was recorded with a ∼4 kBq ^{137}Cs source placed close to the detector, for calibration. The other two are from the rear detector under the same conditions. The gamma-ray spectrum of ^{137}Cs is dominated by a 661.7 keV line due to an internal transition in the daughter nuclide ^{137m}Ba [11]; this is chosen for its similarity to the 551 keV of a positron-annihilation photon. The ^{228}Th dataset contains 62093 events, roughly one day of exposure.

4.1.1 Calibration and cuts

For each event in the calibration sets, the area and height of the pulse was extracted from the first 8 μs (4000 samples) of the event trace. These quantities were calculated relative to a baseline found using a method called *kappa-sigma clipping*, an algorithm commonly used in astronomy that returns an iterative truncated mean. The parameters of this algorithm are some value κ and the number of

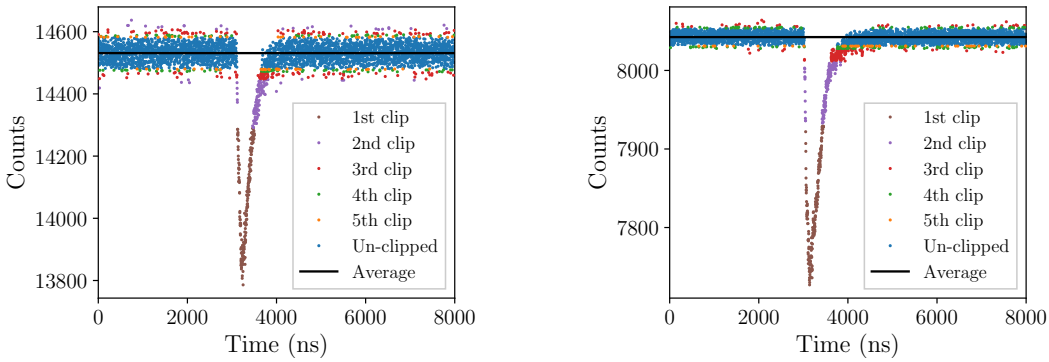

(a) **(b)**

Figure 4.1: Traces of a calibration event in the forward NaI detector (a) and a ^{228}Th event in the SiPM (b). The data clipped in each iteration of the κ-σ algorithm is shown, as is the remaining data and its average. The digitizer records a value between 0 and 16384 counts, corresponding to a measurement in a two-volt range. The position of 0V can be set freely within the output range; in (a) it is in the middle (8192), and in (b) it is at 90% of the maximum (14746). Some small offset of the actual baseline from these values is present.

iterations n. In the first iteration, the mean μ and standard deviation σ of the data are calculated, and outliers are removed from the data (specifically, an entry x is removed if $|x - \mu| > \kappa\sigma$). In subsequent iterations, μ and σ are recalculated for the "clipped" data from the previous iteration, and the process is repeated. After the last iteration is complete, the mean of the remaining "un-clipped" data is returned. In astronomy, this is used as an average that is insensitive to large outliers, such as cosmic-ray hits on a sensor; here, it is used to exclude pulses (of which there may be more than one) and find bins that are part of the baseline. This analysis uses $\kappa = 2$ and $n = 5$ to find the baselines of events in all datasets. The algorithm is visualized in figure 4.1, which shows a calibration event in the forward NaI detector and a ^{228}Th event in the SiPM. Pulse height is determined as the maximum absolute deviation from this baseline, and area as the absolute value of the summed deviation over all samples. Area was found to resolve the ^{137}Cs peak with higher resolution and so was used to measure energy.

Histograms of pulse area for each of the calibration sets can be found in figure 4.2, showing clear peaks due to the presence of the ^{137}Cs source, and secondary peaks at double the pulse area of the main peak, due to events containing two pulses. Fits are used to find the center of each peak; that is, the average pulse area corresponding to a 661.7 keV gamma ray. The acquisition threshold for the rear detector was set too high, and truncates the dataset unfortunately close to the peak, but this can be recalibrated later on the peak from pair production itself.

With this calibration, the energy measured in the NaI detectors for each event

in the ^{228}Th dataset can be determined, and used to determine which of these events are signal-like, background-like, or neither. The few percent of events in which any of the channels has saturated are discarded. Histograms of these energies are shown in figures 4.3 and 4.4. As shown in figure 4.4c, there is a

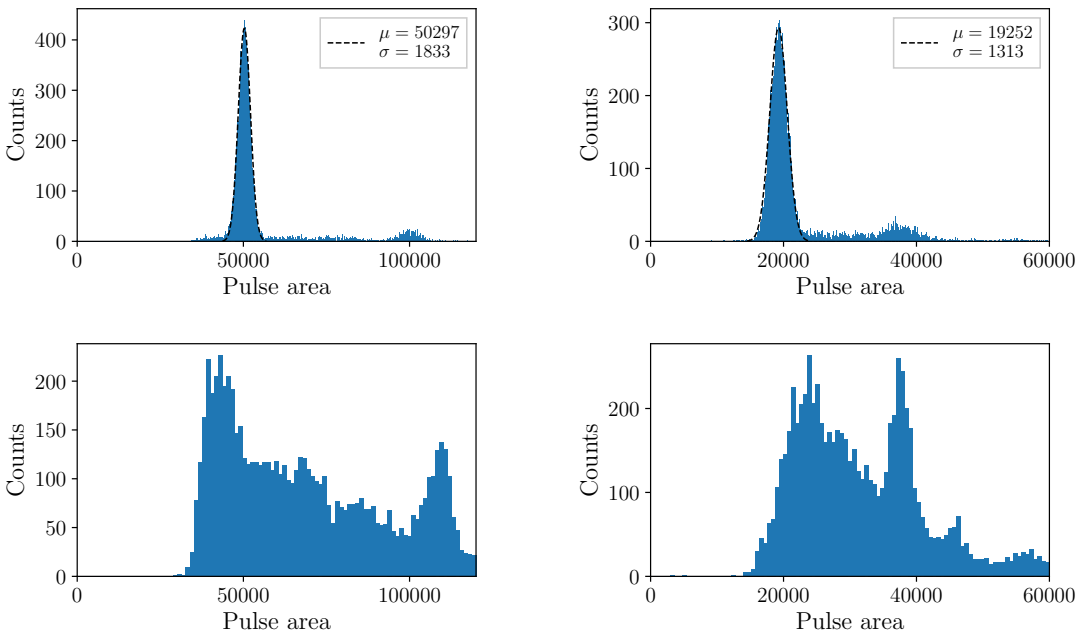

Figure 4.2: Histograms of pulse area for the four calibration datasets; each set contains between 7996 and 8246 events. The sets with the ^{137}Cs source present are on top, and sets with no source are below; sets from the forward detector are on the left, while those from the rear are on the right. It can be seen that any backgrounds around the large peaks are more than an order of magnitude smaller than the peaks themselves, so backgrounds are ignored and a Gaussian is fitted to the upper 90% of each peak.

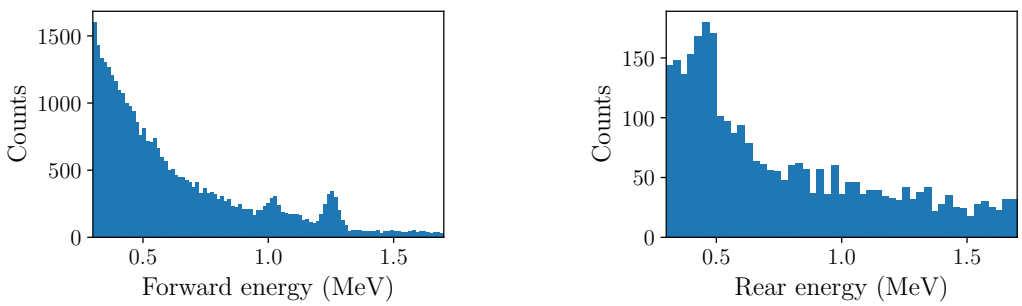

Figure 4.3: Histograms of event energy in either NaI detector during the ^{228}Th run, showing a wide energy window including locations of interest. One can see a large peak in the rear histogram slightly below $m_e c^2 = 0.511$ MeV. In the forward histogram, there is a dubious peak slightly above this location, a broad peak around $2m_e c^2$ (where Compton scatters are expected), and an unidentified peak at 1.25 MeV, which may be due to Compton scattering in an object outside of the xenon. Further analysis is shown in figure 4.4.

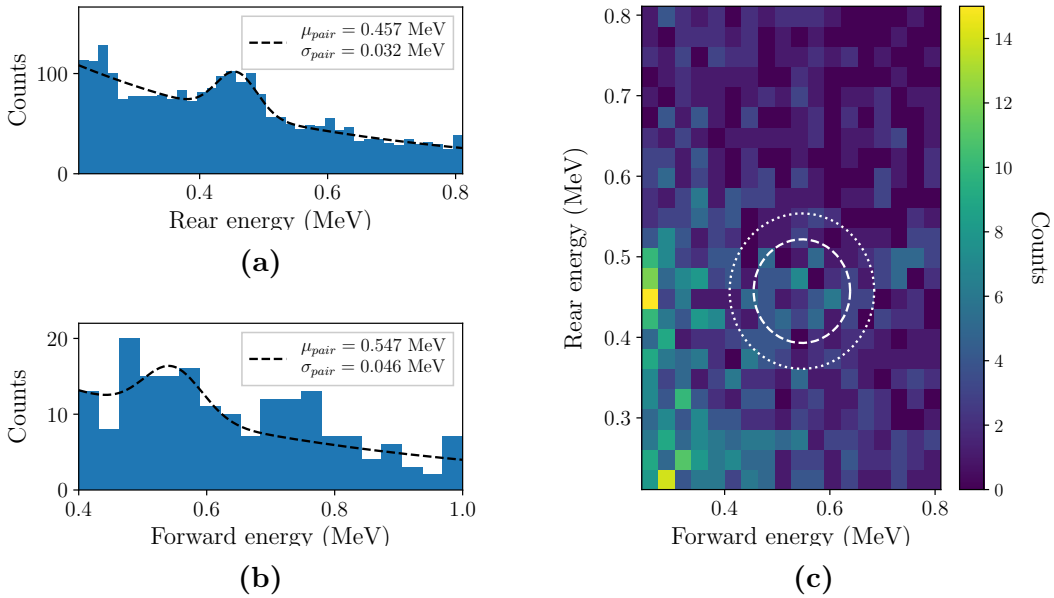

Figure 4.4: Further analysis of the pulses recorded in the NaI detectors with the ^{228}Th source in place. Histogram (a) is a zoomed-in version of the histogram in figure 4.3, fitted with a Gaussian peak and exponential background to event energies in the rear NaI detector. The peak is identified as being at $m_e c^2$, up to some initial miscalibration. The same thing is shown for the forward NaI detector in (b), except that, to reduce the large background, only events with a rear energy within 2σ of the peak are included. The peak in (b) is less well-defined than that in (a), but is shown more clearly in (c). This histogram shows the energies in both detectors, with the combined Gaussian distribution from the fits shown via a dashed line at 2σ and a dotted line at 3σ. The excess of events in the dashed ellipse (69 events in 0.0185 MeV^2) relative to the dotted annulus (60 events in 0.0232 MeV^2) is large, with a significance of 2.9σ.

large excess of events with NaI detection energies near the region expected for pair production, and this excess is interpreted as such; the energy is recalibrated according to the two Gaussian fits, and events falling within the dashed ellipse are tagged as $\beta\beta$-like pair production. This recalibration can easily be applied to the desired Compton-scatter photons; events are tagged as β-like if they satisfy two conditions. The first is that the scattered photons have twice the energy of photons due to pair production (with μ_{pair} and σ_{pair} for the forward detector shown in figure 4.4b), and the second is that the signal in the rear detector is consistent with no pulse at all (86% of events have $E_{\text{rear}} < 0.05$ MeV, while only 0.6% have 0.05 MeV $< E_{\text{rear}} <$ 0.10 MeV). Specifically, an event is tagged as β-like if it has $|E_{\text{forward}} - 2\mu_{\text{pair}}| > 2\sigma_{\text{pair}}$ and $E_{\text{rear}} > 0.05$. When these cuts are applied to the dataset, we are left with 2040 β-like events and 64 $\beta\beta$-like events.

4.1.2 Tests of background rejection

The general goal of this experiment, as described in section 3.3, is to evaluate the ability of the prototype two-phase TPC (and thereby the two-phase framework in general) to facilitate discrimination between β-like and $\beta\beta$-like events. In this analysis, this ability is evaluated using machine learning; the full dataset divided into a "training" set and a "testing" set, each consisting of β-like and $\beta\beta$-like events tagged using the analysis of the signal in the NaI detectors described in section 4.1.1. A classification algorithm is trained on the training set to independently tag each event as one of the two types using only SiPM traces (over the window $[2.5\mu\text{s}, 5.5\mu\text{s}]$, containing S1), and the accuracy of the trained algorithm is evaluated using the testing set, to control for over-fitting.

Fisher discriminant analysis

As mentioned in section 3.2, the DEAP-3600 dark-matter experiment, while operating with a different noble liquid detector medium (argon) and under different signals and backgrounds, achieves background rejection via pulse-shape discrimination in a way that may be applicable to EXO. In DEAP, the ratio of the amount of scintillation light detected early in the event window to that for the full event window (known as F_{prompt}) is used to efficiently tag events as either nuclear recoils, potentially due to interactions with dark matter, or electron recoils due to background radiation. The basis of this technique is that nuclear recoils preferentially produce a short-lived excited state of argon, while electron recoils produce more of a long-lived state [31], in much the same way that $\beta\beta$ events in LXe may result in a higher fraction of the short-lived state than β events (see section 3.2).

To investigate the potential of this sort of analysis, we define a tag similar to F_{prompt}. A vector of quantiles $\vec{q}$ is defined, the traces of the four SiPM channels are summed, and the integral of this sum is calculated for each event. For each event, a vector of times $\vec{t}$ is calculated, where each component is the time taken for the integrated signal to rise to the corresponding quantile:

$$\int_0^{t_i} (f(t) - b)\,\mathrm{d}t = q_i \int_0^{\infty} (f(t) - b)\,\mathrm{d}t \tag{4.1}$$

where $f(t)$ is the sum of the SiPM traces and b is the baseline defined in section 4.1.1. For example, if $\vec{q} = (0.5)$, $\vec{t}$ has one component, which is simply the time at which half of the total light has been detected. A Fisher discriminant is then constructed to determine if the the event that $\vec{t}$ belongs to is β-like or $\beta\beta$-like.

Suppose that $\vec{q}$ has two (or more) components; for example, $(0.2, 0.8)$. We can then define some direction $\hat{a}$, and a scalar test statistic for each event, $x(\hat{a}) = \vec{t} \cdot \hat{a}$.

The usefulness of x in distinguishing the two sorts of events can be measured by the separation $S(\hat{a})$, the mean distance in x between the two event types divided by their combined standard deviation in x:

$$S(\hat{a}) = \frac{|\langle x_\beta \rangle - \langle x_{\beta\beta} \rangle|}{\sqrt{\sigma^2_{x,\beta} + \sigma^2_{x,\beta\beta}}} \tag{4.2}$$

It can be shown that $S(\hat{a})$ is maximized when $\hat{a}$ is along a particular direction [37]:

$$\hat{a} \propto \left(V^{ij}_\beta + V^{ij}_{\beta\beta}\right)^{-1} \cdot \left(\langle \vec{t}_\beta \rangle - \langle \vec{t}_{\beta\beta} \rangle\right) \tag{4.3}$$

When $\hat{a}$ is along this optimal direction, it is called a *Fisher discriminant*.

A Fisher discriminant was calculated from a training set containing 1640 β and 44 $\beta\beta$ events and tested on a testing set containing the remaining 400 β and 20 $\beta\beta$ events. Over a large number of trials with randomly selected training and testing sets, the number of quantiles in $\vec{q}$ and their distribution (concentrated near 0 and 1, uniform, or concentrated near 0.5) were varied, and the best resolution in the testing set was interpolated to be at $\vec{q} = (0.08, 0.21, 0.79, 0.92)$. 1000 more trials were run using this $\vec{q}$, and gave a mean separation of 0.70 ± 0.13. Results of a representative example from these trials are shown in figure 4.5.

For this trial, if we choose an x_{cut} so that the rate of type-I error (α, the probability of failing to reject a background event) equals the rate of type-II error (β, the probability of rejecting a signal event) as equal, we get $\alpha = \beta = 0.21$ (meaning that we would lose $\beta = 21\%$ of $\beta\beta$ events, but reduce our backgrounds by $1 - \alpha = 79\%$). This can be compared to figure 2.4 on page 21; the goal of "improved" background

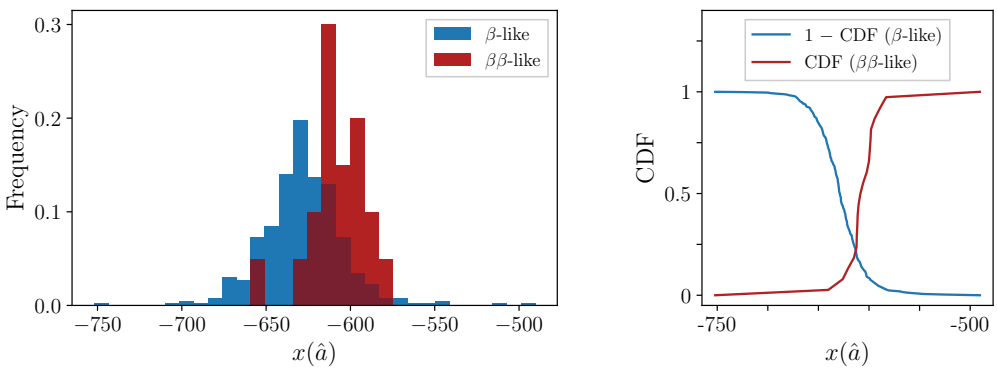

(a) **(b)**

Figure 4.5: Results of a trial of the Fisher discriminant analysis described above, showing the test statistic x for β-like and $\beta\beta$-like events in the testing set. Histograms of x are shown in (a), while (b) shows the CDFs and error rates; one minus the CDF of β-like events equals α, and the CDF of $\beta\beta$-like events equals β. The two distributions have a separation of 0.70.

rejection over the design value for nEXO calls for $\alpha = 0.49$, $\beta = 0$ and would improve sensitivity by 25%, while "perfect" rejection is at $\alpha = \beta = 0$ and would improve sensitivity by a factor of 4.3. In nEXO, $\alpha = \beta = 0.21$ would improve sensitivity by 20% [23].

Neural network analysis

A second analysis was carried out, similar to the first but replacing the Fisher discriminant with a neural network. Many sorts of neural networks exist, with varied applications, but in general they consist of a network of nodes connected by directional weights. Some nodes provide input data to the network, some collect return values from the network, and the rest read values from nodes that are connected to them, weight these input values and collate them into an output, and pass the output on to other nodes. The nodes and weights can be arranges in a variety of ways; this analysis makes use of a sort of network called a *multi-layer perceptron*, implemented using TensorFlow [38].

Neural networks can be can be classified as either *feedforward* networks, in which connections do not form loops (one cannot follow connections from node to node until one returns to the node one started at), and *recurrent* networks, which do contain loops. A perceptron is a sort of feedforward network with a specific structure; the nodes are arranged into layers, with connections only between each layer and the next. A single-layer perceptron has only an input layer and an output layer, with connections running directly between, while a multi-layer perceptron has one or more "hidden" layers between the input and the output [37, 38].

This analysis uses the same input data as with the Fisher discriminant: a randomly selected training set containing 1640 β and 44 $\beta\beta$ events, and a testing set containing the remaining 400 β and 20 $\beta\beta$ events. The classification algorithm is a multi-layer perceptron. Its input layer contains 1500 nodes corresponding to the 1500 samples in the event trace over the interval $[2.5\mu\text{s}, 5.5\mu\text{s}]$. The traces are rescaled for input into the network, with the baseline (calculated as described in section 4.1.1) subtracted and the area set to unity. The output layer consists of two nodes which return values indicating the network's classification of the input on the spectrum from β-like (0) to $\beta\beta$-like (1); these two values sum to 1. One hidden later is used; the addition of a second hidden layer was consistently detrimental to the separation. The nodes in the hidden layer have a rectified linear response function (the response function maps the node's weighted input to its output; a rectified linear function is linear and clamped to be positive) and those in the output layer have a softmax response (each node has an exponential response, normalized to sum to 1 over both nodes), as is common in pattern recognition [38].

A neural network is trained by defining a "loss" function, some measure of the badness of the output, and attempting to set the weights of the network in a way that minimizes the loss. The training of classification networks tend to define loss as the negative of the log-likelihood, called *cross-entropy*. By default, the loss function weights all events equally; in this analysis, the events are reweighed so that the few $\beta\beta$-like events have the same total weight as the many β-like ones. Various iterative algorithms can be applied to the problem of minimizing the loss; in this analysis, the Adam algorithm [39] was found to most consistently produce high separation out of the optimizers available in TensorFlow. Iteration once over each item in the training dataset is called an *epoch*.

At this point, the free parameters remaining in the analysis are the number of nodes in the hidden layer and the number of epochs over which to train the network. Several thousand networks were trained on the data with this number of nodes and the number of epochs uniformly distributed over $[1, 120]$. The results of these trials are shown in figure 4.6. Above 58 nodes or 55 epochs, the separation was seen to saturate, so these parameters were used in the final network, which was built and trained using the TensorFlow calls shown in figure 4.8.

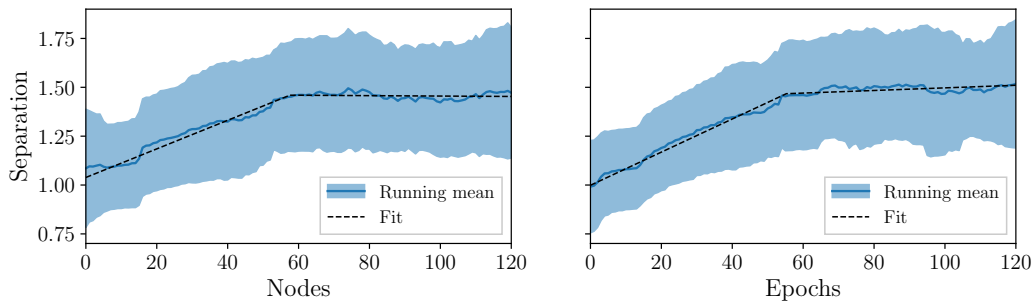

Figure 4.6: Plots of separations achieved by neural networks containing a number of hidden nodes and training for a number of epochs both distributed over $[1, 120]$. The running mean and standard deviation over 10 nodes or epochs is shown, along with a piecewise linear fit. The discontinuities, identified as saturation, occur at 58 nodes and 55 epochs.

In each of 1000 trials, this network was trained on a randomly selected training set, and tested on the corresponding testing set. The mean separation achieved in the testing sets was 1.44 ± 0.25, more than double the separation achieved by the Fisher discriminant analysis. Results from a representative example of these trials are shown in figure 4.7. If we choose a cut so that we have the same $\beta = 0.21$ as in the Fisher discriminant analysis on page 44, this trial gives $\alpha = 0.06$, rejecting 94% of the background, much higher than the 79% rejected by the Fisher discriminant. These rates would improve the sensitivity of nEXO by a factor of 1.9 [23].

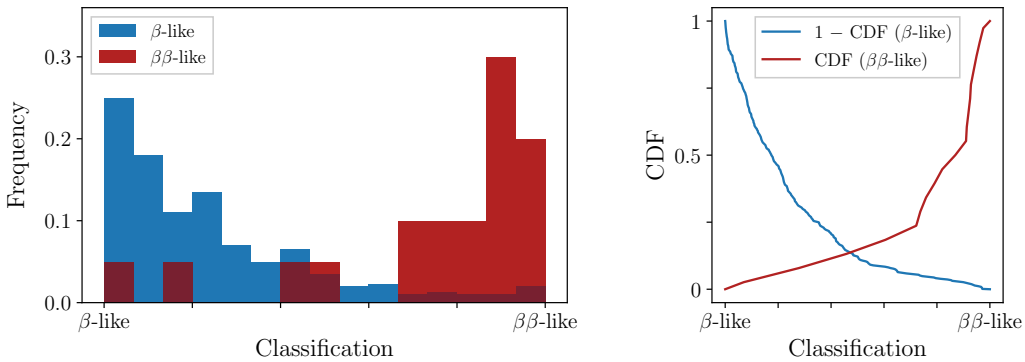

(a) **(b)**

Figure 4.7: Results of a trial of the neural network analysis described above, showing the network's classification of β-like and $\beta\beta$-like events in the testing set. Histograms of the classification output are shown in (a), while (b) shows the CDFs and error rates; one minus the CDF of β-like events is equal to α, and the CDF of $\beta\beta$-like events is equal to β. The two distributions have a separation of 1.45.

```
network = tensorflow.keras.Sequential([          # Sequential class = perceptron
        tensorflow.keras.layers.Dense(           # the hidden layer
                units = 58,                      # the number of hidden nodes
                input_shape = (1500, ),          # the shape of each trace
                activation = 'relu'              # 'rectified linear unit'
        ),
        tensorflow.keras.layers.Dense(           # the output layer
                units = 2,                       # the number of output nodes
                activation = 'softmax'
        )
])
network.compile(
        optimizer = 'adam',
        loss = 'sparse_categorical_crossentropy'
)
network.fit(                            # train the network
        training_data,                  # traces
        training_labels,                # 0 for Compton, 1 for pair
        epochs = 55,
        class_weight = weights          # 1/1640 for Compton, 1/44 for pair
)
```

Figure 4.8: The Python code at the core of the analysis, which builds and trains a neural network for classifying events. The `Sequential` class describes a perceptron, consisting of sequential layers of nodes; `training_data` and `training_labels` are array of traces and the corresponding event-type labels for each event; `weights` is a map connecting each possible label to a weight (specifically the reciprocal of the number of events with that label) to normalize the loss function.

4.2 Analysis of ionization signals

4.2.1 Simulating ionization signals

A proof-of-concept Monte-Carlo simulation was created to evaluate the potential of separating β events from $\beta\beta$ events based on the S2 pulse shape. Specifically, the goal was to determine if the time resolution of a two-phase detector was sufficient to usefully resolve and count the Bragg peaks in the ionization profile.

The first stage of the simulation creates a series of β and $\beta\beta$ events in LXe, and simulates the resulting scintillation and ionization. This was carried out using Geant4 [32], a toolkit developed by CERN for simulating the passage of particles through matter. Events consisting of single 1.6 MeV electrons and pairs of 0.8 MeV electrons in a volume of LXe were simulated using Geant4's generic particle source. The physics parameters used by the central nEXO_MC simulation [23] were used. This stage of the simulation produced, for each event, a description of the track of the ionizing particle(s), consisting of a list of ∼8000 individual ionization deposits and their times, locations, and energies, as well as the total amount of scintillation produced.

The second stage simulated the transport of each ionization cluster from its initial location in the LXe to the readout, using a detector geometry based on that described in section 3.3. It first used a simple model to generate S1 (which was only used for time projection); each energy deposit was assumed to produce scintillation proportional to its energy, and the resulting number of photons incident on the solid angle of the photodetector was modelled as a Poisson process. A second Poisson process converted the ionization energy deposits into clusters of ionization electrons, initially located at the site of the deposit. The drift of the electrons to the top of the liquid was then simulated, under the applied electric field and Gaussian diffusion, finding a time and position of each electron's entry into the gas layer. The positions of the electrons in the horizontal plane were then binned into 512 bins arranged in the $\sim$1mm^2 footprint of the ionization. The probability of a photon being detected, based on the solid angle and reflectivity of the photodetector, was calculated for initial photon positions located vertically at the phase boundary, at the top grid, and in the middle of the two, and horizontally at the center of each bin. For each electron, the EL photons produced were placed stochastically above that bin's center, and were registered as hits in the detector with probabilities interpolated between the precalculated grid, yielding the S2 portion of the trace. Dimensions and physical constants used in the simulation can be found in table 4.1, and a simulated event is shown in figures 4.9 and 4.10.

Parameter	Value
TPC diameter	97.4 mm
LXe depth	88.0 mm
EL layer thickness	4 mm
EL potential difference	8 kV
Electron drift velocity (liquid)	2.0 mm/µs
Electron drift velocity (gas)	1.6 mm/µs
W_S	13.8 eV
W_I	15.6 eV
Longitudinal diffusion	$9.0 \cdot 10^{-3}$ mm^2/µs
Transverse diffusion	$9.0 \cdot 10^{-4}$ mm^2/µs
Digitizer rate	500 MHz

Table 4.1: Dimensions of the simulated TPC, as well as values of physical constants used by the simulation from [18, 30]. W_S and W_I are the *work functions* for scintillation and ionization, the average energy deposit required for an ionizing particle to produce one scintillation photon or electron-ion pair.

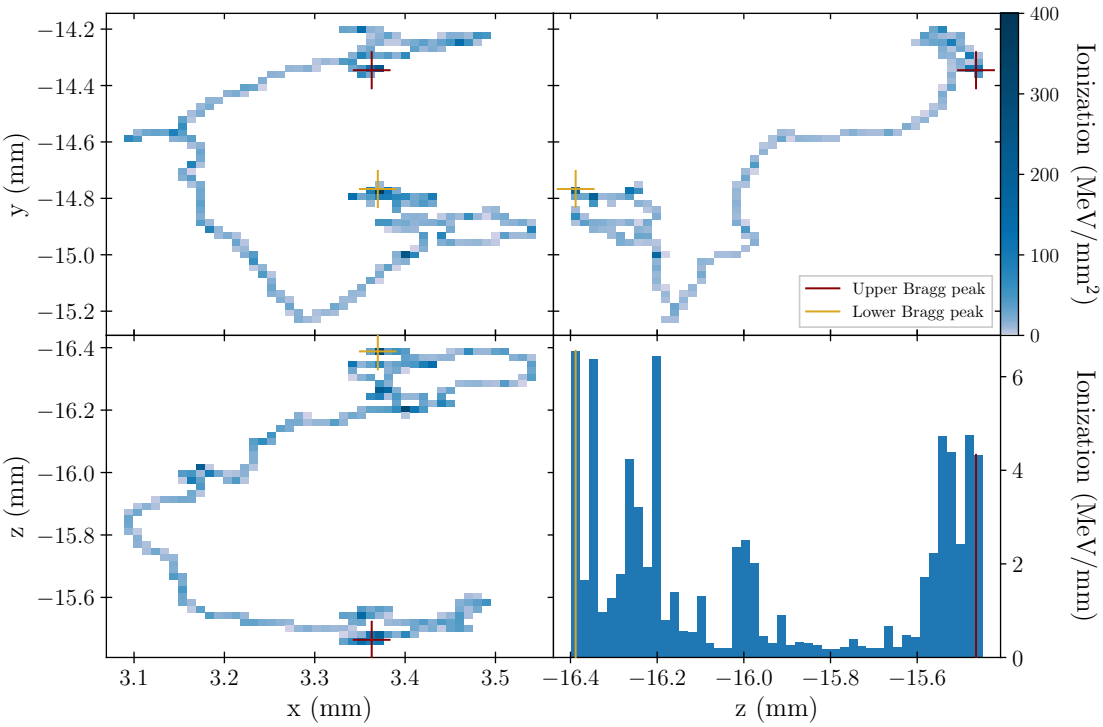

Figure 4.9: An ionization cluster simulated by Geant4 for a pair of 0.8 MeV electrons in LXe. Histograms show the two-dimensional distribution of ionization in each of the xy, yz, and xz planes, and the one-dimensional profile along the axial direction z. The locations of the two Bragg peaks are denoted by crosses in the 2D histograms, and vertical lines in the 1D. A non-Bragg peak also appears at z = -16.0mm in the 1D histogram; this is an impostor due to an electron moving mostly in the xy plane. The origin is defined as the center of the liquid volume.

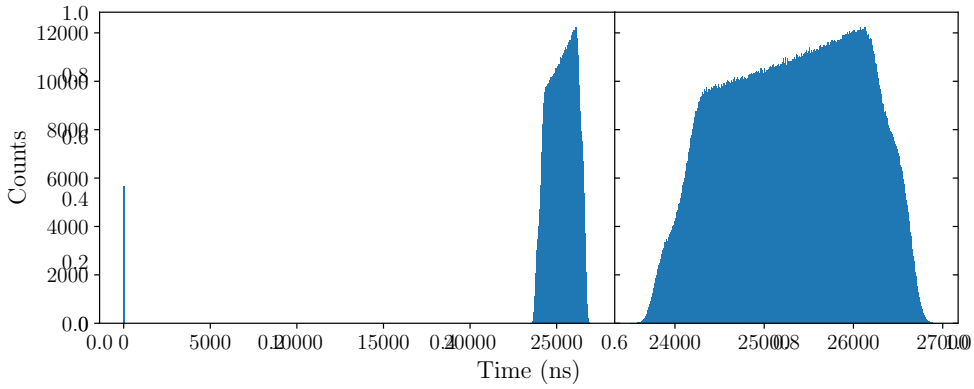

Figure 4.10: Simulated trace of the light produced by the event shown in figure 4.9, with the inset showing detail in S2. Counts are placed in 2ns time bins. The prompt scintillation pulse (S1) defines $t = 0$, while the broader ionization pulse (S2) arrives many microseconds later.

4.2.2 Reconstruction of the simulated events

Even in the raw trace of figure 4.10, useful detail is visible. The basic structure of the S2 pulse is a fast rise, as the ionization enters the gas and begins electroluminescing, followed by a slightly rising plateau, when all the electrons are in the EL region and moving towards the photodetector, and finally a fall to zero as the electrons exit the high-field region. Substructure is visible in the rise and fall; both begin relatively fast, then slow briefly before returning to their original speed. This corresponds to the density of the ionization entering or exiting the EL region (the vertical profile in figure 4.9), which is at first high, at the upper Bragg peak, then low in the middle, then high again at the lower Bragg peak.

This is hardly rigorous; the event shown here was selected for illustrative purposes, because its S2 corresponds so clearly to its ionization profile. No obvious Bragg peaks are visible in most of the simulated events, at least to the naked eye; a full analysis of the simulated signals must consider the transfer function of the detector. This function has two components: a Gaussian broadening due to diffusion in the LXe, and a roughly square broadening as electrons cross the electroluminescent region.

The electroluminescence response is dealt with first. As electrons enter, transit, and exit the EL region, they emit photons at a constant rate, giving a square pulse slightly modulated by the solid angle of the detector changing as the electron moves. This transfer function is easily removed from the signal; convolution with a square pulse wider than the signal is simply integration, and deconvolution is differentiation. To reconstruct the ionization profile as it was upon entry into the gas, the pulse is modulated by the reciprocal of a quadratic fit to the plateau,

flattening the plateau and removing the dependence on the solid angle of the detector. The rise and fall are aligned by their midpoints, and are differentiated and averaged, giving the ionization profile.

This ionization profile is that as the electrons reach the phase boundary, after having diffused in the medium. The diffusion is Gaussian, so some Gaussian transfer function must be deconvoluted from the signal to obtain the initial ionization profile. This is accomplished using the Richardson-Lucy method, an iterative algorithm for deconvolution with noisy signals [40]; the width of the transfer function used is calculated based on S2 and the known drift velocity and diffusion coefficient, with the drift time taken to be the time between S1 and the moment when S2 rises halfway to the plateau. This reconstructed profile, for the event shown in figures 4.9 and 4.10, is shown in figure 4.11.

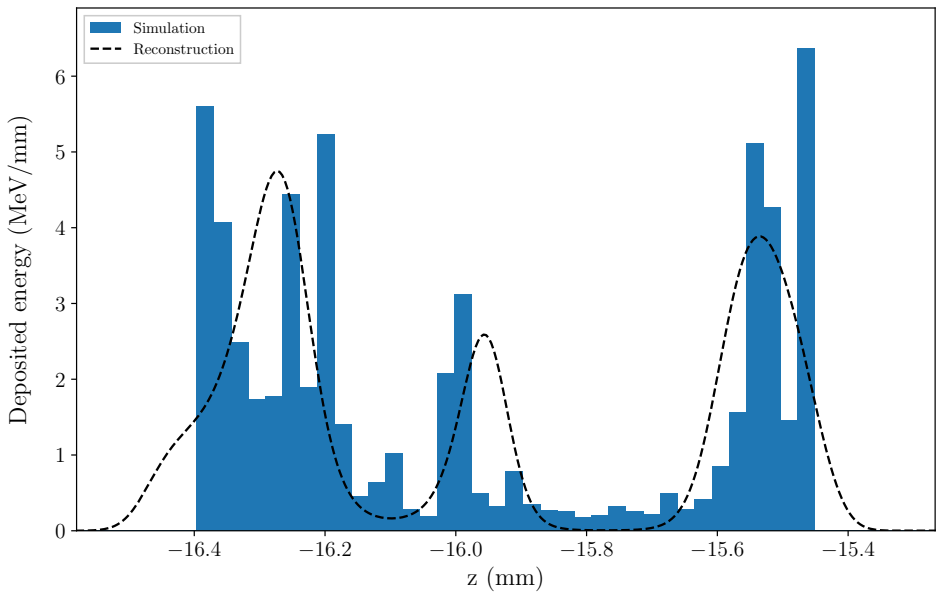

Figure 4.11: The axial profile of the simulated ionization shown in figure 4.9, overlaid with a profile reconstructed by the analysis described above. The reconstruction is based only on the trace of the simulated TPC signals and the known drift velocity and diffusion coefficient.

The reconstruction successfully resolves the two Bragg peaks, as well as the smaller non-Bragg peak. To evaluate resolution, we can integrate the simulation and the reconstruction to find the cumulative distribution of each along z. The inverse of this is the quantile function $Q(p)$, giving the z-coordinate with a fraction p of the initial ionization below it. The deviation in $Q(p)$ between the simulation and reconstruction can be calculated, and identified as the resolution of the reconstruction. Figure 4.12 shows these quantile functions for the profiles shown in figure 4.11, as well as the distribution of resolutions over many simulated events.

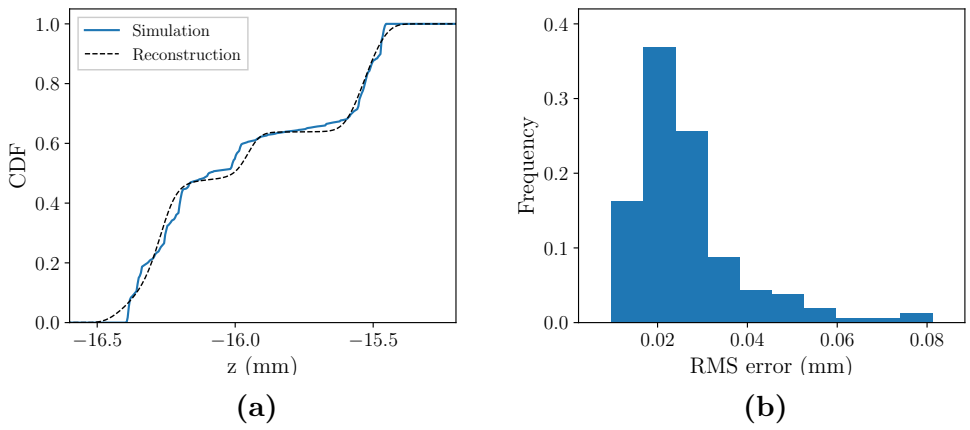

Figure 4.12: The positional resolution of the reconstruction of simulated events. The cumulative distributions of the simulated and reconstructed ionization profiles shown in figure 4.11 are shown in (a); the RMS deviation in z between the two (32 μm) is identified as the resolution. The resolutions of all simulated events (77 β, 83 $\beta\beta$) are histogrammed in (b).

Figure 4.12b shows that most of the reconstructions have resolutions similar to or better than that shown in 4.11, indicating that significant useful information about the ionization profile remains in S2 after diffusion and detection. Counting of Bragg peaks is not always as straightforward as in the event shown; a large peak can be either a single Bragg peak, two Bragg peaks with the same axial coordinate, or simply an electron moving in the xy plane. Nevertheless, it appears that the initial ionization profile is recorded in S2, likely in a much more explicit way than it is recorded in S1, and that the pulse-shape recognition techniques described in section 4.1 could be augmented by the inclusion of the ionization signal.

Conclusion

In this work, a novel detector and analysis technique was put forward for reducing backgrounds and improving sensitivity in the nEXO experiment which, when commissioned, will lead the search for neutrinoless double-beta decay, an important facet of the question of the neutrino masses. The techniques described here have the potential to replace the simple signals produced by the single-phase architecture with substructured pulses containing useful information that was invisible to EXO-200.

The results obtained in this analysis are promising, surpassing the goal of improved background rejection set by the nEXO collaboration. However, this analysis is far from complete; even the scintillation analysis of section 4.1 uses a partial dataset, with much less exposure than would have been taken in a run intended for final analysis. While some pulse-shape discrimination was achieved, the classification algorithms used were only trained on sets of 40 $\beta\beta$-like events; even the beginner's tutorials for TensorFlow use much larger training sets, with the standard training database of handwritten digits or images of clothing containing thousands of examples of each class. The acquisition of a larger dataset, either containing the full TPC signal or only the scintillation, would likely make better use of the potential of this analysis.

Of course, the initial goal for the machine learning analysis *was* to use the full TPC signal, with S2 conjectured and suggested by simulation to be a significantly better tag for pulse-shape discrimination than S1 due to its more direct representation of the initial event type. The effort to better characterize and fix the issues that prevented this is relatively undeveloped, but it is hoped that over the coming months the prototype detector will return to operation, with clean xenon and clear ionization signals.

Beyond the detector itself, there are analysis techniques that may be used to improve the understanding of the detector and the background rejection obtained. While an electron-positron pair does behave similarly to a pair of electrons, the analogue is not exact, and it could be useful to determine the degree of the difference as far as this experiment is concerned, which could be accomplished via

Monte-Carlo simulation. Additionally, more advanced classification algorithms may be used; a Fisher discriminant defines a cut along some hyperplane, while a *kernel* Fisher discriminant is capable of cuts along more complex, non-linear hypersurfaces. The neural network analysis may also be improved by the use of a more complex network. Convolutional networks have layers in which each node only takes input from a cluster of nodes in the previous layer, which can be useful as a first layer in pattern recognition; deep neural networks, with multiple large hidden layers, are capable of recognizing patterns within patterns, providing sophistication beyond that of a single-layer perceptron.

Despite these shortcomings, the effort to achieve pulse-shape discrimination has seen some measure of success, with moderate background rejection achieved without operating the TPC in two-phase mode (with gas, but without electroluminescence). This itself is interesting; discrimination on the scintillation signal alone may present a more immediate opportunity for application to EXO than full two-phase operation, as it requires no changes to the existing detector design beyond a faster readout, though it does not allow analysis of S2. In either case, it seems clear that these results warrant further investigation of pulse-shape discrimination, which may well contribute to the next step in the search for the nature of the neutrino masses, and with it to our understanding of fundamental physics.

References

[1] David Griffiths. *Introduction to Elementary Particles.* WILEY-VCH, 2nd edition **(2008)**.

[2] ATLAS Collaboration. *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC.* Physics Letters **B716** (2012).

[3] C.D. Ellis and W.A. Wooster. *The Average Energy of Disintegration of Radium E.* Proc. Roy. Soc. **A117** (1927).

[4] E. Fermi. *Versuch einer Theorie der beta-Strahlen.* Zeitschrift für Physik **88** (1934).

[5] F. L. Wilson. *Fermi's Theory of Beta Decay.* American Journal of Physics **36** (1968).

[6] IceCube Collaboration. *Measurement of the multi-TeV neutrino interaction cross-section with IceCube using Earth absorption.* Nature **551** (2017).

[7] F. Reines, C. L. Cowan, Jr., F. B. Harrison, A. D. McGuire, and H. W. Kruse. *Detection of the Free Antineutrino.* Phys. Rev. **117** (1960).

[8] R. Davis, Jr., D. S. Harmer, and K. C. Hoffman. *Search for Neutrinos from the Sun.* Phys. Rev. Lett. **20** (1968).

[9] J. N. Bahcall, N. A. Bahcall, and G. Shaviv. *Present Status of the Theoretical Predictions for the ^{37}Cl Solar-Neutrino Experiment.* Phys. Rev. Lett. **20** (1968).

[10] B. Pontecorvo. *Inverse beta processes and nonconservation of lepton charge.* Soviet Physics JEPT **7** (1958).

[11] Particle Data Group. *Review of Particle Physics.* Phys. Rev. **D98** (2018).

[12] B. Kayser. *Are Neutrinos Their Own Antiparticles?* J. Phys. Conf. Ser. **173** (2009).

[13] A. Giuliani and A. Poves. *Neutrinoless Double-Beta Decay.* Adv. High Energy Phys. (2012).

[14] nEXO Collaboration. *nEXO Pre-Conceptual Design Report.* Cornell University Library, arXiv:1805.11142 **(2018)**.

[15] EXO-200 Collaboration. *Search for Majorana neutrinos with the first two years of EXO-200 data.* Nature **510** (2014).

[16] KamLAND-Zen Collaboration. *Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen.* Phys. Rev. Lett. **117** (2016).

[17] EXO-200 Collaboration. *The EXO-200 detector, part I: Detector design and construction.* J. Inst. **7** (2012).

[18] E. Aprile and T. Doke. *Liquid Xenon Detectors for Particle Physics and Astrophysics.* Rev. Mod. Phys. **82** (2010).

[19] EXO-200 Collaboration. *Observation of Two-Neutrino Double-Beta Decay in Xe-136 with the EXO-200 Detector.* Phys. Rev. Lett. **107** (2011).

[20] EXO-200 Collaboration. *Improved measurement of the $2\nu\beta\beta$ half-life of* ^{136}Xe *with the EXO-200 detector.* Phys. Rev. **C89** (2014).

[21] EXO-200 Collaboration. *Search for Neutrinoless Double-Beta Decay with the Complete EXO-200 Dataset.* Phys. Rev. Lett. **123** (2019).

[22] EXO-200 Collaboration. *Measurement of the Drift Velocity and Transverse Diffusion of Electrons in Liquid Xenon with the EXO-200 Detector.* Phys. Rev. **C95** (2017).

[23] nEXO Collaboration. *Sensitivity and discovery potential of the proposed nEXO experiment to neutrinoless double-β decay.* Phys. Rev. **C97** (2018).

[24] M. Moe. *Detection of neutrinoless double-beta decay.* Phys. Rev. **C44** (1991).

[25] nEXO Collaboration. *Imaging individual barium atoms in solid xenon for barium tagging in nEXO.* Nature **569** (2019).

[26] A. Bolozdynya. *Two-phase emission detectors and their applications.* Nucl. Inst. Meth. **A422** (1999).

[27] XENON Collaboration. *First Dark Matter Search Results from the XENON1T Experiment.* Phys. Rev. Lett. **119** (2017).

[28] S. Kang, S. Scopel, G. Tomar, and J. Yoon. *Present and projected sensitivities of Dark Matter direct detection experiments to effective WIMP-nucleus couplings.* Astroparticle Physics **109** (2019).

[29] J. Thomas and D Imel. *Recombination of electron-ion pairs in liquid argon and liquid xenon.* Phys. Rev. **A36** (1987).

[30] E. Aprile, A. Bolotnikov, A. Bolozdynya, and T. Doke. *Noble Gas Detectors.* WILEY-VCH **(2006)**.

[31] DEAP-3600 Collaboration. *First results from the DEAP-3600 dark matter search with argon at SNOLAB.* Phys. Rev. Lett. **121** (2018).

[32] S. Agostinelli et al. *Geant4 - a simulation toolkit.* Nucl. Inst. Meth. **506** (2003).

[33] M. Berger, J. Hubbell, S. Seltzer, J. Chang, J. Coursey, R. Sukumar, D. Zucker, , and K. Olsen. *XCOM: Photon Cross Sections Database.* NBSIR **87-3597** (2010).

[34] A. Compton. *A Quantum Theory of the Scattering of X-rays by Light Elements.* Phys. Rev. **21** (1923).

[35] D. Green. *Positron Cooling and Annihilation in Noble Gases.* Phys. Rev. Lett. **119** (2017).

[36] E. Özmutlu. *Sampling of Angular Distribution in Compton Scattering.* Int. J. Radiat. Appl. Instrum. **A43** (1992).

[37] G. Cowan. *Statistical Data Analysis.* Oxford University Press **(1998)**.

[38] M. Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems.* www.tensorflow.org **(2015)**.

[39] D. Kingma and J. Ba. *Adam: A Method for Stochastic Optimization.* Cornell University Library, arXiv:1412.6980 **(2014)**.

[40] W. Richardson. *Bayesian-Based Iterative Method of Image Restoration.* J. Opt. Soc. Am. **62** (1972).