SEARCH FOR NEUTRINOLESS DOUBLE BETA DECAY WITH EXO-200 AND nEXO

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Michael Jewell
May 2020

Abstract

Neutrinoless Double Beta Decay is a theorized nuclear decay in which a nucleus decays via the emission of two electrons but no neutrinos. Detection of this lepton number violating process represents the best test of the Majorana nature of neutrinos, which could potentially provide an elegant mechanism by which neutrinos obtain mass. Both the recently decommissioned EXO-200 and the future experiment nEXO search for this hypothesized decay using liquid xenon time projection chambers filled with xenon enriched in the isotope ^{136}Xe. In the first part of this work, the results from the final neutrinoless double beta decay search of EXO-200 with the complete dataset is presented. This includes a description of the advanced analysis techniques used to maximize the topological discrimination between double beta signal and gamma backgrounds. The second part of this work presents the current state of R&D efforts towards designing the planned tonne scale experiment, with focus on the characterization of charge-detecting tiles which are proposed to make up the anode plane of nEXO.

Acknowledgments

There are many people to thank for helping guide and support me throughout my graduate career. With all of their help not even a global pandemic could prevent me from receiving my PhD.

I would like to start by thanking my advisor, Giorgio, for giving me the opportunity and resources to pursue this work. Giorgio's constant enthusiasm and active participation in my research was the motivation I needed to push myself each day. In addition, thank you to all of my fellow EXO-200 and nEXO collaborators who created the perfect collaborative environment for a student such as myself to prosper. I would also like to thank the entire Varian staff, most importantly Marica and Sha, for always being there to offer a helping hand.

To the current and former members of the Gratta group, thank you for making the lab a place I looked forward to going everyday and to consistently being there to push me along. Whether it was in offering a never ending supply of coffee or helping to fix the xenon purity, I couldn't imagine having gotten here without your help. I would specifically like to thank the post-doc council, while your power is still ill-defined your advice was invaluable.

Thanks to my friends and family, especially my parents, who have supported me throughout the years. Without your constant encouragement none of this would have been possible. Finally, special thanks to Jenn for always being there to keep me sane. Our many California adventures were the fuel I needed to get to the finish line.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The Standard Model (SM) of Particle Physics encapsulates our current understanding of elementary particles and their interaction with three of the four fundamental forces. This model has been rigorously tested by experiment and with few exceptions has successfully explained their observations. One such exception to this amazing success is the existence and origin of neutrino mass as the SM is built with the assumption of massless neutrinos.

The neutrino is a neutral fermion that interacts weakly with normal matter. Neutrinos are one of the most abundant particles in the universe with more than $>10^{14}$ neutrinos passing through the average person each second [1]. These neutrinos originate both from extraterrestrial sources such as the Sun [2], cosmic rays interacting in the atmosphere [3] and supernova explosions [4] as well as terrestrially from nuclear reactors [5] and even the Earth itself [6]. But neutrino interactions with matter are extremely rare and most of these neutrinos pass right through the earth without interacting. So while neutrinos are ever present, their elusive nature has limited our ability to fully characterize their basic properties most notably their mass.

Further probing the nature of neutrinos will provide important details needed to refine our understanding of the universe. This knowledge could potentially offer a telescope into the early universe via the detection of relic neutrinos [7], a probe of the dynamics of supernova explosions [8] and an explanation for why matter dominates our universe [9]. There are even applications outside of fundamental physics such as

nuclear non-proliferation efforts via reactor monitoring with neutrinos [10].

The existence of the neutrino was first postulated by Pauli in 1930 [11] to account for the apparent lack of energy conservation in β-decay, which was at the time thought to be a simple two-body decay process were a neutron decayed into a proton by emitting an electron. In the case of a two-body decay the energy of the emitted electron was expected to be fixed at the unique energy of the decay. Experimental observations of β-decay instead revealed a continuous spectrum with the electron always having less than the expected energy. The introduction of a nearly massless third particle provided a means of accounting for this seemingly missing energy. In order to conserve electric charge the new particle needed to be neutral and to have avoided detection it must interact only weakly. At the time of Pauli's theory, direct confirmation was thought to be nearly impossible as the small mass and weak interaction made their detection extremely difficult.

The development of nuclear reactors in the 1940s provided a means to overcome the minuscule interaction rate by providing high fluxes ($> 10^{13}\mathrm{s}^{-1}\mathrm{cm}^{-2}$ [12]) of neutrinos produced by β-decay of reactor fission products. This in conjunction with large target volumes and a unique event signature provided a technique to produce a measurable neutrino signal. Using this method an experiment by Reines and Cowan successfully detected neutrinos produced by a nuclear reactor, proving its existence over 20 years after Pauli's initial postulate [5].

Originally neutrinos were assumed to be massless particles and are included in the SM as such but this picture came into question with the Homestake Experiment which showed a deficit of electron neutrinos originating from the Sun [2]. This surprising result was followed by a number of experiments detecting neutrinos from a wide range of sources [3, 13, 14] which together demonstrated the existence of neutrino flavor oscillations resulting from the flavor states being mixtures of mass states.

The establishment of massive neutrinos was unexpected and requires extending the SM to account for this non-zero neutrino mass. A number of ideas have been proposed but one such possibility arises from the potential Majorana nature of the neutrino [15, 16]. Since neutrinos are neutral it is possible they are their own anti-particle, in contrast to the other charged leptons which are so called Dirac particles. While the

Majorana nature of neutrinos is experimentally unconfirmed, it is theoretically well motivated, providing an elegant extension of the SM to account for the abnormally small but non-zero neutrino mass. [17]

The best chance to explore the Majorana nature of neutrinos is through the observation of neutrinoless double beta decay ($0\nu\beta\beta$), a theoretical decay with a half-life in excess of $\sim 10^{25}$ yrs. A number of experiments are searching for this decay including the recently decommissioned EXO-200 and the proposed next generation counterpart nEXO which are the topics of this dissertation. Details about the theory of neutrino mass and the possibility of detecting the Majorana nature of neutrinos via $0\nu\beta\beta$ are discussed in Chapter 2. An outline of the EXO-200 experiment is provided in Chapter 3, with details relevant to the analysis results presented in Chapter 4-6. These results include a measurement of the transport properties of electrons in Liquid Xenon (Chapter 4) needed for the the development of advanced machine learning based discriminators (Chapter 5) used in the final $0\nu\beta\beta$ search with the complete EXO-200 dataset (Chapter 5). Finally, the next generation tonne scale experiment nEXO is described in Chapter 7 with a detailed description of the characterization of a charge tile proposed as the ionization readout for nEXO found in Chapter 8.

Chapter 2

Massive Neutrinos

2.1 Neutrino Oscillations

Currently, there are three known neutrino flavor states (ν_e, ν_μ or ν_τ) governing neutrino interactions via the weak force. These flavors mirror the three generations of charged leptons and their existence has been experimentally tested with measurements of the Z boson resonance at electron-positron colliders, which show results consistent with three light ($m_\ell \ll m_Z/2$) neutrino states [18]. In addition, the SM assumes that neutrinos are massless and in turn the flavor states are equivalent to the neutrino mass states which govern neutrino propagation. But for massive neutrinos the situation is more complicated and the flavor states are composed of a mixture of mass states, resulting in the process of neutrino oscillations.

Neutrino oscillation refers to the phenomenon where neutrinos produced in a particular flavor state will develop components of the other states while propagating

through space. This arises from the fact that the three neutrino flavor states are linear superposition of the three neutrino mass eigenstates (ν_1, ν_2, ν_3) which have different masses (m_1, m_2, m_3). A transformation between these bases can be done using the mixing matrix U, known as the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) Matrix [19, 20]. Using this matrix the relative mixture of mass states needed to create a pure flavor state can be written as

$$|\nu_\ell\rangle = \sum_i U^*_{\ell i}|\nu_i\rangle. \tag{2.1}$$

Here $|\nu_\ell\rangle$ are the flavor states (ℓ=e,μ,τ) and $|\nu_i\rangle$ are the mass states (i=1,2,3). The PMNS matrix can be parameterized in terms of three mixing angles ($\theta_{12}, \theta_{13}, \theta_{23}$) one Dirac Phase ($\delta$) and two Majorana Phases (α_2, α_3) [21]

$$U = \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{13}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & e^{i\alpha_2} & 0 \\ 0 & 0 & e^{i\alpha_3} \end{bmatrix} \tag{2.2}$$

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$. Because the Majorana nature of the neutrino is currently an open question, it is customary to factor out the Majorana phase contribution and show it separately from the Dirac contributions.

Since the masses of each $|\nu_i\rangle$ are not the same, each state propagates with a different energy (E_i) which in turn leads to a variation in the relative phase between each mass state. In this case a neutrino produced in a pure flavor state, which is the coherent sum of mass states, develops components of the other flavors because the phases of each mass component de-cohere resulting in a non-pure state.

This decoherence can be shown by taking the propagation of each $|\nu_i\rangle$ through 1D space which are described by the standard free wave solution

$$|\nu_i(x,t)\rangle = e^{-i\phi_i}|\nu_i(0,0)\rangle \tag{2.3}$$

where $\phi_i = E_i t - p_i x$ (assuming c=$\hbar$=1). Then the probability that a neutrino starting out in a pure flavor state α will be found at some later time with flavor β is given by

$$P(\alpha \to \beta) = |\langle \nu_\beta | \nu(t) \rangle|^2 = \left| \sum_i \sum_j e^{-i\phi_i} U_{\beta j} U^*_{\alpha i} \langle \nu_j | \nu_i \rangle \right|^2 = \left| \sum_i e^{-i\phi_i} U_{\beta i} U^*_{\alpha i} \right|^2 . \quad (2.4)$$

If for simplicity it is assumed that there are just two neutrino states (instead of the three observed states), then the matrix U is a 2x2 matrix with a single mixing angle θ

$$U = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} . \quad (2.5)$$

Allowing for the oscillation probability to be expressed in terms of the mixing angle as

$$P(\alpha \to \beta) = \sin^2(2\theta) \sin^2\left(\frac{(\phi_i - \phi_j)}{2}\right) \approx \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2_{ij} x}{4E}\right) . \quad (2.6)$$

The final form in Equation 2.6 assumes the neutrinos are ultra-relativistic ($E_i >> m_i$ and $t \approx x$), in which case the energy of each state is approximately the same ($E = E_i \approx E_j$) while the difference in the energies is proportional to the mass squared difference ($\Delta m^2_{ij} = m^2_i - m^2_j$).

Equation 2.6 shows the result of the decoherence. As the mass states propagate their relative phases go in and out of phase and in turn the probability of finding flavor β oscillates with respect to that phase difference. For monoenergetic neutrinos detected at a fixed location from a source this oscillation depends only on the mass squared difference (Δm^2_{ij}) and the mixing angle. Thus an oscillation experiment is sensitive to both the mixing angle and the mass squared difference between different mass states.

2.2 Neutrino Mass

A number of oscillation experiments have now successfully observed and measured the relevant oscillation parameters, as presented in a global analysis of all experimental

results [22]. This allows for the determination of the parameters appearing in the PMNS matrix, the results of which are summarized in Table 2.1. As can be seen, the

Table 2.1: Best fit to the PMNS matrix parameters based on a global fit to oscillation experiment data with 1σ error bars included [22].

Parameter	NH	IH
$\theta_{12}[^\circ]$	$34.6^{+1.0}_{-1.0}$	$34.6^{+1.0}_{-1.0}$
$\theta_{23}[^\circ]$	$48.9^{+1.9}_{-7.4}$	$49.2^{+1.5}_{-2.5}$
$\theta_{13}[^\circ]$	$8.8^{+0.4}_{-0.4}$	$8.9^{+0.4}_{-0.4}$
$\Delta m^2_{21}[10^{-5}\text{eV}^2]$	$7.60^{+0.19}_{-0.18}$	$7.60^{+0.19}_{-0.18}$
$\Delta m^2_{31}[10^{-5}\text{eV}^2]$	$2.48^{+0.05}_{-0.07}$	$-2.38^{+0.05}_{-0.06}$
$\delta[^\circ]$	241^{+115}_{-68}	268^{+61}_{-58}

observation of neutrino oscillations has allowed for the precise measurement of the three mixing angles and the mass splittings. While the mass splittings imply that neutrinos have mass, their measurement does not indicate the absolute scale of this mass. Thus oscillation experiments alone are unable to directly measure the absolute neutrino mass.

Complementary to oscillation experiments are experiments which attempt to measure spectral distortions near the end point of β-decay. These experiments are sensitive to the average neutrino mass (m_β) and recent results from KATRIN establish an upper limit on the neutrino mass [23]

$$m_\beta = \sqrt{\sum_i |U_{ei}|^2 m_i^2} < 1.1\text{eV}. \tag{2.7}$$

In addition cosmological measurements of baryon acoustic oscillations are sensitive to the total neutrino mass (m_{tot}) and results from the Planck Collaboration have set an upper limit on this sum [24]

$$m_{tot} = \sum_i m_i < 0.23\text{eV}. \tag{2.8}$$

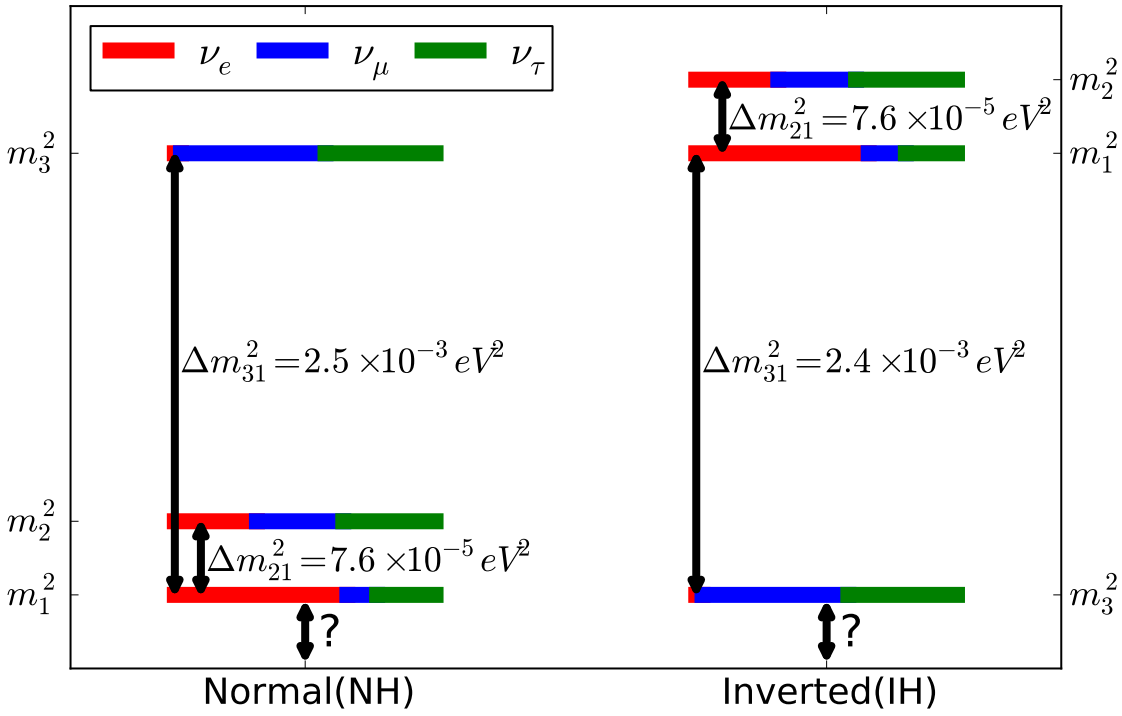

Figure 2.1: Depiction of the possible neutrino mass orderings given the known mixing parameters shown in Table 2.1. Also shown is the relative mixture of each flavor state needed to make each mass state.

Along with the overall mass scale the exact ordering of the neutrino mass states is currently unknown. As shown in Table 2.1 current oscillation experiments have accurately determined the mass splittings Δm^2_{21} and Δm^2_{32}. In the case of Δm^2_{21} the overall sign is known from Solar neutrino experiments which measure the matter effect in the Sun and are sensitive to the overall sign [13]. For Δm^2_{32} only the absolute value is known, as it derives from vacuum oscillation experiments which are not sensitive to the overall sign. This leaves open the possibility of two ordering schemes for the neutrino masses. Either m_3 is the heaviest ($m_1 < m_2 < m_3$) in which case the masses follow the Normal Hierarchy(NH) or m_3 is the lightest state ($m_3 < m_1 < m_2$) in which case the masses follow the Inverted Hierarchy(IH). These potential scenarios are summarized in Figure 2.1 which shows the possible orderings of each mass state along with their relative mixture of flavor states.

In total the results from oscillation experiments, β-decay end-point measurements

and cosmological constraints indicate that neutrinos have mass but their masses are extremely small. Compared to the mass of the lightest charged lepton (the electron) their masses are more than five orders of magnitude smaller which presents a problem when trying to add neutrino mass back to the SM.

In the simplest extension to the SM it is assumed that neutrinos derive their mass from the Higgs in a similar fashion to the other charged leptons. This is achieved by adding a Dirac mass term for the neutrinos to the SM Lagrangian given by

$$\mathcal{L}_D = -m_D \overline{\nu_R} \nu_L + h.c. \tag{2.9}$$

Here ν_L and ν_R are the left- and right-handed chiral projections of the neutrino fields respectively and m_D is the Dirac Mass. This term requires the addition of right handed neutrino and left handed anti-neutrino fields to the SM which have not been experimentally observed. Assuming that these unobserved fields are sterile and simply do not participate in Weak Interactions the standard procedure of adding mass from the Yukawa coupling to the Higgs field can be followed. The issue with this extension arises when trying to produce neutrinos with masses not ruled out by experiment. In order to achieve this, the required coupling strength is significantly weaker than that of the other SM particles [25]. While this disparity is not forbidden, it is considered "unnatural" and this extension to the SM is disfavored for its inability to naturally explain the smallness of the neutrino masses.

Instead, this disparity may imply a new mechanism for generating neutrino mass. One potential solution is the introduction of a Majorana mass term ($\mathcal{L}_M$) to the SM Lagrangian

$$\mathcal{L}_M = -\frac{1}{2} m_R \overline{(\nu_R)^c} \nu_R - \frac{1}{2} m_L \overline{(\nu_L)^c} \nu_L + h.c. \tag{2.10}$$

Here $m_R(m_L)$ are the right(left) Majorana masses and ν^c is the charge conjugate of ν. The introduction of this term violates charge conservation so is forbidden for charged particles but for neutral neutrinos this term is in general allowed. It does introduce lepton number violation but this was only an assumed global symmetry of the SM and is not a required symmetry by any conservation law.

The most general neutrino mass term in the Lagranian can than be written as

the sum of the Dirac and Majorana mass terms in Equation 2.9 and Equation 2.10 respectively

$$\mathcal{L} = \mathcal{L}_D + \mathcal{L}_M = -m_D\overline{\nu_R}\nu_L - \frac{1}{2}m_R\overline{(\nu_R)^c}\nu_R - \frac{1}{2}m_L\overline{(\nu_L)^c}\nu_L + h.c. \tag{2.11}$$

This term can than be reorganized as

$$\mathcal{L} = -\frac{1}{2}\overline{(n_L)^c}Mn_L \ , \ n_L = \begin{bmatrix} \nu_L \\ \\ (\nu_R)^c \end{bmatrix}. \tag{2.12}$$

For simplicity the single flavor assumption can be taken, in which case the matrix M is a 2x2 matrix which has two mass eigenstates with masses given by the eigenvalues $m_\pm$

$$M = \begin{bmatrix} m_L & m_D \\ \\ m_D & m_R \end{bmatrix} \ , \ m_\pm = \frac{1}{2}\left(m_R + m_L + \sqrt{(m_R + m_L)^2 - 4(m_L m_R - m_D^2)}\right). \tag{2.13}$$

Equation 2.12 represents the most general neutrino mass Lagrangian allowed by the SM. Neutrinos being Dirac is than a specific case where $m_L = m_R = 0$ resulting in the Majorana term vanishing. Another case arises when it is assumed that $m_R \gg m_D \gg m_L$. In this case the masses of the eigenstates in Equation 2.13 reduces to

$$m_+ = m_R \ , \ m_- = \frac{m_D^2}{m_R}. \tag{2.14}$$

The Dirac mass (m_D) is expected to be $\sim 10^2$ GeV since it derives its mass through the Higgs and in turn should be constrained to similar order as the other charged leptons. On the other hand m_R is not protected by the SM symmetries and instead could follows symmetries of a much higher energy theory, for example at the grand unified theory (GUT) scale (10^{14}-10^{16} GeV). Under these conditions there is than a heavy mass state (m_R) and a light mass state (m_-) with mass in the range 1-100 meV,

consistent with the current upper bound on the neutrino mass from experiment. This is known as the see-saw mechanism as it provides a natural explanation of the neutrino mass by adding a massive sterile state (m_R) which suppresses the Dirac mass creating a relatively light mass state [26].

The Majorana nature of neutrinos is then motivated by the ability to elegantly explain the mass of the neutrino through the see-saw mechanism. It is possible to test this Majorana nature by than searching for lepton number violating processes (LNVP) since the addition of a Marjorana Mass term introduces lepton number violation into the SM. One of the most promising LNVP is that of neutrinoless double beta decay ($0\nu\beta\beta$) which will be described in detail in the following section.

2.3 Double Beta Decay

Double beta decay ($2\nu\beta\beta$) is an extremely rare second order nuclear decay by which a nucleus decays by simultaneously converting two neutrons to two protons via the emission of two electrons and 2 antineutrinos as shown in Figure 2.2a.

$$(A, Z) \rightarrow (A, Z + 2) + 2e^{-} + 2\bar{\nu}_e$$

In principle this process can occur for any nucleus which has larger mass than its daughter (M(A,Z) $>$ M(A, Z+2)) but because $2\nu\beta\beta$ is second order it is heavily suppressed. For most nuclei it is than expected that the first order process of single beta decay (β-decay) will dominate.

It is possible for the $2\nu\beta\beta$ process to dominate in certain even-even nuclei where β-decay is energetically forbidden due to the increased binding energy from the pairing force. This can be seen by examining the Semi-Empirical Mass Formula (SEMF) [27, 28] which approximately gives the mass of a nucleus containing A nucleons and Z

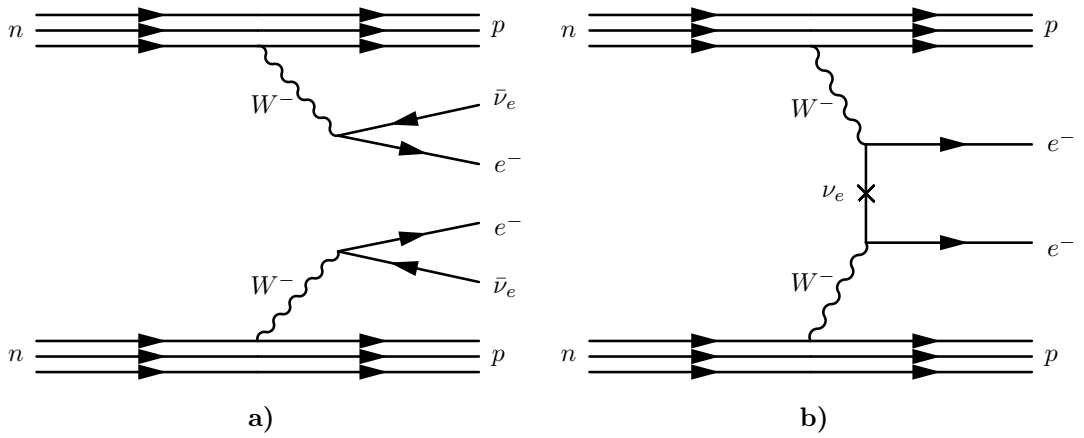

Figure 2.2: Feynman Diagram for (left) $2\nu\beta\beta$ and (right) $0\nu\beta\beta$ assuming light neutrino exchange.

protons:

$$M(A,Z) = \alpha A - \beta Z + \gamma Z^2 + \frac{\delta}{A^{1/2}} f(A,Z)$$
$$\alpha = m_n - a_1 + \frac{a_2}{A^{1/3}} + a_4 \ , \ \beta = 4a_4 + \frac{a_3}{A^{1/3}} + (m_n - m_p) \ , \ \gamma = \frac{4a_4}{A} + \frac{a_3}{A^{1/3}} \ , \ \delta = a_5 \tag{2.15}$$

Here m_p and m_n are the mass of the proton and neutron respectively while a_{1-5} are empirical parameters determined via fits to experimental data. Of importance for $2\nu\beta\beta$ is the last term, known as the pairing term, which depends on the function:

$$f(A,Z) = \begin{cases} 0, & \text{Odd A} \\ +1, & \text{even A, odd Z} \\ -1, & \text{even A, even Z} \end{cases} \tag{2.16}$$

This term accounts for the tendency of the nucleons to form spin-0 couples which are more tightly bound than unpaired nucleons. For odd A there is always an unpaired nucleon and this term vanishes, but for even A nuclei there is an alternating sign which reflects the breaking of nucleon pairs when going from even Z to odd Z states. As a result, the isobaric curves for even A nuclei are composed of two offset parabolas,

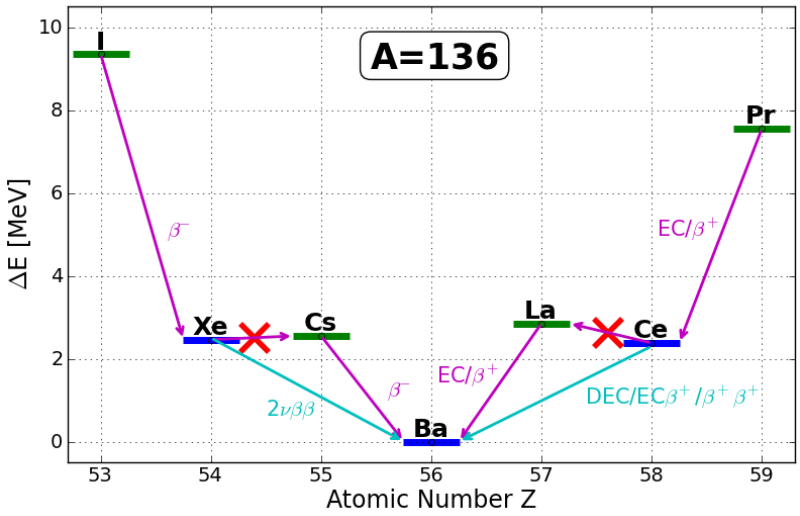

Figure 2.3: Difference in mass with respect to ^{136}Ba for nuclides with A=136. [29, 30] As highlighted, the mass of ^{136}Xe is larger than that of ^{136}Cs preventing β-decay. Instead ^{136}Xe will decay via $2\nu\beta\beta$ to ^{136}Ba with $Q_{\beta\beta}$ =2457.83 $\pm$ 0.37 [31].

with the less tightly bound odd Z nuclei lying higher than the even Z nuclei.

As a result the expected β-decay daughter for certain even-even nuclei can potentially have more mass than the parent (M(A,Z) $<$ M(A,Z+1)) in which case β-decay is energetically forbidden. Instead, the $2\nu\beta\beta$ process will be the main mode of nuclear decay for these nuclei. An example of this is shown for the A=136 isobar in Figure 2.3 using experimentally measured masses [29, 30]. Here it is seen that ^{136}Xe is unable to β-decay to ^{136}Cs and instead decays via $2\nu\beta\beta$ to ^{136}Ba.

The $2\nu\beta\beta$ decay process was initially proposed in 1935 [32] with a predicted half-life in excess of 10^{17} yrs. It was first observed geochemically in 1950 via the detection of excess ^{130}Xe in a 1.5×10^9 yrs old tellurium ore sample arising from the decay of ^{130}Te to ^{130}Xe [33]. The first direct detection was than achieved in 1987 using a Time Projection Chamber(TPC) to detect the decay of ^{82}Se to ^{82}Kr [34]. Since then, $2\nu\beta\beta$ has been directly observed in a handful of isotopes with half lives ($T^{2\nu\beta\beta}_{1/2}$) ranging between 10^{18} and 10^{21} years. The current direct measurements for a selection of isotopes are briefly summarized in Table 2.3 based on recommendations found in [35].

It should be noted that $2\nu\beta\beta$ applies when the nucleus has a proton deficit but an analogous situation arises for nuclei with a proton excess such as ^{136}Ce shown in Figure 2.3. In this case the nucleus can simultaneously convert two protons into two neutrons via double electron capture(DEC), double positron emission($\beta^+\beta^+$) or electron capture plus positron emission(ECβ^+):

$$\beta^+\beta^+ : (A, Z) \rightarrow (A, Z-2) + 2e^+ + 2\nu_e \tag{2.17}$$

$$\text{DEC} : (A, Z) + 2e^- \rightarrow (A, Z-2) + 2\nu_e \tag{2.18}$$

$$\text{EC}\beta^+ : (A, Z) + e^- \rightarrow (A, Z-2) + e^+ + 2\nu_e \tag{2.19}$$

Of these processes only DEC has been observed so far [36–38].

2.3.1 Neutrinoless Double Beta Decay

The process of $2\nu\beta\beta$ is allowed by the SM and is possible even if neutrinos are not Majorana. On the other hand, if neutrinos are Majorana there is the possibility of an alternative decay mode by which the nucleus decays via the emission of two electrons, but unlike $2\nu\beta\beta$ does not produce any neutrinos:

$$(A, Z) \rightarrow (A, Z+2) + 2e^- \tag{2.20}$$

This process, which is shown in Figure 2.2b, is known as neutrinoless double beta decay ($0\nu\beta\beta$) and, as can be seen in Equation 2.20, the lack of anti-neutrinos in the final state results in the violation of lepton number conservation and thus represents physics beyond the SM.

The diagram in Figure 2.2b depicts the commonly assumed mechanism of light-neutrino exchange but there are a number of other possible mechanisms [26]. While the exact mechanism impacts the connection between $0\nu\beta\beta$ and neutrino mass it should be noted that positive detection, regardless of the underlying mechanism, implies that neutrinos are Majorana. This was shown by Schechter and Valle and is known as the black-box theorem [39]. It is depicted in Figure 2.4 which shows that a $\overline{\nu}$ to ν transition is induced by rearranging Figure 2.2b and adding two W^- vertices

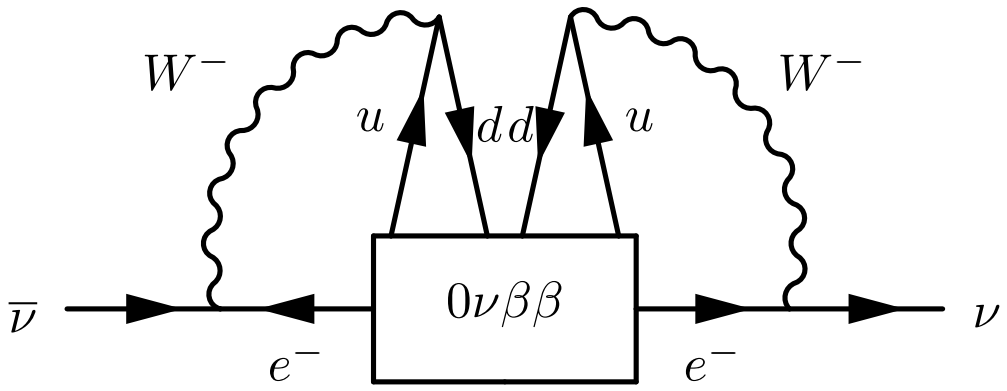

Figure 2.4: Feynman diagram illustrating the black-box theorem of Schechter and Valle [39] which demonstrates any $0\nu\beta\beta$ mechanism implies Majorana neutrinos.

connecting the quark. This transition than implies the existence of a Majorana Mass term in the Lagrangian.

Under the assumption that light-neutrino exchange is the dominant mechanism the half-life is related to the neutrino mass scale:

$$\left(T_{1/2}^{0\nu\beta\beta}\right)^{-1} = G_{0\nu}(Q_{\beta\beta}, Z)|M_{0\nu}|^2\langle m_{\beta\beta}\rangle^2 \tag{2.21}$$

Here $M_{0\nu}$ is the nuclear matrix element(NME), $G_{0\nu}$ is the phase-space factor and $\langle m_{\beta\beta}\rangle$ is the effective Majorana Mass defined as

$$\langle m_{\beta\beta}\rangle = \left|\sum_{i=1}^{3} U_{ei}^2 m_i\right|. \tag{2.22}$$

The value of $G_{0\nu}$ depends on the atomic number Z and the the energy of the decay ($Q_{\beta\beta}$) and can be accurately calculated analytically for different isotopes. The values used here are summarized in Table 2.2 and come from the calculation presented in [40].

In addition, the value of $M_{0\nu}$ needs to be determined for each isotope but in this

Table 2.2: Table of calculated $G_{0\nu}$ and $M_{0\nu}$ values needed to relate the $0\nu\beta\beta$ half-life to the neutrino mass $\langle m_{\beta\beta}\rangle$. This definition of $G_{0\nu}$ is defined such that it absorbs the factor of g_A^4 and m_e^2 which sometimes appear in Equation 2.21.

Iso	$G_{0\nu}[10^{-25}$ y^{-1}eV$^2]$	$M_{0\nu}$				
	[40]	EDF[41]	ISM[42]	IBM-2[43]	QRPA[44]	SkQ[45]
^{76}Ge	0.24	5.55	2.81	4.68	4.64	5.09
^{130}Te	1.42	6.41	2.65	3.70	3.65	1.37
^{136}Xe	1.45	4.20	2.19	3.05	2.02	1.55

case, unlike $G_{0\nu}$, the complex many body problem has no exact solution. [46] Several different approaches have been employed to calculate $M_{0\nu}$ but their results can differ by as much as a factor of four depending on the method and isotope. A summary of results for a selection of models is shown in Table 2.2 for three $0\nu\beta\beta$ candidate isotopes. The resulting uncertainty introduced by the spread in NME calculation is the dominate uncertainty when interpreting a half-life value as a measure of the neutrino mass.

There is an additional uncertainty in Equation 2.21 due to the possible reduction of the axial-vector coupling (g_A). This arises from the fact that the value of g_A needs to be reduced or "quenched", by an empirical observation with no theoretical explanation, in order to match the calculated decay strength to that observed in nature for β-decay [47, 48]. Because g_A enters Equation 2.21 as g_A^4 this effect potentially has a large impact on the conversion between $0\nu\beta\beta$ and $\langle m_{\beta\beta}\rangle$ and could result in the reduction of the reported sensitivity to $\langle m_{\beta\beta}\rangle$. While there are a number of proposed models to explain the need for this quenching there currently exists no general theory for the effect and its unclear how to extrapolate to $0\nu\beta\beta$. Because of this it is common to take the un-quenched free value deriving from neutron decay experiments ($g_A = 1.2724 \pm 0.0023$) [49].

In the case of a positive $0\nu\beta\beta$ detection the relationship between $T_{1/2}^{0\nu\beta\beta}$ and $\langle m_{\beta\beta}\rangle$ in Equation 2.21 would result in a measurement of the absolute neutrino mass. Even in the absence of a positive detection this can be used to further constrain the absolute

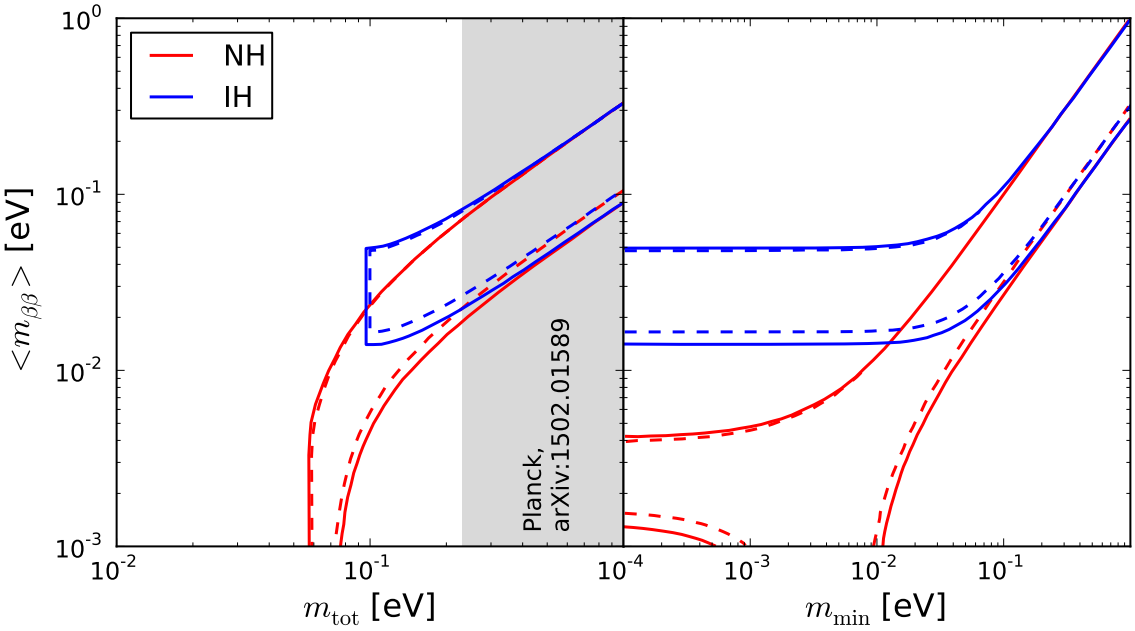

Figure 2.5: Allowed values of $\langle m_{\beta\beta}\rangle$ as a function of the (left) sum neutrino mass and (right) lightest neutrino mass. In both cases the IH is shown in Blue and the NH in red. The dashed lines represent the spread due only to the unconstrained Majorana Phases and the solid lines include the additional error from uncertainty on oscillation parameters given in Table 2.1.

mass scale of neutrinos in the light neutrino exchange model. Using the experimentally measured oscillation parameters in Table 2.1, the allowed values of $\langle m_{\beta\beta}\rangle$ can be plotted as a function of either the sum of the neutrino mass (m_{tot}) or the lightest neutrino mass (m_{min}), as shown in Figure 2.5. The large spread in $\langle m_{\beta\beta}\rangle$ is mostly due to the completely unknown Majorana Phases appearing in Equation 2.2 with a small contribution from the uncertainty in the measured oscillation parameters. The vertical grey exclusion comes from cosmological constraints described in Section 2.21 and the horizontal exclusion comes from $0\nu\beta\beta$ searches [50].

Detection of $0\nu\beta\beta$ is achieved via the observation of the sum electron energy produced by a $2\nu\beta\beta$ source isotope. Similar to β-decay the electron energies in $2\nu\beta\beta$ follow a continuous spectrum as there is some energy carried away by the two neutrinos which go undetected. On the other hand in $0\nu\beta\beta$ no neutrinos are produced

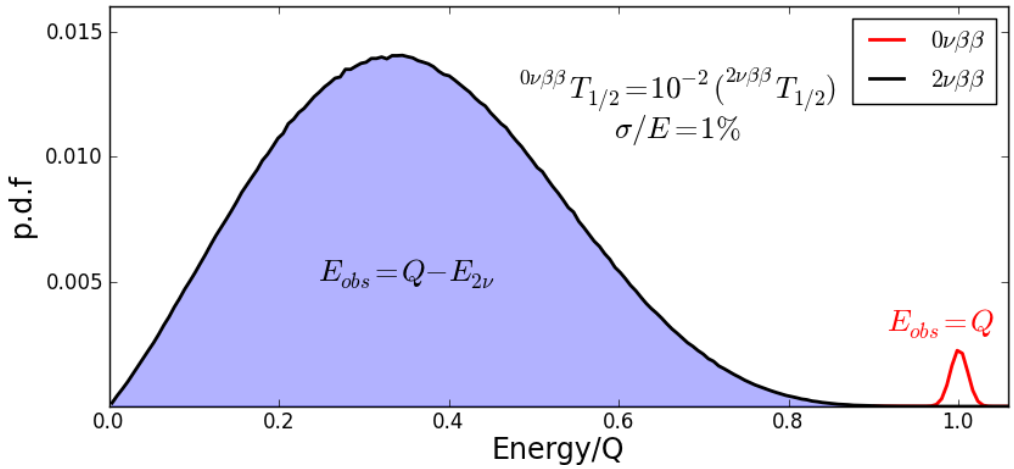

Figure 2.6: Expected energy distribution of the electrons produced in (blue)$2\nu\beta\beta$ and (red)$0\nu\beta\beta$. The spectrum is smeared with a $\sigma/E = 1\%$ energy resolution and for illustrative purposes $T_{1/2}^{0\nu\beta\beta}$ is chosen to be $\times 100$ smaller than $T_{1/2}^{2\nu\beta\beta}$.

in which case the the two electrons possess all of the decay energy. This results in a mono-energetic peak directly at $Q_{\beta\beta}$ as shown in Figure 2.6 where the width of the peak comes from the resolving power of the detection technique. Because $2\nu\beta\beta$ is effectively an irreducible background, this unique signature is fundamental to any $0\nu\beta\beta$ search as it allows for the distinction between the electrons produced via $0\nu\beta\beta$ and $2\nu\beta\beta$.

While $2\nu\beta\beta$ is an inherent background to $0\nu\beta\beta$ searches, any realisable experiment will also be subjected to backgrounds from a number of naturally occurring radioactive decays. Measuring the energy of these interactions can remove many of these backgrounds but some of these decays can mimic the energy signature of $0\nu\beta\beta$ and due to the extremely long half-lives ($> 10^{25}$ yrs) being studied even trace amounts of these backgrounds can potentially obscure a $0\nu\beta\beta$ signal. Thus it is important to limit background levels as much as possible. This can be achieved for example by careful material screening and underground deployment. In addition modern experiments are able to further reduce background contributions by reconstructing not just energy but also position and topology, which can be used for further discrimination between signal-like and background-like events.

A common figure of merit used to quantify the reach of a $0\nu\beta\beta$ search is the Sensitivity, defined as the corresponding half-life of the largest signal that would be hidden under the background of the experiment [51, 52]. Using the law of radioactive decay the half-life Sensitivity of an experiment can be expressed as

$$T_{1/2}^{0\nu\beta\beta} = \ln(2)\,\epsilon N_{\beta\beta}\frac{\Delta t}{n_S}. \tag{2.23}$$

Here n_S is the largest number of counts that would be undetectable, Δt is the experiments integration time, ϵ is the detection efficiency and $N_{\beta\beta}$ is the total number of possible $0\nu\beta\beta$ decay isotopes.

To understand the choices inherent to the design of a $0\nu\beta\beta$ search one can model a detector as a simple counting experiment and examine two limiting cases. First, if the experiment is background dominated and assuming Poisson Statistics n_S is just the statistical error in subtracting the background counts N_B ($n_S = \sqrt{N_B}$). The background counts than come from integrating over the Region of Interest (ROI), which is defined in terms of the detector energy resolution σ. Finally modeling the backgrounds as flat and making the naive assumption that the number of background counts scale linearly with the detector mass allows N_B to be written in terms of a background level (B) expressed per unit mass, energy and time $N_B = B\sigma M_{\beta\beta}\Delta t$. The sensitivity can than be expressed as

$$T_{1/2}^{0\nu\beta\beta} = \ln(2)\left(\frac{N_A}{A_{\beta\beta}}\right)\epsilon\sqrt{\frac{M_{\beta\beta}\Delta t}{B\sigma}}. \tag{2.24}$$

Here $M_{\beta\beta}$ is the total mass of the $\beta\beta$ isotope, N_A is Avogadro's number and $A_{\beta\beta}$ is the molar mass.

In addition to the background dominated experiment there is also the optimal scenario in which the expereiment is background free and the measured signal than n_S no longer depends on σ, $M_{\beta\beta}$ or Δt. Instead n_S is simply a constant (C) allowing the sensitivity to be expressed as

$$T_{1/2}^{0\nu\beta\beta} = \ln(2)\left(\frac{N_A}{A_{\beta\beta}}\right)\epsilon\frac{M_{\beta\beta}\Delta t}{C} \tag{2.25}$$

In this case the growth of the experiment is linear with exposure, instead of growing like $\sqrt{\Delta t}$, which further highlights the importance of background reduction for $0\nu\beta\beta$ searches. There are some ideas which have the potential to achieve nearly background free searches, such as identification of daughter isotopes [53], but currently this goal has yet to be fully realized. While Equation 2.24 and Equation 2.25 make many assumptions, that do not capture the full power of modern experimental techniques, they represent a useful tool to highlight the experimentally controllable parameters which need to be optimized for maximal sensitivity. Namely a large volume of potential $2\nu\beta\beta$ nuceli($M_{\beta\beta}$), a high $0\nu\beta\beta$ detection efficiency (ϵ) and a long integration time (Δt) are needed to maximize the possibility of producing a detectable signal. In addition, the background level needs to be reduced as much as possible (B), while also achieving a precise energy measurement (σ), in order to limit the background contamination. While modern experiments have yet to achieve truly background free searches, many techniques have been employed to lower background levels allowing for nearly background free experiments. This results in sensitivities growing at rates between the two limiting scenarios presented here.

	$Q_{\beta\beta}$ [keV]	$T^{2\nu\beta\beta}_{1/2}$ [yr] [35]	$T^{0\nu\beta\beta}_{1/2}$ [yr]	$\langle m_{\beta\beta}\rangle$ [meV]	Ref.
^{76}Ge	2039 [54]	$1.88^{+0.08}_{-0.08} \times 10^{21}$	9.0×10^{25}	124-244	GERDA [55]
			2.7×10^{25}	226-446	MJD [56]
^{130}Te	2528 [57]	$7.91^{+0.21}_{-0.21} \times 10^{20}$	2.3×10^{25}	86-404	CUORE [58]
^{136}Xe	2458 [31]	$2.18^{+0.05}_{-0.05} \times 10^{21}$	1.07×10^{26}	53-164	KZ-400 [50]
			3.5×10^{25}	93-286	EXO-200 [59]

Table 2.3: Summary of measured $T^{2\nu\beta\beta}_{1/2}$ for a select set of isotopes from recommendations in [35]. In the case of multiple measurements the reported value is the error weighted average. Also included are current limits on $T^{0\nu\beta\beta}_{1/2}$ for each isotope as well as the corresponding range of $\langle m_{\beta\beta}\rangle$ limits. The spread in $\langle m_{\beta\beta}\rangle$ is due to the uncertainty of the NMEs defined in Table 2.2.

As of this writing, a number of $0\nu\beta\beta$ searches have been performed but no positive detection of $0\nu\beta\beta$ has yet been confirmed. In 2001 an experiment claimed a detection

of $0\nu\beta\beta$ in ^{76}Ge with a measured half life of $T_{1/2}^{0\nu\beta\beta} = 2.23^{+0.44}_{-0.31} \times 10^{25}$ years [60]. Recent results using the same isotope have since surpassed this measurement with no evidence of a $0\nu\beta\beta$ signal effectively ruling it out [55, 56]. Table 2.3 gives a brief overview of the current limits for a selection of $0\nu\beta\beta$ source isotopes including the corresponding range of ruled out $\langle m_{\beta\beta} \rangle$. These efforts have so far pushed the $T_{1/2}^{0\nu\beta\beta}$ sensitivity for a number of isotopes to the 10^{25-26} years level. This effort is ongoing and the next generation of tonne scale experiments aim to improve the sensitivity to $0\nu\beta\beta$ to 10^{28} years.

Chapter 3

EXO-200

3.1 Overview

The EXO-200 experiment uses a liquid xenon (LXe) time projection chamber (TPC) to search for $0\nu\beta\beta$ of ^{136}Xe. EXO-200 is deployed at the Waste Isolation Pilot Plant (WIPP) [61] in Carlsbad, New Mexico. The experiment was commissioned in 2011 and began taking data toward the $0\nu\beta\beta$ search in September 2011. Operations were interrupted in February 2014 due to an incident at the underground host facility unrelated to EXO-200 [62]. Upon regaining access to the site, data taking restarted in January 2016 and continued until December 2018 when EXO-200 completed operations and was decommissioned.

Therefore, data taken during the lifespan of EXO-200 are divided into two phases as shown in Figure 3.1. Phase-I consists of data taken prior to the incident at WIPP (September 2011 - February 2014) and Phase-II consists of the data taken after access was restored (January 2016 - December 2018). Data taken during these periods consists of both calibration data and data used for the $0\nu\beta\beta$ search, which is called "low background" (LB) data. In addition to the WIPP incident this division accounts for improvements in detector performance related to detector upgrades to both the electronics and the operating drift field. A system was also installed in Phase-II to suppress radon in the air gap between the copper cryostat and the lead shield.

The final EXO-200 dataset including both phases consists of 1181.3 days of live-time corresponding to an exposure of 234.1 kg·yr with 117.4 kg·yr in Phase-I and 116.7 kg·yr in Phase-II, giving a median 90% confidence level $0\nu\beta\beta$ haf-life sensitivity of $5.0 \cdot 10^{25}$ yr. No statistically significant evidence for $0\nu\beta\beta$ is observed, leading to a lower limit on the half life of $T_{1/2} > 3.5 \cdot 10^{25}$ yr [59].

The EXO-200 detector [63], analysis [64] and results [59, 65, 66] have been described in extensive detail elsewhere. The following summary focuses on the details relevant to the results presented in Chapter 4-Chapter 6.

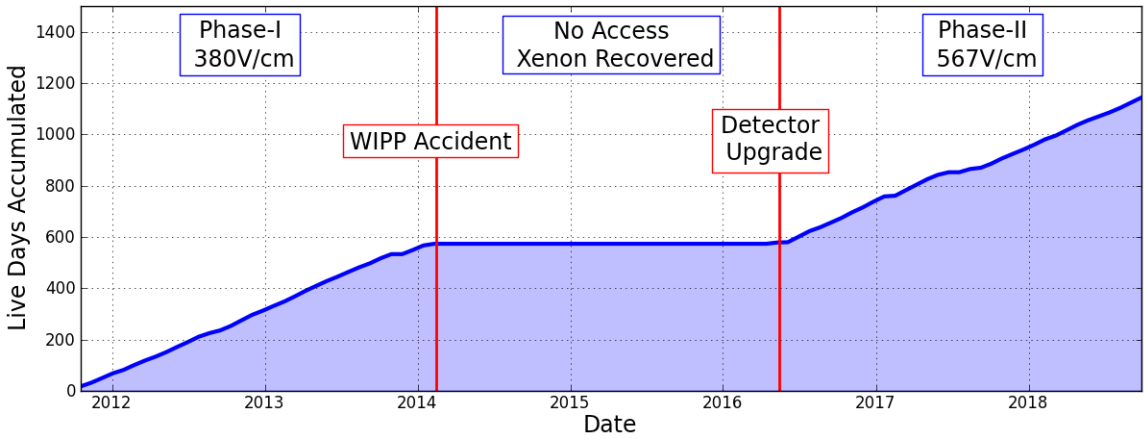

Figure 3.1: Accumulated days of data during EXO-200 operation. Data is divided into two phases separated by an extended outage due to an accident at the host facility.

3.2 EXO-200 Detector

3.2.1 Principle of LXe TPCs

Ionizing particles in LXe deposit energy into two detectable channels. First, a charge channel resulting from free electrons produced in the ionization process which is collected on the appropriate electrodes. Second, a light channel produced via the de-excitation of xenon dimers (Xe_2^*) which emit scintillation light at a wavelength of 178 nm and is collected by light sensors. These excited dimers form as a result of excited Xe atoms produced both directly from incoming radiation and indirectly from the recombination of ionized electrons with the positively charged Xe ions. In addition, some fraction of energy is lost in the form of heat dissipation but current LXe TPCs do not detect this channel. [67]

As a result of the recombination process the observed charge and light exhibit an anti-correlation, as the additional scintillation derives directly from the conversion of electrons into light [68]. The relative number of recombining electrons depends on many details of the initial deposit topology and as a result the fraction of electrons which recombine can fluctuate for each event. Estimates of the energy from the

charge or light channel alone are than affected by the amount of recombination, with additional uncertainty introduced by the unknown fraction of recombination. This results in worse energy resolution being obtained from the detection of a single channel than would be expected without recombination [69]. In order to achieve an optimal energy measurement the two signals can instead be combined in a linear combination. This allows for the approximate recovery of the amount of recombination for each event, which reduces the uncertainty on the energy measurement and in turn increases the observed energy resolution. Because of this, LXe TPCs require the simultaneous detection of both charge and light.

In addition, the recombination process is heavily dependent on the transport properties of electrons in LXe, as the probability for an electron to recombine is related to the proximity of electron-ion pairs. This in turn results in a strong dependence on the applied electric field, where larger drift fields suppress the number of recombining electrons by separating the negative electrons and positive ions more quickly. As a result, the magnitude of the applied electric field is an important design feature of any LXe TPC.

3.2.2 Detector Description

In order to limit the background introduced by the components near to the LXe volume, the EXO-200 detector is composed of a radiopure copper vessel which is just 1.37 mm thick. The TPC is ∼44 cm long and ∼40 cm and is filled with LXe, enriched to $80.672 \pm 0.014\%$ in the isotope ^{136}Xe [64]. The detector includes two identical back-to-back drift regions known as TPC1 and TPC2. These TPCs share a common cathode located at the center of the detector. In order to maintain the operating temperature of 167 K the TPC is submerged in HFE-7000 cooling fluid [70] contained in a copper cryostat. The cryostat is located in a class 1000 clean room. [63].

As discussed in Chapter 2, limiting backgrounds is key to a successful $0\nu\beta\beta$ search. While backgrounds can be mitigated in the analysis using observables such as energy and position (described in Section 3.5), it is also possible to limit backgrounds by careful selection of low background materials and the use of both passive and active

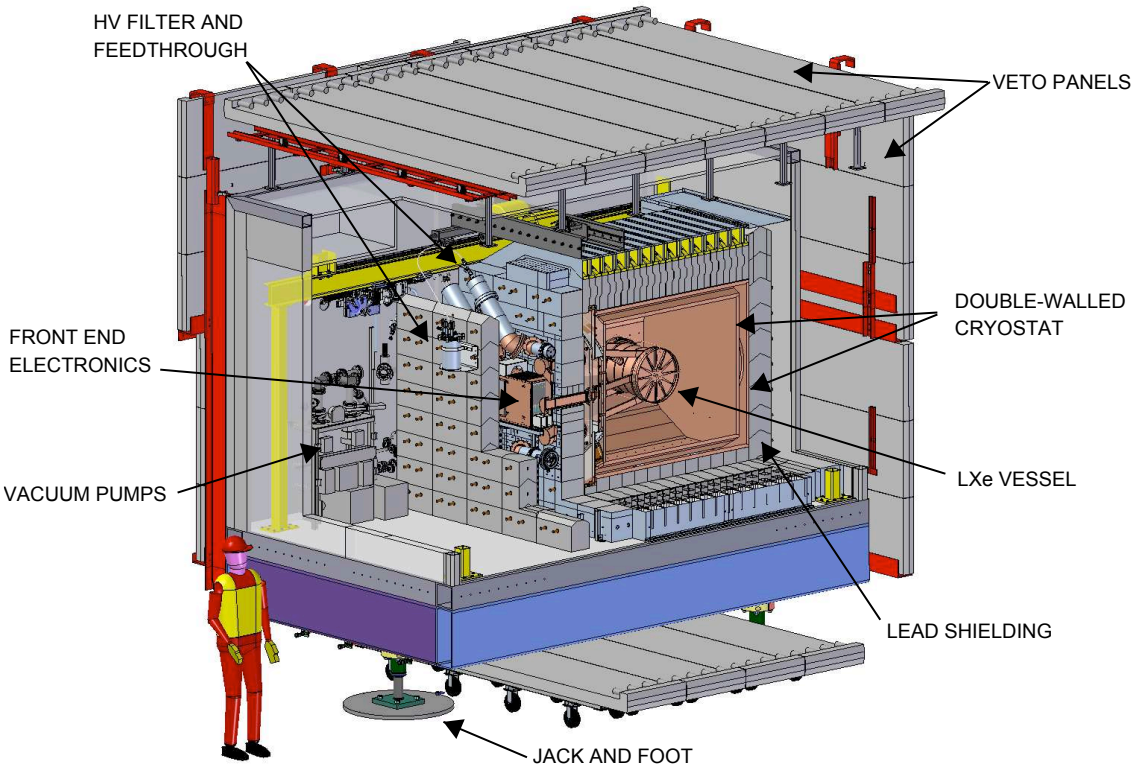

Figure 3.2: Diagram of the EXO-200 detector in the cleanroom.

shielding. To ensure minimal contributions to the background rate from construction material, all potential materials underwent extensive radiopurity screening as described in [71]. In addition, the underground location provides an overburden of 1624^{+22}_{-21} meters of water equivalent [72] which offers shielding from cosmogenic backgrounds. An active muon veto composed of scintillation panels on 4 sides of the clean room identifies > 94 % of cosmic ray muons passing through the TPC [59]. Further passive shielding is achieved with $\sim$25 cm of lead shielding in all directions, 50 cm of HFE-7000 at $\sim$167 K and 5.4 cm of copper. A diagram of the EXO-200 setup with the key features highlighted is shown in Figure 3.2.

As described in Section 3.2.1, interactions in LXe produce both a charge signal in the form of free electrons and a light signal in the form of 178 nm scintillation photons. The electrons produce via ionization are drifted from their initial location toward the anodes by a uniform bulk electric field, where they are detected at either end by a

pair of wire planes. Each wire plane lies in front of an array of large-area avalanche photodiodes (APDs) [73] that are used to observe light produced via scintillation. The APDs have a diameter between 19.6 mm and 21.1 mm, with an active diameter of 16 mm, and are hexagonally packed onto two platters to produce a packing ratio of 48%. The platters are vacuum coated with aluminum and MgF_2 to reflect scintillation light which does not collect an APD surface. Each APD platter contains 234 APDs which are grouped together into gangs of five to seven devices in order to reduce the number of connections required, resulting in 37 readout channels per plane.

The front-most wire plane in each pair (V-wires) is located 192.4 ± 0.5 mm from the cathode and serves both as a shielding grid and to detect induced signals as electrons are drifted from the interaction location to the second wire plane (U-wires) where charge is collected. The U-wires are separated from the V-wires by 6 mm. Each APD plane is 6 mm behind the U-plane. The U and V grids at either anode are oriented at a 60° angle with respect to each other. Each grid consists of 114 wires separated by 3 mm, which is chosen to maximize the fiducial volume while limiting the high voltage required to achieve full transparency on the V-Wires. In order to reduce the additional material near the TPC, the wires are than grouped into readout channels containing 3 wires each. This grouping gives 38 readout channels per wire plane that are each 9 mm in pitch and is chosen to maximize the topological discrimination while simultaneously minimizing the channel count. A cutout view of the detector can be seen in Fig. 3.3, which shows the cathode, APD plane and wire planes.

The operating field used to drift electrons was 380 ± 20 V/cm (567 ± 30 V/cm) in Phase-I (Phase-II) which was achieved by holding the cathode at a potential of -8 kV (-12 kV), the V-wire grid at -780 V (-1045 V) and the U-wire grid at ground. These potentials are chosen so that the average of the non-uniform field between the U- and V-wire grids (778 V/cm (1170 V/cm)) is approximately twice the main drift field, ensuring full transparency of the V-wire grid and full collection of the U-wire grid for ionization.

The magnitude of the electric field in both regions is determined using an electrostatic simulation of the full EXO-200 detector performed with COMSOL [74]. The

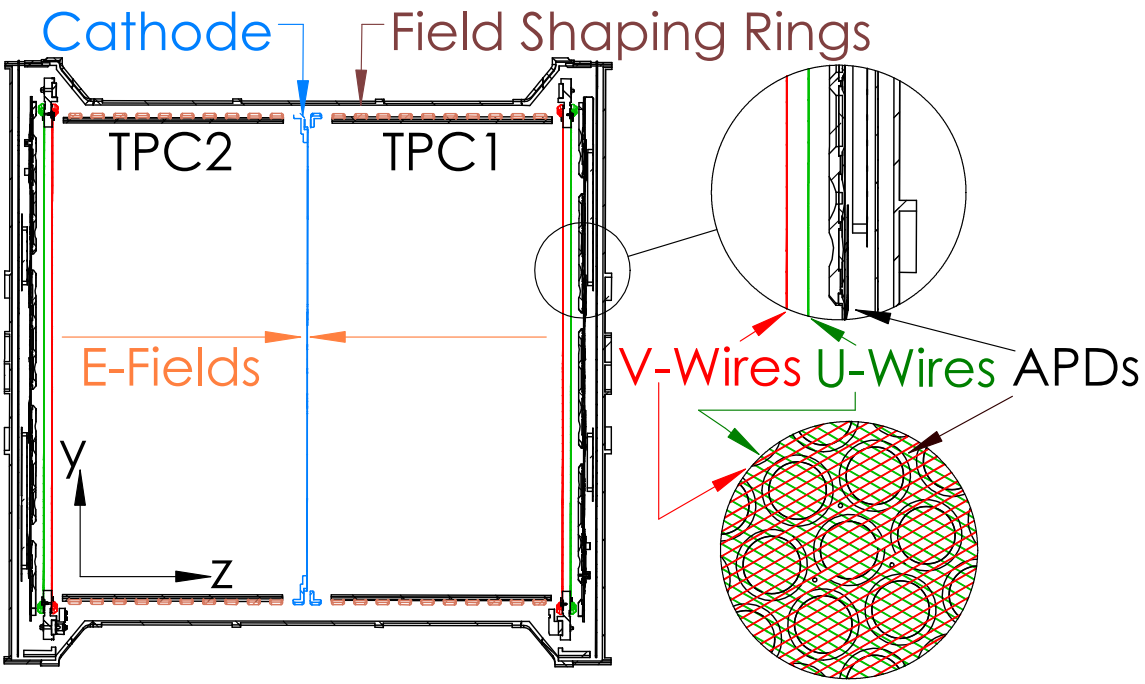

Figure 3.3: Cut out schematic of the EXO-200 LXe TPC. Shown in the center is the shared cathode plane (blue) and on either side the APD (black), V-wire (red) and U-wire (green) planes are shown. Also included on either side of the cathode are the copper field shaping rings and their support structure. The top inset shows a zoomed view of the wire and APD planes while the bottom inset shows the same view in the X/Y plane to illustrate the relative angle between the wire planes.

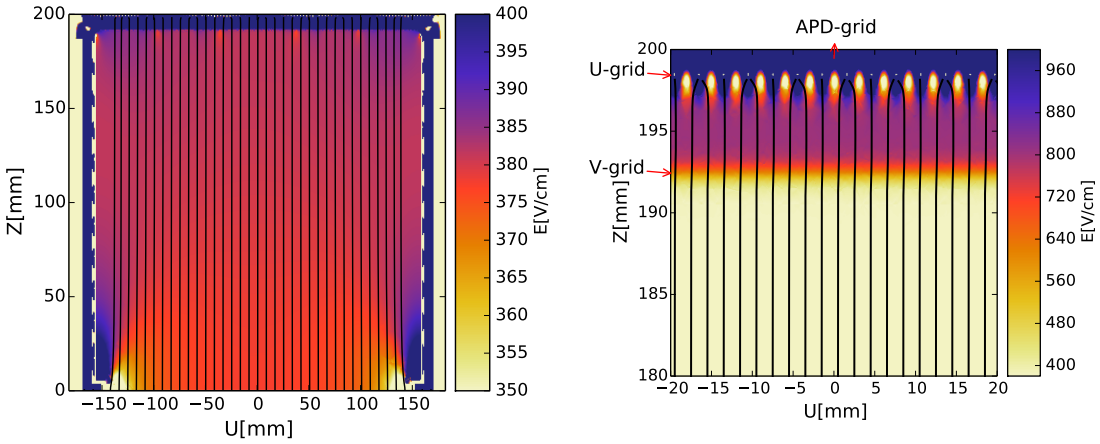

Figure 3.4: Magnitude of the electric field in the U-Z plane for (left) one entire drift space of the TPC (cathode is at Z=0 mm) and (right) a small segment near the anode. Also shown in black are the electric field lines that electrons would follow to the collection plane assuming no diffusion.

error on the bulk field denotes the maximum spatial variations around the mean field. The COMSOL simulation is only performed for the Phase-I configuration. The Phase-II field and associated error are derived assuming the field scales proportionally. Figure 3.4 shows a 2D segment of the resulting COMSOL model. At the edge of each TPC the field is graded in 10 steps by copper field shaping rings to produce a more uniform field along the z-axis within the LXe bulk.

The basic structure of the electronic readout chain used in EXO-200 is shown in Figure 3.5. Signals created by the detectors response to both charge and light are fed through kapton insulated flex cables to the front end electronics cards (FECs), where they are amplified by charge sensitive pre-amplifiers. Next the amplified signals are shaped using two integration stages with shaping time τ_1 to filter out high frequency noise and two differentiation stages with shaping time τ_2 to filter low frequency noise. A final differentiation stage with time constant τ_3 comes from input to the FET. A schematic of the first four stages of shaping for the U-Wires is shown in Figure 3.6. The shaped signals than go to to analog-to-digital converters (ADCs) before finally being passed into the data acquisisition system (DAQ), which contains a trigger event

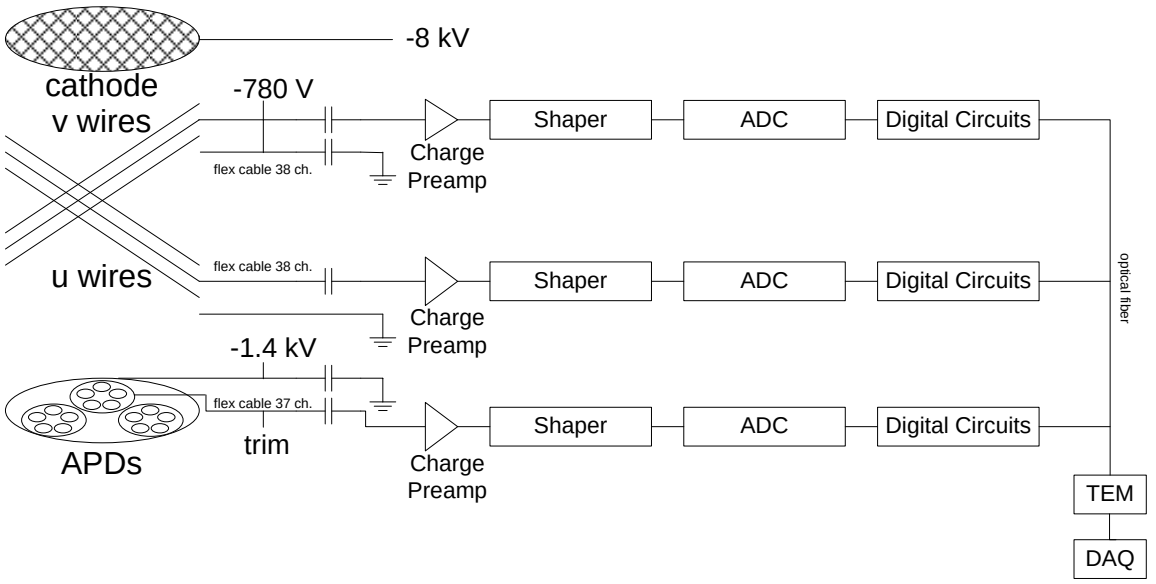

Figure 3.5: Schematic drawing of the EXO-200 electronics.

module (TEM).

Calibration of the detector is periodically performed by positioning γ-ray sources at several locations around the detector via a copper guide tube that wraps around the outside of the LXe TPC. A set of the standard source locations (denoted S2, S5, S8, S11 and S17) are chosen to produce events at a variety of points in the detector allowing for the study of position dependent responses. These locations are shown in Figure 3.7 and details of their positions are summarized in Table 3.1. Four sources (^{137}Cs, ^{60}Co, ^{228}Th and ^{226}Ra) are used to span the energy range of interest for the double beta decay search. Calibrations using the ^{228}Th positioned at the S5 location are done 3 times each week. This source has a prominent γ line at 2615 keV which is used for monitoring the energy resolution near $Q_{\beta\beta}$.

Electrons drifting toward the anode through the LXe volume can be captured on electro-negative impurities resulting in a Z-dependent attenuation of the observed charge (N_e)

$$N_e(t_d) = N_0 \exp\left(-t_d/\tau_e\right) \tag{3.1}$$

where N_0 is the initial number of electrons, t_d is the electron drift time and τ_e is the

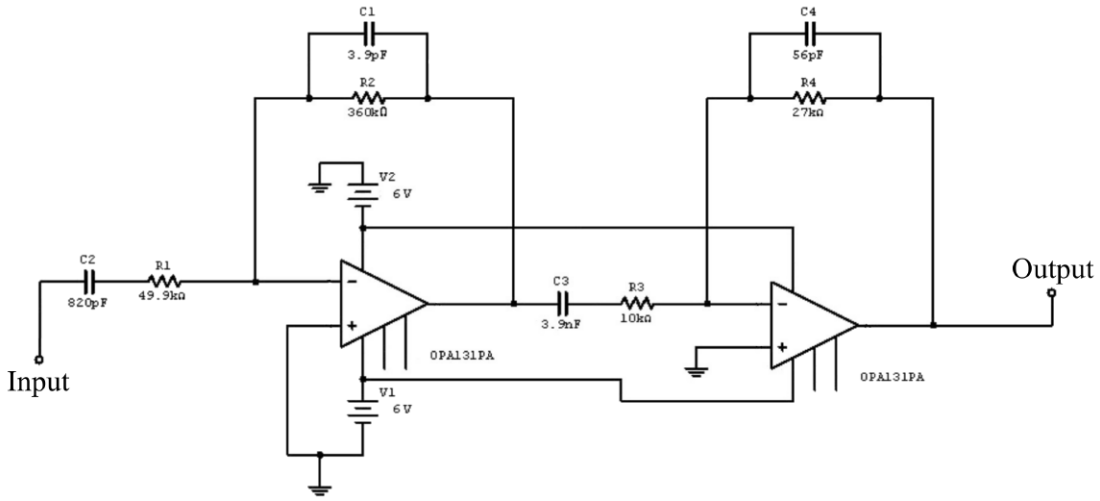

Figure 3.6: Schematic drawing of the first four shaping stages for the U-Wires. Taken from [75].

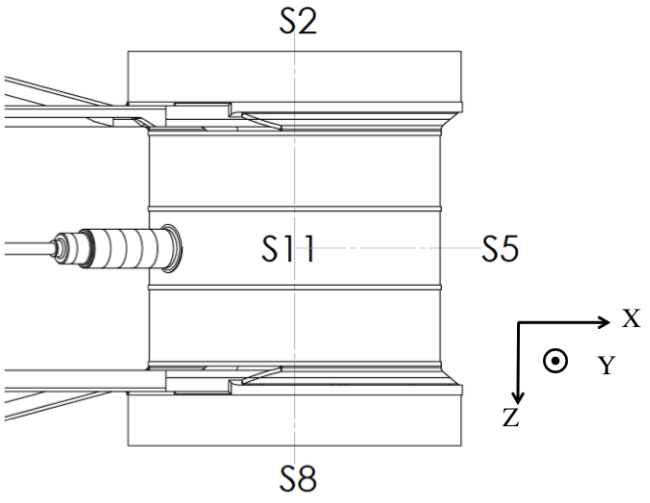

Figure 3.7: Diagram of the calibration source locations around the TPC. The third cathode position (S17) is located directly behind S11 so is not shown here. Taken from [64]

Table 3.1: Summary of the five standard source locations as well as their nominal positions in the X,Y and Z plane. The center of the cathode is the origin of the TPC coordinate system located at (0,0,0).

		Position [cm]		
Name	Location Type	X	Y	Z
S2	Anode	0	0	-29.5
S5	Cathode	25.5	0	0
S8	Anode	0	0	29.5
S11	Cathode	0	25.5	0
S17	Cathode	0	-25.5	0

electron lifetime in the detector. To minimize this attenuation the xenon is continuously circulated through purifiers to remove impurities and maximize the electron lifetime [63]. In addition, the electron lifetime is measured using calibration data taken weekly using the S5 cathode position to produces events distributed along the z-axis. This measurement allows for an accurate correction (described in Section 3.5.1) to the charge energy to normalize the detector response and achieve an accurate measure of the charge energy. [64].

3.3 Detector Simulation

Analysis of EXO-200 data relies on Monte Carlo (MC) simulations to model the detector response to ionizing radiation. For a given source, the MC simulation first produces energy deposits within the detector using a Geant4-based application [76] which employs a detailed model of the detector geometry [64]. The energy deposits are binned into cubic voxels with 0.2 μm edges known as pixelated charge deposits (PCDs). These voxels are than used to simulate both the charge and light response of the detector.

The light signal of each voxel is simulated using an empirically measured response

function to model the position dependence in the observed light signal. This response function is determined using calibration data and gives the relative distribution of signal among the two TPCs. The signal on each TPC is then evenly distributed among its APD channels, which is approximately true because only the sum APD response is used in reconstructing an event. The unshaped signals are than generate assuming a step function model.

In order to simulate the charge response, each voxel is tracked as it is drifted from the interaction location to the collection wires using a simulation of the electric field in the TPC. For every time step, the charge induced on each readout channel is determined from the Shockley-Ramo theorem [77, 78]. Early analyses [64], used a two-dimensional (2D) electrostatic modeling of the electric field in the detector and assumed azimuthal symmetry. This model worked well for these early analyses which relied on basic observables but as the analysis grew complex and examined finer details in the the detector response a more detailed model was required.

The main limitation is that the approximation of azimuthal symmetry does not accurately reflect the 60° orientation of the U-wires relative to the V-wires, which has a substantial effect on the trajectory of the charge deposits in the collection region. While modeling this trajectory is not necessary for the classic multiplicity definition, described in Section 3.5.2, it does impact the distribution of charge among individual U-wire channels. As a result, after incorporating diffusion [79], this 2D simulation gave poor agreement to data in the multiplicity and amplitude of signals induced on U-wires neighboring the collection channel. In order to improve the agreement, a full three-dimensional (3D) simulation of the electric field was implemented. This substantially improved the agreement of these signals between data and simulation. To optimize the speed of the simulation, only a segment spanning 4mm x 4mm x 25mm, with edges oriented parallel to the U- and V- wires, was used to approximate the field throughout the detector, assuming translational symmetry. This does not fully account for imperfections in the electric field near the edge of the detector, but provides a good description within the volume used for analysis.

To model diffusion, at each time step of size dt, additional random displacements, dx and dy, are added to the position in the x and y direction. As described in

Section 4.3, these displacements are drawn from independent Gaussian distributions with $\sigma^2 = 2D_T dt$ (where D_T is the transverse diffusion coefficient). To ensure sufficient granularity of the charge distribution, the initial 0.2 μm voxels are further split into equal deposits consisting of a maximum of ~600 electrons (~10 keV) prior to simulating their drift. MC studies indicated that this level of granularity is sufficient to accurately model the diffusion process, and the resulting reconstructed charge deposits are not affected by finer pixelization.

The result of both the charge and light simulations described above are unshaped waveforms for each readout channel. These waveforms are sampled at 1 MHz and contain 2048 samples with the light signal centered at the interaction time. Each waveform is then shaped with the known transfer functions described in Section 3.4.1 and the pulse amplitudes are scaled using the gain of the detector readout electronics. Additionally, noise samples from data are added to each waveform in order to produce a noise spectrum more closely resembling data, which is non-white. These noise samples are randomly taken during normal low background operation by periodically triggering the detector at 10 Hz. Waveforms for both MC and data are processed with the same analysis code to reconstruct signals on the U- and V-wires which is described in Section 3.4.

3.4 Event Reconstruction

A major advantage of TPCs is the ability to reconstruct the position and topology of interactions which deposit energy in the LXe. This is possible in EXO-200 by combining the reconstructed signals from the three readout planes described in Section 3.2 which are recorded when the detector is triggered. Here we describe the most advanced and likely final reconstruction algorithm which was used in the analysis of the complete EXO-200 dataset presented in [59]. Earlier and simpler versions of reconstruction were used for the previous analyses presented in [65, 66, 80]

Upon triggering, the DAQ records the output of all 152 charge and 74 APD channels in the detector. For each channel the recorded waveform consists of 2048 samples taken at a sampling rate of 1 MHz. These waveforms are then input into the

reconstruction algorithm which proceeds in three stages. First, a signal finding stage searches each waveform to determine if a real signal above noise is present. Second, each signal found in the first stage is fit using a known signal model to extract the relevant parameters. Finally, the signals of each channel type are grouped together to form clusters which represent individual interaction sites in the detector. For each cluster the Charge energy can be estimated from the U-Wire amplitude. The energy in the scintillation channel for each event is then determined using the APD amplitude. However, because of the large value of the speed of light, the scintillation signals from different clusters in a signal event cannot be resolved and are added together. The 2D position in the plane perpendicular to the drift (U/V) is determined using the U- and V-signals, which are related to the orthogonal X/Y coordinate frame by a linear transformation. The third position dimension parallel to the drift (Z) is then determined using the difference in time between the APD and U-Wire signals, which give the time of the initial interaction and the collection of the charge signal respectively.

3.4.1 Signal Finding

Signals are found on each channel by applying a matched filter to each waveform which compares the observed waveform with an expected signal model. Accurate templates are generated for each channel by applying the known transfer functions to the expected unshaped pulses. For APDs the unshaped signals are approximated by a step function, while for U/V-Wires a simulation of a point deposit is used to generate the unshaped signals.

As a charge deposit drifts, it follows the field lines until it ultimately deposits on a single channel. At each stage during its drift, the charge induces a signal on all of the wire channels which is inversely related to the distance between the deposit and the given channel. For a single point deposit, only one channel ultimately collects charge while all other channels see the charge from a distance. This results in a large "collection" signal on the hit channel while all other channels observe small "induction" signals which have amplitudes that fall off sharply as a function of the

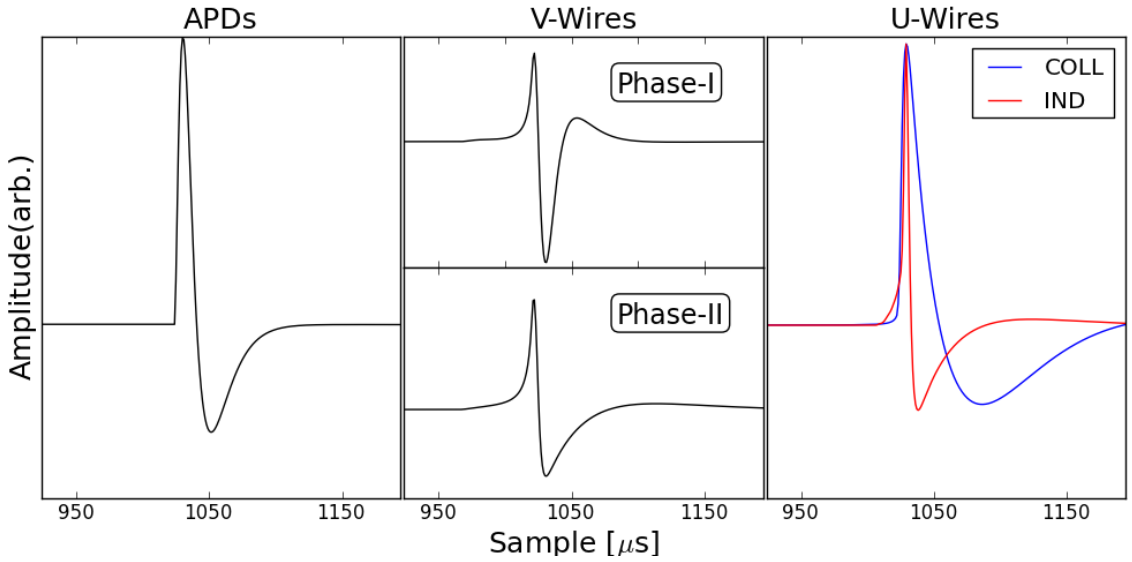

Figure 3.8: Signal models for APD (left), V-Wire (middle) and U-Wire (right) channels. For V-Wires the Phase-I and Phase-II models are shown separately due to the change in the nominal shaping times which changed the model shape. For U-Wires both the collection and induction models are shown for comparison.

distance from the collecting channel. In order to account for this, two separate signal models are used to characterize U-wire signals as shown in Figure 3.8. These models represent the basis from which any U-wire signal can than be represented, because a realistic deposit in LXe produces thousands of electrons, resulting in signals which are really a superposition of the two models. The case is simpler for V-wires, which do not observe collection, allowing signals to be represented by only one induction like signal model.

As described in Section 3.2.2, the signals which are recorded by the DAQ are the signals after shaping. In order to produce accurate signal models which match the shaped signals in data, the transfer function for each channel is required. The nominal time constants needed for the transfer function can be estimated using the known electronic configurations such as that shown for the U-Wires in Figure 3.6. This only accounts for the first four stages, the fifth stage comes from the input to the FET which for the U-wires uses a 1 pF capacitor and a 60 MΩ. This capacitor is only accurate to within $\pm$0.5 pF, resulting in a large uncertainty for this final

differentiation stage. In order to mitigate the impact of this uncertainty, this final shaping constant is individually measured for each U/V-wire channel by using charge injection calibrations. These more accurate times are found to improve the modeling of detector response. While APDs in principle have the same issue, the shaping for these channels is more aggressive and the nominal times are sufficient. The nominal shaping constants for each of the channel types are shown in Table 3.2.

	Stage Type				
Channel Type	Integration		Differentiation		
APDs	3	3	10	10	300
U-Wires	1.5	1.5	40	40	∼60
V-Wires(Phase-I)	3	3	10	10	∼60
V-Wires(Phase-II)	1.5	1.5	40	40	∼60

Table 3.2: Shaping times (in µs) used to define transfer functions in creating signal templates. The U-Wire and APD parameters are the same in both phases but the V-Wire shaping times were changed in Phase-II in an attempt to improve the signal to noise ratio. The longest differentiation stage for the the U/V-Wires is the nominal value but the signal templates are generated with the values extracted from calibration data.

After applying the matched filter to each waveform, the presence of a signal is determined by searching for a peak in the filtered waveform which is above a given threshold. This threshold is defined independently for each individual waveform using the observed baseline and is defined as 5(4) times the mean absolute deviation (MAD) for wire (APD) channels.

The matched filter is optimized to identify the presence of a signal but its broad response means it is unable to resolve the position of multiple peaks on the same channel which are too close in time. As mentioned, this is not necessary for APD signals. On the other hand charge drifts relatively slowly (∼2 mm/μs) and multiple deposits collecting on the same channel can produce distinguishable pulses. In order to search for additional peaks a second signal finding stage is employed for all channels

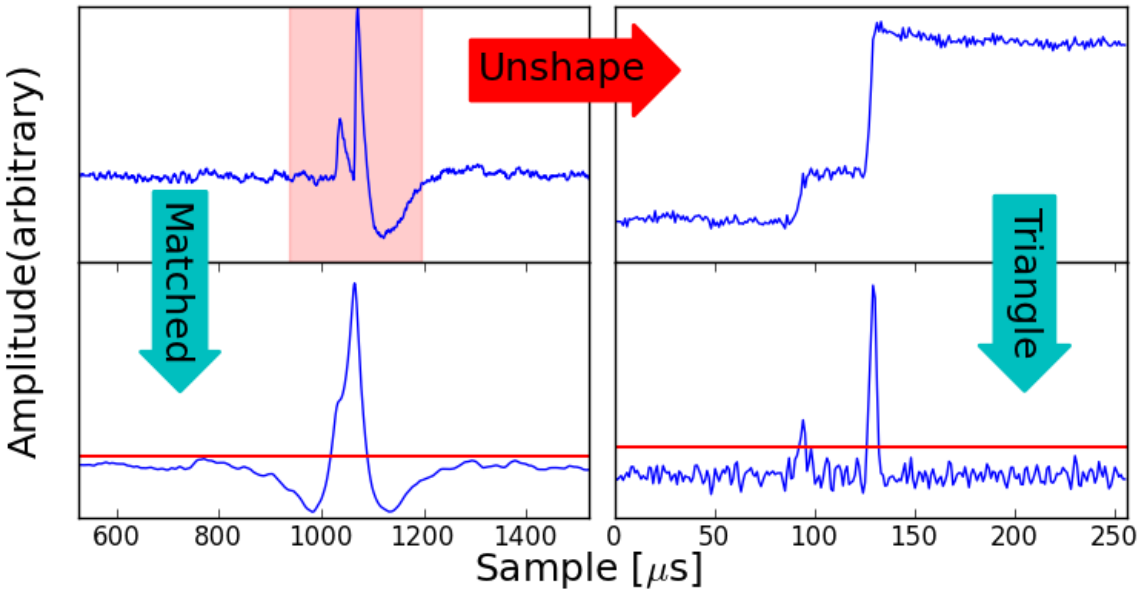

Figure 3.9: Example of signal finding procedure for a U-Wire waveform with two signals (upper left). The result of the matched filter (bottom left) shows a clear pulse but is too broad to distinguish the two peaks. Instead the unshaped signal (upper right) is passed through a triangular filter (bottom right) at which point the two signals can be disentangled. The red lines in the matched/triangular filter show the threshold defined by the MAD of the baseline.

with a found signal. First, this additional stage un-shapes each wire waveform which had a signal found by the matched filter using a 256 us interval around the signal. The resulting waveform is than reshaped with a 2 us triangular filter and this waveform is searched for additional peaks using the same approach as for the matched filter. This procedure is depicted for a U-Wire waveform with two distinct signals in Figure 3.9.

3.4.2 Signal Fitting and Parameter Extraction

The signal finding stage gives a rough estimate of the timing and amplitude of each found signal but a better estimate of these parameters is achieved by performing a χ^2 fit of each waveform with a found signal to its respective signal model. For V-Wire channels in Phase-I and APD channels the fit is done using a $\pm$40 us window around the signals output from the signal finding stage. For V-Wire channels in Phase-II

and U-Wire channels the window is extended to +140 us to account for the larger undershoot from the longer differentiation times. In the case of multiple signals on the same waveform, the windows are defined separately and combined if they overlap.

As mentioned in Section 3.4.1 there are two distinct signal classes for U-Wire channels. The signal finding algorithm is optimized to identify an excess above noise but is not well suited to distinguish between these two classes. Instead all found U-Wire signals are fit with both an induction and collection template and a discrimination algorithm is used to determine which class is more appropriate. This algorithm works by applying a set of 1D cuts on the expected signal characteristics which are optimized using MC data. A signal which passes all cuts is considered to be induction-like.

Previously it was assumed that these induction tagged signals did not correspond to charge collecting on the channel and thus were ignored in subsequent reconstruction stages including the charge energy calculation. This assumption is true for point-like deposits but in general is not valid. In reality, energy deposits in LXe consist of clouds of electrons produced from ionizing radiation traveling through the detector. Each electron in the cloud will than drift towards the anode and collect on a U-Wire channel. The observed signals on each channel will then be the summation of the signal induced by each electron. In this case a U-Wire signal can generically have a component from both collection and induction, with one component vanishing only in the specific cases.

Induction signals on the U-wires are typically small compared to collection signals, having amplitudes between 5-10% of the depositing channel. This can be seen in Figure 3.10 which shows the simulated collection signal on the hit channel as well as the induction signal on the neighboring channel resulting from a $\sim$2 MeV point-like energy deposit. Because of this, they are quickly overshadowed by deposits with more than $\sim$50 keV of energy. In the case of an induction signal produced from a large neighboring collection pulse, small charge deposits diffusing away from the main deposit can produce collection signals of similar scale as the induction signal. This can than result in the loss of energy if the pulse is assumed to be purely induction like.

This effect can be demonstrated in MC where the true energy deposited in the

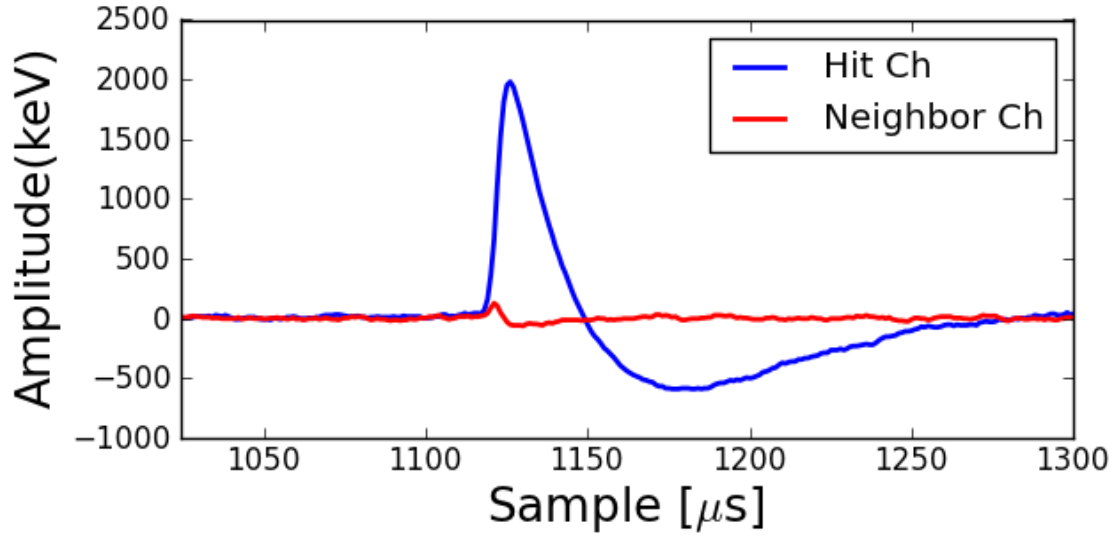

Figure 3.10: Example collection signal on a U-wire and the resulting induction signal on the neighboring U-wire channel produced from a ∼2 MeV point-like charge deposit in MC simulation.

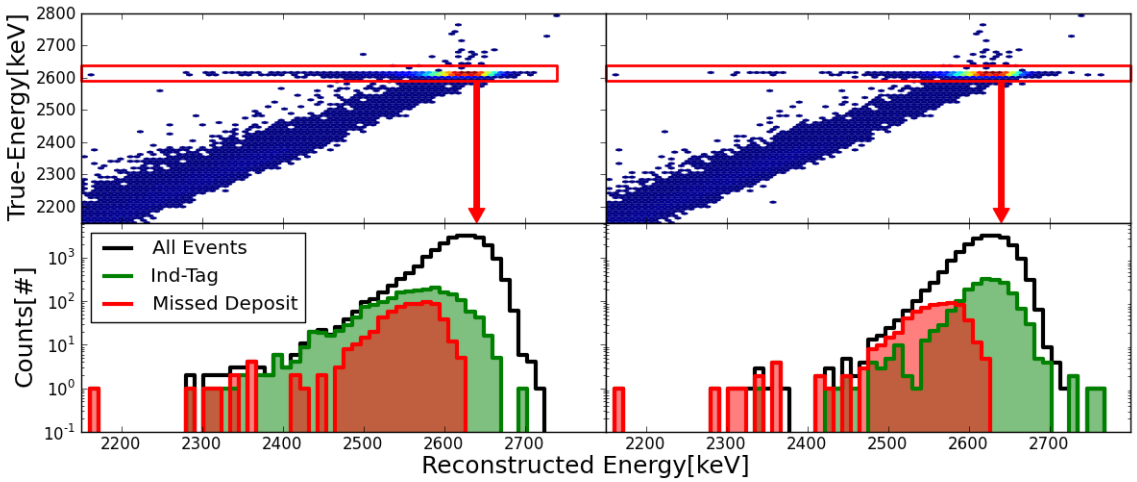

Figure 3.11: Distribution of true vs reconstructed energy from a simulation of the ^{228}Th calibration source (top) before refitting the induction signals (left) after refitting the induction signals (right). The bottom panel shows the distribution of reconstructed energies at the 2615 keV peak with the contributions from induction tagged events highlighted in green and the signals missed by reconstruction highlighted in red.

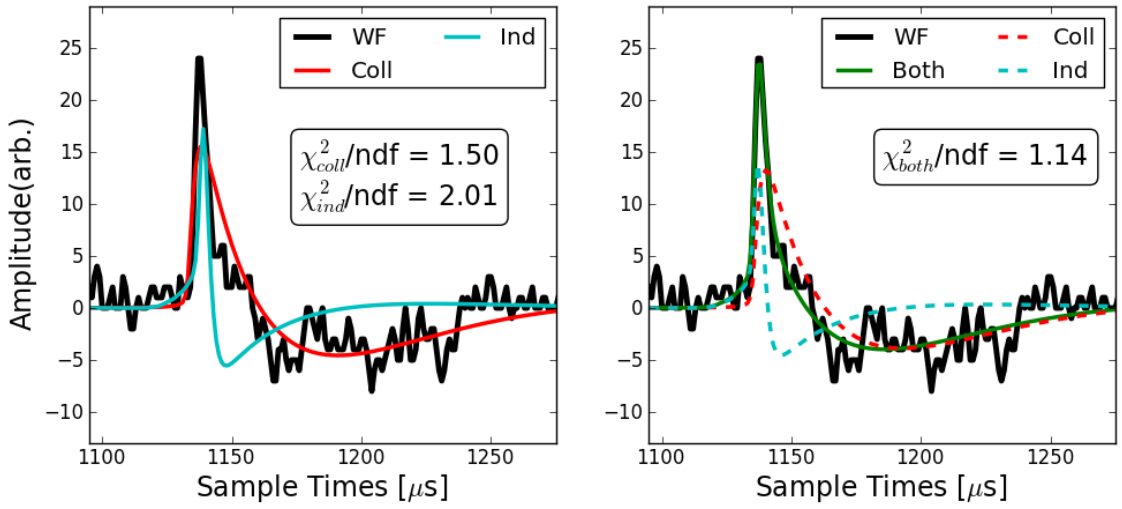

Figure 3.12: Example waveform from MC with real noise, which collects a small energy deposit and is neighboring a large collection signal. As can be seen by the χ^2 values, using a pure collection/induction model results in poor fit (left) but using a mixed model which is a superposition of the two models improves the fit (right). Dashed lines represent the individual collection/induction components in the sum fit.

TPC is known. Figure 3.11 shows the true vs reconstructed energy of MC events from a simulation of the ^{228}Th calibration source. Here a clear non-Gaussian tail in the reconstructed energy is observed at the 2615 keV peak. An examination of these events found that the tail was in large part composed of events containing at least one tagged induction signal. Further comparison to the results of the simulated charge propagation showed that these signals tend to correspond to small energy deposit with similar energy to that which was "missing".

To recover this previously lost energy each signal which is tagged as induction-like is refit using a χ^2 fit that allows for both a collection and induction component simultaneously. An example fit for one such mixed signals is shown in Figure 3.12 to demonstrate the improvement in the fit result. This shows both the initial fits assuming either a pure induction/collection signal as well as the fit after allowing for both components simultaneously. After implementing the joint fit the previously observed non-Gaussian tail is significantly reduced, as shown in Figure 3.11. The remaining tail is mostly due to small deposits falling under the signal finding threshold

that are missed during signal finding, which are highlighted in red.

3.4.3 Matching and Clustering

All found U-Wire, V-Wire and APD signals are then grouped together to form "charge clusters", which approximately correspond to distinct energy deposits within the TPC. This is accomplished by first using the timing, amplitude and location of each signal to determine which signals of the same channel type (U,V or APD) derive from the same energy deposit in the TPC. The resulting groups of U- and V- signal are then combined into clusters with 2D position in the U/V plane using a matching algorithm to determine the combinations which are most likely to be caused by the same energy deposit. This is achieved with a set of probability density functions (PDFs) that described the time, energy and position expected from the U-Wires given a set of V-Wire signals generated by a common energy deposit. Each of the resulting 2D clusters are then assigned a 3D position by matching them with an APD grouping to determine their Z-position.

In the case that signals from the two wire planes and the APD plane are found, the above procedure will result in a set of charge clusters with all three coordinates reconstructed. Due to the different signal finding threshold of channel type, there is the possibility that at least one signal type will be missing, resulting in partially reconstructed clusters which are missing at least one dimension. To determine the threshold of each channel type the reconstruction efficiency can be determined as a function of energy. Because the U-Wires have the lowest threshold they can be used as a "standard candle" for determining the V-Wire and APD efficiencies. This is done by finding the ratio of one charge cluster events with both a U-Wire and V-Wire/APD signal to all events containing a U-Wire signal. For these events the charge energy can be estimated with the U-Wire signal. For U-Wires there is no additional reference energy and instead MC with noise sampled from data is used. The efficiency is than determined as the ratio of events with a found U-Wire signal to those without and using the known simulated energy as the reference energy. The resulting efficiency curves of each channel type are shown in Figure 3.13 for Phase-I and Phase-II.

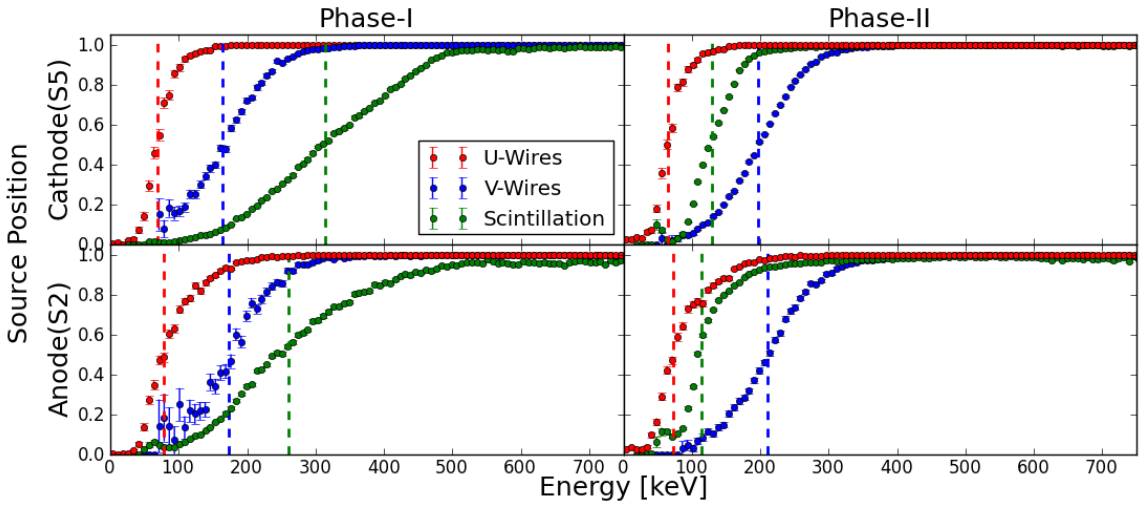

Figure 3.13: Reconstruction efficiency in (left)Phase-I and (right)Phase-II for the three channel types for calibration data taken at (top) the S5 Cathode position and (bottom) the S2 Anode position. The Scintillation and V-Wire efficiency are determined using data with the U-Wires acting as the energy reference. For U-Wires the efficiency is determined with MC.

The threshold for each channel is determined as the 50% efficiency mark. For the U- and V-Wires the thresholds are found to be ∼80 keV and ∼200 keV and are approximately the same in both data taking phases with small differences likely due to time dependent noise variations. On the other hand, the scintillation efficiency is found to be ∼300 keV in Phase-I and ∼150 keV in Phase-II. This improvement is largely due to the upgraded electronics which reduced the noise on the APD channels prior to the start of Phase-II. In addition, both the curves using calibration data at the anode and cathode are shown separately to highlight position dependent effects. This is most noticeable in the scintillation curve which is slightly worse near the cathode which is expected due to the lower probability for photons produced near the cathode to reach the APD plane.

In the case of multiple deposits it is possible for at least one found charge cluster in the event to contain a U-Wire and APD signal but no V-Wire signals. This is due to the fact that the sampling rate is unable to disentangle the light signal from individual deposits and the observed light signal is the sum signal produced by all

deposits. Thus the light threshold is dictated by the sum energy of all charge clusters. On the other hand, each single cluster is subject to the individual U- and V-Wire thresholds and due to the to the higher reconstruction threshold for U-Wire signals a charge cluster with energy < 300 keV has some chance to be missing a V-Wire signal.

Because these clusters lack a V-Wire signal, their position in the V-Wire plane is unknown and only a 2D position in the U/Z plane can be determined. Events containing such a cluster are considered "non-3D" as they lack a fully reconstructed position. However, these clusters and their associated event have an accurate energy measurement because the V-Wires are not used in the charge energy determination. In addition, the relevant charge energy corrections described in Section 3.5.1 are accurately applied as they only rely on the known Z-position. The light-map correction requires both an X and Y position and so this correction is done using only the position of the fully 3D reconstructed charge clusters.

3.5 Background Discrimination

As described in Section 2.3.1 a key component to a $0\nu\beta\beta$ search is the ability to determine if events interacting in the detector are coming from $0\nu\beta\beta$ or one of the backgrounds. This discrimination is achieved in EXO-200 via a combination of energy, position and topology.

3.5.1 Energy

The main observable used to separate $0\nu\beta\beta$ signal is the total energy deposited in the detector. Estimating the energy of an event in LXe can be done independently using either the charge (E_Q) or light (E_S) response of the detector. In addition, the light and charge response are anti-correlated. This allows for a linear combination of the two observables to be formed which gives an optimal estimate of the energy.

Charge Energy

The charge energy is estimated for each charge cluster using the sum amplitude of all U-Wire collection signals contained in the cluster. To remove variations in the relative response of each U-Wire channel, a gain correction is applied to each channel. The gains are measured using the pair-production peak produced of the 2615 keV γ from the ^{228}Th calibration source. In addition, a correction is applied to the energy of each cluster which contains a mixed Induction and Collection signal in order to recover small deposits which are otherwise missed, as described in Section 3.4.2.

As described in Section 3.2.2, electro-negative impurities in the detector can cause the observed charge energy to be attenuated as a function of the time spent drifting electrons from their initial location to the anode. In order to stabilize the detector response, the electron lifetime in the detector is monitored weekly using calibration data and the measured electron lifetime can be used to correct the charge energy of each cluster based on the expected attention given by Equation 3.1. In addition, a selection cut is applied to all data to require that only data taken during periods for which $\tau_e \geq 2$ ms are used. This ensures that losses are $\lesssim 5\%$ over the full drift length of $\sim$100 µs, to minimize the impact of uncertainties in extracting and applying the lifetime correction.

Finally a correction is applied to E_Q to account for the imperfect shielding of the U-Wire grid by the V-Wire grid. This results in a residual Z-dependence in the detected U-Wire amplitude. This dependence can be modeled as

$$E_d = \frac{E_0}{1 + p_0 \exp\left[\left(|Z| - Z_{max}\right)/p_1\right]} \tag{3.2}$$

where E_d is the detected energy, E_0 is the true energy of the cluster, Z_{max} is the maximum drift distance and p_0, p_1 are free parameters determined using a fit to calibration data.

Light Energy

The light energy is determined for each event from the summed APD response. To achieve the optimal light estimate a correction is applied to the observed light energy

based on the 3D position of the event in the detector. This accounts for variations in distribution of light among the APD channels caused by differences in the solid angle and gain among different channels. Using data taken with the ^{228}Th calibration source an empirical correction function is generated by dividing the detector into 3D voxels and fitting the light response of each voxel at the 2615 keV γ line.

In addition a "denoising" algorithm is applied to both Phase-I and Phase-II data to optimally estimate the scintillation energy of each event. The algorithm was initially developed for Phase-I to address the observation of excess correlated noise among the APD channels which limited the scintillation resolution [81]. While this correlated noise is mostly removed by the electronics upgrade prior to the start of Phase-II, the denoising algorithm is still able to provide a slight improvement to the energy resolution in Phase-II because the algorithm optimally weights the contribution of each APD to the sum energy to maximize the signal to noise ratio.

Rotated Energy

As discussed in Section 3.2.1 the recombination of electrons with positive ions produces an anti-correlation between the charge and light response in LXe. Random fluctuations in the number of recombining electrons, which is still not well understood in LXe, then results in degraded resolution when measuring the energy with only one of these two signals. In order to obtain the optimal resolution a linear combination of the charge and light energy can be used.

This new energy estimate is equivalent to a rotation in the 2D charge/light energy plane as shown in Figure 3.14 and is defined as the Rotated Energy (E_R). The linear combination is then defined with respect to a rotation angle (θ):

$$E_R = E_Q \cos\theta + E_S \sin\theta \tag{3.3}$$

In order to calculate θ, ^{228}Th calibration data is used to scan and fit for the angle resulting in the best resolution at the 2615 keV gamma line. The improvement in resolution using the linear combination can be seen in Figure 3.14 which shows the observed energy spectrum as well as the resulting resolution using each of the three

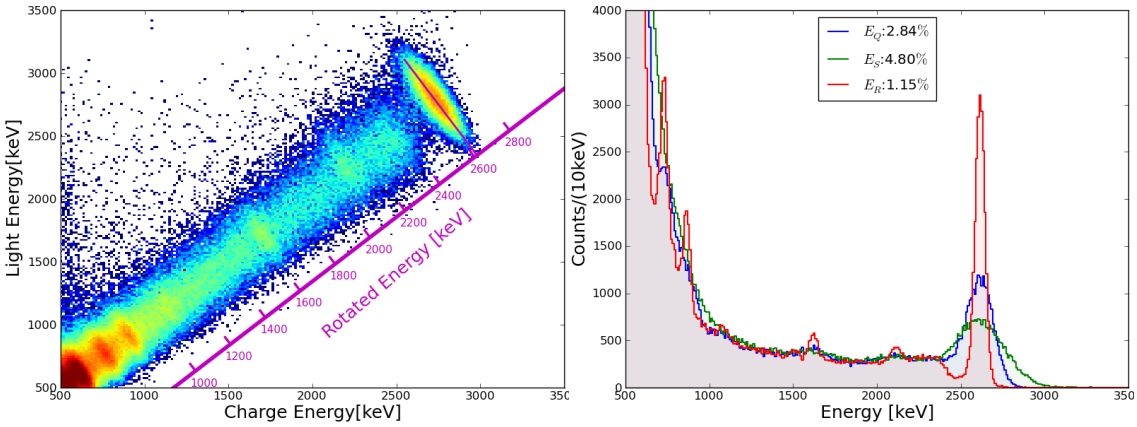

Figure 3.14: (left) The light/charge response of the detector during a ^{228}Th calibration run showing the anit-correlation between the two signals. (right) A comparison of the energy spectrum using each of the three energy estimates as well as the resulting energy resolution at the 2615 keV peak.

energy estimates.

This optimal θ and the resulting resolution are measured periodically over the course of data taking to capture potential variations arising from changes in the detector response. The achieved resolution at the 2615 keV line for both Phase-I and Phase-II is shown in Figure 3.15. The resolution improvement between the phases is in part due to the electronics upgrade, which reduced the electronic noise on the APDs. The increase in the drift field also played a role, as it reduced the number of electrons which recombine. As a result a larger fraction of the overall energy is shifted into the better measured charge channel, resulting in an improved energy estimate.

To obtain the expected resolution at $Q_{\beta\beta}$, the resolution is parameterized as a function of the energy(E)

$$\sigma^2 = \sigma^2_{elec} + bE + cE^2 \tag{3.4}$$

where σ^2_{elec}, b and c are determined by fits to calibration data. In order to account for spatial variations in the detector resolution and differences in the distribution of events taken during calibration, the average detector resolution is determined by uniformly weighting all calibration data from several position and accounting for the detector livetime. The resulting average resolution curve using Equation 3.4

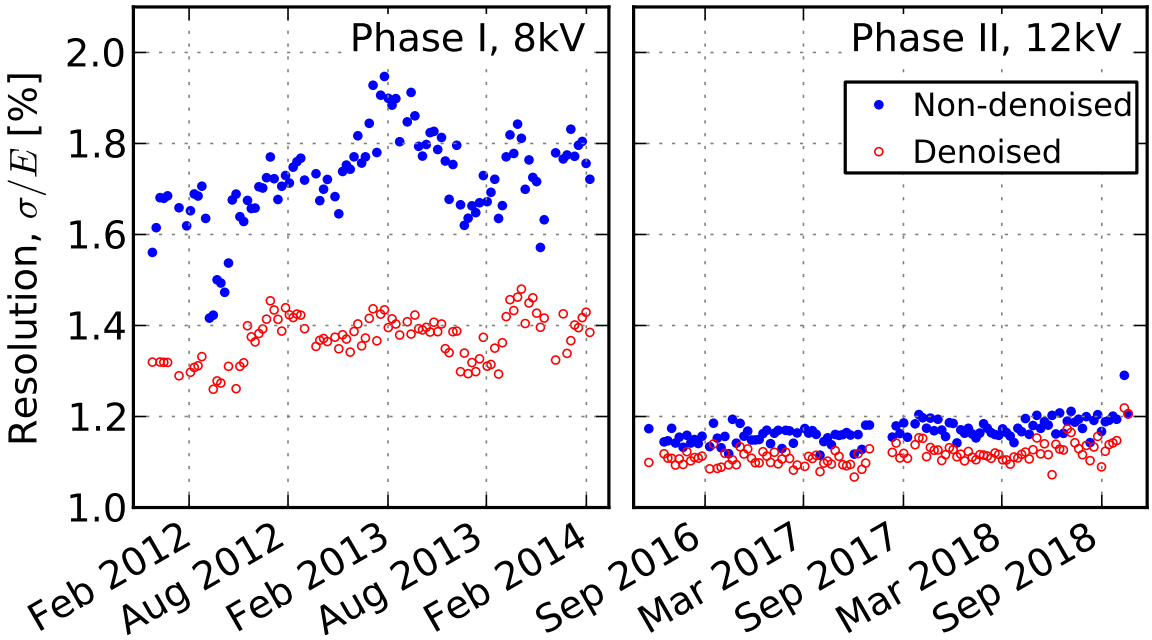

Figure 3.15: Observed energy resolution for SS events at the 2615 keV of the ^{228}Th during weekly calibrations at the S5 cathode position. Both the non-denoised(blue) and denoised(red) resolutions are shown.

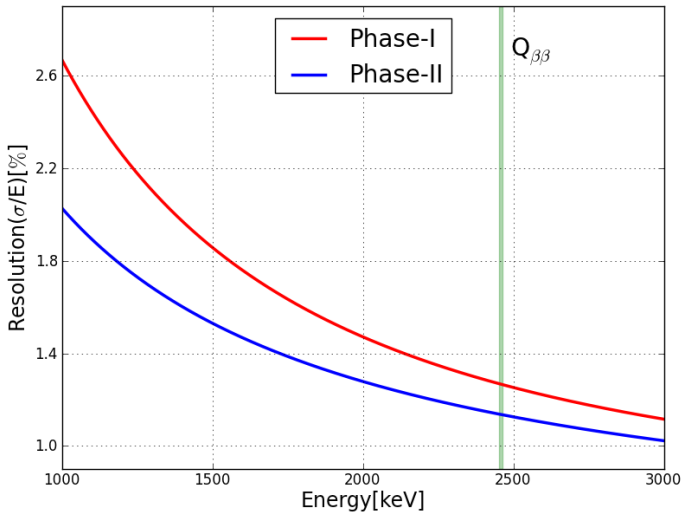

Figure 3.16: Average resolution curves for SS events in Phase-I (red) and Phase-II (blue). The observed resolution at $Q_{\beta\beta}$ for Phase-I (Phase-II) is $1.35 \pm 0.09\%(1.15 \pm 0.02\%)$.

is shown in Figure 3.16. The observed resolution at $Q_{\beta\beta}$ for Phase-I (Phase-II) is $1.35 \pm 0.09\%(1.15 \pm 0.02\%)$.

Light to Charge Ratio

While the 1D projection onto E_R gives the energy estimate which optimizes the resolution, there is additional information contained by the location of the event in the 2D energy plane. Specifically, it is possible to identify α's interacting in the detector due to their relatively short range ($\sim \mu$m) in LXe. These tracks contain clouds of very high density electron-ion pairs, resulting in higher rates of recombination and hence reducing the charge yield and increasing the light yield.

In order to remove these events a "diagonal" cut is applied which sets an upper bound on the expected light signal for a given charge energy. This cut is depicted in Figure 3.17 which shows the light/charge response during low-background operation. A clear population of α events exists off the main diagonal with disproportionately high light energy. These events are removed by the diagonal cut.

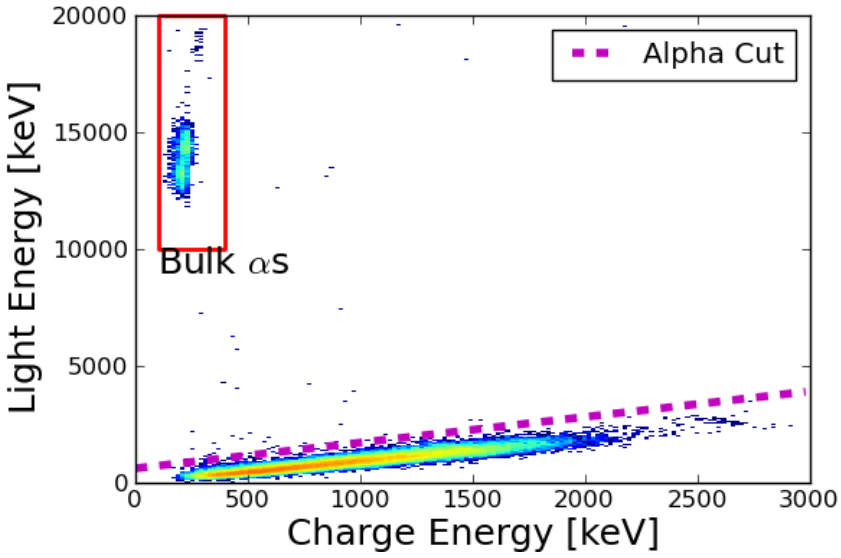

Figure 3.17: Cut on the ratio of charge to light designed to remove α's such as the highlighted population of bulk α's in the upper left corner.

3.5.2 Topology

Other than $2\nu\beta\beta$ (which can mostly be removed from its energy signature), the main backgrounds which limit the search sensitivity are the result of γ's produced by natural decay chains with long lived parent isotopes, such as ^{238}U and ^{232}Th. While careful material screening was done to limit their presence in the materials used to construct the detector, low levels of these isotopes still remain. As a result the main source of this γ background are the external detector materials, such as the copper TPC vessel and the cryostat. Many of these γ's can be removed based on energy but there is still some population of events which are expected to deposit energies too close to $Q_{\beta\beta}$ to distinguish with energy along.

In order to improve background rejection, it is possible to take advantage of the differences in the physical processes which govern γ and β interactions with LXe. At energies near $Q_{\beta\beta}$ the dominating process for γ interactions is Compton scattering. Due to the large detector size relative to the γ's mean free path, a γ will tend to deposit its energy in more than one location within the LXe. On the other hand, β's have a range of $\sim$6 mm at $Q_{\beta\beta}$ [82] and thus tend to deposit their energy in a single

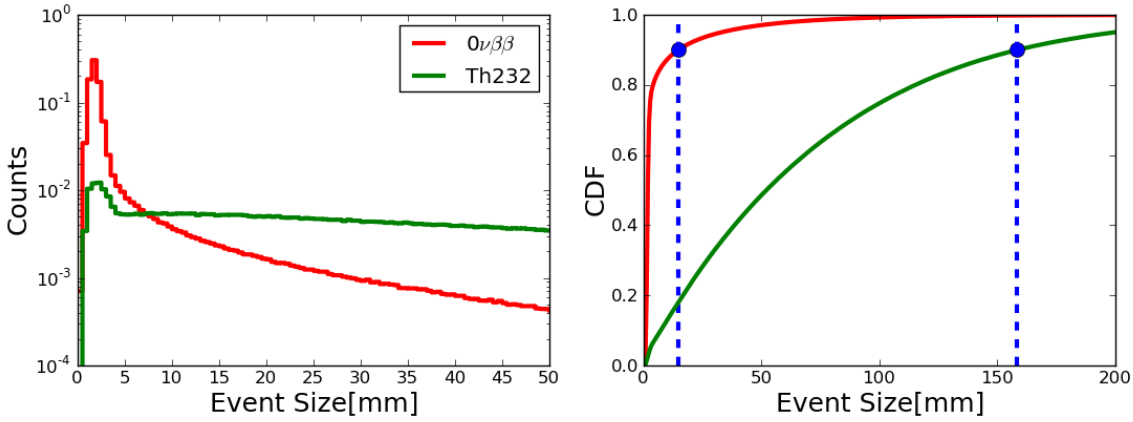

Figure 3.18: Comparison of the simulated event sizes for signal like $0\nu\beta\beta$ and a prominent γ background(left) and the corresponding CDF for both distributions with the 90% point marked by a vertical line(right). While both signal and background have a peak at sizes < 5 mm the γ background has a significant tail extending to large event size due to Compton scatters.

site. This results in γs having a larger spatial extent than that expected for β like interactions. This is depicted in Figure 3.18 which shows the distribution of the true spatial extent in MC for both $0\nu\beta\beta$ and a prominent γ background, where the true extent is calculated as the largest distance between simulated energy deposits in the LXe. Additionally the cumulative distribution function (CDF) of each distribution is shown with the 90% mark denoted. This shows the significantly longer tail of the γ like data set, from which it can be observed that only $\sim$15% of γ's have sizes under 10 mm.

In order to separate events based on their spatial extent an additional observable, called multiplicity, can be created by making use of the TPCs ability to separate individual energy deposits as described in Section 3.4.1. This allows events to be classified into two categories of multiplicity. Events with one deposit location are classified as single-site (SS) and are expected to be predominately due to β/signal like interactions. On the other hand events with more than one deposit can be classified as multi-site (MS) and are expected to be predominately due to γ backgrounds.

To first order the multiplicity determination is based on the number of reconstructed charge clusters. In order to account for scatters perpendicular to the X/Y

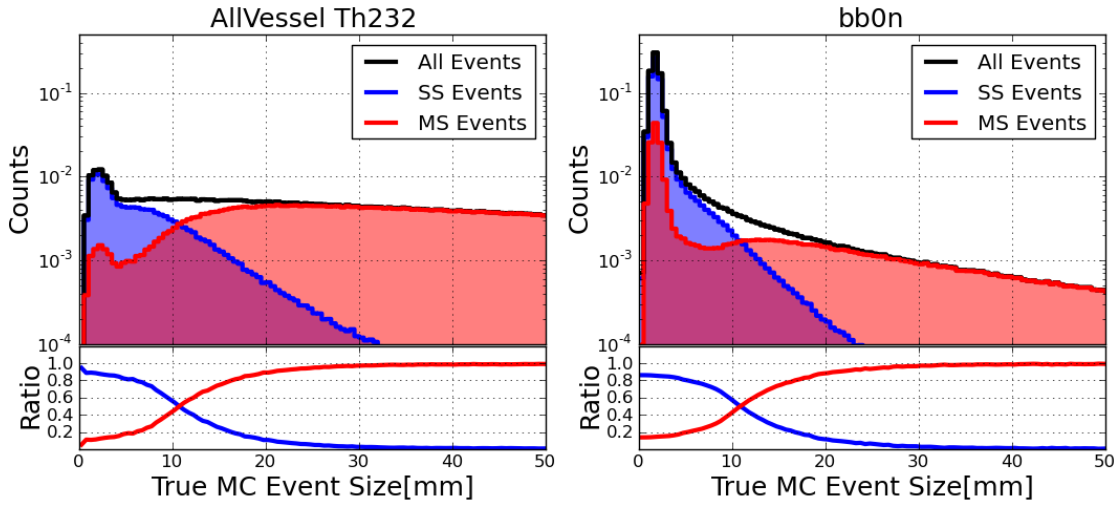

Figure 3.19: Distribution of the true spatial extent of events in MC with the individual contribution from SS/MS events highlighted(top) and the ratio of the SS/MS distributions to the total number of events(bottom) for a γ background(left) and $0\nu\beta\beta$ signal(right).

plane which are too close to be distinguished, an additional criteria requires that a single charge cluster must have no more than two U-Wire channels seeing collection signal to be classified as SS. This criteria is based on the expectation that the majority of β-like events with energy near $Q_{\beta\beta}$ have extents < 10 mm, as shown in Figure 3.18. Because the U-Wire pitch is 9 mm an event extending over more than two U-Wires would be in excess of 18 mm, which is inconsistent with this expectation.

The ability to separate signal and background using multiplicity can be demonstrated in MC as shown in Figure 3.19. This plot shows the event size distribution for events classified as SS or MS for both $0\nu\beta\beta$ signal and a prominent γ background. In addition, the ratio of events classified as SS/MS to the total number of events is shown in the bottom panel to demonstrate the efficiency of multiplicity to identify events based on their spatial extent. From Figure 3.19 it can be seen that multiplicity alone is a powerful tool for separating events, with near 100% efficiency at identifying events in excess of 20 mm, accounting for nearly 80% of the γs from this background component.

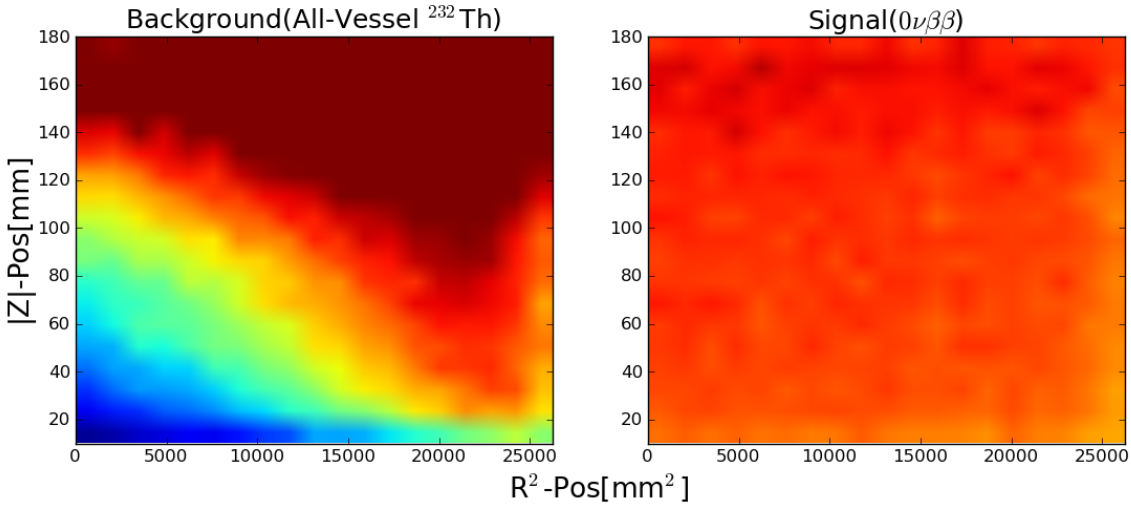

Figure 3.20: Distribution of events within the detector in MC simulations for an external background source(left) and $0\nu\beta\beta$ events(right). A clear non-uniformity can be seen in the distribution of backgrounds.

3.5.3 Position

Another benefit of TPCs is the ability to reconstruct the 3D position of an event. This information is important for optimizing the energy resolution via position dependent corrections but can also be used to separate signal and background.

Position based discrimination is possible due to the non-uniform position distribution of many γs backgrounds. This arises from the fact that γs with energy near $Q_{\beta\beta}$ have a mean free path of $\sim$9 cm in LXe, compared to the detector which is approximately 40 cm in height and width. External γs will then tend to interact more frequently near the edge of the TPC, resulting in a non-uniform position distribution clustered more closely to the detector's surface. On the other hand internal sources, such as $2\nu\beta\beta$ and $0\nu\beta\beta$, which are dispersed in the LXe produce interactions uniformly distributed within the volume of the detector. This effect can be seen in MC by comparing the distribution of events within the detector volume for both an external γ background (^{232}Th from the TPC Vessel) and signal ($0\nu\beta\beta$) as shown in Figure 3.20.

To capture this effect a new observable can be defined as the shortest distance between any charge cluster in an event and the edge of the detector. A depiction of

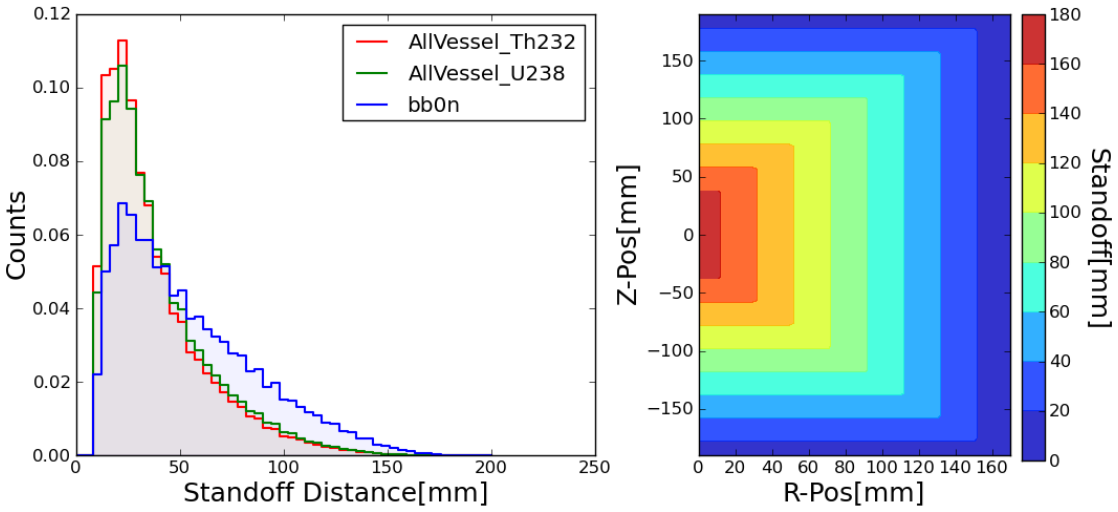

Figure 3.21: Depiction of the standoff variable definition (right) and distribution of standoff in MC for two prominent external γ background sources as well as the $0\nu\beta\beta$ signal (left).

this variable, named "standoff distance", is shown for different regions in the detector in the right of Figure 3.21. In addition Figure 3.21(left) shows the expected standoff distance distribution for two prominent external γ backgrounds as well as $0\nu\beta\beta$ signal in MC. As expected the γ backgrounds have more events at small standoff(near the edges), then evenly distributed sources such as $0\nu\beta\beta$.

In order to make full use of position information, an accurate position reconstruction is necessary. The position resolution can be determined using MC where the true position of an event is known based on the location of simulated energy deposits. Comparing the relative difference between the reconstructed position and the true position, shown in Figure 3.22, results in a resolution of 1.82 mm, 1.42 mm and 0.41 mm for U,V and Z, respectively. The superior resolution in Z is due to the 1 µs sampling time, which approximately corresponds to an effective segmentation of $\sim$2 mm (assuming a drift of $\sim$2 mm/µs) compared to the 9 mm segmentation of the U/V-wires. In addition, the slightly better resolution in V over U is due to the higher channel multiplicity of V-Wire signals, which allows for an improved constraint on the position.

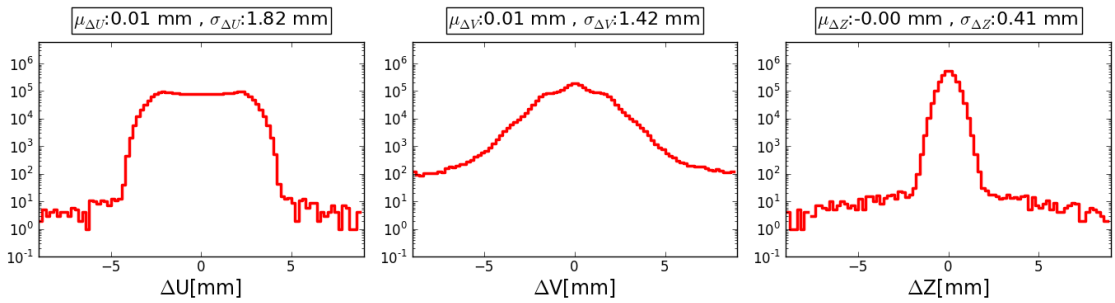

Figure 3.22: Distribution of the difference between reconstructed and true position along the (left)U, (middle)V and (center)Z axes estimated using MC. The mean and width of each distribution are shown in the legend. The resolution along each axis is 1.82 mm, 1.42 mm and 0.41 mm for U,V and Z respectively.

Chapter 4

Electron Transport Properties in LXe

4.1 Overview

As described in Section 3.5.2, a powerful tool of LXe TPCs is the ability to reconstruct event topology to differentiate between different types of interactions in the detector. For example, topology is used in $0\nu\beta\beta$ searches to discriminate between a $0\nu\beta\beta$ signal, consisting of two spatially unresolved electrons, and γ backgrounds, that tend to produce multiple energy deposits from Compton scattering at energies near $Q_{\beta\beta}$. This

discrimination, in addition to the energy of the deposit, helps to eliminate sources of background radiation that could otherwise obscure a signal and results in an improved sensitivity to $0\nu\beta\beta$ [59, 66].

Such topological classification requires a detailed understanding of the diffusion process for electrons as they drift under the influence of the electric field in LXe. Diffusion effects that are modest in EXO-200, which has a relatively short drift length ($\lesssim$ 20 cm), become more important in the ton-scale LXe detectors being developed for both $0\nu\beta\beta$ and dark matter searches, which will have drift lengths $\sim$1 m [83, 84]. Accurate measurements of the electron diffusion in LXe are therefore important for understanding background discrimination in these next-generation experiments.

Electron diffusion in LXe has previously been studied at fields higher than those of EXO-200. Early work using short drift lengths ($\lesssim$5 mm) measured the transverse and longitudinal diffusion coefficients for drift fields between 700 V/cm $< E_d <$ 7500 V/cm [67, 85]. In addition, a recent measurement using data from the XENON10 detector determined the longitudinal diffusion coefficient at E_d = 730 V/cm [86]. As described in Section 3.2.2, EXO-200 typically operates with a bulk electric field of 380 V/cm (567 V/cm) in Phase-I (Phase-II), allowing the measurement of the transverse electron diffusion at lower electric fields.

Summarized here are the results, presented in [79], of a measurement of the transverse diffusion constant and drift velocity of electrons in LXe at drift fields between 20 V/cm and 615 V/cm using EXO-200 data. At the operating field of 380 V/cm EXO-200 measures a drift velocity of $1.705^{+0.014}_{-0.010}$ mm/μs and a transverse diffusion coefficient of 55$\pm$4 cm^2/s.

4.2 Low Field Data Set

In order to study the electric field dependence of the electron transport proprieties in LXe at low fields, a dedicated campaign was taken using the ^{228}Th calibration source at fields ranging from 20 V/cm to 615 V/cm. With the exception of the higher field point, data taken at each field included deployment of the source at both the S2 anode and S5 cathode positions summarized in Table 3.1. This allowed for the

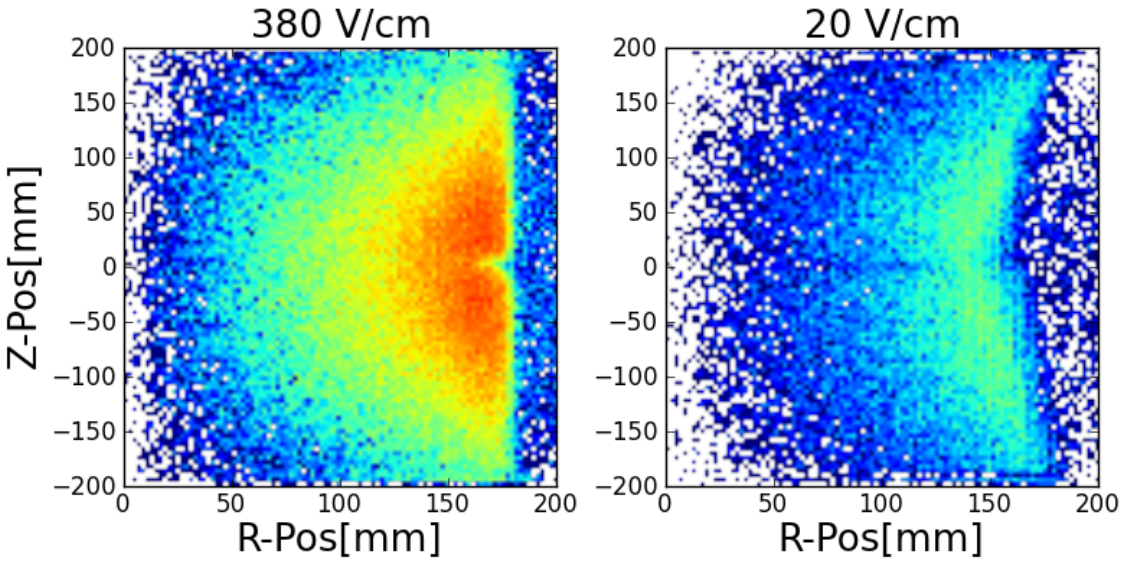

Figure 4.1: Spatial distribution of events in the detector for a ^{228}Th calibration run at the S5 cathode position for data taken at 380 V/cm (left) and 20 V/cm (right). A clear deficit of events at large radius near the cathode is observed in the low-field run, which can arise from imperfections in the electric field near the edge of the detector.

study of potential bias introduced from the non-uniform position distribution within the detector.

When operating at these additional fields, the bias applied to the V-wire grid and cathode was chosen to maintain a constant factor of 2 higher mean field in the collection region between the U- and V-wires relative to the bulk of the detector, in order to maintain full transparency of the V-wire grid to electrons. The additional data acquired during this campaign as well as the nominal biases are summarized in Table 4.1.

To limit effects of poor purity, only data for which the electron lifetime is > 2 ms is used in this analysis. In addition, the radial fiducial volume cut was tightened, relative to that in Section 6.5.2, to remove regions of non-uniform electric field. As shown in Figure 4.1 a deficit of events at high radius near the cathode was observed for data taken at lower fields. This indicates that significant radial non-uniformities were present in the drift field. As described in Section 3.3, the electric field model used to drift electrons in simulation uses a segment of the full 3D model to approximate the

Field[V/cm]	Cathode [kV]	V-Wires [kV]	S5 Run [#]	S2 Run [#]
615	12.93	1.1	6481	N/A
187	4.07	0.47	6454	6447
75	1.75	0.31	6453	6451
38	0.98	0.25	6460	6461
20	0.60	0.23	6466	6467

Table 4.1: Summary of calibration data taken at fields outside of the nominal 380 V/cm used in Phase-I, including the nominal bias of the V-Wire and Cathode.

field in the detector. This segmented 3D approximation provides a good description within the fiducial volume and improves the modeling of electron drift between the wire grids but it does not fully account for imperfections in the electric field near the edge of the detector. To avoid introducing systematic errors due to these non-uniform regions which are not simulated, the radial fiducial volume cut was tightened to remove these regions from the analysis. As described in Section 4.5, checks were performed to ensure no radial dependence was present within the fiducial volume used at each field.

4.3 Diffusion Model

For an initial δ-function charge deposit of N electrons centered at position $\vec{x} = (0, 0, 0)$ at time $t = 0$, the charge density, $n(\vec{x}, t)$, at later time t and position $\vec{x}$ can be determined by solving the 3-dimensional diffusion equation [87]:

$$n(\vec{x}, t) = \frac{N}{4\pi D_T t \sqrt{4\pi D_L t}} \exp\left[\frac{-(x^2 + y^2)}{4D_T t}\right] \exp\left[\frac{-(z - v_d t)^2}{4D_L t}\right] \tag{4.1}$$

where D_T (D_L) is the "transverse (longitudinal) diffusion coefficient", describing diffusion in the directions perpendicular (parallel) to the electric field. Equation 4.1 describes a Gaussian charge distribution that diffuses while drifting in the $+Z$-direction with velocity v_d due to an applied electric field as shown in Figure 3.4. Different

conventions for defining D_T exist in the literature; here we use the same convention as in previous measurements of D_T for liquid noble gases [88] where $\sigma^2 = 2Dt$ in each degree of freedom. Because each degree of freedom can be treated independently, the transverse diffusion process can be modeled as a random walk along the X-axis and Y-axis where at each time step dt a random step is sampled from independent Gaussian distributions with variance $\sigma^2 = 2D_T dt$. Considering only the distribution in the 2-dimensional plane transverse to the electric field for an initial point like distribution at $t = 0$ gives a radial variance at time t:

$$\langle R(t)^2 \rangle = \langle x(t)^2 \rangle + \langle y(t)^2 \rangle = 4D_T t \tag{4.2}$$

4.4 Measurement of Drift Velocity

In order to measure the diffusion constant, the transport of electrons in LXe must be accurately modeled. The simulation described in Section 3.3 requires a realistic electron drift velocity to produce signals with the correct shape and timing structure. Previous works have reported the drift velocity measured by EXO-200 at the standard operating field of 380 V/cm [64]. Here, the velocity is measured over the larger range of electric fields described in Section 4.2.

The drift velocity is determined from the drift time for events originating from a known location in the detector. To perform the measurement in the region of the detector between the cathode and the V-wire plane (i.e., the "bulk velocity"), events that drift the full length of the TPC are used. The drift time is converted to an average velocity using the separation between the cathode and V-wire grid of 192.4 mm. Two separate analysis techniques were used to check for consistency.

4.4.1 Analysis with α Particles

The first analysis measures the maximum drift time for α events emitted from surface contamination on the cathode and V-wire plane. Due to the short range of an α in LXe, the separation between the initial charge distribution and the originating

surface is small relative to the uncertainty in the reconstructed Z-position [64]. Energy deposits from α events have higher charge density relative to γs and βs causing substantially more scintillation from recombination relative to ionization. This allows α events to be selected using a cut on the ratio of scintillation light to charge yield, shown in Figure 3.17, which is tuned using peaks from known α backgrounds [89]. To obtain sufficient statistics for this measurement requires several days of "low background" data, for which a calibration source is not present. Approximately $\sim$100k α events were observed in the both the Phase-I and Phase-II low background data-sets. Low background data were not taken at any of the other fields considered here.

As shown in Figure 4.2, the distribution of drift times from αs has a large concentration of events arising from surface contamination on both the cathode and V-wire planes. The distribution of drift times from both surfaces are separately fit to Gaussians to determine the mean drift time for events originating on each. The electric field is less uniform between the U- and V-wire planes than in the bulk, as shown in Figure 3.4. For this reason it is not possible to define a single drift velocity in this region, but the mean drift time between the U- and V-wire planes can still be determined and is subtracted in the measurement of the bulk drift velocity.

The dominant systematic with this technique is due to the measurement of the distance between the cathode and the V-wire plane. This separation is determined using the known geometry at room temperature as well as the expected thermal contraction of the TPC materials at LXe temperatures. After cooling the TPC, there is some uncertainty on the resulting drift length and the position of the cathode relative to the wire planes. As shown in the inset of Figure 4.2, a $\sim$1 µs difference in the drift time is seen between the drift spaces of TPC1 and TPC2 for α events. This difference is consistent with an offset of the cathode position from center by $\sim$0.5 mm. In addition, the offset between TPCs is similar in Phase-I and Phase-II showing no significant change after refilling and temperature cycling the detector between phases. This potential offset is taken as an additional systematic error on the total distance between the V-wire plane and cathode in each TPC.

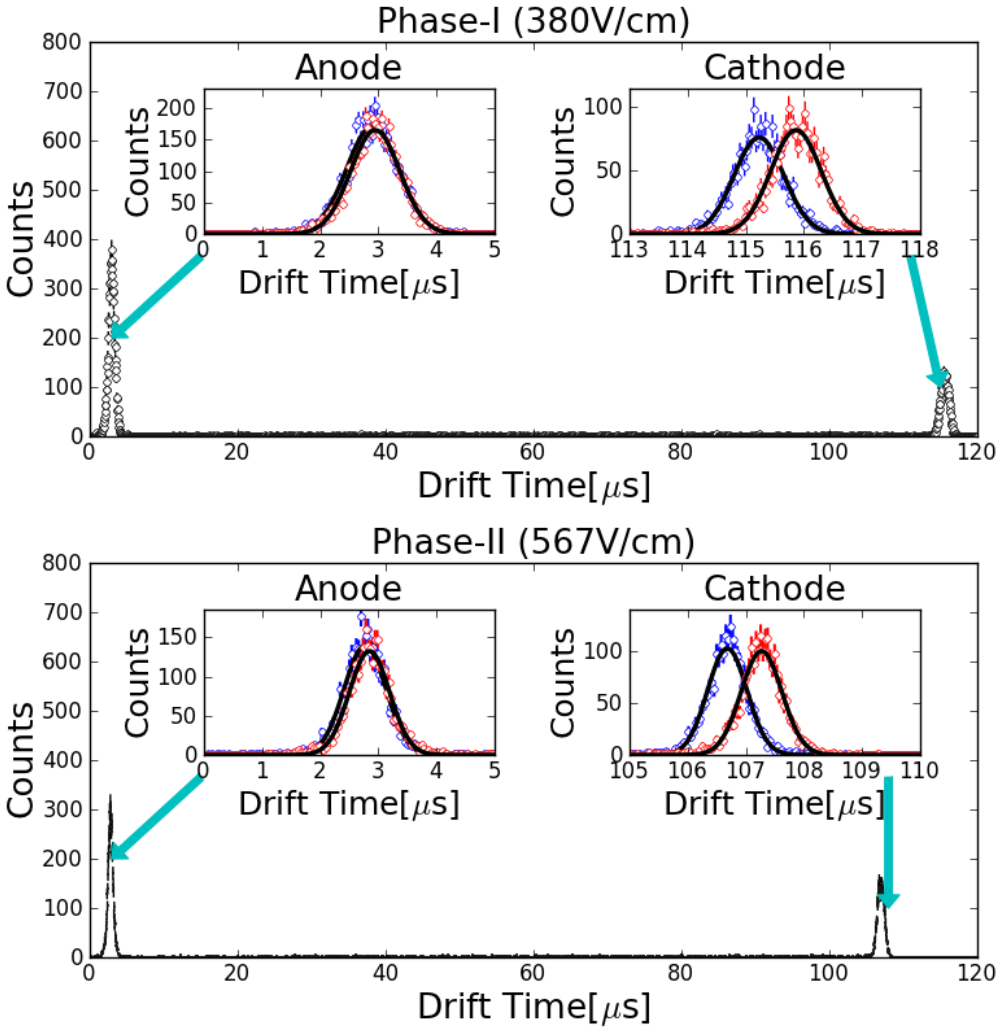

Figure 4.2: Combined distribution of drift times for α events from both drift spaces. The inset shows events near the cathode separately for TPC1 (blue) and TPC2 (red). Also included are the best Gaussian fits used for determining the maximum drift time in both TPCs. The offset between the TPCs corresponds to a $\sim$0.5 mm uncertainty in the cathode position as explained in the text.

4.4.2 Analysis with γ's

In addition to using α events from surface contamination, γ-ray source calibration data can be used to measure the drift time. Data taken with a ^{228}Th source at the cathode was used in order to maximize the number of events traveling the full length of the detector. Due to the finite size of charge deposits and position-dependent variations in the drift length, the distribution of drift times has a finite width edge at its maximum, as shown in Fig. 4.3. A simulation is used to calculate the position along this edge that corresponds to the average drift time for events traveling the full length of the detector in MC.

Following the same procedure as for the measurement using α events, the portion of the total drift time corresponding to the transit time between the V- and U-wires must be subtracted to obtain the drift time in the main drift region alone. This V-to-U drift time is directly measured using the α events at the two fields for which low background data were acquired (380 V/cm and 567 V/cm), as described in Section 4.4.1. For the remaining fields, this drift time is estimated using the known spacing between the V- and U-wires and the measured value of the drift velocity at the mean field in this region. Since the field in the U-V gap was chosen to be twice the field in the main drift region for each dataset, and since the field was also varied by a factor of two between different data sets, the measurement of the drift velocity in the bulk region for each higher field point can be used to estimate the drift velocity in the collection region at the point below it. While this approximation neglects the non-uniformity of the electric field between the U- and V- planes, the correction itself is small, and the resulting error gives a negligible contribution to the total error.

From Figure 4.3 it can be seen that drift time distribution in data are not fully reproduced by MC. In particular, the tail of the distribution in data are slightly broader than in the simulated distribution. This difference may arise from the approximate modeling of the electric field geometry near the TPC surfaces in simulation. To ensure that any differences in the detailed modeling of the edge of the distribution are accounted for, the width of the edge in the data distribution is added as a systematic error in the measurement of the maximum drift time.

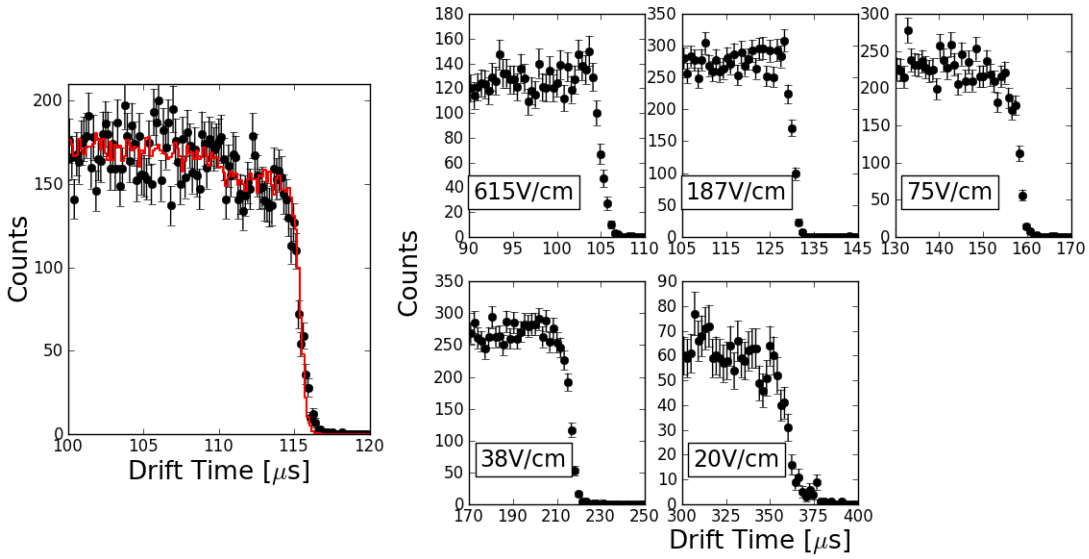

Figure 4.3: Distribution of drift times for events with the longest drift times for MC (red line) and data (black points) for data taken with a ^{228}Th source positioned near the cathode. Shown for both data taken at the nominal Phase-I field of 380 V/cm (left) and data taken at the alternate fields (right). From the comparison between data and MC for the Phase-I field, the observed edge width in data is broader then MC due to model a simplified model of the complicated cathode geometry. This width is taken as a systematic error in the measured velocity.

4.4.3 Summary

The measured drift velocity from α events and source data at each field are shown in Fig. 4.4. Both techniques agree within errors at the fields at which data are available for both (inset). Figure 4.4 also shows previous measurements of the drift velocity of electrons in LXe [90–94]. The EXO-200 measurements differ from previous measurements at the same fields by up to $\sim$10%. The difference between the EXO-200 measurements and the early measurements by Miller et al. [90] and Gushchin et al. [91], using substantially shorter drift distances, cannot be fully explained by the temperature differences between the measurements. Previous measurements have found the drift velocity to vary by $\sim$ 0.012 mm/μs K^{-1} at similar fields [95]. In the current measurement, data were taken at 167.0 $\pm$ 0.2 K, while the Miller et al. [90] and Gushchin et al. [91] measurements were obtained at lower temperatures (163 K and 165 K respectively). These temperature differences are consistent with slower drift velocity at higher temperatures. However, they would account for only a $\lesssim$ 3% variation between the measurements shown in Fig. 4.4. The more recent measurements from XENON10 [92], XENON100 [93] and LUX [94] are taken at 177 K, 182 K and 174 K respectively. These measurements take advantage of the $\gtrsim$10 cm drift lengths and higher chemical purity available with modern LXe TPCs and are in better agreement with the EXO-200 measurements presented here. The residual disagreement between EXO-200 and these more recent measurements is comparable to the expected deviation due to temperature differences.

In addition, early studies suggested that the purity of the LXe can affect the drift velocity of electrons in LXe [96, 97]. Although only high purity data were included for the measurement shown in Fig. 4.4, the drift velocity was also measured for data with lower purity. For this data electron lifetimes varied between 80 μs and 600 μs. The drift velocity for this data were measured using α events following the method outlined above and agreed within error to the measurement at purity >2 ms included in Fig. 4.4.

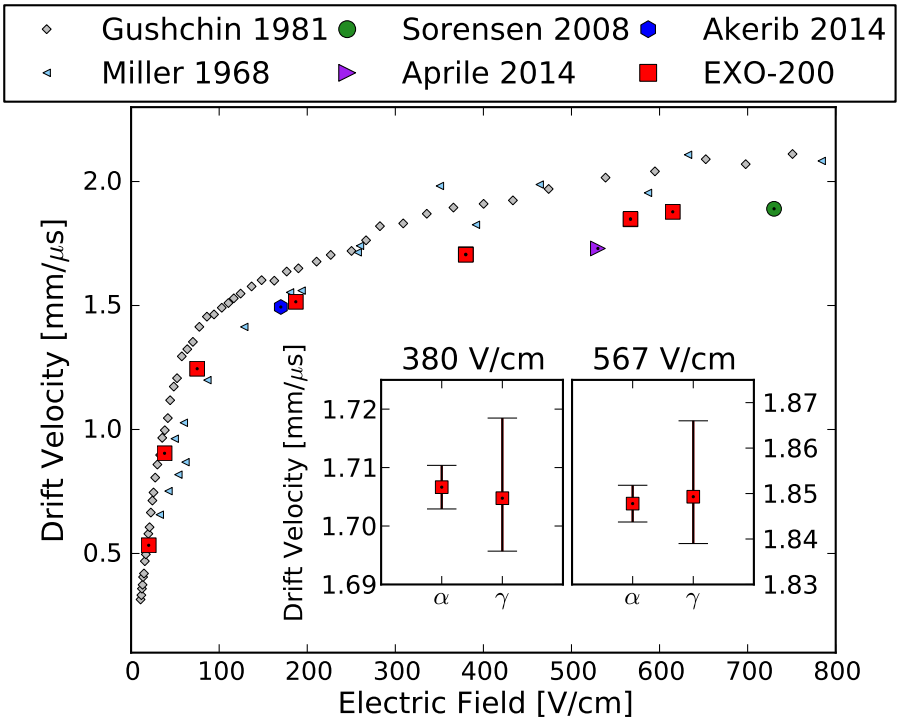

Figure 4.4: Drift velocity measured using EXO-200 for γs in ^{228}Th source data and αs in low background data (167 K, red). Also included are results from Miller et al. (163 K, cyan) (Miller 1968: [90]), Gushchin et al. (165 K, grey) (Gushchin 1981: [91]) and the XENON10 (177 K, green) (Sorensen 2008: [92]), XENON100 (182 K, purple) (Aprile 2014: [93]) and LUX (175 K, blue) (Akerib 2014: [94]) collaborations. LXe operating temperatures are included to account for possible variations due to temperature [95]. The inset shows a comparison of the two methods of measurement at 380 V/cm and 567 V/cm in EXO-200.

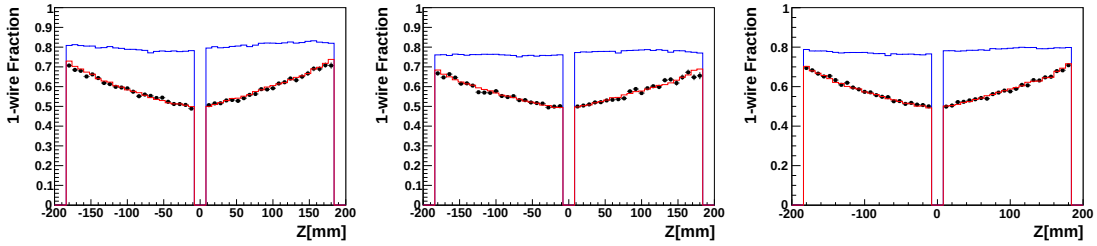

a) ^{228}Th **b)** ^{60}Co **c)** ^{226}Ra

Figure 4.5: EXO-200 comparison of 1-wire fraction between Data (black points) and MC without diffusion (blue line) and with diffusion (red line) for operation at 380 V/cm. The sources shown are (a) ^{228}Th, (b)^{60}Co and (c) ^{226}Ra. For MC with diffusion a coefficient at the best fit of 55 cm^2/s was used for all sources.

4.5 Measurement of Transverse Diffusion

At the standard operating field of 380 V/cm, data were taken using ^{226}Ra, ^{228}Th and ^{60}Co calibration sources, as described in Section 3.2.2. Runs were performed with the sources positioned near both the anode and cathode to account for any systematic errors that may arise from different position and energy distributions. Variations in the measured diffusion coefficient for different data sets were at most ∼10%. For data taken at operating fields other than 380 V/cm, only the ^{228}Th source was used. A minimum of ∼100k events were used for the measurement at each field.

MC was generated using a range of diffusion coefficients. For each simulated coefficient, the distribution of the fraction of clusters containing only 1 readout channel instead of 2 readout channels (i.e., the "1-wire fraction") for different Z-positions was determined and compared to that observed in data. A maximum likelihood fit is performed to determine the diffusion coefficient that gives the best agreement between data and MC. The distributions of 1- and 2-wire events for each source are binned into 6 Z-bins spanning the length of both TPCs. Only events inside the fiducial volume described in Section 4.2 are used to avoid field non-uniformity near the edges of the detector. Fits are performed separately for each source and source position, and results from the fit to each dataset agree within errors. The final measurement at each field is then performed using a combined fit to all sources.

Figure 4.5 shows the Z-dependence of the 1-wire fraction for data and MC for each of the three sources at the standard field. This includes MC with a diffusion coefficient near the best fit value of $D_T = 55 \pm 4$ cm^2/s as well as that generated without diffusion. A linear dependence in the 1-wire fraction versus Z-position can be seen in source calibration data, consistent with increasing event size at longer drift lengths. For the MC without diffusion, the distribution does not depend on drift length. As shown in Fig. 4.5, incorporating diffusion into the MC results in good agreement in the distribution of the 1-wire fraction with data.

The simulation of the EXO-200 detector response has been studied in detail to optimize its agreement with data, but some differences remain. Inaccuracies in the simulation of 1- and 2-wire event distributions are estimated by measuring the differences in these distributions between data and MC for events within 10 mm of the V-wire plane, where the effect of diffusion is negligible. Differences in the 1-wire fraction for each Z-bin between 5-20% was observed at the fields considered here. The largest of these errors was associated with the fields farthest from the standard operating field, where the MC is less studied.

To account for the error associated with systematic difference between the data and MC, an additional scaling of the 1-wire and 2-wire event distributions in each bin is allowed to vary. These scalings are constrained in the likelihood fit by Gaussian errors determined from the differences between data and MC for the short drift time events. An additional independent scaling of the overall ratio between 1-wire and 2-wire events was also included and constrained using the same technique. These scalings and constraints allow the effect of any errors in the MC simulation of the expected wire distributions to be propagated as systematic errors to the measurement of the diffusion constant. No correlation between the constraints were observed so the constraints were separately fit for each data set and the likelihood was summed over all data sets to determine the best fit diffusion constant and error at each field.

In addition, there is an error associated with inaccurately reproducing the signal amplitude of charge channels in MC. At the standard field of 380 V/cm, the relative signal amplitude of each channel is empirically scaled to match the amplitude observed in data to within 10%. At lower electric fields, the signal amplitude is reduced due

to the lower charge extraction efficiency. This variation in charge yield with electric field has been well studied in previous experiments [67, 68] and more recently in EXO-200 [98]. The relative summed signal amplitude is measured for events in the ^{228}Th peak for each field and a correction to the expected charge yield in simulation is applied by scaling all charge signals by the observed factor. This does not fully account for the detailed microphysics of this process, and future work could use simulation packages such as NEST [76, 99] to account for the varying charge yield with field on an event-by-event basis. Since only an empirical scaling is used here, an additional systematic error is estimated by performing the fit to the diffusion constant for both the scaled and un-scaled U-wire signals. The shift in the central value of the fit is then included as an additional systematic error to conservatively account for any error in the simulated charge yield with field. The resulting error from the signal magnitude systematic is 5% at the standard field compared to the 8% total error. For the measurements at lower fields, this error increases due to the larger variations in charge yield. For the lowest field, this systematic results in a 46% error compared to the 48% total error.

The effect of self-repulsion of charge clouds is ignored in the full simulation of charge propagation and diffusion described in Section 3.3. This effect was studied separately to determine its impact on the measurement and is qualitatively different from diffusion since it depends on the density of the charge distribution. This introduces a dependence on the event energy and an additional non-linear z-dependence in the size of the charge distribution versus drift time. However, the variation with cluster energy for the diffusion coefficients measured here is less than 10% over the full energy range and is consistent with errors. This indicates that the data are not sensitive enough to show the effect of self-repulsion of charge deposits.

As a cross-check, a standalone simulation including the effects of both self-repulsion and diffusion was performed, using a similar technique as in [88]. This simulation tracked the drifting electrons while accounting for both diffusion as well as repulsion from each electron to calculate the expected spread of the initial deposit. The screening of polarized LXe ions was ignored to give a conservative estimate of the error from the repulsion process. Results of this study confirm that repulsion is a negligible

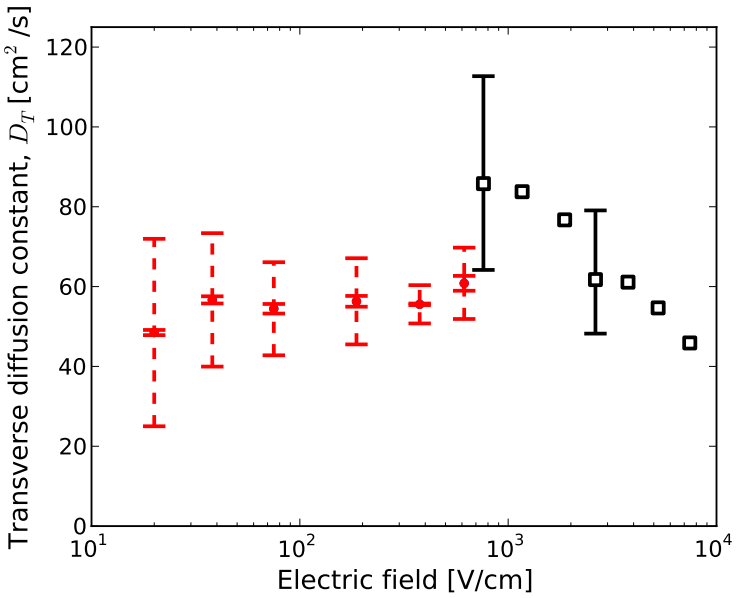

Figure 4.6: Measured diffusion coefficient versus electric field in EXO-200 (red points). Solid error bars are statistical and dashed errors include the systematic errors discussed in the text. Also plotted are previous measurements made by T. Doke and collaborators (black points) [67, 69, 85, 100]. Error bars for the previous measurements from [69] are shown, where errors were provided only for two of the seven points.

effect, contributing at most 3% to the overall spread of the charge distribution.

As described in Section 4.2, the fiducial volume at each field was defined to minimize systematic effects from field non-uniformities. To check for any residual systematic effects, the diffusion measurement was compared for different radial segments in the detector. At the standard operating field of 380 V/cm, measurements using high- and low-radius segments of the detector were consistent, and these measurements constrained any residual variation between the segments to $< 2.6\%$, compared to the 8% total error. For the lowest field at 20 V/cm, where the field non-uniformities were greatest, the maximum variation was $< 5.5\%$ compared to the 48% total error. This shows the resulting systematic effect of residual non-uniform fields in the fiducial volume is small relative to the systematic error from other sources.

Figure 4.6 shows the measured diffusion coefficient from EXO-200 data from the fields used, as well as results of a previous measurement [67, 69, 85, 100]. The error

bars for the previous measurement are taken from [69], where errors in the measurement of eD_T/μ are shown by the authors for only two of the seven points. The diffusion coefficient is calculated by the authors using the transverse spread of the electron clouds and the drift velocity measured in [69], assuming the relationship to the mobility $v_d = \mu E_d$, where μ is the electron mobility, v_d is the electron drift velocity, and E_d is the electric field [101]. The resulting measurements of the diffusion coefficient shown in [67, 69, 85, 100] do not include the error bars on these two points from the original reference, but they are propagated to the measurement of D_T shown in Fig. 4.6 here for comparison. Although there is no direct overlap in the fields considered, the results of the current study are consistent within errors near 600 V/cm, but suggest a 25% lower central value than the previous measurement. These results use a maximum drift length of 198 mm, which is substantially larger than the 2.7 mm drift length used in the previous measurement [69]. This longer drift length limits the effect of systematic errors on the measurement of the diffusion coefficient. Finally, the detector MC and event reconstruction available for EXO-200 data allows this analysis to account for systematic effects in the measurement of the charge distribution versus drift length in detail.

4.6 Conclusions

A measurement of the drift velocity, v_d, and the transverse diffusion coefficient, D_T for electrons in drift fields between 20 V/cm $< E_d <$ 615 V/cm has been performed using data from the EXO-200 detector. These results contribute to the understanding of the transport properties of electrons in LXe. Ton-scale LXe particle detectors, which are currently in development for the search for $0\nu\beta\beta$ and dark matter, will be more sensitive to effects such as diffusion because of their longer drift lengths, making this measurement relevant for the development of future LXe technology.

Chapter 5

Machine Learning Based Discriminator

5.1 Overview

As described in Section 3.5, γ backgrounds and $\beta/\beta\beta$ signals can be distinguished by their topological signature due to differences in physical processes governing the interactions. Through the use of multiplicity, EXO-200 is able to efficiently separate a majority of the γ backgrounds with minimal loss to the $0\nu\beta\beta$ signal, as shown in Figure 3.19 for a simulation of the expected signal and background.

A perfect topological discriminator would identify all γ's which Compton scatter in the detector. While the multiplicity definition described in Chapter 3 is able to identify events which have clearly separated deposits within the detector, it is not optimized for the more subtle cases such as distinguishing bremsstrahlung or resolving overlapping Compaton scatters. These deficiencies arise from limitations related to detector parameters such as the wire pitch of 9 mm and the digitization time of 1 us, which are built into the multiplicity definition. This is reflected in the efficiency curves shown in the bottom of Figure 3.19. Here the efficiency for identifying backgrounds begins to noticeably drop at roughly twice the U-Wire pitch, which follows directly from the requirement that SS events include no more than two U-Wire collection signals. As a result, the event size distributions shown for signal and background in Figure 5.1 still show clear separation even after the multiplicity classification. In the case of SS events, the remaining set of γ's with sizes larger than $\sim$6 mm can arise from multiple scatters which are too close to be separated by reconstruction. Similarly, some $0\nu\beta\beta$ events are classified as MS due to bremsstrahlung radiation which can produce multiple distinct scatters in the detector.

As shown in Figure 3.19, a clear separation still remains between signal and background for both SS and MS events, implying that further discrimination is still possible for both classes of events. The challenge is then finding features of the detector response which can be used to extract topological information not captured by multiplicity.

5.2 Boosted Decision Tree

As discussed in Chapter 3 the observed response on the wire channels is the joint response induced by each individual electron. Because of this it is expected that information about the extent of electron clouds should be encoded into the shape, size and distribution of signals on the wire planes. To improve discrimination it is then possible to extract relevant observables from the wire signals which correlate to the topology of the deposits of electrons. Extensive studies were done to identify potential observables for both SS and MS events as will be described in the following

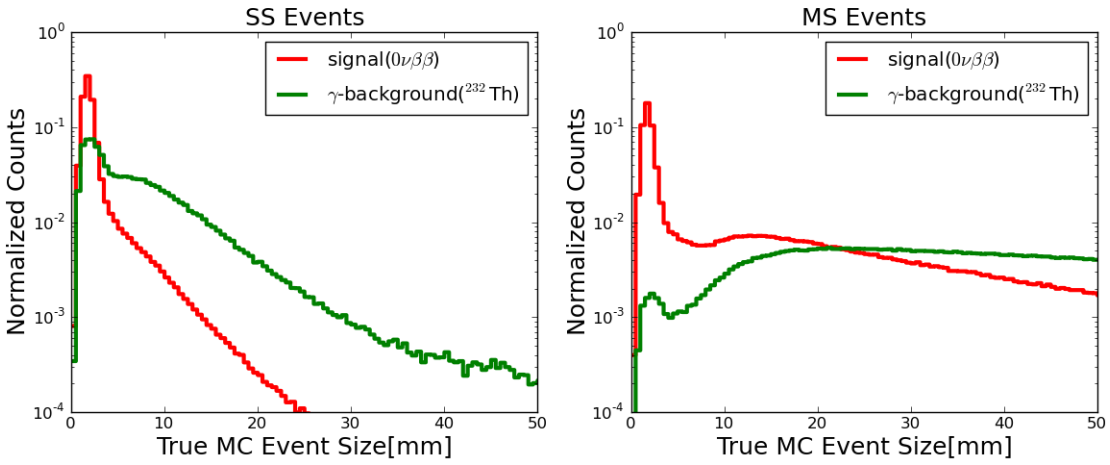

Figure 5.1: Distribution of the true spatial extent of events in MC for both SS(left) and MS(right) events. A clear separation still remains between signal and background, implying that further discrimination is still possible given additional information.

sections.

In order to take advantage of the additional discrimination offered by these finer details, a Boosted Decision Tree (BDT) was developed using the TMVA software package [102]. This allowed for the combination of the identified observables into a final joint discriminator which could be used in the final $0\nu\beta\beta$ fit. The BDT method is a supervised learning algorithm which combines a set of individual decision trees (DT), each of which consist of a series of binary cuts on the input variables to classify an event as either signal(+1) or background(-1). Boosting is then performed so that the set of DTs is created such that events which were incorrectly classified in earlier DTs are weighted more strongly in later stages. This ensures that the later DTs are optimized to correctly classify the previously missed cases. The final output is then the average response of all DTs, weighted by their classification performance. This technique was first used by EXO-200 to study the $2\nu\beta\beta$ decay of ^{136}Xe to an excited state [103]. Further details about the BDT technique can be found in [102, 104].

5.2.1 SS Discriminator

In the case of SS events, further discrimination is possible if one can separate true single scatter events from events consisting of multiple deposits which are too close for reconstruction to disentangle and are mis-identified as single scatters. This population of events can be observed in Figure 3.19 in the region less than ~ 20 mm but before the peak at ~ 3 mm where the efficiency of correctly identifying γ's as MS decreases. In order to identify these events, observables related to spatial extent of individual deposits were studied.

The first observable studied was the number of U-Wire channels which observed signal, which correlates to the extent in the 2D plane perpendicular to the drift axis (transverse plane). Due to the multiplicity definition which allows for SS events to extend over at most two U-Wire channels, mis-identification can occur when events fall within roughly twice the wire pitch ($\sim$18 mm) of each other in the transverse plane. In order to improve upon the multiplicity definition, the number of U-Wire channels which detect collected electrons can be used to further discriminate signal and background. Because the majority of $0\nu\beta\beta$ events are expected to have extents less than the 9 mm wire pitch, $0\nu\beta\beta$ events will tend to collect on a single U-Wire. On the other hand, γ backgrounds resulting in multiple spatially unresolved deposits will tend to be larger and produce more two U-wire events than expected for $0\nu\beta\beta$. This is shown in Figure 5.2 which compares the distribution for the number of U-Wire channels for SS events in MC between $0\nu\beta\beta$ and the ^{232}Th background component.

In addition, scatters which are separated along the the drift axis can be mis-identified if the scatters fall within $\sim$7 mm. This is due to the 3.5 µs matching time used to identify U-Wire signals caused by the same energy deposit. Identification of events which have large longitudinal extents is possible by using the rise time of the U-Wire signals (defined as the time between the 5% and 95% points in the sum U-Wire signal for each cluster). Because the time required to collect each electron depends on the electron's starting position, the rise time is correlated to the time it takes for the electrons in a deposit to fully collect. Background-like events, which tend to extend further along the drift axis, will then take longer to collect. As a result these events tend to have on average longer rise times as shown in Figure 5.2.

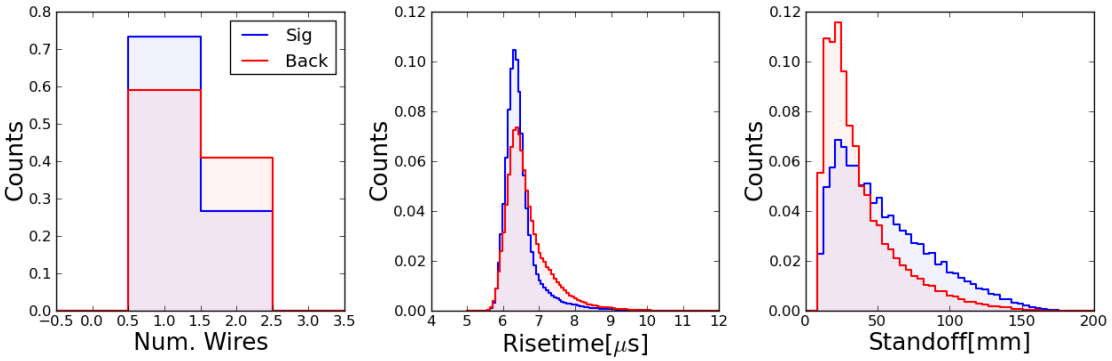

Figure 5.2: Comparison between the simulated SS event distribution for signal($0\nu\beta\beta$) and background(prominent γ's) for each of the observables studied to improve the discrimination between signal and background in the SS channel. Description of the individual variables can be found in the text.

To make use of these additional observables, a BDT algorithm was developed to discriminate between SS $0\nu\beta\beta$ events and γ backgrounds for the $0\nu\beta\beta$ search [66]. The final BDT was created by combining the topological variables above along with the position based discriminator standoff distance which was described in Chapter 3. This allowed for a single discriminator to be produced which captured the discrimination power offered by both position and topology, simplifying the implementation into the final analysis framework described in Chapter 6. While other inputs were studied, this set of obsevables was chosen in order to optimize the discrimination power while maintaining good agreement ($< 20\%$ residuals) between simulation and data.

The BDT algorithm was trained to maximize the separation between SS $0\nu\beta\beta$ events and the most prominent γ backgrounds arising from ^{238}U, ^{232}Th and ^{60}Co. Training was done using a subset of 1M simulated events for both signal and background and its performance was validated on a statistically independent test set to check for over training. Further validation was done by checking the agreement between data and MC for both the resulting discriminator and each of the inputs. This comparison is shown in Figure 5.3 which compares the predicted and observed distributions for both calibration data with a ^{226}Ra source and $2\nu\beta\beta$ distribution determined by subtracting the backgrounds from the low background data. Agreement for all inputs and the BDT discriminator itself is observed to better than 10% in

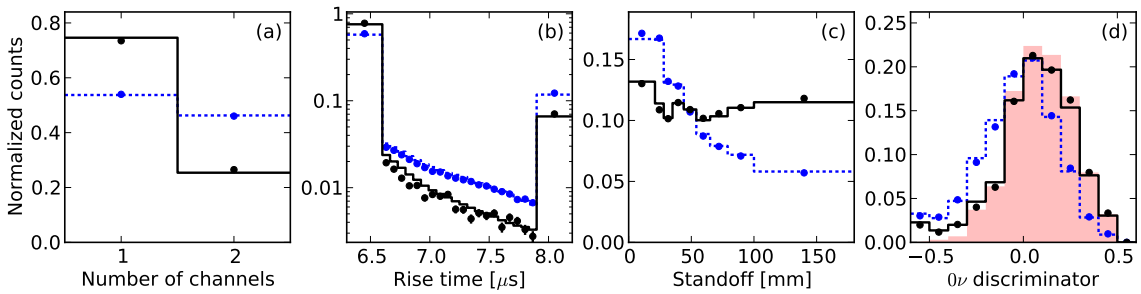

Figure 5.3: Comparison of the SS distributions of the input observables as well as the resulting $0\nu\beta\beta$ discriminator between data(dots) and MC(lines). Shown for both a γ calibration using ^{226}Ra (blue) and background subtracted $2\nu\beta\beta$ (black). Additionally the predicted distribution of the discriminator is shown for $0\nu\beta\beta$ as the red shaded region. All distributions are normalized by area and the edge bins represent overflow.

all bins. The binning shown in Figure 5.3 is chosen to optimize the discrimination power while simultaneously minimizing the systematic uncertainties introduced from a mismatch between data and MC.

A common tool used to visualize and quantify the performance of a binary classifier, such as a BDT, is the Receiver Operating Characteristic (ROC) curve [105]. Calculating the ROC curve is done by scanning over all possible binary cuts on a chosen observable and determining the efficiency of correctly identifying signal as a function of the efficiency to reject backgrounds for each cut. The ROC curve for the individual inputs as well as the resulting BDT based $0\nu\beta\beta$ discriminator are shown in Figure 5.4. Also shown as reference is a dashed line corresponding to the expected curve for a discriminator which has no classification power. Comparing the curves it can be seen that the BDT achieves higher signal efficiency for all background rejections, indicating that it is correctly combining the individual inputs into a more powerful discrimination.

The complete shape of the curve is important for understanding the detailed performance, as the priority weighting of signal efficiency and background rejection may vary with the situation. In order to compare the overall performance of different discriminators, the overall performance can be approximately captured by calculating the area between the entire ROC curve and the dashed line. While the area does not show whether individual points on an ROC curve are superior, this shows the average

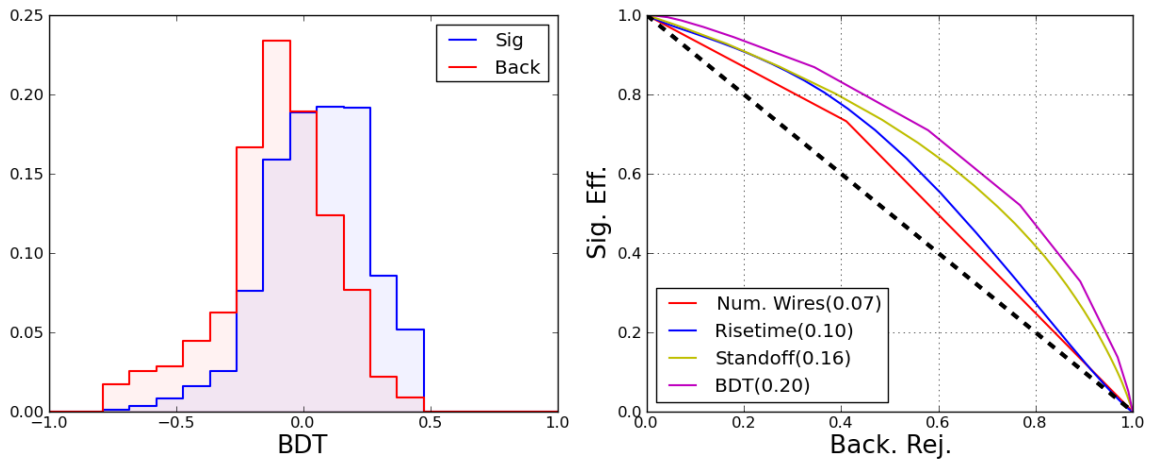

Figure 5.4: (left) Simulated distribution of the BDT discriminator for the $0\nu\beta\beta$ signal and the prominent γ backgrounds. (right) A comparison between the ROC curve of the BDT to that of each of each input observable, with the area between each curve and the dashed line shown in the legend.

performance achieved over all possible cuts with curves covering more area having on average more discrimination. The area covered by the BDT curve as well the input observables is shown in the legend of Figure 5.4 from which it can further be seen that the BDT is superior to each cut individually covering ∼20% more area than the next best discriminator.

While the ROC curves are useful in visualizing the discrimination power, the real metric of interest is the search sensitivity to $T_{1/2}^{0\nu\beta\beta}$. As will be discussed in Chapter 6, the BDT can be added as an additional dimension into the final fit used to determine if a positive $0\nu\beta\beta$ signal is present. Comparing the fits using the BDT to those using just SS/MS classification alone, a ∼15% increase in sensitivity to $T_{1/2}^{0\nu\beta\beta}$ was observed.

5.2.2 MS Discriminator

Further discrimination is also possible for MS events, as there as about 10% of the $0\nu\beta\beta$ events which get classified as MS. This can occur either from bremsstrahlung radiation emitted from the β's or from diffusion causing a single $0\nu\beta\beta$ deposit to grow in size such that it spans more than two U-Wire channels. While these events fall under

the MS definition decribed in Chapter 3, their physical signature looks different than MS events caused by γ's which can Compton scatter multiple times in the detector. This can in part be observed in Figure 5.1 which compares the distribution of the spatial extent for MS events in simulation from both $0\nu\beta\beta$ signal and γ backgrounds. A clear separation between the two distributions remains, indicating the potential to recover the $0\nu\beta\beta$ events classified as MS by leveraging the finer details of the detector response.

Using a similar approach to that for SS events, the first two observables investigated were the number of U-Wire collection channels as well as the rise time. In this case it should be noted that, because each event can be composed of more than one reconstructed deposit, the rise time is defined as the largest rise time among all clusters while the number of wires is defined as the total number of collection wires summed over all clusters. The simulated distribution for both variables is shown in Figure 5.5. While the number of wires continues to show clear discrimination power there is little separation observed for the rise time distributions.

The apparent lack of discriminating power for rise time is due to the fact that these two variables are optimized for estimating the size of a single cluster but don't fully capture information about the relative topology of the set of deposits. For MS events the total spatial extent of the entire event is important for signal and background discrimination, as bremsstrahlung typically produces smaller secondary clusters which are relatively close to the main deposit. This is in contrast to γ's which can scatter multiple times producing many disperse deposits. The rise time observable then presents little separation because it is optimized to constrain just the longitudinal size of an individual cluster. On the other hand, while the number of wires was also optimized for single clusters, it continues to remain a powerful discriminator as it is summed over all clusters. This preserves information about the total transverse size of the entire event. This still results in some lost information about the total spread in the detector, as this definition ignores the relative spacing between individual U-Wires. Thus additional variables were also explored for MS events.

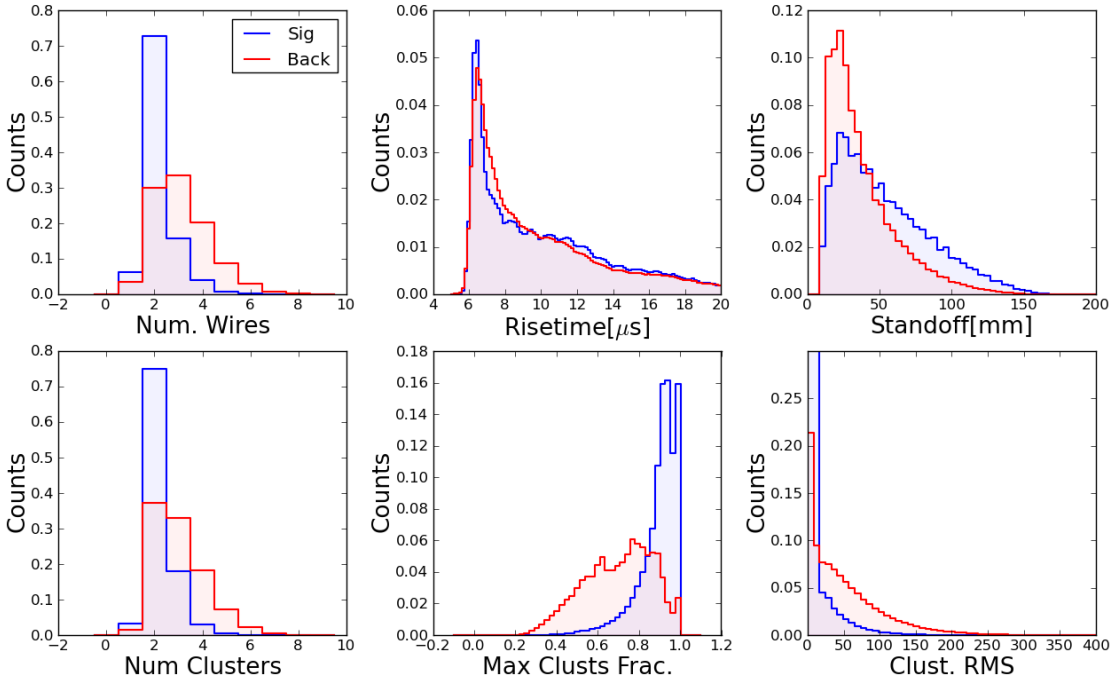

Figure 5.5: Comparison between the simulated MS event distribution for signal($0\nu\beta\beta$) and background(prominent γ's) for each of the observables studied to improve the discrimination between signal and background in the MS channel. Description of the individual variables can be found in the text.

In order to break the degeneracy between the total number of channels and the relative spatial separation between the U-Wire signals, the total number of reconstructed clusters was used. Because of the low probability of bremsstrahlung resulting in more than one additional distinctly separated deposits, the MS $0\nu\beta\beta$ distribution tends to favor fewer reconstructed clusters. This is can be seen in Figure 5.5, where the majority of MS $0\nu\beta\beta$ events contain just two clusters, one from the main electron deposit and one additional deposit from bremsstrahlung. While additional deposits are possible, the probability decreases sharply. This is in contrast to γ's, which result in more clusters on average due to Compton scattering. Further breaking of the degeneracy is then possible by not only counting the number of clusters but by determining the RMS spread of the reconstructed cluster positions. This gives an additional handle in determining how disperse a given event is, allowing for identification of signal-like events with secondary deposits relatively close to the main deposit, as depicted in Figure 5.5.

Further information about the relative distribution of energy deposits is then found in the fraction of energy contained in the largest cluster. Because bremsstrahlung results in low energy radiation, the secondary deposits are expected to contain only small fractions of the total event energy. Most of the energy for a $0\nu\beta\beta$ like event is then found in the large deposit created by the main β track. This can be seen in Figure 5.5 which compares the simulated signal and background distributions for this maximum cluster energy fraction.

Following the same procedure described for SS events, a BDT based discriminator was built to optimally combine the set of MS observables. The resulting distribution for the MS BDT discriminator and the correspondng ROC curves are shown in Figure 5.6. The BDT again shows better overall discrimination power than each of the input variables, although in this case the fraction of energy contained in the largest cluster appears to dominate the MS discriminator. While the BDT for MS events showed improved discrimination it was ultimately not used in the final $0\nu\beta\beta$ search. Instead, as discussed in Section 5.3, an alternative method was explored to produce a more powerful joint discriminator for both SS and MS events using lower level information contained in the U-Wire waveforms.

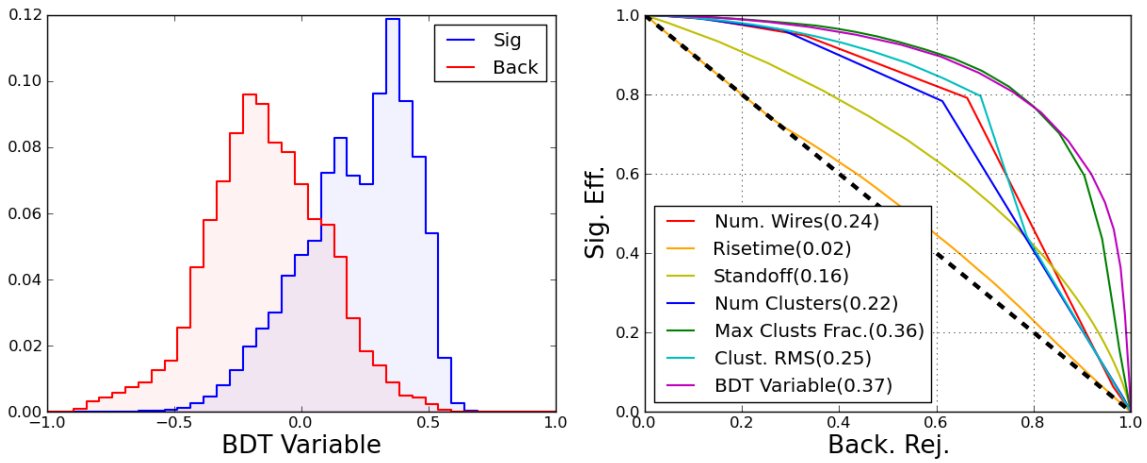

Figure 5.6: A comparison of the resulting ROC curves for the potential MS observables studied to further discriminate between signal and background in the MS channel.

5.3 Neural Network

As shown in Section 5.2 additional discrimination is possible through the use of features of the U-Wire signals not captured by the standard SS/MS definition. The BDT developed for SS events offered a powerful tool to combine these additional features into a joint discriminator which helped improve the $0\nu\beta\beta$ search sensitivity by $\sim$15%. The drawback of this method is that it requires the ability to determine and extract the relevant features from the waveforms. While the reconstruction algorithm is able to estimate many such parameters, there are many simplifying assumptions which may cause complex features to be lost during reconstruction and clustering. As a result the BDT algorithm may still miss the full topological information contained in the wire signals.

Motivated by improvements in the EXO-200 energy reconstruction using Deep Neural Networks (DNN) [106], an alternative method was explored to use lower level information in order to optimize the $0\nu\beta\beta$ discrimination. This method uses a DNN based discriminator that takes images of the U-wire signals as input and outputs how much a given events resembles signal($\beta\beta$) or background(γ). The design of the DNN architecture is inspired by the Inception architecture proposed by Google [107] and

implemented with the Keras library [108] using the Tensorflow backend [109].

Similar to the BDT, the goal of this new discriminator is to capture topological information, which can improve discrimination over the traditional multiplicity, in addition to limiting bias introduced from inaccuracies in the MC. A main difference in the design of the DNN algorithm is to ignore position dependent discrimination. This is in part to force the DNN to focus on topology but is also motivated by the lack of the V-Wires, which would be required for the DNN to fully capture the position of an event. Including V-Wires in the image is currently not possible due to modeling inaccuracies, which result in poor agreement between data and MC. These inaccuracies translate into increased bias in the final discriminator, which outweighs the additional discrimination offered by the extra information. In addition, studies have indicated that the additional topological discrimination offered by the V-Wires is small. Because it is not possible to disentangle the features of the waveforms which are related to position and topology, the new variable can be combined with energy and standoff in a multi-dimensional fit to maximize the search sensitivity, as will be described in Section 6.

5.3.1 Input Image

The input to the DNN algorithm are images of the U-Wire waveforms, built by stacking all neighboring channels in each TPC. To normalize the response of each channel the baseline of each waveform is measured and subtracted. Additionally, each waveform is corrected for its channels measured gain. Each waveform saved by the detector triggering is 2048 µs long but the maximum drift time within the detector is only ∼100 µs. The additional samples are used in reconstruction to measure the baseline and look for coincident events but are not necessary for use in topological discrimination. In order to decrease the image size and save on computational time, the image generated for the DNN uses a smaller 350 µs window around the expected signal location. In order to preserve the relative z-Position this window is defined with reference to the scintillation time, with 24 us coming before and 326 us after.

For this study two methods of generating the U-Wire images were studied. The

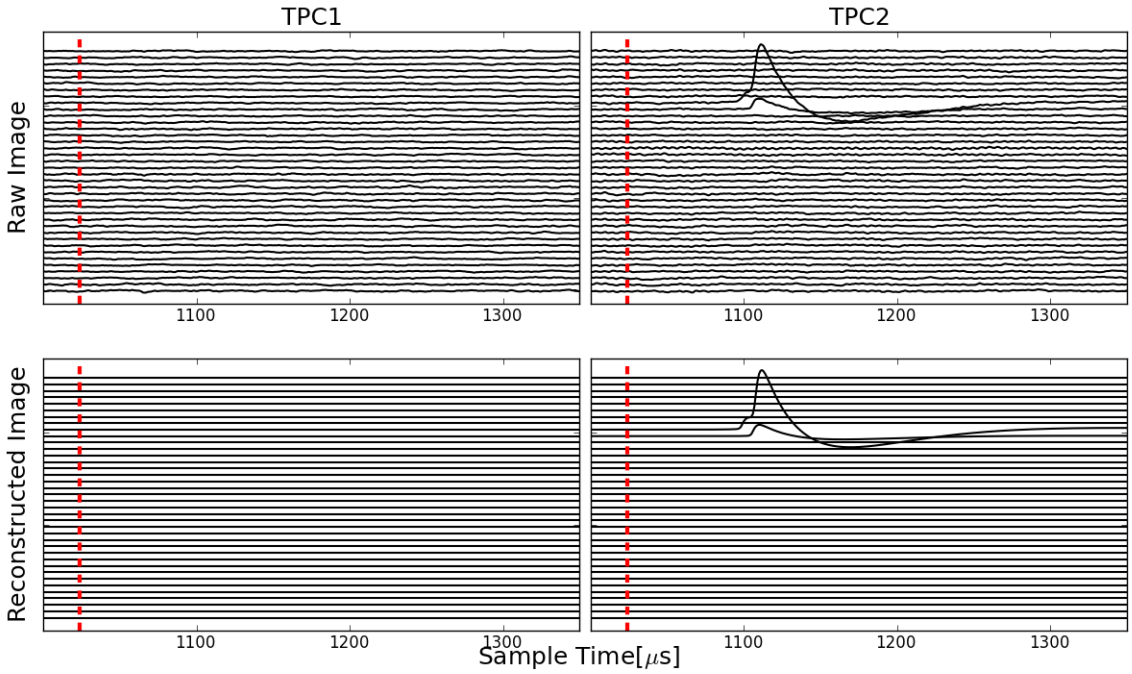

Figure 5.7: Example of a (top) raw and (bottom) reconstructed image for the same event in data. The vertical red dashed line shows the time of the scintillation signal which is used as a reference in defining the window.

initial method used the minimally processed "raw" waveform from each channel which are saved when the detector triggers. While the raw image preserves maximal dimensionality, it is more computationally intensive. In addition the raw image is potentially more sensitive to inaccuracies in the modeling of the detector response. To alleviate this, an alternative method was developed to instead regenerate the waveforms using the fit results of the reconstruction algorithm described in Section 3.4. Unlike the raw image the image built from the fits relies on parts of the reconstruction algorithm but it avoids the simplifying assumptions made in the clustering stage. Because of this the reconstructed image represents an intermediate step, with less complexity than the raw image but more than the BDT inputs. In addition, it also provides more control of the features available to the DNN potentially allowing for better agreement between observation and predication. Example images for both methods are shown in Figure 5.7 for an event taken during low background data taking.

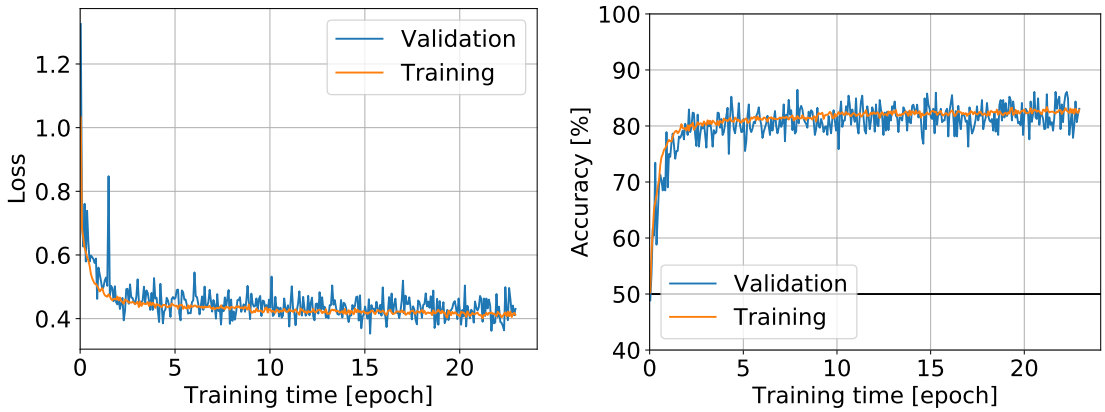

Figure 5.8: The loss (left) and accuracy (right) of the DNN training for the raw DNN in Phase-II for both the training and validation sets. Agreement between the validation and training set indicate the algorithm has not been over-trained. Taken from [110].

5.3.2 Training

Training of the discriminator is done by generating a MC data set for the expected signal and background using the method described in Section 3.3. Signal like events are generated by simulating $0\nu\beta\beta$ uniformly in the LXe volume but with a energy which is randomly sampled from a uniform distribution between 1 MeV and 3 MeV. Background-like events are generated from a simulation of a uniformly distributed γ source in the same energy range. A uniform energy and position distribution is used in both data sets to maximize the importance of topology. This ensures that the DNN algorithm prioritizes features related to topology and does not rely on those correlated to position and energy, which are captured elsewhere in the analysis.

A minimal number of cuts are applied to maximize the exposure of the DNN to different image types. In order to remove events with an ambiguous time window, a cut is applied to remove all events which do not have exactly one scintillation cluster. These events are also removed from the final analysis so the DNN does not need to handle them. Additionally events are required to have at least one found charge cluster and have a total reconstructed energy above 1 MeV to match the analysis threshold used in the final fit. Finally to match the Fiducial volumne in the final

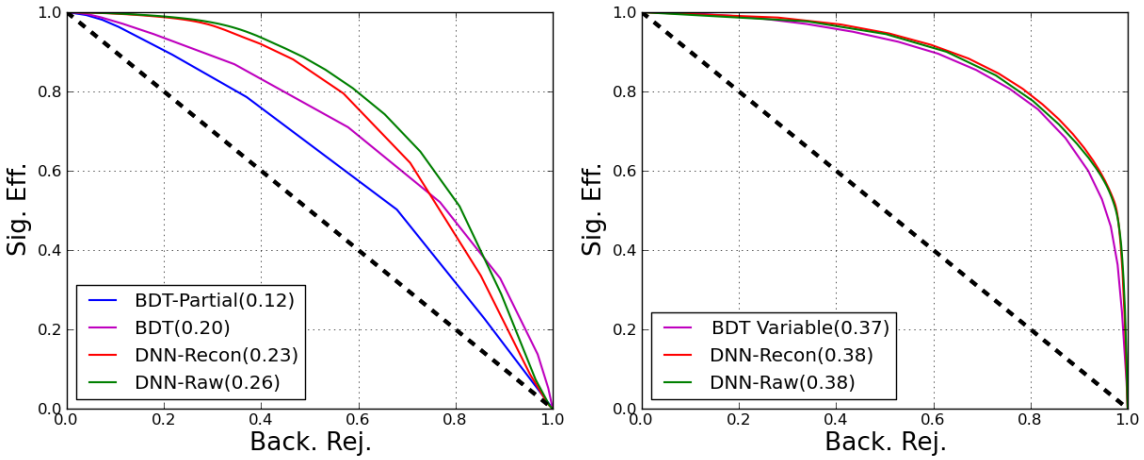

Figure 5.9: The resulting ROC curve for (left) SS and (right) MS events using both DNN methods compared to corresponding BDTs described in Section 5.2.

analysis, events which are outside of the Fiducial Volume are removed. The final data set is populated with events passing the above cuts such that there is an even number of signal and background events.

This total set is then divided such that 90% of events are used for training while the other 10% are set aside as a validation set for monitoring the DNN during training. In order to check for over-fitting, the loss and accuracy for both the training and validation set are compared as a function of the number of full passes through the training set (epoch). As shown in Figure 5.8 both sets of data show similar performance, indicating there is no over-fitting. Training was then done independently using the same procedure for both the reconstructed and raw images. As a result two algorithms are produced, one for the raw image (DNN-Raw) and one for the reconstructed image (DNN-Rec).

5.3.3 Performance and Validation

As done with the BDT, the first step in understanding the performance of the DNN is to examine the ROC curve of the resulting discriminator to see how the discrimination power compares to that of the BDT algorithm. This comparison is shown in Figure 5.9 for both SS and MS events, whose distributions are shown in Figure 5.12. From this

comparison it can be seen that both the DNN-Raw and DNN-Rec discriminators cover more area than the corresponding BDT algorithms, indicating that the DNN is outperforming the BDT in overall discrimination.

A more detailed examination of the SS ROC curves show that the BDT outperforms the DNN for background rejections above ∼0.7. This is due to the fact that the BDT includes knowledge of the standoff distance while the DNN is trained to explicitly ignore position. In order to make a fair comparison, an additional BDT was created (BDT-Partial) which included just risetime and the number of collection signals. Comparing this BDT discriminator to the DNN, the DNN is observed to outperform the BDT for any given point on the ROC curve. While the DNN itself sacrifices the discrimination power offered by standoff, this can be recovered by adding standoff as a third observable to the fit, as will be discussed in Chapter 6. This is also the case for MS events but as shown in Figure 5.6, the relative strength of standoff is less for the MS BDT and thus the effect is less noticeable.

In addition, due to the reduced number of dimensions contained in the reconstructed image relative to those of the raw image, the discrimination power of the DNN-Raw is expected to be marginally better. This is observed for SS events in Figure 5.9 where the DNN-Raw covers ∼10% more area than the DNN-Rec, indicating that there is some information lost when converting from raw to reconstructed images. Likely this is due to the DNN picking up on low energy signals, either induction or collection, not found by reconstruction which help constrain the size of the event. This can be further understood by looking at the ROC curves for SS events which contain one wire and two wire events separately, as shown in Figure 5.10. The relative improvement of the DNN over the BDT is larger for events with two wire channels as opposed to events with one wire, which is consistent with the expectation that the additional U-Wire signal increases the information available to the DNN allowing for a more informed decision. It can also be observed that the difference between the raw and reconstructed DNN algorithms is larger for one wire events, which could arise if the raw DNN is finding signals missed by reconstruction. For MS events the difference between the raw and reconstructed is significantly smaller, with the two methods covering nearly the same area. This indicates that simplifications introduces

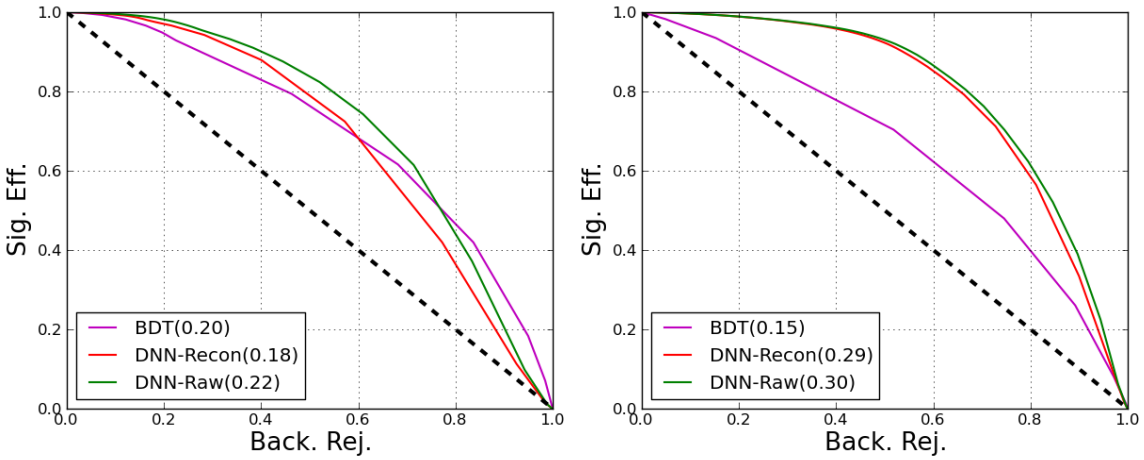

Figure 5.10: Comparison between the ROC curves for the DNN and BDT discriminators shown separately for SS events with (left) one wire channel and (right) two wire channels. A significantly larger difference in performance is observed for two wire events, which is consistent with expectations. The BDT shown here is the full BDT including standoff distance.

by the reconstruction algorithm have less impact on the MS discrimination.

Because the design of the discriminator is based on the hypothesis that the event size is the optimal discriminator for SS events, it is possible to further validate the performance by comparing the resulting discriminators to that which would be achieved by the presumed perfect discriminator. This is done in Figure 5.11 by comparing the ROC curve for the different SS discriminators to that which would be obtained if the true event size was known. While the goal is to achieve performance as close to this ideal case as possible, it is not expected that any discriminator can match the perfect discriminator as the true size is determined using the truth information only accessible in simulation and ignores all detector effects such as diffusion and reconstruction. Thus this comparison both shows the remaining discrimination power still theoretically available and ensures that the discriminators are not outperforming the expected limiting case.

One feature of the DNN algorithm is that unlike the BDT, it is not possible to pinpoint the exact features in each image which the DNN uses to classify events. This is a direct result of the DNN attempting to squeeze the maximal information from

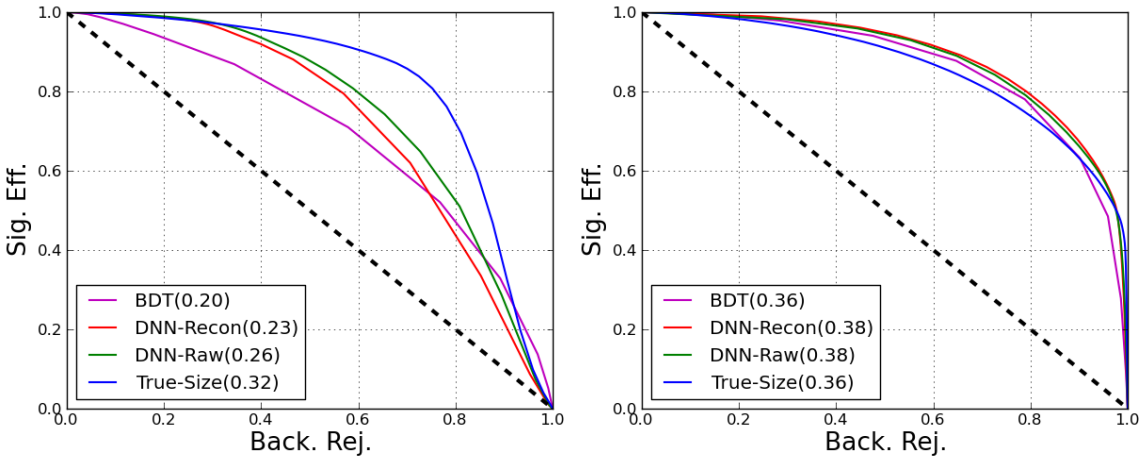

Figure 5.11: Comparison of the ROC curve for the discriminators applied to (left) SS and (right) MS events to the ROC curve which would result from having access to the true event size.

the images, specifically information not easily captured by the simplified set of BDT observables. Instead, validation is done by confirming that the underlying physical parameters, such as event size, are correctly identified. This is possible by correlating the fraction of events which are correctly classified as a function of these physical parameters. Classification in this case is done by choosing a point on the ROC curve, which corresponds to a binary cut on the DNN discriminator. For this comparison the DNN variable corresponding to a signal efficiency of 90% was used.

As indicated in Figure 5.1 the additional discrimination power offered by the DNN is expected to come from features which correlate to the spatial extent of events. Thus the first cross check is to confirm that the improved performance of the DNN is in fact correlated with event size. This is shown in Figure 5.13 for SS events by comparing the classification efficiency to the true simulated event size. As expected, events with larger sizes tend to be classified as background while smaller events are identified as signal. In addition, the roll off in efficiency begins at ∼9 mm compared to multiplicity alone which begins at ∼18 mm as shown in Figure 3.19. This is consistent with the expected resolving power of the U-Wires, which is ultimately limited by the 9 mm pitch. Going one step further and using the results of the SS BDT shown in Section 5.2, the rise time of the U-Wire signals is known to be a feature of the U-Wires

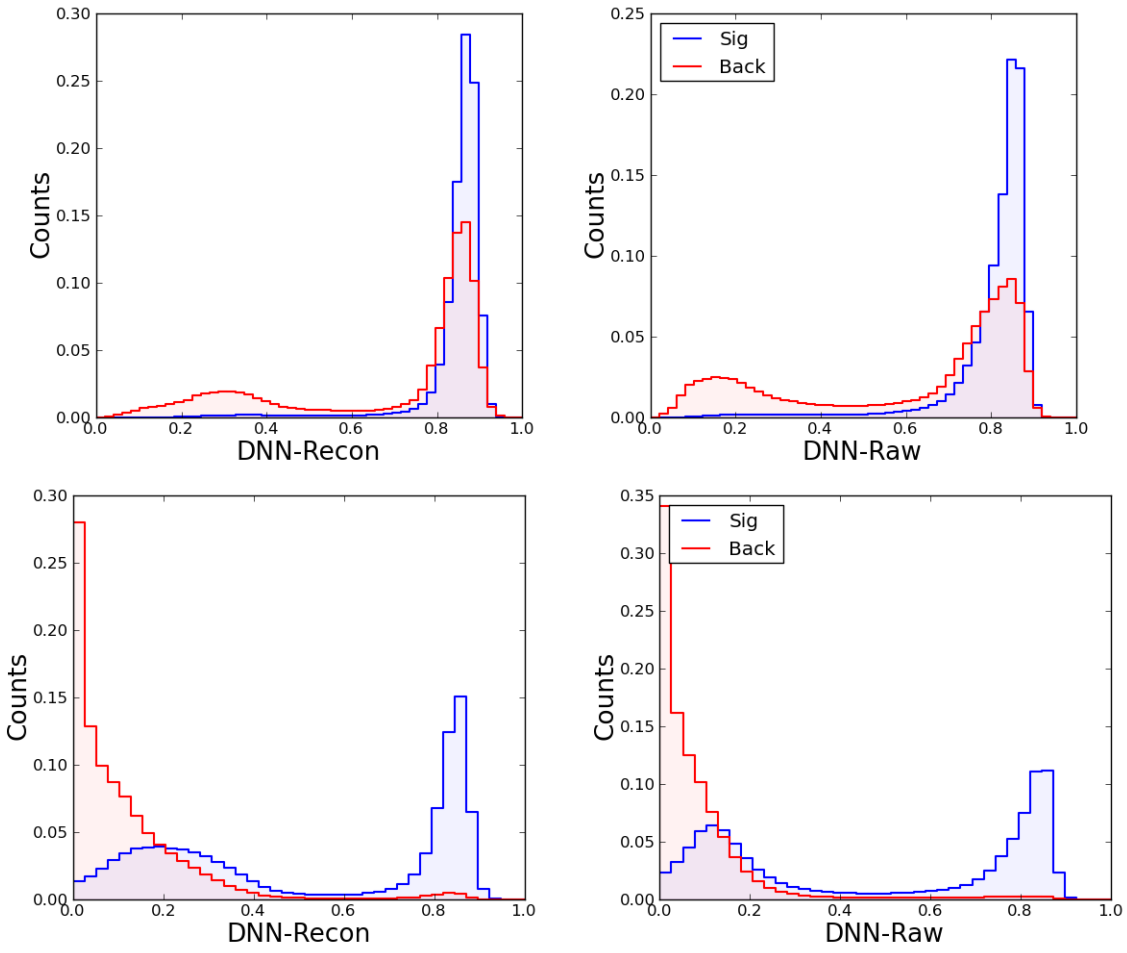

Figure 5.12: Distribution of the DNN discriminator for simulated signal($0\nu\beta\beta$) and background(γ) which are classified by reconstruction as (top) SS and (bottom) MS. This is shown for both the (left) reconstructed and (right) raw DNN algorithms.

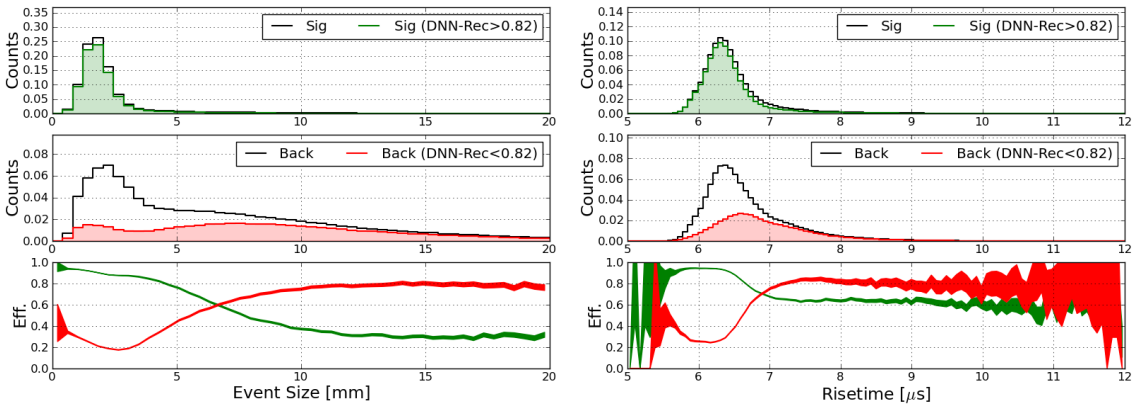

Figure 5.13: The simulated distributions of (left) event size and (right) risetime for SS (top) signal and (middle) background events compared to the distribution of events which are correctly classified by the DNN algorithm. (bottom) The resulting classification efficiency is shown to demonstrate that the DNN discriminator is correctly picking up features related to topology.

which offers discrimination power. Thus by also comparing the classification efficiency as a function of rise time, it can also be seen that the DNN has independently found the rise time as a feature useful in classification. The results of this cross check show that with only images of the U-Wires the DNN is correctly identifying features of the waveforms which correlate to size.

This cross check can also be done for MS events as shown in Figure 5.14. Similar to SS events the DNN shows correlation to the true event size, correctly identifying the larger backgrounds. The observed correlation is much smaller for MS events, with a more gradual slope in the background efficiency curve, indicating that the event size is not the only physical parameter being used by the DNN. A much larger correlation is observed between the maximum cluster fraction and the DNN discriminator though. This is consistent with the comparison of the MS ROC curves to the ROC curve for the true event size shown in Figure 5.11. Both DNN discriminators are shown to outperform the true event size indicating that the event size does not contain the full discrimination power.

The above cross checks indicated that the DNN correctly identified features correlated to topology but further validation is possible by checking for the absence of

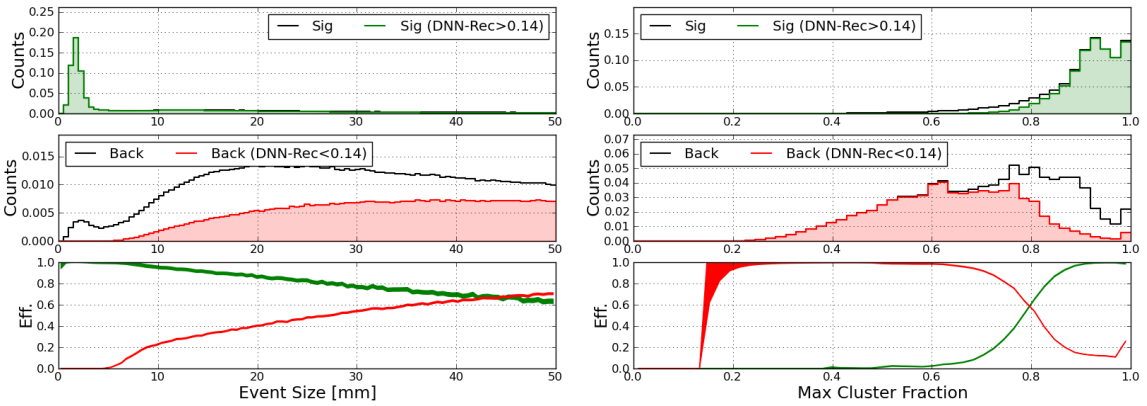

Figure 5.14: The simulated distributions of (left) event size and (right) maximum cluster fraction for MS (top) signal and (middle) background events compared to the distribution of events which are correctly classified by the DNN algorithm. (bottom) The resulting classification efficiency is shown to demonstrate that the DNN discriminator is correctly picking up features related to topology.

correlation. Because the DNN was intended to focus on topology the training set used uniformly distributed signal and background to ensure the DNN did not learn to use any position information in its classification. To confirm this, the correlation of the discriminator to the standoff distance for both SS and MS events is examined as shown in Figure 5.15. For SS events the observed efficiency is flat, showing that the DNN has not unintentionally picked up on features related to position.

In the case of MS events, this cross check is slightly more complicated as standoff distance is defined as the smallest standoff over all clusters. This results in bias for MS events at large standoff to have smaller event sizes (more signal like) as each additional cluster increases the probability that the event would have at least one cluster which falls at a smaller standoff. This results in a slight correlation in the MS efficiency, with worse identification efficiency for backgrounds at large standoff. Due to the bias in the standoff definition, this correlation is the result of the underlying correlation between standoff and event size for MS events and does not represent the DNN using position to classify events.

Up to this point, validation has focused on consistency checks on the simulation and have not involved the performance on real data. The final validation step is

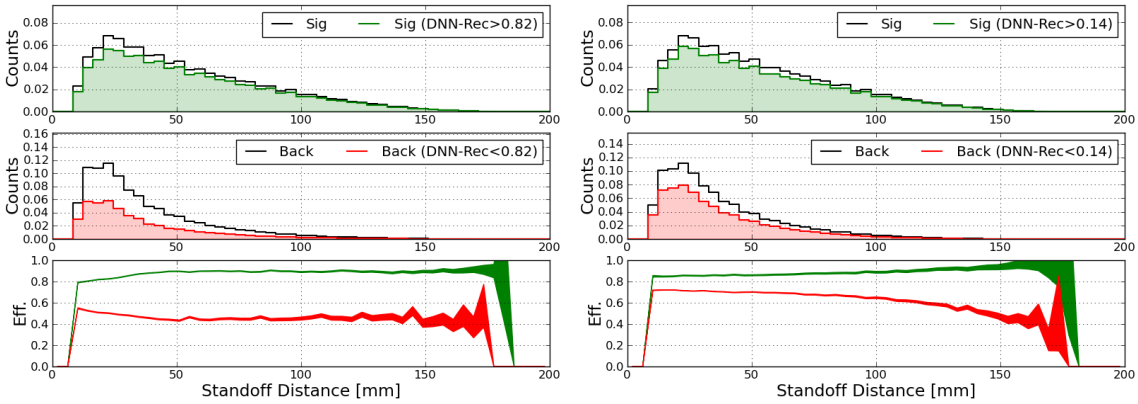

Figure 5.15: The simulated distributions of standoff distance for (left) SS and (right) MS for (top) signal and (middle) background events, compared to the distribution of events which are correctly classified by the DNN algorithm. (bottom) The resulting classification efficiency is shown to demonstrate that the DNN discriminator is correctly ignoring position in determining signal likeness.

then to compare the results of the simulation to what is observed when applying the DNN algorithm to data. Because the DNN is trained entirely with simulation, model inaccuracies in the MC can then be imparted onto the DNN discriminator. As a result discrepancies can then arise between the predicted distributions of the DNN and those observed in data, which can cause sysematic bias in the final analysis. While the same concern exists for the BDT as well, the DNN uses the lower level information and so it is harder to pinpoint exactly which parameters of the model might effect the agreement between data and MC.

In order to quantify this agreement, the observed distributions for data taken during calibrations are compared to that predicted by simulation. Calibration sources are ideal because they have sufficient activity that the events coming from other sources, like the $2\nu\beta\beta$ of ^{136}Xe, are negligible. Comparisons for both the raw and reconstructed DNN discriminators are shown in Figure 5.16 for the three γ calibration sources deployed at the main cathode location (S5). For both the raw and reconstructed DNN the residuals for are all within $\sim$25%($\sim$50%) for SS(MS) events, showing relatively good agreement between data and MC. The bins with the worst agreement are the central bins $(0.3 - 0.7)$ which are sparsely populated but show as much as a 50%

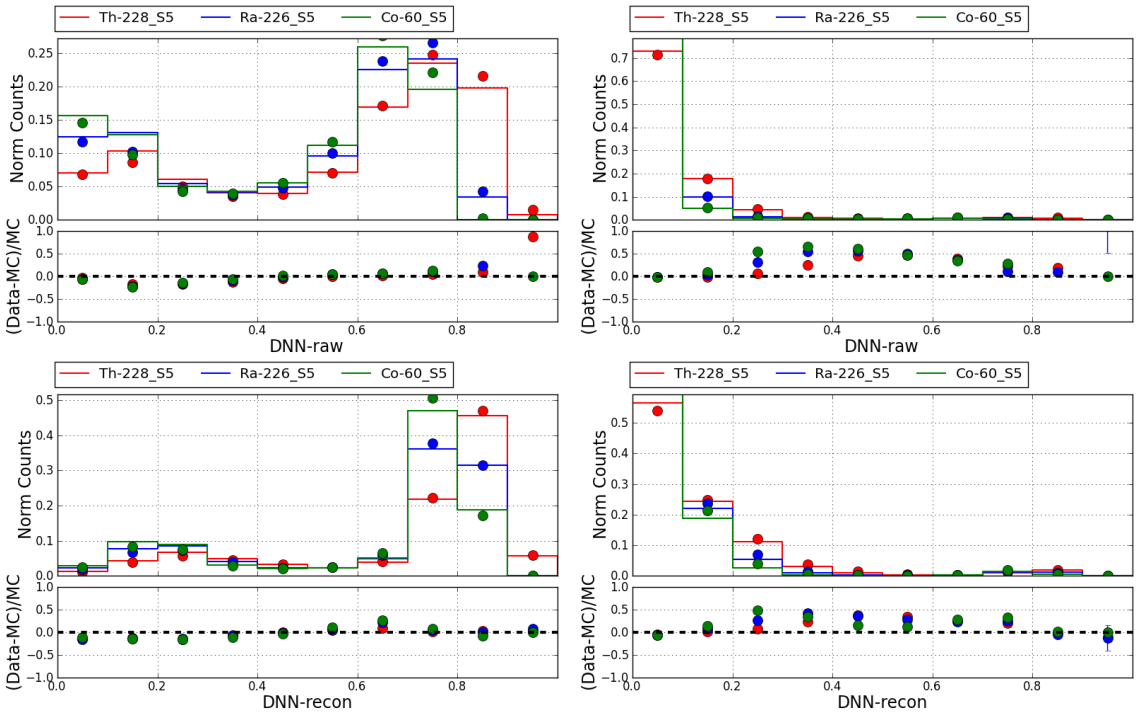

Figure 5.16: Comparison between data and simulation for the (top) DNN-Raw and (bottom) DNN-Rec discriminators for (left) SS and (right) MS events in calibration data taken at the S5 Cathode position.

excess in data.

While the calibration sources offer the simplest comparison between data and simulation, these sources result in mainly γ's. As a result, calibration data mainly allows for validation of background-like events. In order to extend this performance validation to signal like $\beta\beta$ events, the expected distribution of the DNN variables can be compared to the observed $2\nu\beta\beta$ spectrum in low background data. This comparison is somewhat more difficult than for calibration data because the low background spectrum contains both $2\nu\beta\beta$ as well as backgrounds. Estimating the $2\nu\beta\beta$ spectrum requires taking the low background data and fitting the observed energy spectrum with the expected background components, as will be described in Chapter 6. Using these fit results, the $2\nu\beta\beta$ spectrum is estimated by subtracting the best fit background model from the observed spectrum. The resulting comparison between data and simulation is shown in Figure 5.17 for the raw and reconstructed DNN. From this it

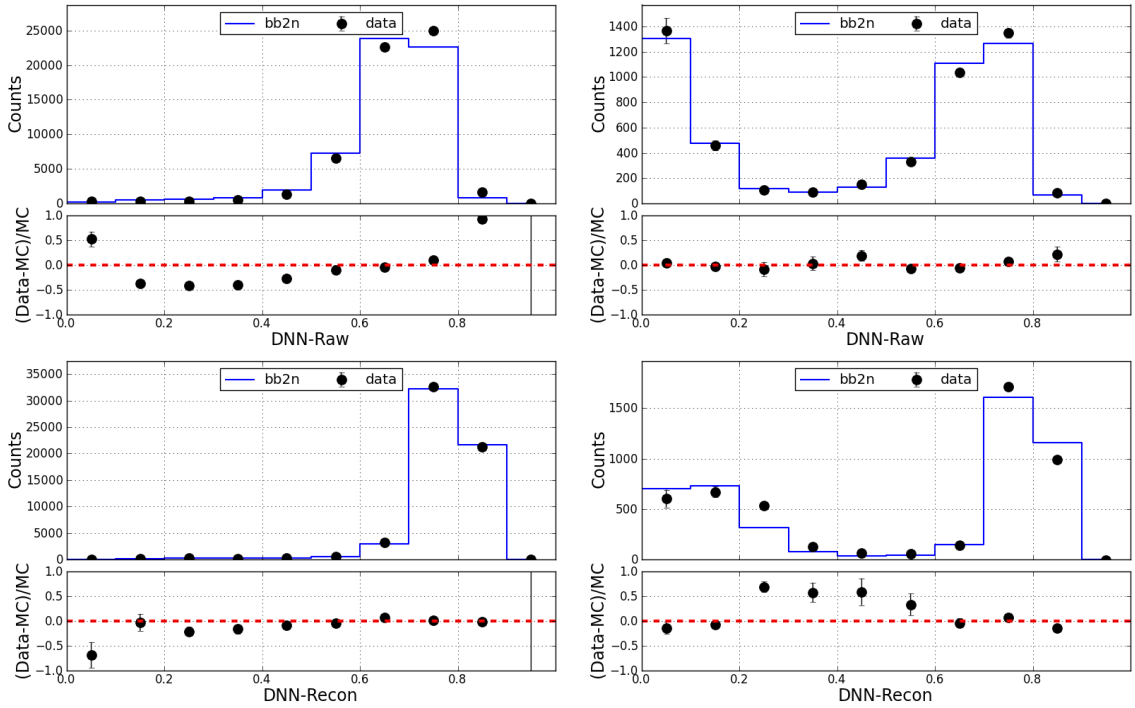

Figure 5.17: Comparison between data and simulation for the DNN-Raw discriminator for (left) SS and (right) MS events in the background subtracted $2\nu\beta\beta$ spectrum.

can be seen that both the DNN-Raw and DNN-Rec discriminators have generally good agreement with the residuals generally within ∼50%, with the observed exceptions occurring for the sparsely populated bins. While the calibration data generally saw good (<25%) agreement for the SS DNN variables, there is worse agreement in the DNN-Raw distribution for SS events with residuals exceeding 25% in multiple bins. In addition, an overall slope in the residuals for the DNN-Raw is observed, indicating a systematic shift with more signal like events appearing in data than predicted in MC.

While examining the residuals disagreement between the distributions in data and MC give a general indication on the accuracy of the simulation they do not directly show the impact to the final analysis. In addition, some optimization is possible by varying the number of bins as shown in Figure 5.18 for the reconstructed DNN. Changing the binning does not fix the underlying issue but it can in principle remove

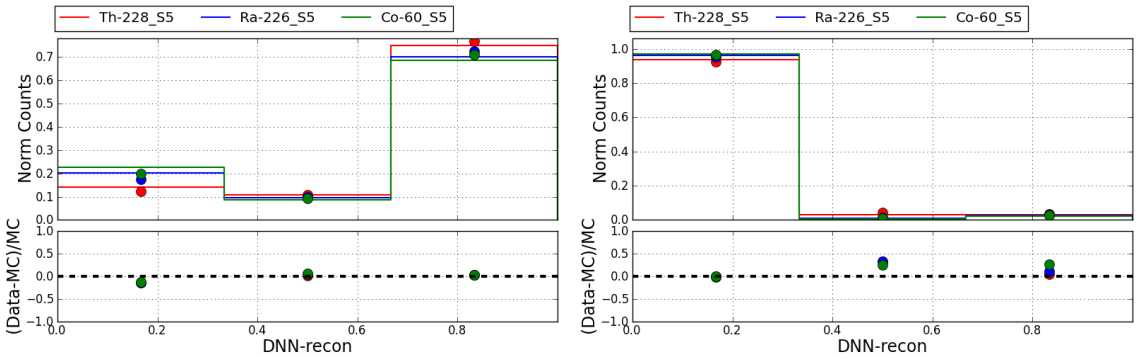

Figure 5.18: Comparison between data and simulation for the DNN-Rec discriminator for (left) SS and (right) MS events in calibration data with reduced number of bins.

potential bias introduced by a mismatch between the shape of the data and MC distribution, at the cost of potentially sacrificing discrimination power. This trade off between agreement and discrimination is not fully captured in the above agreement plots but as will be discussed in Chapter 6 its impact on the $0\nu\beta\beta$ search is determined by an unskewing procedure. Extensive studies were then done to determine the optimal discriminator which included determining which method (BDT, DNN-Raw or DNN-Rec) and what binning to use to maximize the sensitivity to $T_{1/2}^{0\nu\beta\beta}$.

Chapter 6

Results from Final $0\nu\beta\beta$ Search

6.1 Overview

As described in Chapter 3, EXO-200 took data towards the $0\nu\beta\beta$ search between September 2011 and December 2018. During this period the results toward the $0\nu\beta\beta$ search have incrementally been updated to reflect both the additional exposure and improvements to hardware and analysis. This chronology is shown in Table 6.1, which lists the corresponding exposure, limit and sensitivity for each publication. While the details of the previous analysis are described elsewhere [64], the final analysis details, utilizing the entire EXO-200 exposure, will be explained in the following chapter [59].

Date	Exposure[kg · yr]	Sensitivity[$\times 10^{25}$ yrs]	Limit[$\times 10^{25}$ yrs]	Ref.
2012	32.5	0.7	1.6	[80]
2014	100.0	1.9	1.1	[65]
2018	177.6	3.7	1.8	[66]
2019	234.1	5.0	3.5	[59]

Table 6.1: Summary of the published EXO-200 results.

As with previous analyses, a blinding procedure was implemented to prevent the tuning of the analysis toward any particular outcome. This was accomplished by hiding or "masking" all SS $0\nu\beta\beta$ candidate events with energies between $Q_{\beta\beta} \pm 2\sigma$. All of the optimizations presented below were done using this masked dataset, with the data un-blinded only upon finalizing the analysis strategy. This blinding procedure was applied to both the newly taken data and data which had been previous analyzed. While it is impossible to remove all bias from the previously examined data, re-blinding prevents iterative improvements aimed at removing or adding individual events.

6.2 Fit Model

In order to determine the rate of $0\nu\beta\beta$ observed in data, the observed spectrum in LB data is fit to an expected signal and background model. This is done by performing a binned maximum likelihood fit using the MINUIT package of RooFit [111]. The following describes the background model, fit procedure and fit parameters.

6.2.1 Background Model

An accurate background model is required to determine whether a $0\nu\beta\beta$ signal is present in the data. The background model is created using the simulation, described in Section 3.3, to model the expected detector response from different background sources. The sources included in this model are

- **LXe decays**: $2\nu\beta\beta$ (^{136}Xe), ^{137}Xe and ^{135}Xe
- **Th-Like (2615 keV γ from ^{208}Tl)**: near (TPC vessel) and far (inner cryostat)
- **U-like (2448 keV γ from ^{214}Bi)**: near (TPC vessel), far (air gap radon) and LXe (active, inactive, cathode)
- **^{60}Co (2506 keV sum γ peak)**: near (TPC vessel) and far (calibration sources)
- **n-capture components**: ^{1}H in HFE, ^{63}Cu and ^{65}Cu in the outer cryostat, inner cryostat, and TPC vessel and ^{136}Xe de-excitation γs
- **Other**: ^{40}K and ^{65}Zn in TPC vessel

For EXO-200 the main backgrounds originate from trace amounts of long lived isotopes in the material used to construct the TPC vessel and its components. The most notable of these isotopes are ^{238}U and ^{232}Th which have γ lines near $Q_{\beta\beta}$ in their decay chain, at 2448 keV and 2615 keV respectively. In addition, the copper can also contain ^{65}Zn, ^{40}K and ^{60}Co. The ^{40}K and ^{65}Zn components are of less importance for the $0\nu\beta\beta$ search because their corresponding γ's and β's have energies below 2 MeV. They are still included in the model to ensure an accurate fit of the $2\nu\beta\beta$ spectrum.

The decay of ^{60}Co results in two γ's with energy 1173 keV and 1333 keV. While neither γ is near $Q_{\beta\beta}$, the sum energy of 2506 keV falls in the 2σ ROI. Most of these sum events should be reconstructed as MS but due to the finite resolving power of the detector there is some probability that they are reconstructed as SS, if they are spaced too closely. For the decay chains with events near $Q_{\beta\beta}$ (^{238}U, ^{232}Th and ^{60}Co), the model includes both a near and far component to account for changes in the spectral shape arising from different source locations. Backgrounds produced from the prompt γ's produced by the neutron capture on the hydrogen in HFE and the copper of the TPC vessel and cryostat are also included in the model.

In addition, there are also internal backgrounds from sources in the LXe volume. The most prominent of these sources being the backgrounds produced by the xenon itself. This includes the previously discussed $2\nu\beta\beta$ decay of ^{136}Xe as well as the decay of ^{137}Xe and ^{135}Xe, which are produced by neutron capture on ^{136}Xe and ^{134}Xe respectively. Apart from the xenon isotopes, the main background in the LXe volume is the decay of ^{222}Rn which results in the 2448 keV γ from ^{214}Bi. This includes ^{222}Rn decays which occur at the cathode as well as those in the inactive and active volumes.

6.2.2 Maximum Likelihood Fit

The simulation results for both the $0\nu\beta\beta$ signal and each of the background sources described in Section 6.2.1 are used to create binned probability density functions (PDFs). A maximum likelihood fit is then performed to simultaneously fit each of these PDFs to the observed data. This is done by minimizing the negative log likelihood function (NLL)

$$\mathrm{NLL} = \sum_i \left[m_i^{SS} - d_i^{SS} \ln m_i^{SS} + m_i^{MS} - d_i^{MS} \ln m_i^{MS} \right] + G \tag{6.1}$$

where $d_i^{SS(MS)}$ is the number of SS(MS) counts observed in the i^{th} bin and $m_i^{SS(MS)}$ is the expected number of SS(MS) events from the fit model in the i^{th} bin. The number of counts in the fit model is calculated by summing over all the signal and background

components of the fit model

$$m_i^{SS} = N \sum_j n_j s_j f_j^{SS}(\boldsymbol{x}) \tag{6.2}$$

$$m_i^{MS} = N \sum_j n_j (1 - s_j) f_j^{MS}(\boldsymbol{x}) \tag{6.3}$$

where N is the overall normalization, $f_j^{SS(MS)}(\boldsymbol{x})$ is the SS(MS) PDF for the j^{th} component, n_j is the expected number of counts in the j^{th} component and s_j is the fraction of events which are SS for the j^{th} component. Each PDF is a function of the multi-dimensional set of observables $\boldsymbol{x}$, which can include energy and any additional set of observables such as standoff or the $0\nu\beta\beta$ discriminators described in Chapter 5. The last term in Equation 6.1 represents the Gaussian constraints used in the fit, defined as

$$G = \frac{1}{2}(\boldsymbol{\rho} - \boldsymbol{\rho_0})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\rho} - \boldsymbol{\rho_0}) \tag{6.4}$$

where $\boldsymbol{\rho}$ is the vector of constrained parameters, $\boldsymbol{\rho_0}$ is the set of expected value for the parameters and $\boldsymbol{\Sigma}$ is the covariance matrix for the set of parameters.

Each PDF component of the fit model is defined by two component specific parameters (n_j and s_j). For all but the ^{222}Rn component the total counts are left unconstrained. In the case of ^{222}Rn, a constraint is added to restrict the fractional contribution of the three ^{222}Rn background sources to 10% of that measured using an analysis of α's in the data as described in [64]. The SS-Fraction of each component is determined by the simulation but is constrained by the observed agreement between data and MC for the calibration sources, described in Section 6.2.4. Based on the observed correlations between the SS fraction residuals of the three γ calibration sources, an 85% correlation is added to the SS-Fraction constraint for each of the γ-like background PDFs. While the total number of neutron-captures ($\sum_{nc} n_j$) is unconstrained, the fractional contribution of each individual neutron capture component to the total number of neutron captures

$$c_j = \frac{n_j}{\sum_{nc} n_j} \tag{6.5}$$

is constrained to be within 20% of that predicted in simulations of the neutron capture rates [72].

In addition, the fit also includes three global parameters. First there is the overall normalization (N), which accounts for potential differences in the data selection between data and MC. This parameter is constrained by the overall uncertainty in the signal efficiency, which is evaluated in Section 6.5 and summarized in Table 6.5. A second normalization parameter is added to all SS PDFs to account for possible discrepancies in the signal efficiency specific to the ROI itself. This is done by subdividing each PDF into a section inside and outside the 2σ region around $Q_{\beta\beta}$. A separate normalization (N_{ROI}) is applied to the section of each SS PDF which is inside the ROI. This normalization is constrained by the uncertainty on the number of background events in the ROI, which will be discussed in Section 6.3. Finally, a linear scaling factor (β) is applied to the energy of all β-like PDFs ($E_\beta = \beta E$). Because the energy scale in data is calibrated with γ sources, this factor accounts for the possibility that the energy scales of β-like events and γ-like events are different. While β is expected to be within $\sim 1\%$ of unity, it is conservatively left unconstrained and allowed to freely float in the fit. A summary of the fit parameters and their respective constraints are shown in Table 6.2

Table 6.2: Summary of the fit parameters and their respective constraints.

	Name	Constraint
PDF	n_i	^{222}Rn constrained to 10% of activity determined from α's [64]
	s_i	Constrained within uncertainties in Section 6.2.4
Neutron	f_j	Constrained to 20% of MC predictions [72]
Global	N	Constrained with uncertainties in Table 6.5
	N_{ROI}	Constrained with uncertainties in Table 6.3
	β	None

6.2.3 Search Sensitivity

In order to optimize the analysis, a figure of merit is needed to compare different fit configurations. As discussed in Section 2.3.1, one of the main metrics used to quantify a $0\nu\beta\beta$ is the search sensitivity. For $0\nu\beta\beta$ searches it is common practice to define this sensitivity as the median upper limit an experiment would report if it were repeated a large number of times. In this way, the sensitivity represents a figure of merit which is not impacted by statistical fluctuations present in data.

Evaluating the sensitivity is done by using Monte Carlo to produce a large number ($\sim$1000) of toy datasets, which are used to simulate the expected distribution of upper limits. Assuming the background model described in Section 6.2.1 is correct, the expected contribution of each background can be estimated by performing a fit to blinded LB data. The result of this fit can then be used to generate toy sets by randomly sampling each background PDF based on its relative contribution in the best fit. A fit is performed to each toy dataset and the 90% upper confidence limit on the number of $0\nu\beta\beta$ counts is evaluated. The experimental sensitivity is then defined as the median of the resulting distribution of limits.

6.2.4 Data and Simulation Agreement

Because the fit uses a background model created by simulating the expected detector response of each PDF, reliable fit results depend on the model's completeness and the simulation's ability to reproduce the detector response for each component. Quantifying this accuracy is done by measuring the differences in the predicted and observed detector response for known sources. In the case of γ-like backgrounds, this is done by using the three γ calibration sources and their corresponding simulations. For β-like events this is more difficult, as no pure sample of β events exists. Instead, the comparison is done between the simulated $2\nu\beta\beta$ spectrum and the spectrum in low background data after all other backgrounds are subtracted out. This subtraction is done by first fitting the low background data to the background model and then using the fit results to remove the relative contribution of each background PDF from the observed spectrum.

6.3 Background Uncertainty

The first parameter of interest in quantifying the simulation accuracy, is the ability to accurately predict the number of background events in the 2σ ROI. Three main factors contribute to this uncertainty in the model. The first, is the uncertainty introduced by variations in the energy resolution due to the non-uniform spatial and temporal response in the detector. Calibration data is used to quantify these variations, which are observed to be $\sim$0.1% near $Q_{\beta\beta}$. This results in a relative change in the number of background events in the 2σ ROI of 1.5(1.2)% in Phase-I (Phase-II).

In addition, inaccuracies in the simulation can result in differences in the spectral shape of the PDFs. These shape differences can then result in a systematic bias in the best fit results, which translate to a bias in the best fit number of $0\nu\beta\beta$ events. Quantifying the accuracy of the simulation in reproducing the true spectral shape of each PDF is done by first examining the differences between the observed and predicted spectra. This is shown for both Phase-I (Phase-II) in Figure 6.1(6.2) for the energy dimension and Figure 6.3 for the DNN and standoff dimension.

In order to translate this observed agreement into an uncertainty on the final number of $0\nu\beta\beta$ events, a new set of PDFs is created based on the observed disagreement between data and MC. This is accomplished by scaling each original PDF by the observed differences, such that the residual difference between data and MC is removed. The scaling for each β-like PDF is done with the residuals estimated from the $2\nu\beta\beta$ agreement and for the γ-like PDFs the corresponding γ source is used. A "toy" dataset can then be built with these scaled PDFs by randomly sampling each of them based on the relative contributions estimated from fits to the LB data using blinded data. This is repeated many times and each resulting toy set is fit with the un-scaled PDFs to determine the number of events in the 2σ ROI. The systematic bias introduced by the residual shape disagreement is estimated from the fractional difference between this result and the result when fitting to toy sets built from the unscaled PDFs.

The bias introduced by the shape uncertainty depends on the set of observables and their respective binning in the fit. For each fitting configuration which was

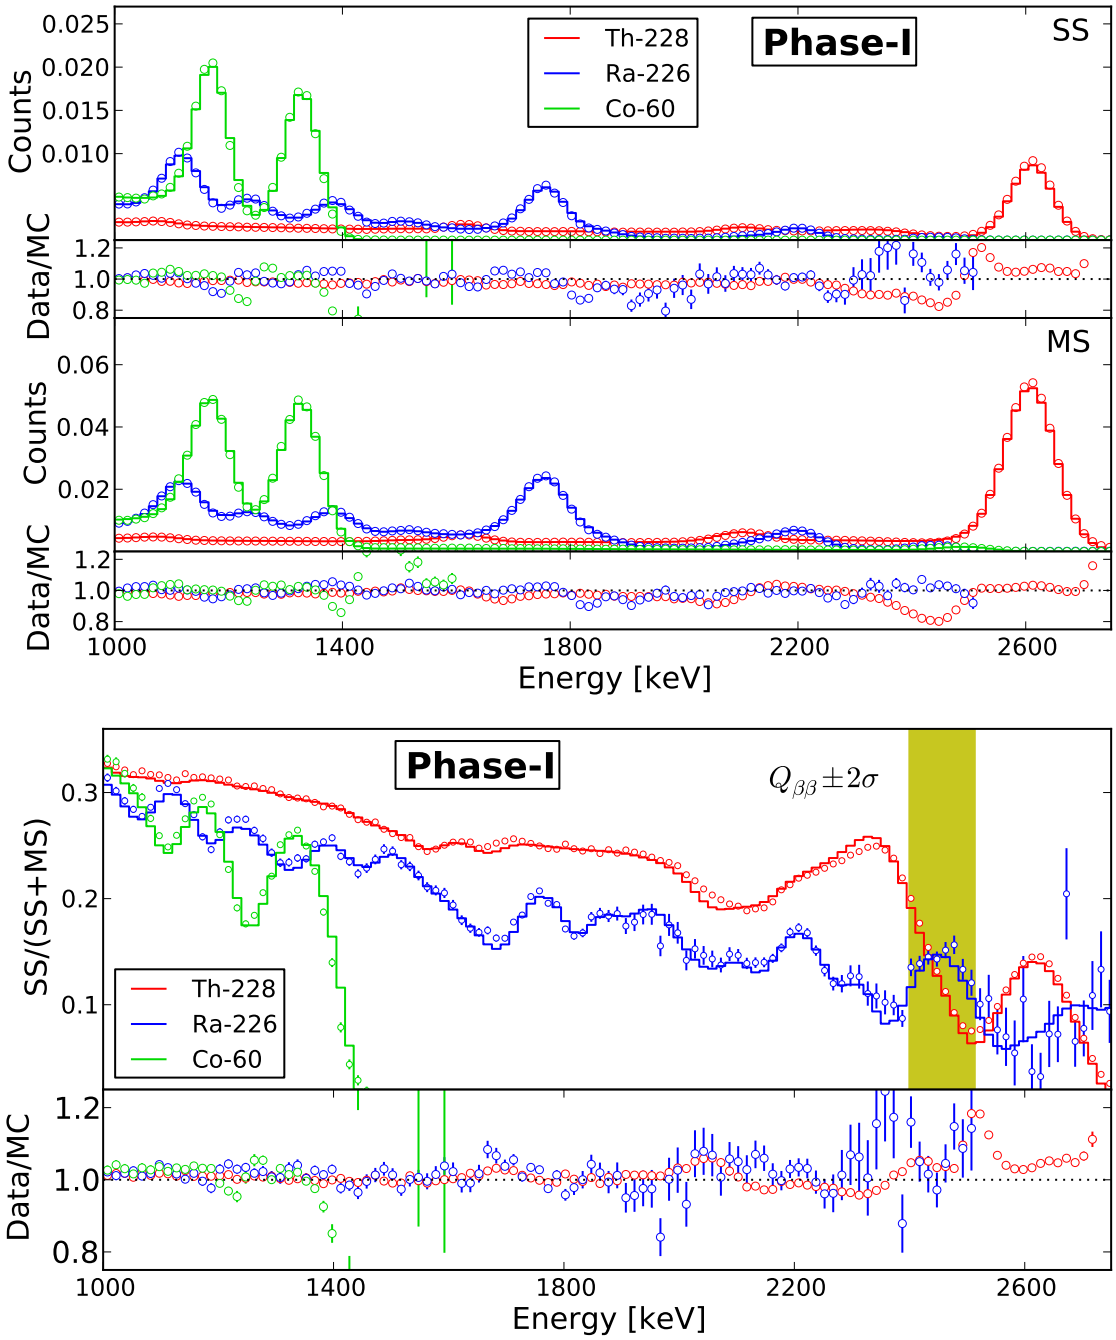

Figure 6.1: Comparison of the (top)energy and (bottom) SS-fractions(SS/(SS+MS)) distributions between data and MC for Phase-I calibration data taken at the S5 position. The bottom panel shows the ratio between data and MC.

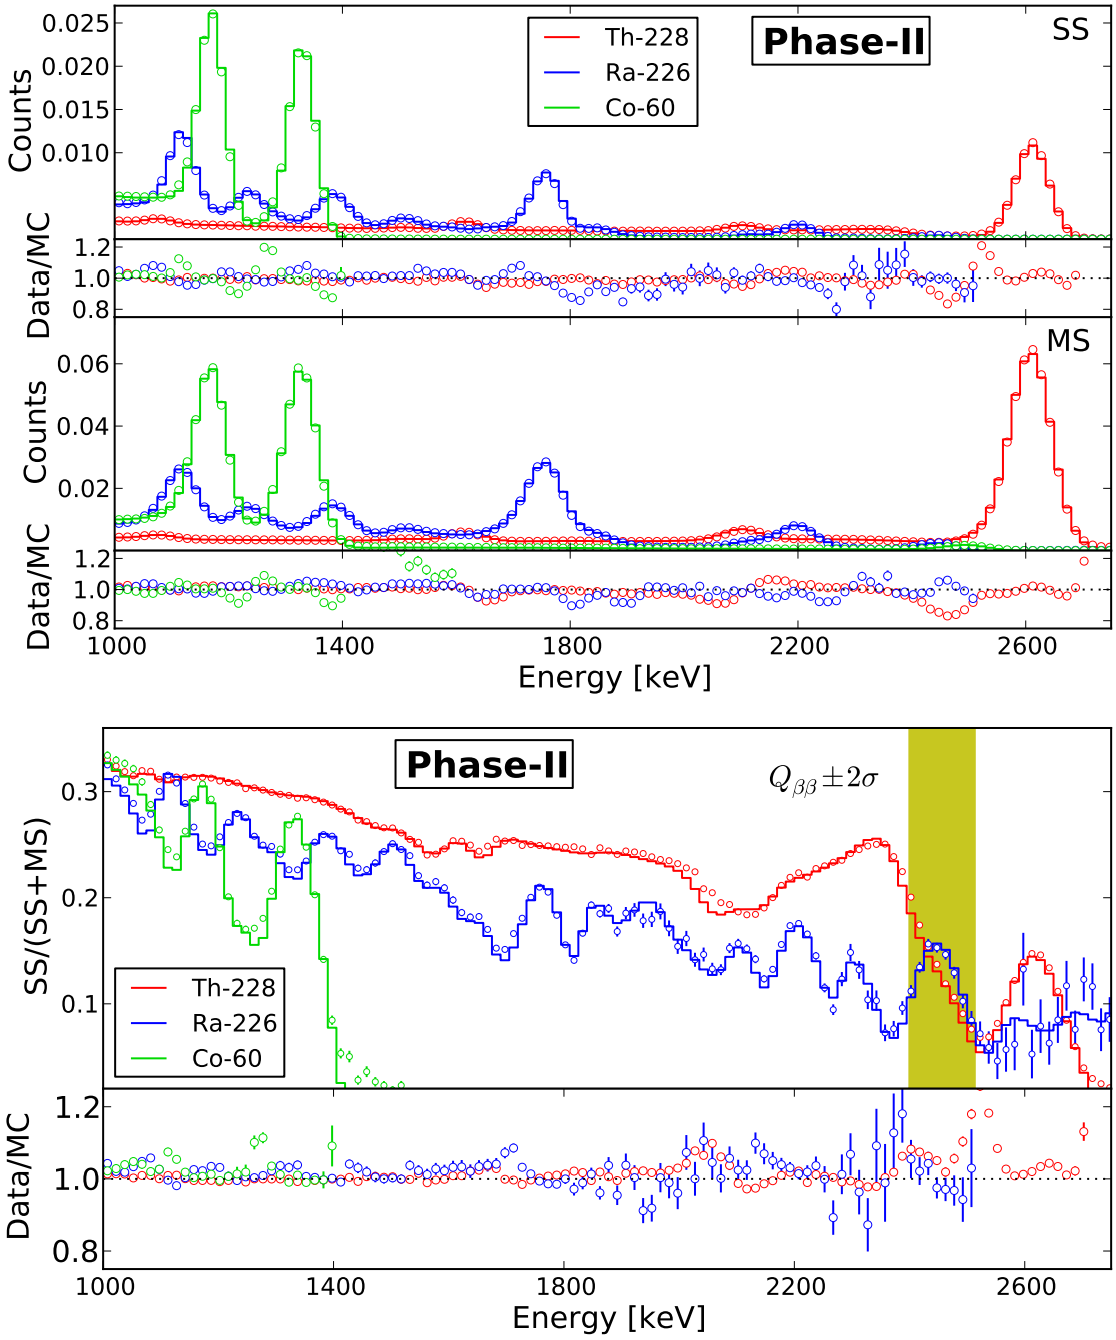

Figure 6.2: Comparison of the energy distribution between data and MC for Phase-I calibration data taken at the S5 position. The bottom panel shows the ratio between data and MC.

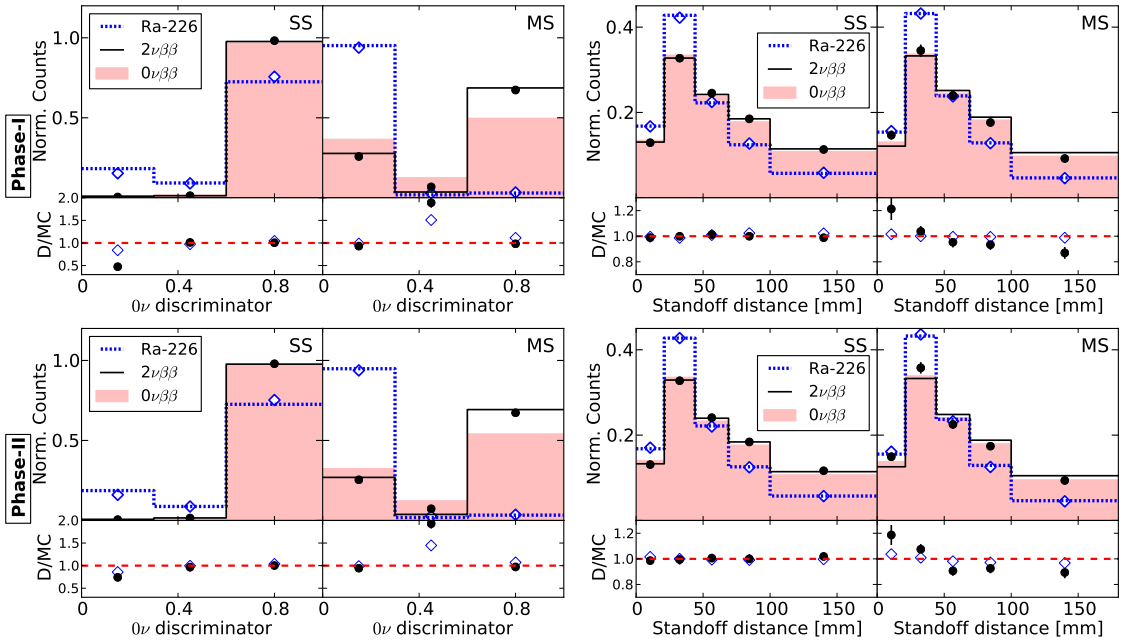

Figure 6.3: Comparison between data(dots) and MC (lines) for the DNN discriminator (left) and standoff distance (right) for Phase-I (top) and Phase-II (bottom). Shown for the distributions of ^{226}Rn (blue) and background subtracted $2\nu\beta\beta$ (black). In addition, the simulated distribution expected for $0\nu\beta\beta$ is shown in red. The difference between the $0\nu\beta\beta$ and $2\nu\beta\beta$ DNN distributions are due to the higher rate of bremsstrahlung at higher electron energies.

tested, this bias was reevaluated for sensitivity studies. In the case of the final fit configuration, described in Section 6.6, this error is estimated to be 5.8%(4.4%) for Phase-I (Phase-II).

The final contribution to the background uncertainty comes from the completeness of the background model. This mainly derives from the unknown locations of the far components for ^{238}U, ^{232}Th and ^{60}Co. Because there are many possible degenerate source locations, only a representative location is used in the model. Studies are then performed to determine the impact of the assumed location on the $0\nu\beta\beta$ search. This is done by fitting the blinded LB data using background models with alternative locations for the far components of each of these sources. The change in the relative number of events in the 2σ ROI can then be used to determine the contribution of the background model assumptions to the overall ROI normalization. The uncertainty is determined to be 4.0%(4.6%) for Phase-I (Phase-II).

The final uncertainty in the background counts is estimated by summing, in quadrature, the individual contributions described above. As summarized in Table 6.3, this results in a 7.1(6.5)% systematic uncertainty in Phase-I (Phase-II). These uncertainties are used to constrain the ROI normalization term (N_{ROI}) in the LB fit.

Source	Phase-I	Phase-II
Spectral shape agreement	5.8%	4.4%
Background model	4.0%	4.6%
Energy Scale and Resolution	1.5%	1.2%
Total error	7.1%	6.5%

Table 6.3: Summary of systematic uncertainty on the number of background ROI events.

6.4 Fraction of SS Events

In addition to the agreement of the spectral shape of each PDF, the fit also relies on the predicted fraction of events which are SS. This is quantified by comparing

the predicted and observed SS fractions, as a function of the event energy, for the three calibration sources. Figure 6.1(6.2) show this comparison for the calibration sources deployed at the S5 position in Phase-I (Phase-II). All bins are shown to have agreement to within 25%. The weighted average of the residual difference, using the inverse of the uncertainties as the weights, is used to evaluate the uncertainty for each source type. Taking into account variations based on source type and source position the systematic uncertainty is observed to be 5.8%(4.6%) in Phase-I (Phase-II).

In previous analyses, the SS fraction in the 2σ ROI was measured to be $\sim$24% [66] but as seen in Figure 6.1 and Figure 6.2 the observed SS fraction is $\sim$14%. This change is due to the relaxed reconstruction cut, described in Section 6.5.4, which is applied prior to the SS fraction calculation. As a result of this relaxed cut, nearly all of the previously removed MS events are recovered relative to previous calculation. For $0\nu\beta\beta$, which are mostly SS, only 10% of the total events are impacted by this cut but for γ's near $Q_{\beta\beta}$ this cut previously removed $\sim$50% of events.

6.5 Livetime and Efficiency

As discussed in Section 2.3.1, maximizing both the livetime and efficiency are important features of any $0\nu\beta\beta$ search as they ensure that any potential $0\nu\beta\beta$ events are detected. While this alone would imply an analysis with no selection cuts, some cuts are still needed to remove periods of data which may have poor signal to noise ratios. These regions may contribute to higher background rates which can potentially hide a $0\nu\beta\beta$ signal. It is then important to quantify and model these cuts to ensure that the observed efficiency is both optimal and reproduced by the simulation.

For the $0\nu\beta\beta$ analysis two types of selection cuts are applied to the data. First, data periods which are deemed to be poor quality are removed from the data set. This removes periods which have high background rates that can potentially limit the search sensitivity. Because these cuts remove segments of time and not specific events, they only impact the overall detector livetime and do not change the $0\nu\beta\beta$ detection efficiency. In addition, there are also cuts applied to specific types of events. These cuts remove events which are either likely caused by backgrounds or poorly modeled

in simulation. Unlike vetoes, these cuts do not impact the detector livetime but instead factor into the overall $0\nu\beta\beta$ detection efficiency. A summary of these cuts and the resulting livetime and efficiency are presented below.

6.5.1 Vetoes

To reduce the background rate from cosmogenic activation of the copper vessel and HFE fluid, data taken 1 ms before and 25 ms after a trigger of the muon veto system are cut from the analysis. For muon events which pass through the detector, the main background results from the capture of spallation neutrons on ^{136}Xe [72] . This produces ^{137}Xe, which has a half life of 3.8 min and decays via a β emission with a total energy of 4173 keV [112]. Because this decay is accompanied by a prompt γ cascade, potential ^{137}Xe events can be identified by searching for events which are coincident with a muon and have energies consist with that expected for the γ cascade. A veto is than applied to reject all data taken for 5 half-lives (19 min) after such events, which is estimated to reduce the the background from ^{137}Xe by $23 \pm 8\%$ [66]. An additional 1 min veto is included for all muon events reconstructed in the TPC to remove muon-induced spallation products, as well as minimize the impact of events which are missed by the ^{137}Xe tagger.

	Phase-I		Phase-II	
Source	Days	%	Days	%
Muon	34.02	5.9	29.42	5.2
Environment	10.07	1.8	22.36	3.9
^{137}Xe	21.06	3.7	16.62	2.9
Total	64.61	11.3	67.46	11.8

Table 6.4: Summary of livetime lost due to the vetoes described in the text in both data taking phases.

In addition, data taken with poor environmental conditions, such as during mining activities, are removed from the analysis. These periods have been observed to contain

bursts of electronic noise in the detector, which can result in poor data quality. These vetoed periods result in <12% loss to the livetime, as summarized in Table 6.4. Taking into account all vetoes, the integrated livetime in the detector is 1144.52 days, corresponding to 573.56(570.96) days in Phase-I (Phase-II).

6.5.2 Fiducial Volume

While it is important to maximize the number of ^{136}Xe nuclei being observed, it is also important to limit systemic effects which can potentially bias the final result. For this reason, a fiducial volume (FV) is defined for the final analysis which excludes certain regions of the detector that are not well modeled by the detector simulation. Because the wire planes form a regular hexagon, events used in the final analysis are required to lie within a hexagon in the xy plane with an apothem of 162 mm. In addition, events are also required to be more than 10 mm from the cylindrical PTFE reflector. These radial cuts account for regions near the edge of the detector, where field lines are not perfectly parallel to the z-axis and instead have components which point toward the center of the detector. These radial non-uniformities can alter the trajectory of drifting electrons resulting in a distortion of the detector volume. This effect is small for the field used in the main $0\nu\beta\beta$ search but is more pronounced at lower fields used in the transport analysis of Chapter 4, as shown in Figure 4.1. As discussed in Section 3.3, the simulation used to drift electrons selects a central segment of the full electric field model to approximate the field in the detector. This allows for an accurate model in the bulk and at the wire planes but does not model imperfections near the edge of the detector. In addition, events are required to be more than 10 mm from both the V-Wire and cathode planes due to the excess background rates observed in these regions.

The FV does not impact the $0\nu\beta\beta$ detection efficiency, as it is defined only for events which are fully contained in the FV. An accurate measure of this volume is still needed to estimate the number of ^{136}Xe atoms. This FV contains $3.31 \cdot 10^{26}$ atoms of ^{136}Xe, with an equivalent mass of 74.7 kg, The uncertainty in determining this volume is 3.1% for both phases and is determined by combining the position

reconstruction accuracy with the difference between the predicted and observed event rates in calibration data [64].

6.5.3 Coincidence Cut

A coincidence cut is also included to remove events which occur too closely in time to each other. This cut aims to reduce the potential background caused by short-lived intermediate states, such as ^{214}Po which has a half life of 164.3 µs and is produced by the β decay of ^{214}Bi. In order to remove these correlated events, all events within a given time window of another event are removed from the analysis. Previous analyses [65, 66, 80] used a 1 s window, which being three orders of magnitude longer than the half-life of ^{214}Po is extremely conservative. Because random coincidences are also removed by this cut, this resulted in a 7% lose to the $0\nu\beta\beta$ efficiency. Recent studies [113] indicated that this window could instead be shortened to 100 ms, with minimal increase in the background rate in ROI. This in turn reduces the random coincident rate, resulting in just a $\sim$0.5% efficiency loss.

6.5.4 Reconstruction Cut

In order to pass the fiducial cut described in Section 6.5.2 an event is required to be fully contained within the FV. In the case of MS events, this means that all clusters be within the FV. As discussed in Section 3.4.3, it is possible for at least one found charge cluster in a MS event to contain a U-Wire and APD signal but no V-Wire signals. In this case, it is not possible to determine whether the cluster is fully contained in the FV. For this reason previous analyses [65, 66, 80] have included a requirement that all charge clusters have their 3D position reconstructed.

While $0\nu\beta\beta$ is predominately SS, roughly $\sim$10% of events are expected to be MS due to bremsstrahlung. Many of these events are than removed by the reconstruction cut because bremsstrahlung tends to result in low energy secondary deposits. This results in a $\sim$10% loss in the $0\nu\beta\beta$ detection efficiency. In order to recover these events the reconstruction requirement was relaxed such that events are only required to have 60% of their total energy associated with fully reconstructed clusters. While

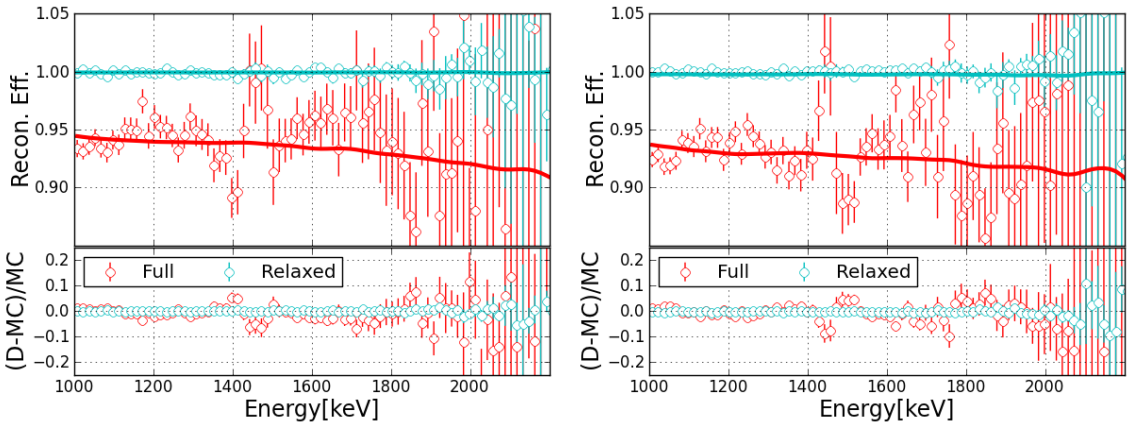

Figure 6.4: Shown in the top panel is the efficiency, estimated as the fraction of events passing the cut, as a function of energy in both background subtracted LB data (points) and $2\nu\beta\beta$ MC(lines). The bottom panel shows the residual difference between data and MC. Included for both Phase-I (left) and Phase-II (right) and for the full (red) and the relaxed (cyan) reconstruction requirement.

these events are inherently MS, it is possible to distinguish them from the dominant γ backgrounds using a powerful discriminator, such as those described in Chapter 5.

The efficiency of this cut can than be determined in simulation as the ratio of events which pass the cut to the total number of events. Using the results of the $0\nu\beta\beta$ simulation, this relaxed cut is observed to recover nearly all MS events. The resulting loss in detection efficiency is determined to be 0.3(0.6)% for Phase-I (Phase-II). The uncertainty of this estimate is determined by comparing the efficiency predicted for $2\nu\beta\beta$ events in simulation to that observed in background subtracted LB data. Using fits with and without the reconstruction cut, the efficiency in data is estimated as the ratio of the two fit results. The resulting comparison is shown in Figure 6.4 for both data taking phases. The corresponding uncertainty is estimated from the error weighted average of the residuals, which gives 0.04% for both phases. Figure 6.4 also shows the comparison between data and MC for the full 3D cut resulting in a $\sim$0.4% average residual. To be conservative, the full 3D result is taken for the uncertainty on the reconstruction efficiency.

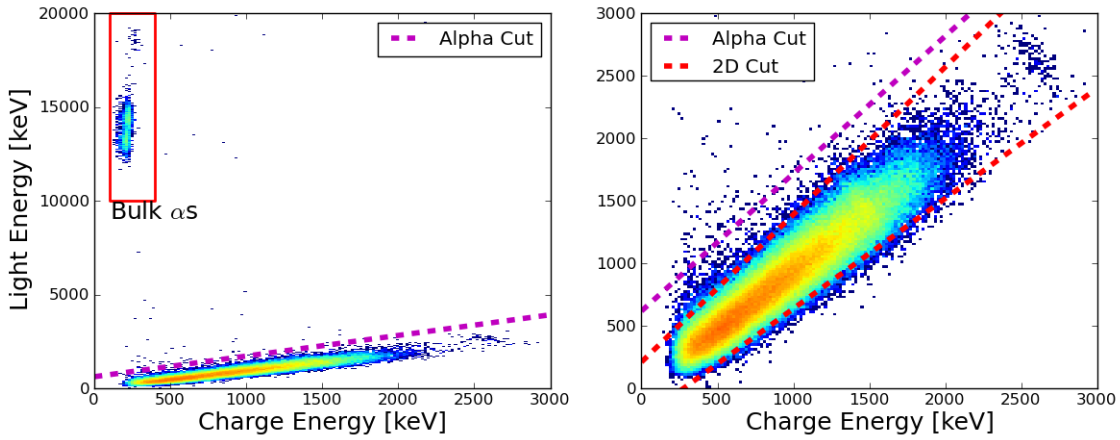

Figure 6.5: Cut on the ratio of charge to light designed to remove αs such as the highlighted population of bulk αs in the upper left corner(left) and the revised cut which additionally removes poorly reconstructed events on either side of the main diagonal (right). Note the different coordinate scale in the two plots.

6.5.5 2D Cut

As discussed in Section 3.5.1, a cut on the ratio of the light and charge energy can be used to remove α's from the final data set with negligible impact to the overall $0\nu\beta\beta$ detection efficiency.

Following the results from previous analysis [66], the diagonal cut was further tightened to also remove poorly reconstructed events, such as the case where the scintillation light is partially blocked. This is done by requiring that the orthogonal distance between the calibrated 2D energy of an event and the rotated axis fall within a given window. Using both calibration and $2\nu\beta\beta$ data, this distance is found to be approximately Gaussian and is measured in σ's of the energy resolution along the direction orthogonal to the rotated axis. This cut as well as the original α cut are depicted in Figure 6.5 for low background data.

As done for the other cuts, the impact to the $0\nu\beta\beta$ detection efficiency is evaluated in simulation by the fraction of events which survive this cut. The efficiency and the corresponding uncertainty of depends on the threshold used to define the distance from the center of the gaussian. Comparisons between the predicted and observed

distribution indicated agreement within 1% for all thresholds above 2.2σ. In order to have a negligible impact on the sensitivity a cut of 2.5σ is ultimately chosen. As a result, the final efficiency is evaluated to be 98.6(98.4)% in Phase-I (Phase-II), with a corresponding 0.9% uncertainty for both phases.

6.5.6 Other

In addition to the cuts described above, all events which have more than one reconstructed scintillation cluster are removed from the analysis. This serves a similar purpose to the coincidence cut described in Section 6.5.3, but targets correlated events faster than 1 ms, which can fall in the same event frame. Based on the observed rate of such events, this cut is estimated to have negligible impact to the overall efficiency and carries a $\sim$1.0% uncertainty.

Finally, after un-blinding the full dataset, all SS candidate events having energies between $Q_{\beta\beta} \pm 2\sigma$ were examined as a consistency check. During this examination, one event with an un-denoised energy inside this region was found to be mis-reconstructed by the de-noising algorithm. As a result, the event was pushed to an un-physical region of the 2D energy plane causing it to be removed by the 2D cut. This mis-reconstruction was identified as a dead/empty voxel in the Phase-II light map used by the de-noising algorithm and the overall efficiency loss was estimated to be $< 1.0\%$. To account for this, a 1.0% efficiency loss is conservatively added to the Phase-II detection efficiency. An additional 1.0% uncertainty is also added in quadrature to the overall systematic uncertainty.

6.5.7 Detection Efficiency for $0\nu\beta\beta$

The detection efficiency of a $0\nu\beta\beta$ event is determined by multiplying the resulting efficiency for each of the above cuts. This is summarized for both Phases in Table 6.5. The overall efficiency, accounting for all vetos and selection criteria are determined to be $97.8 \pm 3.1\% (96.4 \pm 3.1\%)$ in Phase-I (Phase-II). The slightly lower efficiency in Phase-II is due to events which are mis-reconstructed by the denoising algorithm. This is an improvement over the previous analysis, which had a detection efficiency

of 82.4(80.8)% in Phase-I (Phase-II) [66]. This increase in efficiency is due to the relaxation of the reconstruction cut described in Section 6.5.4 and the coincidence cut described in Section 6.5.3.

	Phase-I		Phase-II	
Source	Efficiency[%]	Error[%]	Efficiency[%]	Error[%]
Fiducial volume	-	2.8	-	2.6
0.1 s Coincidence	99.5	< 0.2	99.5	< 0.2
Partial 3D cut	99.7	< 0.4	99.4	< 0.4
Light-to-charge ratio	98.6	0.9	98.4	0.9
De-noising mis-rec	-	-	99.0	1.0
Other [64]	-	< 1.0	-	< 1.0
Total	97.8	3.1	96.4	3.1

Table 6.5: Summary of the final signal efficiency and the corresponding systematic errors introduced by event selection requirements for events above threshold which occur in the fiducial volume.

6.6 Fit Dimensions

While the main observable of the fit is energy, the fit is defined to generically extend to an arbitrary number of dimensions. This allows for additional observables to be incorporated into the analysis in order to optimize the search sensitivity to $0\nu\beta\beta$. As discussed in Chapter 5, a number of additional dimensions were explored based on the ability to improve the discrimination between signal and background. While adding observables which offer additional discrimination tends to improve the sensitivity, there is the potential to introduce additional uncertainties if the new dimensions are poorly modeled in the simulation. In order to finalize the set of observables used in the final fit, an extensive study was done to determine which configuration maximized

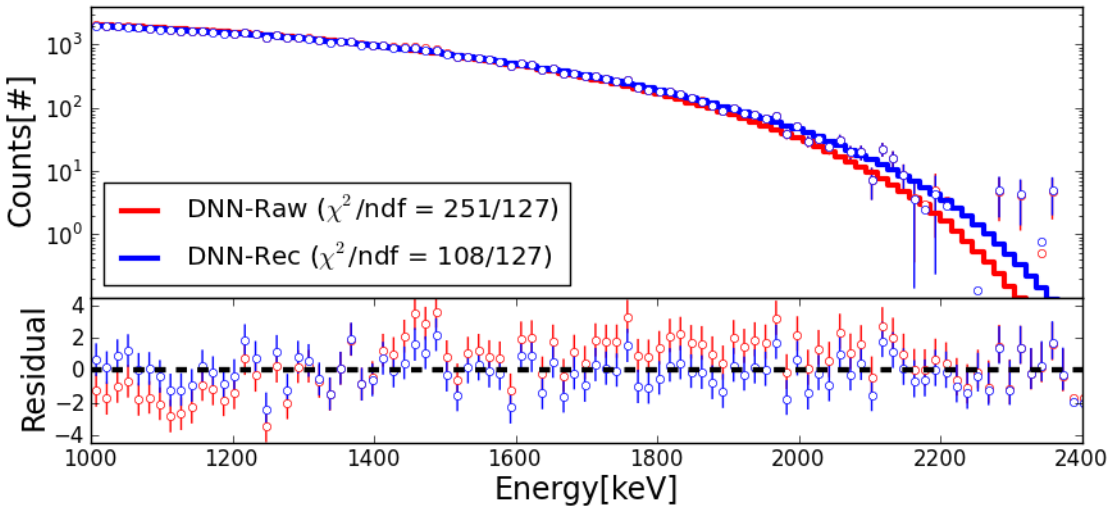

Figure 6.6: Phase-II background subtracted fit to LB data using DNN-Rec and DNN-Raw, with residuals shown in the bottom panel. The χ^2/ndf, calculated using statistical uncertainties only, is shown for each configuration in the legend.

the sensitivity while limiting uncertainties. This is done by evaluating the uncertainty associated with N_{ROI} for each configuration and calculating the corresponding sensitivity to $0\nu\beta\beta$.

Because two DNN-based discriminators were developed, the first step in picking the final configuration was to determine which DNN discriminator to use. While the sensitivity resulting from a fit with DNN-Raw achieved ∼3% better sensitivity than DNN-Rec, assuming comparable shape errors, the agreement between data and MC was worse for the DNN-Raw variable then for DNN-Rec, as shown in Section 5.3.3. As expected, this additional shape error results in the fit attempting to optimize agreement in the DNN dimension, at the cost of introducing bias along other dimensions of the fit. Specifically, the best fit β using DNN-Raw deviates from 1.0 by ∼1.5% to compensate for the excess of data at high signal likeness observed in Figure 5.17. As a result, the χ^2 for the fit using DNN-Raw is worse than the fit using DNN-Rec, as shown by the projection of the best fit onto the energy axis in Figure 6.6. A clear increase in the residuals is observed for the fit using DNN-Raw, resulting in a χ^2 which is more than 2× worse. Thus even though the DNN-Raw results in ∼3% better sensitivity, under the ideal condition, the worse agreement seen in the DNN-Raw

dimension introduces additional bias to the final result. For that reason the DNN-Rec was the preferred observable for the final fit. In addition, as discussed in Chapter 5, the DNN observables showed worse agreement for finer binning. To limit the additional bias introduced by this worse agreement, a binning was chosen to maintain the uncertainty on N_{ROI} to $< 10\%$. This resulted in the DNN dimension being composed of three bins.

Based on these results, two fit configurations were considered for the final analysis. First, a fit using a BDT in both the SS and MS dimension was studied as a natural extension to the previous BDT analysis. Second, a fit using the DNN-Rec and standoff in both SS and MS was studied as a potential improvement. In addition, a fit using just energy along with the standard SS/MS definition was used as a reference point. The results of this study are shown in Table 6.6, from which it can be seen that the configuration using the DNN-Rec observable achieved the best sensitivity. While the BDT analysis provided a $\sim$3% worse sensitivity, the slight margin indicates that the discrimination power of the DNN-Rec observable can be mostly accounted for by careful construction of the BDT.

Dimensions	N_{ROI} Error[%]	Sensitivity[cts]	$T^{0\nu\beta\beta}_{1/2}$ Sensitivity[$\times 10^{25}$yrs]
E	4.6%	13.00	2.7
E+BDT	7.3%	10.63	3.4
E+DNN-Rec+SD	6.4%	10.20	3.5

Table 6.6: Summary of the sensitivity to $T^{0\nu\beta\beta}_{1/2}$ for different fit configurations in Phase-II.

6.7 Final Result

Prior to un-blinding, the 90% confidence upper limit is evaluated for the final fit configuration for both phases of data using the approach described in Section 6.2.3. This is done independently for each phase and found to be be $3.3(3.4)\times 10^{25}$ yrs for Phase-I (Phase-II). The combined sensitivity is estimated by randomly combining MC toys

from each phase and adding their respective profiles, resulting in a total sensitivity of 5.0×10^{25} yrs. Comparing to an analysis which uses only energy and multiplicity, this configuration achieves a ∼25% improvement to the search sensitivity. In addition, this analysis approach represents a ∼10% improvement to the BDT analysis presented in [66], when applied to the same data set. The better sensitivity derives from improvements in detection efficiency, energy resolution and $0\nu\beta\beta$ discrimination.

After un-blinding the entire dataset, the observed LB spectrum for Phase-I and Phase-II are fit separately using the model defined in Section 6.2. Results of the fit for Phase-I (Phase-II) are shown in Figure 6.7(6.8). The relative contribution attributed by the fit to each background source in the 2σ SS ROI are summarized in Table 6.7 along with a comparison to the total number of observed events in the same energy range. In total, 63.2 background events were expected and 65 events were observed.

(counts)	^{238}U	^{232}Th	^{137}Xe	Total	Data
Phase-I	12.6	10.0	8.7	32.3±2.3	39
Phase-II	12.0	8.2	9.3	30.9±2.4	26

Table 6.7: Best-fit background contributions to $Q_{\beta\beta} \pm 2\sigma$ versus observed number of events in data.

A scan of the profile likelihood for different number of $0\nu\beta\beta$ events is done to obtain a confidence interval on the number of $0\nu\beta\beta$ events. This is done independently for each phase and the resulting profiles are added, while taking into account differences in livetime and signal efficiency, to determine the combined limit. These profiles are shown in Figure 6.9, giving a best fit of 2.7 $0\nu\beta\beta$ events in the combined profile. This result is consistent with the null hypothesis of no $0\nu\beta\beta$ to within 1σ and has a 90% confidence upper limit of 20.15 $0\nu\beta\beta$ events. This dataset consists of 1144.52 days of livetime and 3.31×10^{26} atoms of ^{136}Xe with a 97.1% livetime weighted average efficiency. Combining this information, a lower limit of $T^{0\nu\beta\beta}_{1/2}$ $>3.5\times10^{25}$ yrs is obtained with Equation 2.23.

In addition, the profile of each individual phase results in lower limits of $T^{0\nu\beta\beta}_{1/2}$

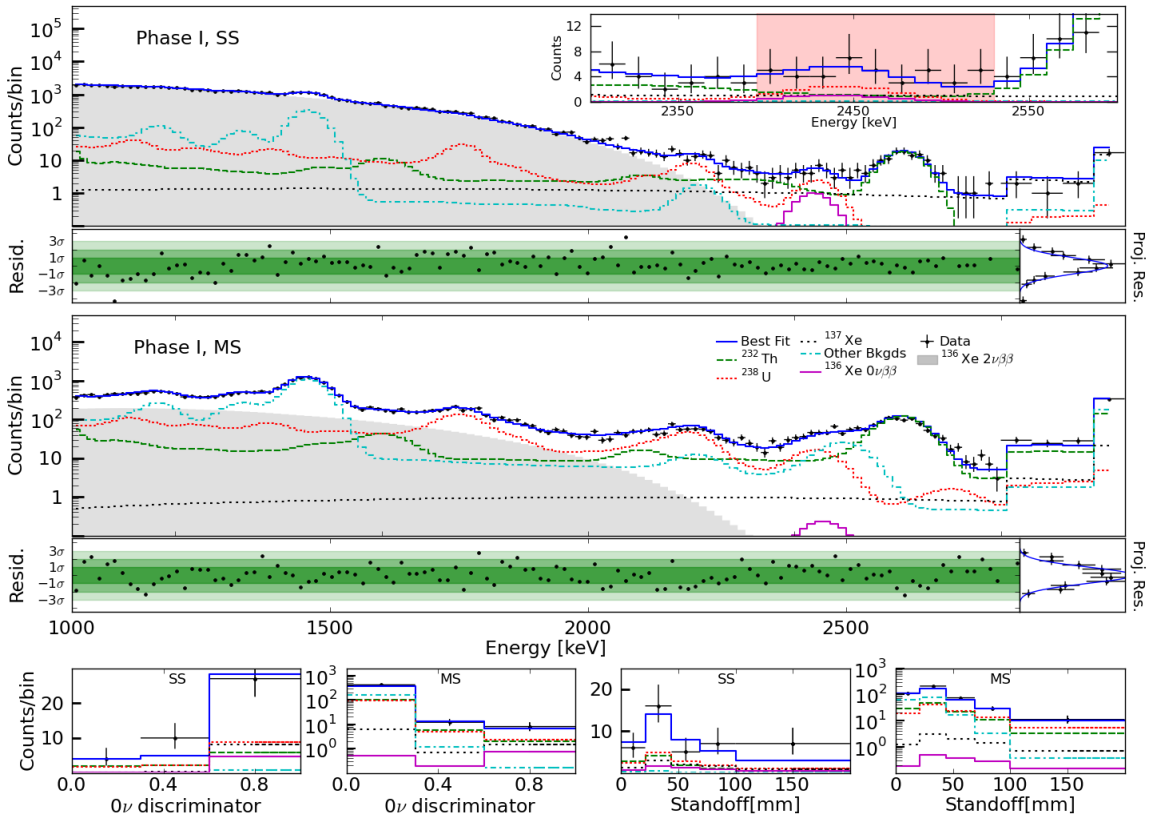

Figure 6.7: Best fit to the Phase-I low background data for the (top) SS (middle) MS energy spectrum using bins of 15 keV and 30 keV below and above 2800 keV, respectively. The inset shows a zoomed in view around the best-fit value for $Q_{\beta\beta}$. Projection of events in the range 2395 keV to 2530 keV on the (bottom left) DNN (bottom right) standoff fit dimension for SS and MS events. The best-fit residuals are shown for the energy dimensions and typically follow normal distributions, with small deviations taken into account in the spectral shape systematic errors.

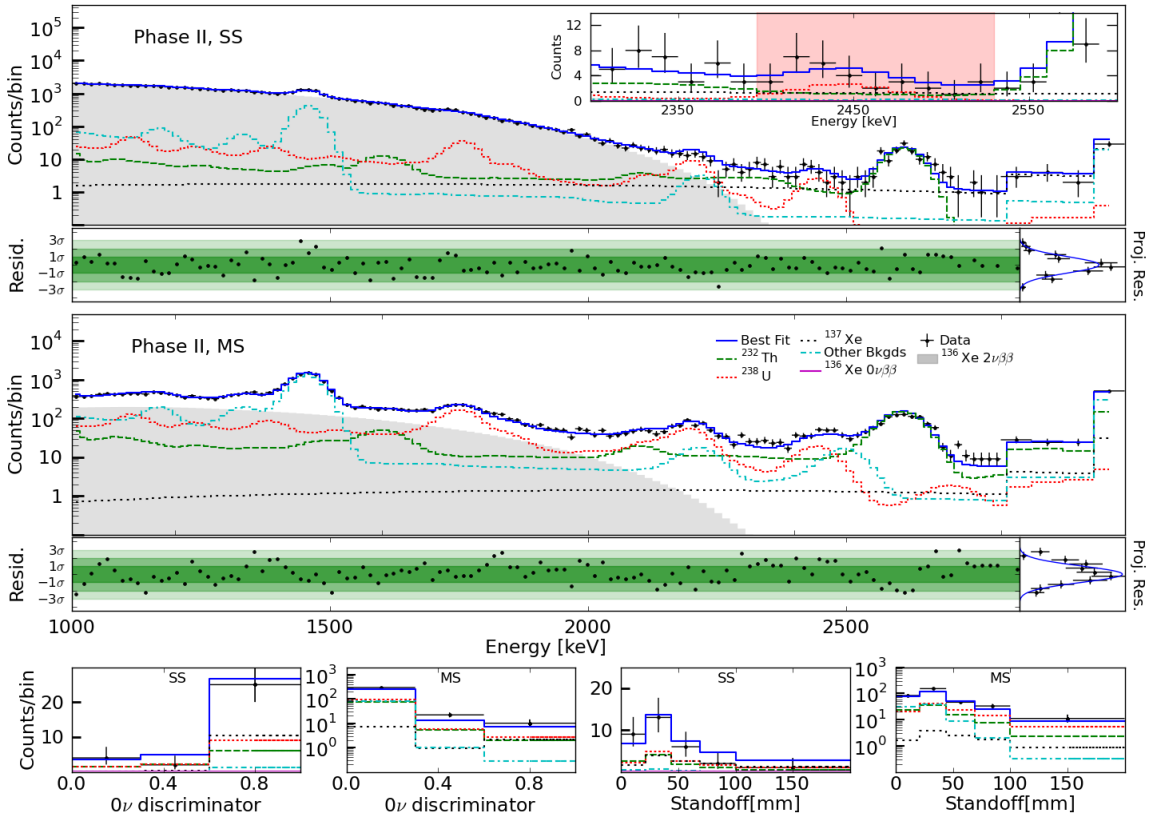

Figure 6.8: Best fit to the Phase-II low background data.

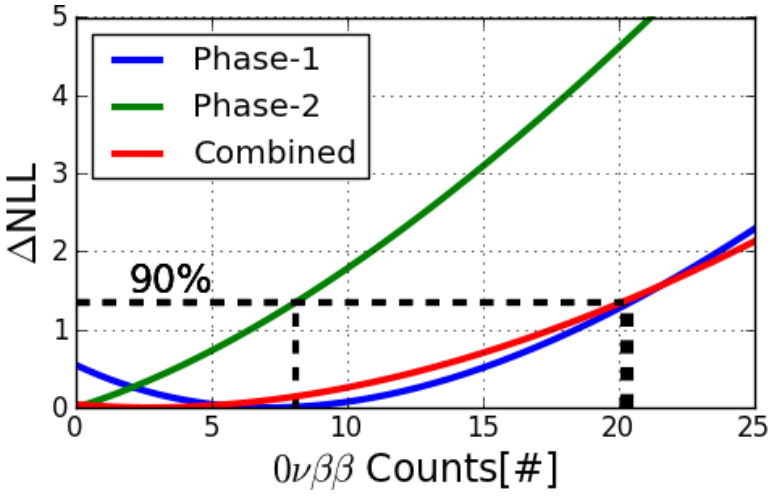

Figure 6.9: Profile likelihood of the number of $0\nu\beta\beta$ events for each phase as well as the combination of the two phases. The dashed lines represent the 90% confidence upper limit on the number of $0\nu\beta\beta$ events in each case.

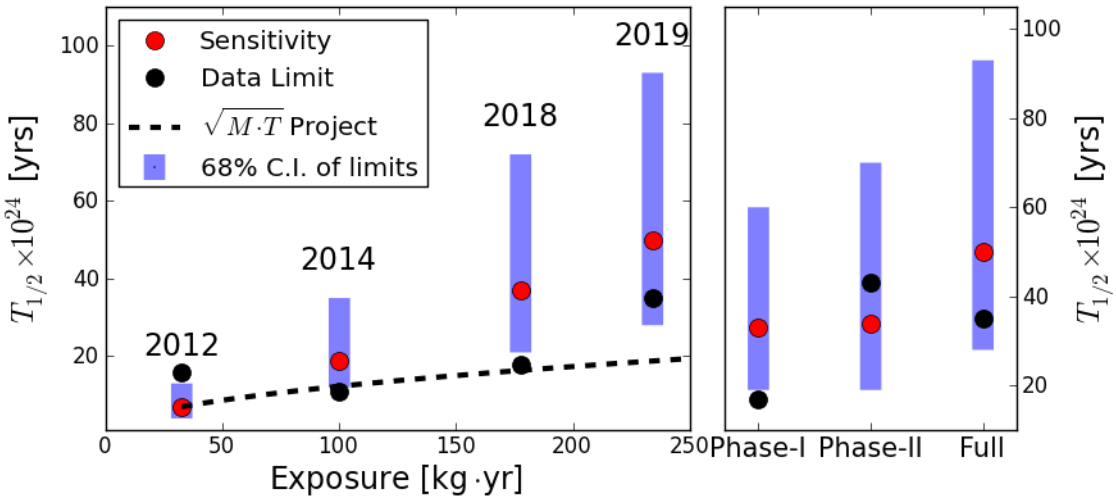

Figure 6.10: Plot showing the evolution of EXO-200 limits and sensitivities for (left) each publication and (right) the individual phases in the final result. Also shown as a dashed line is the expected increase in sensitivity after the 2012 publication of given only an increase in livetime.

$> 1.7 \times 10^{25}$ yrs (4.3×10^{25} yrs) for Phase-I (Phase-II). While the phases have approximately equal livetime, the limit in Phase-I is marginally worse due to an upward fluctuation of observed events. This can be understood by comparing the limit to the expected sensitivity, which as defined in Section 6.2.3 is independent of these statistical fluctuations. Figure 6.10 shows this comparison, from which it can be seen that the observed limits of each phase fluctuate around the median limit within the expectation from the sensitivity distribution. Also shown in Figure 6.10 is the expected $\sqrt{t}$ increase in sensitivity after the first publication of data [80] assuming only the increase in livetime. As a result of numerous improvements to both hardware and analysis, the actual increase in sensitivity for the final result is ~60% larger than this extrapolation.

Finally, this result can be translated int an upper limit on the Majorana neutrino mass using Equation 2.21. This results in a limit of $\langle m_{\beta\beta} \rangle$ <(93-286) meV, using the nuclear matrix elements of [41–45] and the phase factor from [40]. Given this result, Figure 6.11 shows the remaining regions allowed for $\langle m_{\beta\beta} \rangle$ assuming no quenching to g_A and the standard see-saw mechanism.

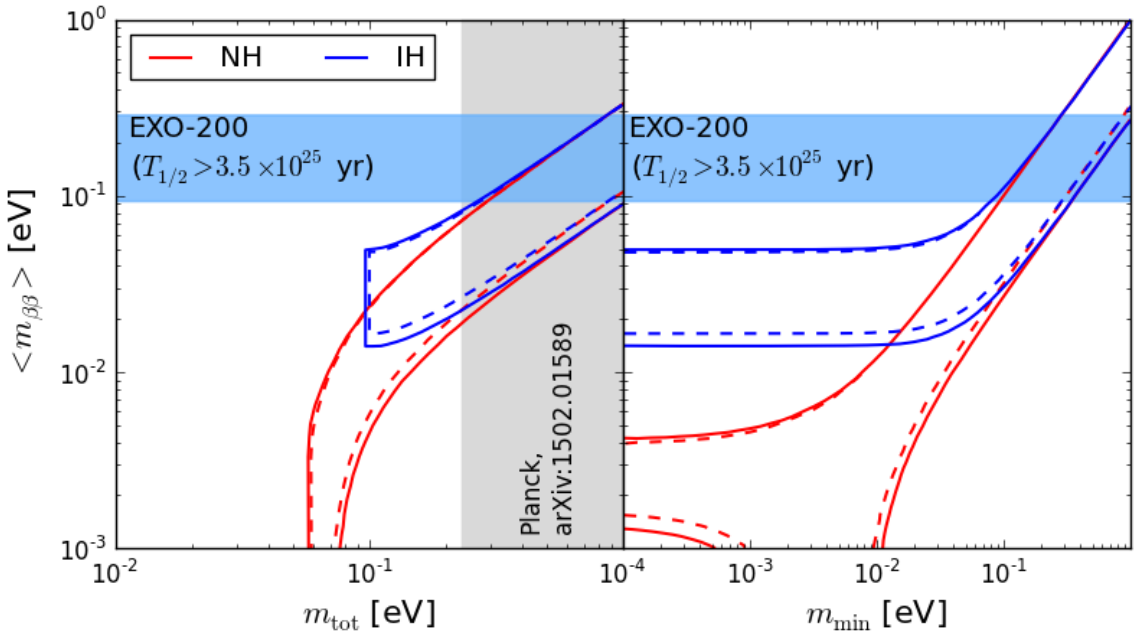

Figure 6.11: Allowed values of $\langle m_{\beta\beta} \rangle$ as a function of (left) sum neutrino mass and (right) lightest neutrino mass for both the IH and NH. The upper limit derived from the final EXO-200 result is shown as a shaded blue band.

Chapter 7

nEXO

7.1 Overview

In order to probe longer half-lives of $0\nu\beta\beta$, a larger and more sensitive detector is required. Following the success of EXO-200, the nEXO experiment is a proposed next-generation LXe TPC that will continue the search for $0\nu\beta\beta$. A detailed Conceptual Design is described in [83] but the following provides a brief overview relevant to Chapter 8.

Similar to EXO-200, the nEXO experiment will consist of a single phase liquid xenon TPC with much of the design inheriting directly from EXO-200. A conceptual sketch of the nEXO detector is shown in Figure 7.1. While the underground facility has yet to be determined, this sketch is shown for the underground Cryopit at the Sudbury Neutrino Observatory Laboratory (SNOLAB) in Sudbury, Ontario, Canada. The TPC is held in a cryostat filled with HFE-700 and the cryostat is held inside a giant tank filled with purified water, which acts as a shield for γ's and fast neutrons. In addition, the water tank is instrumented with PMTs so that it can be used as a cosimic-ray muon veto.

Also shown in Figure 7.1 is a cross-section of the TPC design. The TPC will

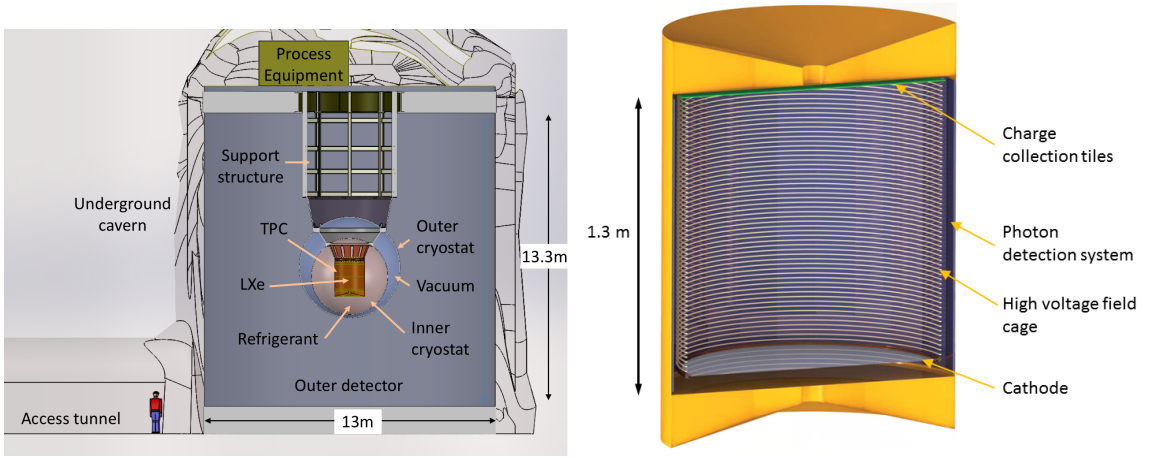

Figure 7.1: (left) Conceptual sketch of the nEXO experiment deployed in the SNOLAB Cryopit. (right) Cross section of the nEXO TPC highlighting the key features.

than be filled with ∼5 tonnes of LXe, enriched to ∼90% in the isotope ^{136}Xe. Instead of two back-to-back TPCs with a shared cathode, as used in EXO-200, the detector will consist of a single TPC with just one cathode and anode. This eliminates the need for a central cathode, producing a single homogeneous detector which reduces radioactive backgrounds from the center of the TPC and allows for nEXO to take full advantage of γ-attenuation in LXe.

Charge produced by ionizing radiation will be collected and read out by an anode plane composed of an array of silica tiles, patterned with crossed metallic strips. Studies to optimize the design of the tile dimensions are on going, but currently each tile is proposed to be 10×10 cm^2 and have a channel pitch of 3 mm. In addition, this design does no include a shielding grid. As a result, signal is induced on the anode through an electrons entire drift. The choice to use arrays of tiles is motivated by the ability to avoid a substantial tensioning frame, which can withstand temperature cycling while supporting the wire grid. Such a design is complex and could result in a large material cost. In addition, this results in a modular design, allowing for each tile to have its own front end and digitization electronics.

The light readout of nEXO will use an array of VUV-sensitive Silicon PhotoMultipliers (SiPMs) in place of the APDs used in EXO-200. Since the charge-tiles are

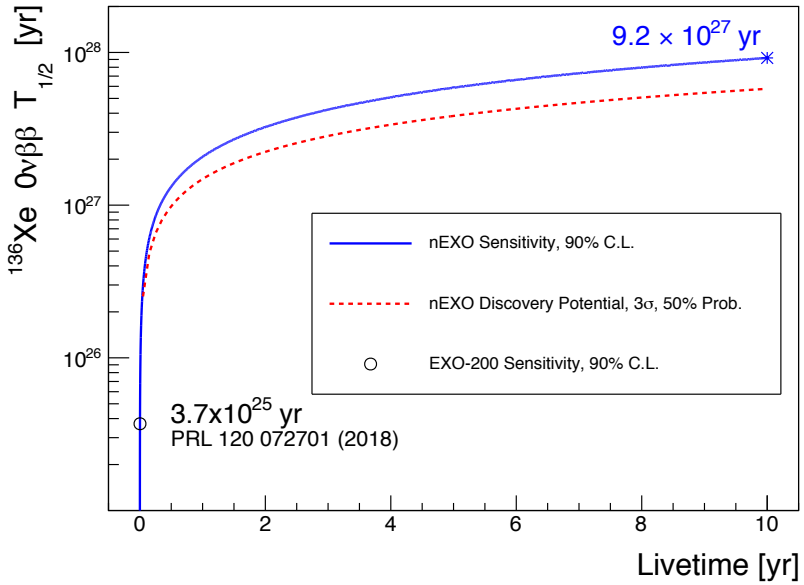

Figure 7.2: Projected Sensitivity (blue line) and Discovery Potential (dashed red line) of nEXO as a function of the detector livetime. After 10 years of data taking nEXO will achieve a sensitivity of 9.2×10^{27} years. Taken from [114].

opaque, no SiPMs are installed behind the cathode. Instead, the SiPMS are installed behind the field shaping rings and arranged around the edge of the TPC in a barrel configuration. This design optimizes the coverage and provides sufficient light collection efficiency to achieve the goal of 1% energy resolution at $Q_{\beta\beta}$.

Using the currently proposed design [83], nEXO is expected to achieve a half-life sensitivity of 9.2×10^{27} years after 10 years of operation, as shown in Figure 7.2, corresponding to a discovery potential of 5.7×10^{27} years [114]. This represents more than a $\times 100$ improvement over the current limit by EXO-200 presented in Chapter 6. Following the procedure outlined in Section 2.3.1, $\langle m_{\beta\beta} \rangle$ sensitivity between 5.7 meV and 17.7 meV, as shown in Figure 7.3.

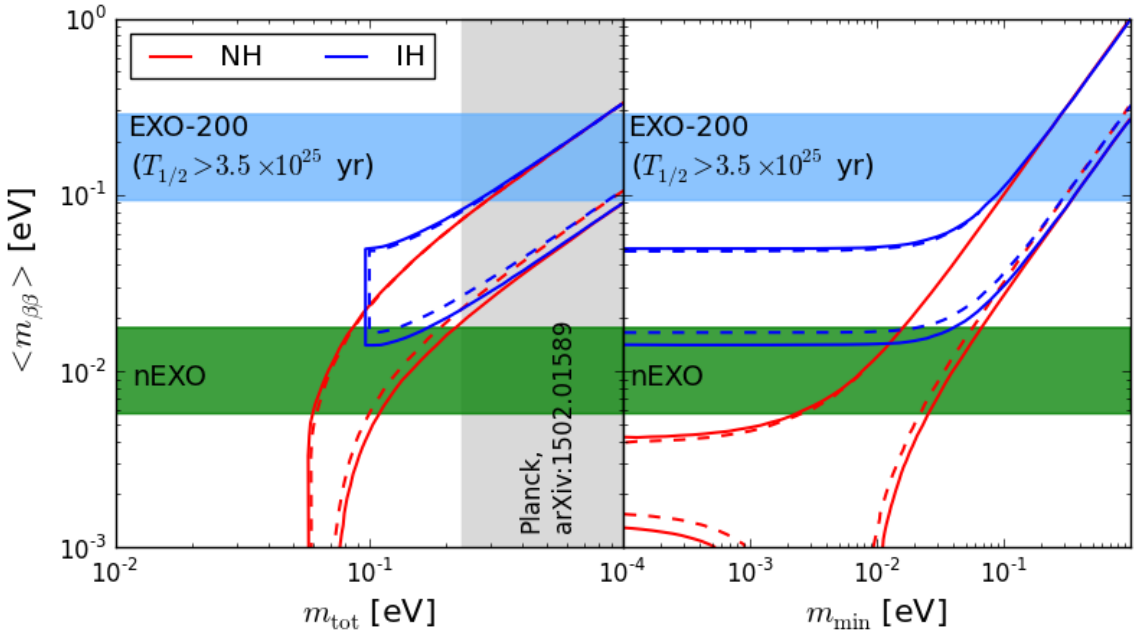

Figure 7.3: Projected sensitivity of nEXO to $\langle m_{\beta\beta} \rangle$ with 10 years of data as a function of the (left) sum neutrino mass and (right) lightest neutrino mass for both the IH and NH. The upper limit derived from the final EXO-200 result is shown as a shaded blue band and the nEXO result in shown in green.

Chapter 8

Characterization of Ionization Tile for nEXO

8.1 Overview

In EXO-200 the ionization signal is measured by two planes of crossed wires where one plane is used as a shielding grid and the other as the charge collection grid [63]. The EXO-200 TPC is approximately 36 cm in diameter while the diameter of nEXO will be over 1 m. At this larger diameter, a rather substantial tensioning frame that can be temperature cycled to ∼165 K would be required, which would pose a challenge to the required radioactivity budget of nEXO. In addition, the larger its diameter, the more vulnerable a crossed-wire design is to ambiguity in reconstructing the position of multiple energy deposits in the detector. Wires are also susceptible to microphonic pickup from environmental noise. For these reasons it was suggested [115] to explore using anode pads as an alternative readout.

As described in Chapter 7, the nEXO collaboration is investigating a segmented anode composed of an array of tiles to avoid these issues. Modular by design, an array of tiles can cover a sizable area, allowing them to be employed by large detectors. Each tile consists of a dielectric substrate covered with an array of conductive strips for collecting charge. In addition, charge tiles can be made using only materials that are either known to be obtainable with extremely low radioactive contamination (such as fused silica), or employed in minimal amounts (i.e. the thin conductive strips). Finally, no mechanically-robust tensioning system is required and because the width of each strip is small compared to the size of the tile, a Frisch grid is not required. A grid-less, tiled anode design is beneficial for an experiment such as nEXO, where a wire tensioning support structure, designed for cryogenic temperature cycling, would contribute significantly to the radioactive backgrounds.

Summarized here are the results, presented in [116], of the characterization the proposed charge tile with a dedicated LXe test setup. The setup consists of a ∼9 kg LXe TPC with a mounted 10 cm × 10 cm tile used to readout the ionization signal. Good agreement is achieved between the measured ionization spectrum of a ^{207}Bi

source and simulations that include the microphysics of recombination in xenon and a detailed modeling of the electrostatic field of the detector. An energy resolution σ/E=5.5% is observed at 570 keV, comparable to the best intrinsic ionization-only resolution reported in literature for liquid xenon at 936 V/cm.

8.2 Stanford Test Stand

8.2.1 Prototype Anode Tile

A 10 cm × 10 cm prototype ionization readout tile which is 300 µm thick was fabricated by the Institute of Microelectronics of the Chinese Academy of Sciences. The tile substrate is a fused-silica wafer with 60 electrically isolated strips (30 "X" strips and 30 orthogonal "Y" strips). Strips are made by depositing layers of Au and Ti onto the silica wafer surface.

Each strip is approximately 10 cm long consisting of 30 square pads, which are 3 mm across the diagonal and daisy-chained at their corners. This geometry maximizes the metallic cover of the substrate, reducing the risk of charge accumulation, and minimizes the capacitance at the crossing between strips. Layers of 1.5µm thick SiO_2 are used at the crossing points of X and Y strips to provide electrical isolation. The capacitance at each crossing is 80 fF, assuming that the conducting structures are 0.5 µm thick gold. This results in a capacitance of 0.57 pF between pairs of crossed strips and a capacitance of 0.86 pF between adjacent parallel strips. In addition a resistance of ∼5 Ω at LXe temperature is expected along each single strip on the tile. A diagram of the tile mounted on a stainless steel support used for testing is shown in Figure 8.1 with details of the strip geometry and orthogonal strip crossing in the insets.

The design of the tile tested here is representative of what is currently proposed for nEXO, although parameter optimization is still under way. The integration concept for the nEXO charge collection plane composed of many such charge tiles covering the anode surface is shown in Figure 8.2. This represents a departure from the crossed wire plane design adopted by EXO-200 and avoids the need to provide a substantial

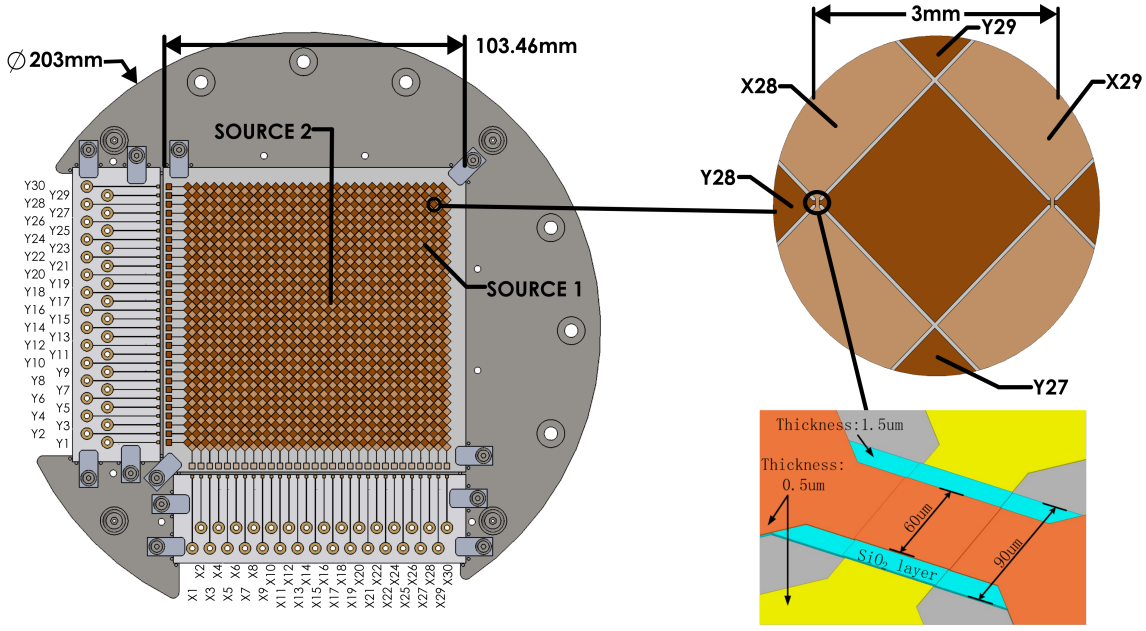

Figure 8.1: Sketch of the prototype charge tile (left), with 60 orthogonal strips, 30 in each direction mounted on a stainless steel support used for testing. The light-colored X strips are vertical and the darker Y strips are horizontal in this view. Each strip is made of 30 3 mm × 3 mm square pads connected at two opposite corners. Also shown (top right) is a detail view of strips on the anode tile. The X and Y strips cross each other at the pad junctions (bottom-right). The metal X and Y layers are separated by a thin layer of SiO_2. This arrangement maximizes overall pad coverage while limiting the capacitance. The projected locations of the two ^{207}Bi sources located on the opposite cathode are indicated. Also shown are the ceramic interface boards mounted on two sides of the tile for testing purposes. Pads on the ceramic interface allow wire bonding to the tile. The ring terminals are connected to feedthrough leading to pre-amplifiers outside of the LXe.

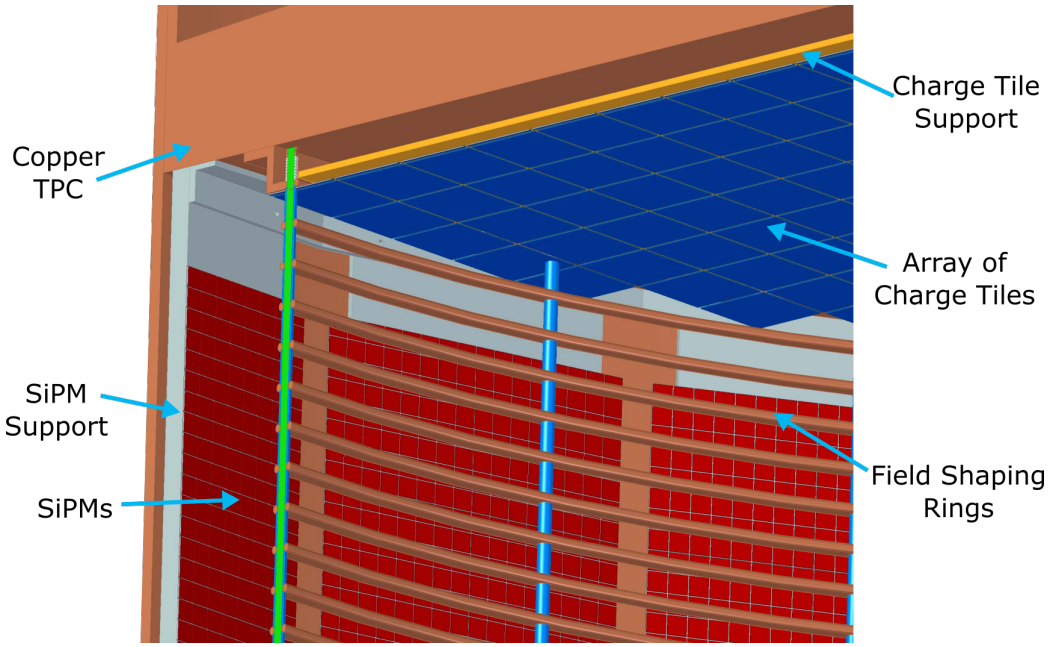

Figure 8.2: Diagram of the integration scheme of charge tiles into the nEXO LXe TPC. The anode is composed of a mosaic of many adjacent charge tiles. Also shown on the outer edge of the copper TPC are UV-sensitive silicon photomultipliers (SiPMs), the proposed method for detecting the LXe scintillation light.

tensioning frame that is both cryogenic compatible and meets the required radioactivity budget. This approach does not use a Frisch grid but offers the advantage of additional channel segmentation to reduce possible ambiguity in reconstructing the position of individual charge clusters in events with multiple charge depositions.

8.2.2 Test Cell

In this test setup, the anode mounting plate is inserted into a liquid xenon time projection chamber (TPC) as shown in Figure 8.3 and Figure 8.4. This TPC is used to characterize the performance of prototype anode tiles in LXe. The TPC is built from a 12 in ConFlat spool-piece and two flanges. The body of the TPC is 304 stainless steel with a 30 cm outer diameter and a 13 cm height. Both TPC anode and cathode mounting plates are approximately 20 cm in diameter. The prototype tile is mounted at the center of the anode plate as shown in Figure 8.1. The drift length can be varied from 18.2 mm to 33.2 mm by changing spacers behind the anode, but was set to the longer length of 33.2 mm for the data presented here.

On the tile, each strip ends with a square pad where the signal can be read out. For this test each channel is wire-bonded to an adjacent ceramic interface board which is also mounted on the anode plate (Figure 8.3). Ring terminals crimped to Kapton-insulated copper signal wires are bolted to through holes on the ceramic interface boards and bring the signals to electrical feedthroughs on top of the xenon cell (Figure 8.4). This solution allows for easy reconfiguration of the strip-readout channel assignments and avoids the use of solder, which could deteriorate the chemical purity of the LXe.

The center of the TPC cathode, 150 mm in diameter, is a photo-etched 302 stainless steel hexagonal mesh, 127 µm thick with 95% optical transparency and 3 mm hexagon size (see Figure 8.4). Below the cathode is a 6 in diameter optical window made with UV grade synthetic silica. Clamped against the face of the optical window and located outside of the xenon volume submerged in the HFE is a cryogenic photomultiplier (PMT), EMI 9921QB, optimized for VUV sensitivity. The PMT collects the xenon scintillation light (peaked at 178 nm) and provides a trigger signal for each

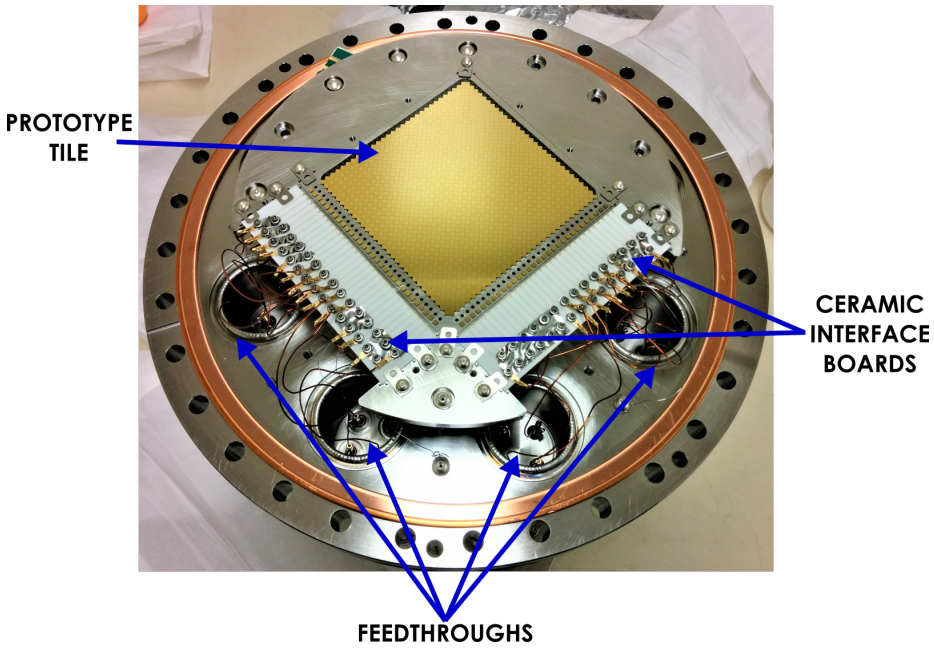

Figure 8.3: Photograph of the anode tile mounted to the top flange of the TPC.

event. The cathode is typically biased at -3110 V resulting in an average drift field of 936 V/cm.

The TPC is filled with ∼9 kg of LXe, which is sufficient to fully submerge the anode tile. Operation of the detector is typically performed at ∼900 torr and ∼168 K. Mimicking the cryogenic scheme used in EXO-200 and proposed for nEXO, the TPC is submerged in HFE-7000 cooling fluid [70] to maintain good temperature uniformity. The HFE-7000 is contained in a cryostat and cooled via a large copper heat exchanger ("cold plate") immersed in the HFE-7000 and cooled with liquid nitrogen flowing through internal tubing. The cold plate is shaped as a semi-cylindrical surface that wraps around the TPC chamber, purposely azimuthally asymmetric to force convection of the HFE-7000 fluid. The temperature of the cold plate is measured by three thermocouples placed at different locations on its body. One of them regulates the flow of the liquid nitrogen via a PID feedback loop managed by a LabVIEW application [117]. The temperature of the cell is recorded using additional thermocouples attached at three locations on the outside of the TPC. The pressure is also monitored at two locations in the Xe system using two Model 121A MKS Baratrons

[118]. The HFE-7000 dewar and TPC are shown in Figure 8.4.

Before liquefaction the TPC is initially pumped down to $\lesssim 10^{-6}$ torr, which results in a small enough oxygen contamination to achieve the required LXe purity for drifting electrons over the TPC length. The empty cell and the HFE-7000 are cooled down and held at the base operating temperature of 168 K. When filling the detector, the xenon is passed through a SAES MonoTorr model PF3C3R1 getter [119] to remove electronegative impurities and transferred into the cold cell, where it enters through a line mounted on the side and liquefies. Continuous recirculation through the purifier is not currently needed since there is no noticeable degradation in the electronegative purity during the typical data taking run (∼2 days). A limit of ∼150 µs on the electron lifetime is estimated by observing the variations in the location of the 570 keV peak as a function of drift time in data and simulation. This lifetime is sufficient to drift and collect charge in the current setup, which has a maximum drift time of 33.2 mm. In nEXO the expected electron lifetime is >10 ms but the ultimate drift length in nEXO is expected to be ∼625 µs [83, 114] which is significantly larger than that achieved here. The approximate ratio of lifetime to drift length is still comparable making this representative of the expected purity effects in nEXO. A MKS 1479A Mass-Flow Meter [120] mounted along the room-temperature portion of the fill line is used to determine the volume of xenon flowing into the cell.

A ^{207}Bi source is plated onto a platinum-rhodium wire, 127 µm in diameter and woven onto the cathode mesh in two locations as shown in Figure 8.4. Projected onto the anode tile the first source is approximately located at the crossing of channels X28 and Y24 and the second at the crossing of channels X17 and Y17, as shown in Figure 8.1. Each source emits two primary gamma rays of 570 keV and 1064 keV. The event rate (before any veto) measured as the PMT trigger rate is 960 Hz. Each source contributes approximately half of the activity.

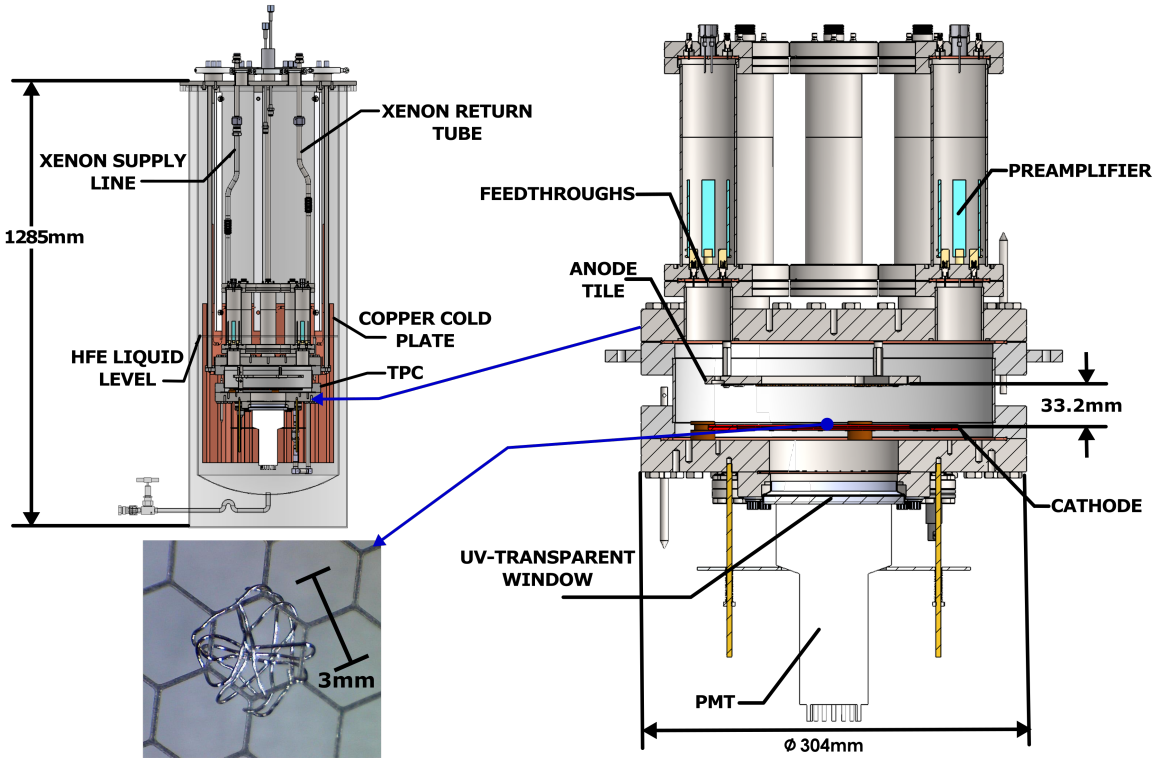

Figure 8.4: Sketch of the TPC mounted inside the HFE-7000 dewar (top left) showing the key cryogenic components. Also shown is a cut out view of the TPC (right) and a photograph of one of the Bi^{207} sources woven into the cathode mesh (bottom left).

8.2.3 Data Taking

Each charge-readout channel from the tile is instrumented with a custom charge-sensitive preamplifier, based on the architecture shown in Figure 8.5. This architecture was chosen because of its low noise characteristics (200 electrons RMS at a peaking time greater than 3 µs for 24 pF strip capacitance and an operation temperature of ∼168 K), compact footprint, and the ability to drive cables over long distances. Power consumption was not a factor, nevertheless, each preamp meets the electrical specifications at less than 45 mW power consumption. An additional stage of gain is placed after the preamplifier in order to match the dynamic range of the signal processing electronics. The full functionality of the preamplifier is described in [121].

The preamplifiers are installed in eight leak-tight cans filled with nitrogen gas and mounted just above the TPC feedthroughs to minimize the distance from the strips to the front-end electronics, which in our case is approximately 10 cm. The cans are partially submerged in HFE to keep them near the designed operating temperature of ∼168 K. No temperature sensors are installed to measure the temperature inside the cans so the exact operating temperature is not precisely known and the electronic noise is expected to grow proportionally to the temperature. The output of each preamplifier is further amplified 10-fold by a Phillips Scientific 776 amplifier installed ∼2 m away and outside the cryostat. The output signal from the second stage of amplification is then sent to a digitizer were the waveforms are readout and stored on the DAQ computer.

Digitized waveform data is collected at 25 MS/s with two 16-channel 14-bit SIS3316 digitizers [122]. Thirty two channels are recorded: thirty charge channels, the PMT signal, and a pulser signal used for noise measurements. The signal from the PMT is split into two: one signal enters a discriminator to generate an event trigger, the other signal is shaped with an Ortec 672 spectroscopy amplifier with 500 ns shaping time, then digitized. Since there are sixty charge channels on the tile and only thirty digitizer channels available, some strips are ganged together on the ceramic interface board inside the TPC. Data from all 32 readout channels is collected during each event. A prescaler vetoes 4.3 ms after every PMT trigger, to keep the data to a rate

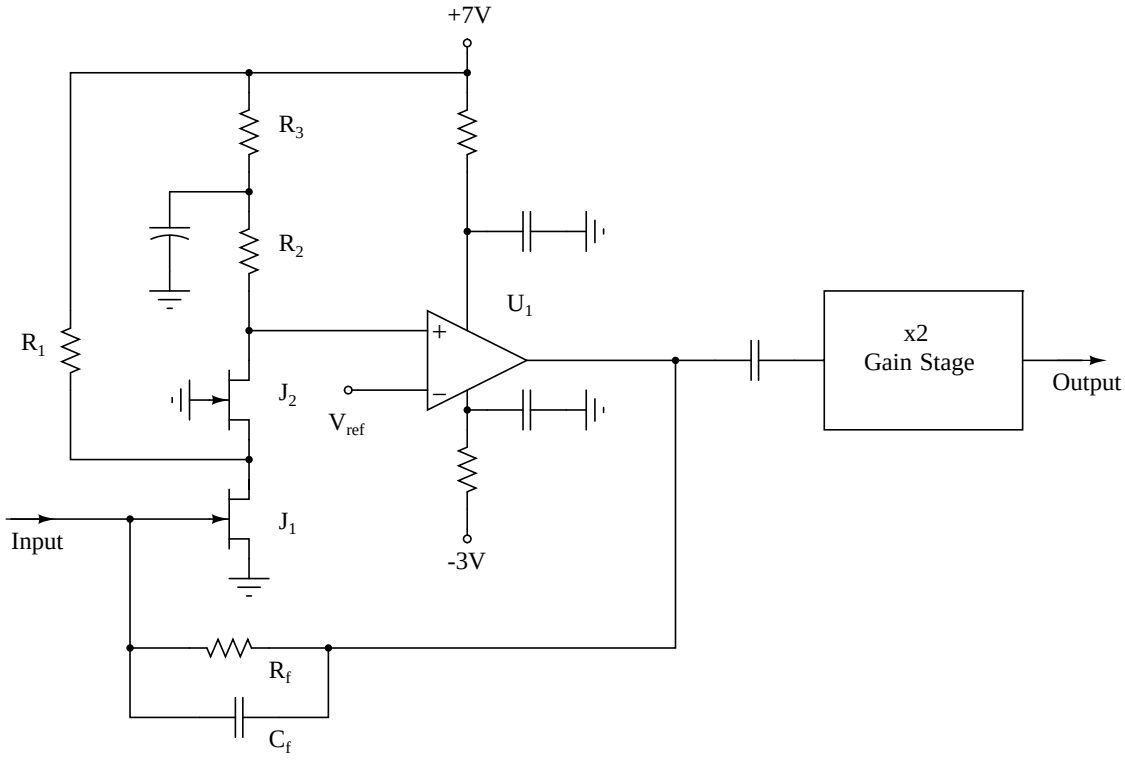

Figure 8.5: Schematic diagram of the charge-sensitive preamplifiers used in this study. The input JFET, J1, is a commercial off-the-shelf BF862, at about 5 mA drain current. The current is set mostly by the series R2 and R3, and is boosted as needed by R1. The transistor J2 acts as a cascode to improve dynamic performance and increase the open-loop gain at medium to high frequencies. The presence of the operational amplifier U1 guarantees a very large value for the open-loop gain of the circuit to minimize potential crosstalk among adjacent strips. The second stage of gain is optional, but in this case was used to match the DAQ dynamic range.

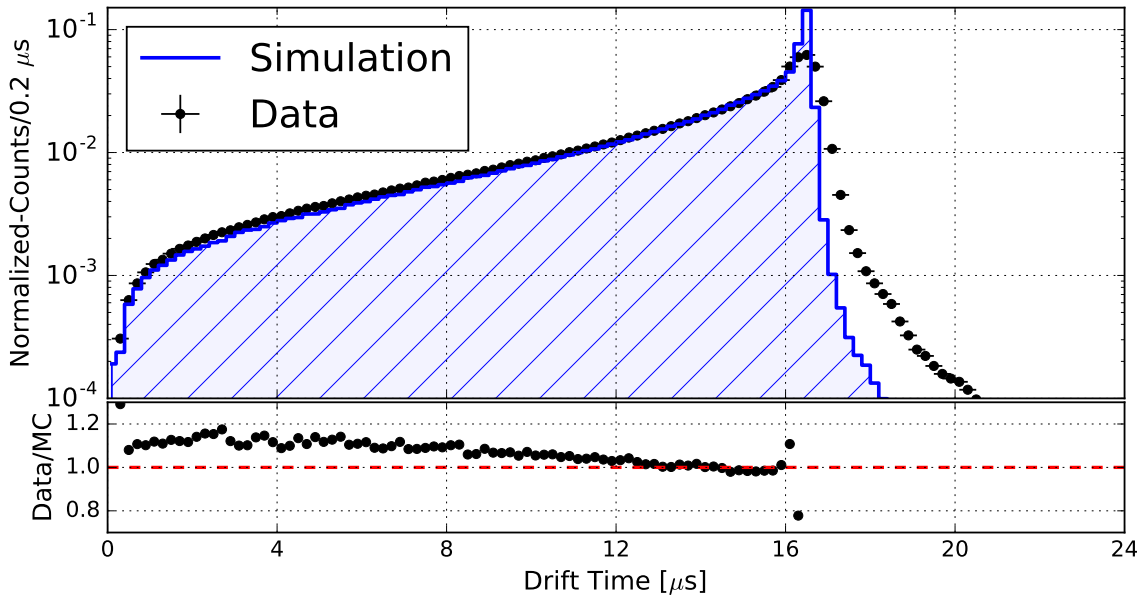

Figure 8.6: Spectrum of drift times in both data (points) and simulation (shaded area) with the ratio of data to simulation shown in the lower panel. The peak at the cathode (16.6 µs) has a larger tail to the right in data that is not reproduced in simulation. This discrepancy is attributed to the detailed geometry and electrostatic properties of the plated source and cathode mesh which are not fully captured in the simulation.

of about 1 GB per minute, the maximum that is manageable for the data acquisition computer.

At a field of 936 V/cm, the drift velocity of electrons in LXe is approximately 2 mm/µs as measured in [90]. This gives a maximum drift time of 16.6 µs at the drift length of 33.2 mm. This is later confirmed by comparing the observed peak drift time in data to that in simulation, as shown in Figure 8.6. The digitized waveforms are 42 µs long (1050 samples) and include 11 µs (275 samples) before the PMT trigger. A sample event is shown in Figure 8.7, which also shows how strip channels are grouped for readout.

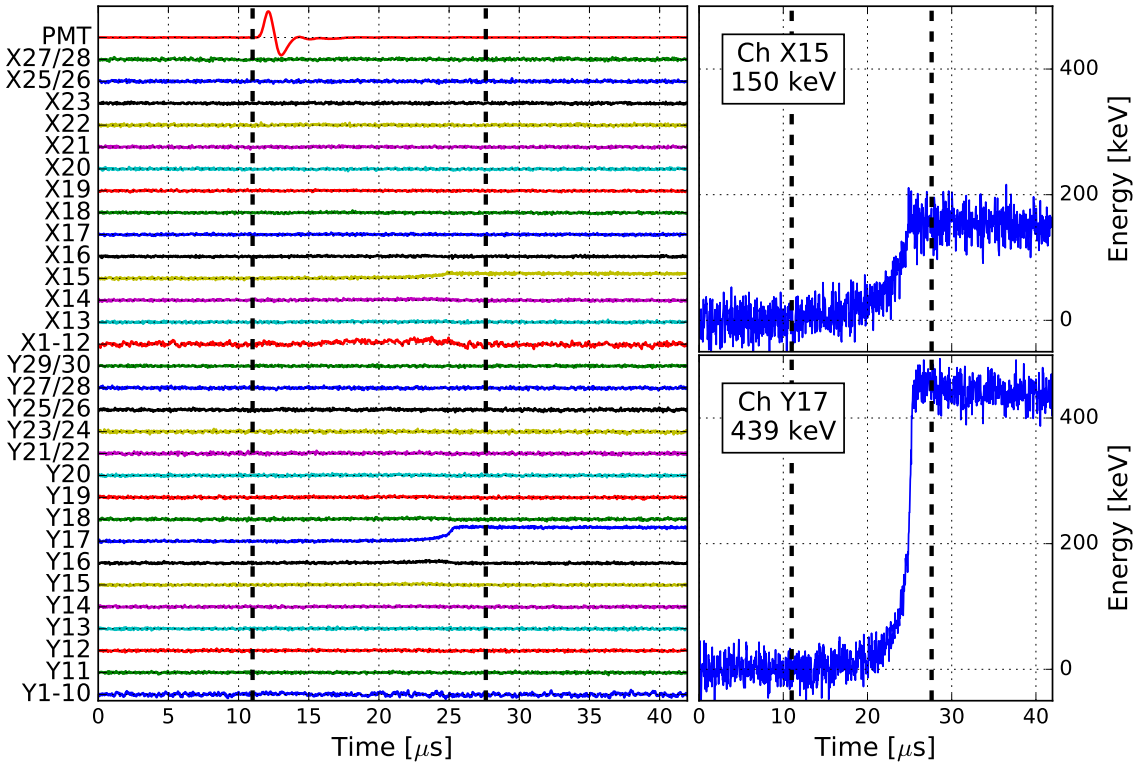

Figure 8.7: A sample event that passed all of the data selection cuts and fell in the Bi-207 peak at 570 keV. On the left the WFs on all charge channels and the PMT are shown, with the grouping of strips in each channel labeled on the vertical axis. Two single-strip charge channels received energy deposits above thresholds (X15 and Y17) resulting in a total energy deposit of 589 keV. The vertical lines at 11 µs represents the trigger from the PMT. The vertical line at 27.6 µs is the maximum expected drift time for an event that started at 11 µs. On the right is a zoomed in view of the 2 charge channels above threshold. Refer to Figure 8.1 for locations of the channels.

8.2.4 ^{207}Bi Source

The ^{207}Bi source attached to the cathode decays via electron capture to ^{207}Pb with a half-life of 32 yrs[123]. This decay is followed by the emission of both γ's and Internal Conversion Electrons. The decay level diagram as well as the most probable emissions lines are shown in Figure 8.8 and Table 8.1 respectively. From this is can be seen that the most probable chain results in the approximately simultaneous emission of γ's with energy 1064 keV and 570 keV.

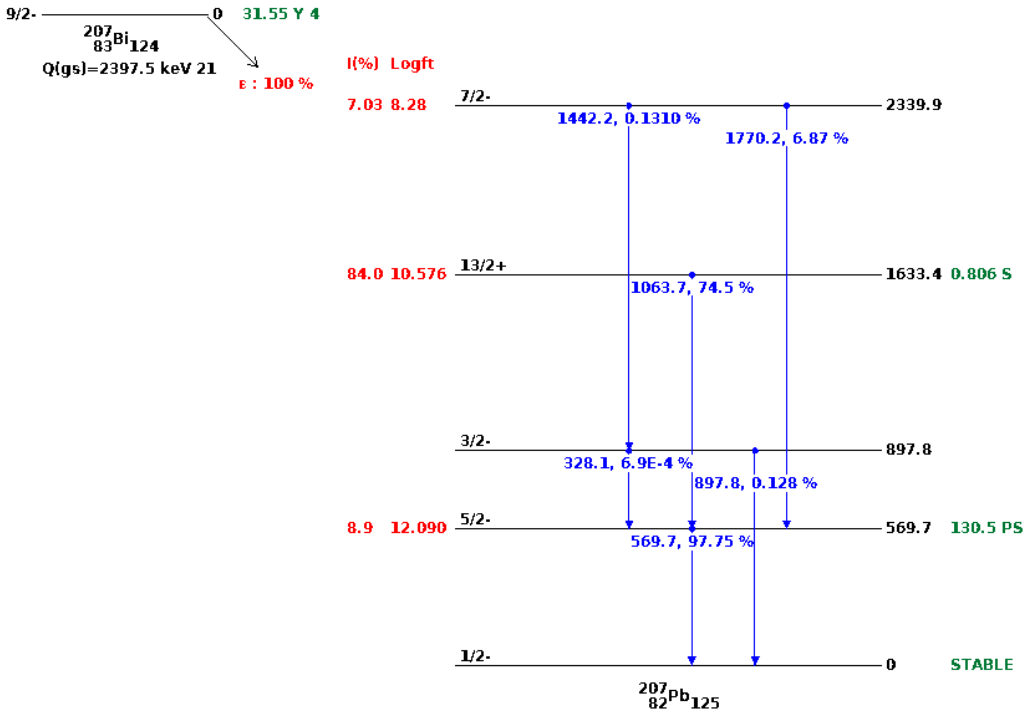

Figure 8.8: ^{207}Bi decay level scheme from NuDat2 [123].

8.3 Detector Simulation

The simulation of the detector is split into two independent stages. The first stage uses a GEANT4-based application [76, 124] in addition to the Noble Element Simulation Technique (NEST) model [99] to parametrize the geometry and determine the number

Table 8.1: Prominent emission lines from ^{207}Bi decay (Intensity>0.1%) [123].

Electrons			Gammas		
Process	Energy[keV]	Intensity[%]	Process	Energy[keV]	Intensity[%]
Auger L	7.97	54.4	XR	10.6	33.2
Auger K	56.7	2.9	XR	72.8	21.4
CE K	481.7	1.537	XR	75.0	35.7
CE L	553.8	0.442	XR	84.5	4.31
CE M	565.8	0.111	XR	84.9	8.27
CE K	975.65	7.08	XR	87.3	3.02
CE L	1047.8	1.84	γ	569.7	97.75
CE M	1059.8	0.44	γ	897.8	0.128
			γ	1063.7	74.5
			γ	1442.2	0.1310
			γ	1770.2	6.87

and location of ionization electrons and scintillation photons produced by particles interacting in the detector. The version GEANT4.10.2.p02 is used for this stage of the simulation. The second stage of the simulation uses the output of the first stage to simulate the signals produced by drifting electrons collected on the strips and the electronics response of the detector.

8.3.1 GEANT4/NEST Simulation

A detailed description of the detector geometry is input into a GEANT4 application, including the TPC Vessel, the PMT window, the internal components (tile, interface boards, cathode, the source, etc.), and the surrounding HFE cryostat.

Particles interacting with the LXe deposit energy by producing both scintillation light (178 nm) and electron-ion pairs (ionization). Due to recombination of a fraction of electron-ion pairs the charge and light yield of individual events of a given energy

and drift field display anti-correlated fluctuations that conserve the total deposited energy [68]. The microphysics that determines the relative amount of energy going into each channel is not currently supported by the standard GEANT4 package. To model this response NEST was used to accurately simulate both the ionization and scintillation response for different detector configurations. The relative light and charge yield varies with the local electric field in the region of the energy deposit, which in turn determines the amount of recombination. To account for this, a map of the electric field magnitude was produced from a cylindrically symmetric COMSOL [74] simulation of the liquid xenon volume of the detector. This electric-field map is used by the simulation to provide the correct magnitude of the electric field to NEST at each point in space. The field map, which neglects the complicated geometry of the cathode mesh by modeling as a flat sheet, is shown in Figure 8.9. The drift field in the region under the tile, assumed to be circular, varies from 890 V/cm to 990 V/cm. This results in an average yield of ∼25 k electrons and ∼15 k photons for gamma's in the 570 keV peak.

Since the cuts described in Section 8.4.2 constrain events to be in the center of the detector, only the ^{207}Bi source located under the center of the anode was simulated to produce the energy spectrum for the single-strip channels shown in Figure 8.17. The complicated geometry of the ^{207}Bi source is approximated as a circular disk of radius 2 mm in the plane of the cathode.

8.3.2 Signal Simulation

Ionized electrons produced in the first stage of simulation that fall within the drift region of the detector are diffused radially according to their drift time, using an electric-field-dependent transverse-diffusion coefficient determined from [79]. Currently no diffusion in the longitudinal direction is included in simulation. The diffused electrons are then binned into voxels with sides of length 530 µm in the x and y directions and 80 µm along the drift direction, z. The x and y dimensions of the voxels were chosen to minimize processing time while preserving signal quality; the z dimension is equivalent to one sample of the 25 MS/s digitizer for a drift velocity

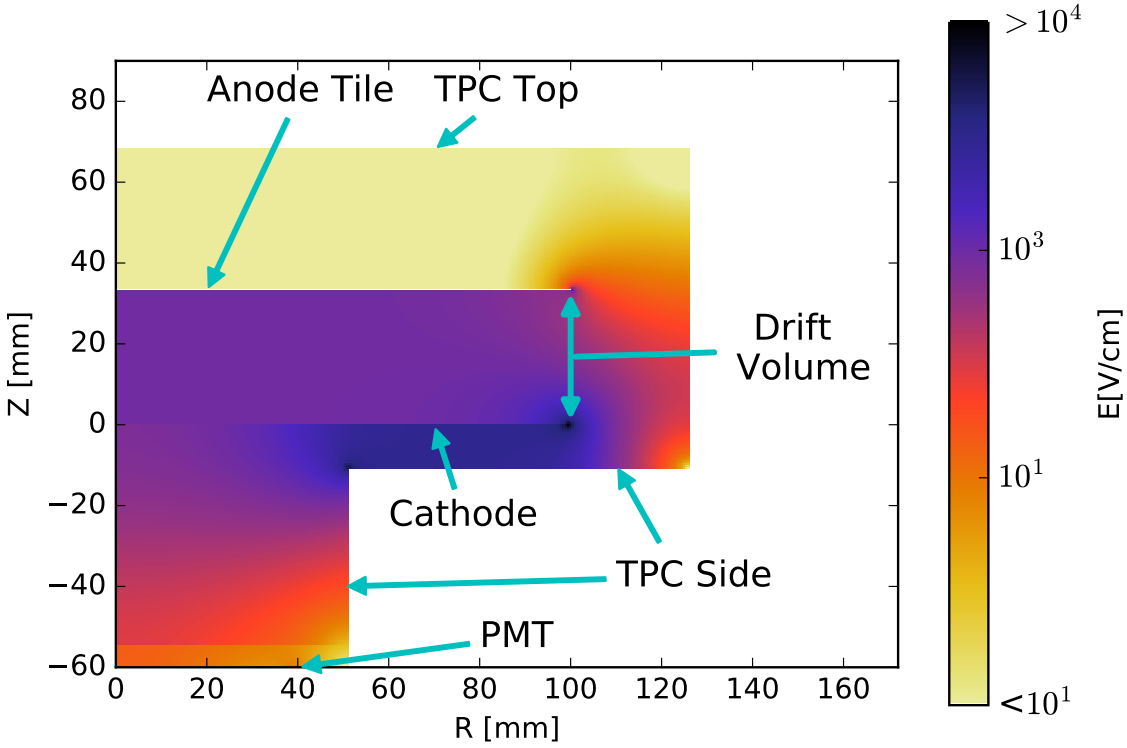

Figure 8.9: Azimuthally-symmetric map of the electric field in the liquid xenon volume. The cathode is at z=0 and the anode is at z=33.2 mm. The cathode is biased at -3110 V while the anode and the TPC walls are at ground.

of 2 mm/µs. Each voxel is tracked as it is drifted from the interaction location to the charge collection tile assuming a uniform drift velocity along the drift axis. The approximation of perfectly parallel field lines is valid in the bulk of the LXe but not near the cathode, where the non-uniform geometry of the mesh as well as the source wire shown in Figure 8.4 cause field distortions. This results in some divergence in the field lines and subsequently the path of electrons, broadening the drift time of events in data near the cathode. From Figure 8.6 this smearing occurs as expected near the cathode where the data sees a much broader peak of maximum drift times. This disagreement between data and simulation motivates the fiducial cut to remove events near the cathode. Diverging field lines are also expected near the edges of the drift region but events in those regions are rejected from the current analysis which only looks at events which hit central strips. For every time step of 40 ns (0.08 mm), the charge induced on each readout strip is determined using an analytical calculation of the charge induced on a square pad in an infinite plane electrode arrangement. The charge induced on a single strip is calculated as the sum of the induced charge on each pad comprising the strip. Corrections are applied to account for electrostatic effects inside the TPC that affect the development of the charge signal, such as cathode suppression and ion screening. A summary of these effects is shown in Figure 8.18 and the details of the calculation are described in Appendix A.

Waveforms for each charge readout channel are then produced using a charge propagation simulation to track electrons from their initial deposition location to their final collection point. The waveforms are sampled at 25 MS/s and contain 1050 samples each, to match the data measured from the LXe setup. In order to make simulated waveforms more realistic, noise waveforms are recorded with solicited triggers throughout data taking. A sample of these solicited waveforms is superimposed on those generated in simulation from drifting charge resulting in waveforms that more accurately resemble what is recorded during data collection. Simulated charge waveforms are then processed using the same reconstruction algorithm as data to find signals and determine event energies. The effect of the electronics used here (Figure 8.5) on the waveform shape is negligible and is not included in the simulation.

Although the first stage of the simulation produces both ionization and scintillation signals, the latter is not propagated to the signal generation stage of the simulation since the light response of the detector is currently used solely as a trigger and not for measuring event energy. This is the result of the single PMT being heavily shadowed by the layout of the charge tile and the amount of scintillation light it detects limits its use for an energy measurement. Work is currently in progress to use an array of high-QE, cryogenic SiPM to detect the LXe scintillation light and overcome this limitation.

Current electronic simulations do not take into account any cross talk which may be present between crossed charge channels. Dedicated studies of this cross talk showed evidence of small levels of signal contamination. The observed signals resulting from these cross talk studies were shown to mimic induction only like signals. For the current analysis which only looks at collection signals the impact of cross talk is minimal, but the treatment of induction signals, particularly for a grid-less design, will be important for the future studies.

8.4 Charge Only Analysis

8.4.1 Event Reconstruction

Once acquired, the digitized data is processed to extract parameters such as waveform energy, PMT amplitude, pulse rise times and delay between PMT trigger and charge collection. Each event contains 30 charge-channel waveforms. Data analysis to extract a good quality energy measurement proceeds in two phases. First the output of each channel is analyzed separately to look for individual signals. Depending on the results of this first pass, channels are then grouped together to determine the total energy deposited in the LXe. In the initial signal finding stage each charge-channel waveform is processed as follows:

1) The first 275 samples of each charge waveform, which are before the trigger, are used to compute the mean and RMS of the waveform baseline. The mean baseline value is then subtracted from the waveform.

2) The waveform is corrected for the exponential decay of the signal, using time constants measured for each charge-sensitive preamplifier. The decay times of the preamplifiers are long compared to the trace lengths, ranging from 200 to 500 µs, making this corrections small.

3) The last 275 samples of the waveform are averaged. Since the baseline has already been subtracted, this plateau is the uncalibrated charge energy, in ADC units.

4) The energy is multiplied by a calibration constant determined from an energy spectrum fit, examples fits for a few channels can be seen in Figure 8.12.

5) An energy threshold is applied by selecting channels with a measured energy that is more than five times the baseline RMS determined in Step 1. The average RMS noise of each channel is shown in Figure 8.10. The energy of all channels above this threshold are added to produce a total event energy.

An example waveform for a channel with an observed signal and the results of the reconstruction algorithm are shown in Figure 8.11.

After waveforms from each calibrated channel in an event are processed, some event-level information is calculated. The energy from all channels above the energy threshold are added to determine the total charge energy for the event. In addition, a sum waveform is produced by adding all charge signals above threshold together. The time to reach 95% of the sum signal height is then estimated. The difference of this risetime and the trigger time are than used to determine the drift time of the event.

8.4.2 Analysis Cuts

A set of data quality cuts are then applied to exclude events if they meet any of the following criteria:

1) Any waveform reaches the maximum or minimum of the ADC range. The range of all charge channels are > 7 MeV, which is well outside the range of energies studied here.

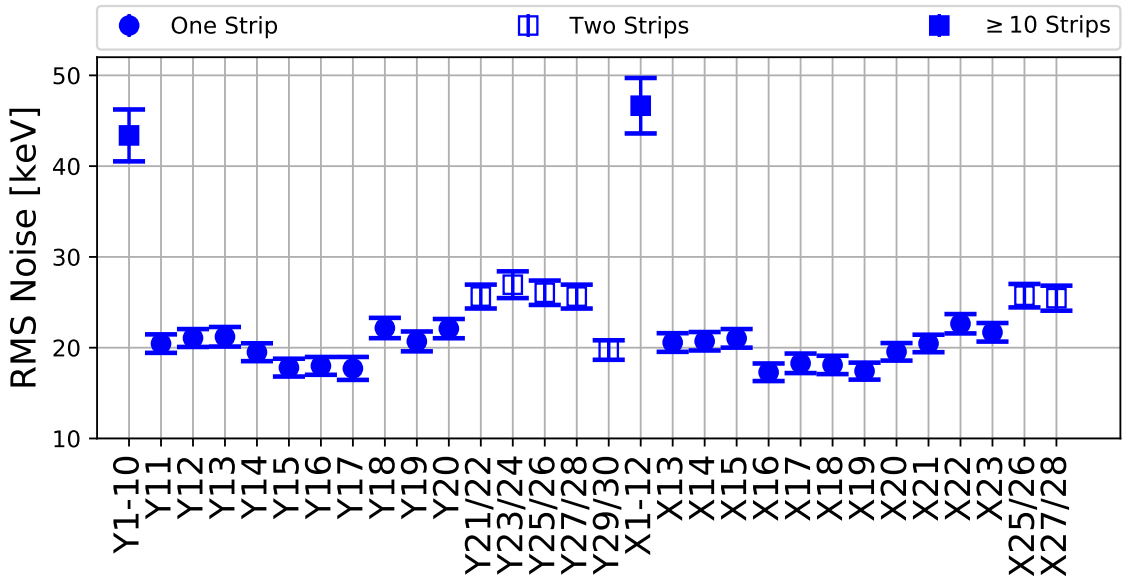

Figure 8.10: RMS noise, measured from 275-sample baseline, for each charge readout channel. Each data point is the mean of the RMS distribution; the error bars are the RMS width of the RMS distribution. Channels with more than one strip have, on average, larger RMS noise because of their larger capacitance. In general, the RMS noise grows with the square-root of the number of strips ganged together in the readout channel.

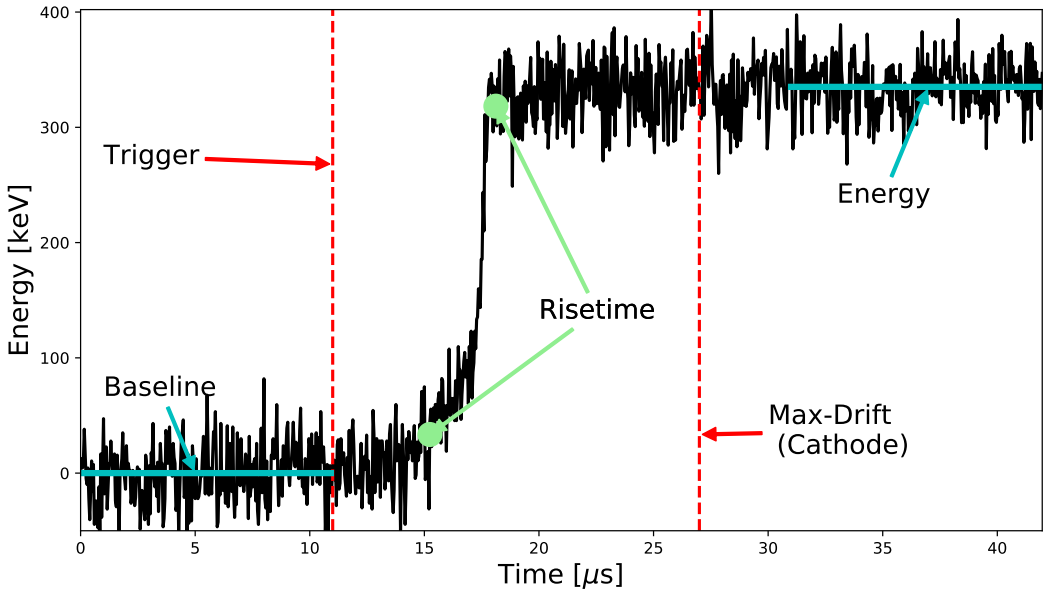

Figure 8.11: Sample signal in data after correcting for the decay time, with an illustration of the reconstruction algorithm. The location of the trigger and the maximum expected drift time are shown as dashed red lines. The samples used for the basline/energy measurement along with the result of this measurement are shown as solid light-blue lines. The 95% risetime is marked in green.

2) The PMT waveform shape has poor agreement with a template shape. This cut eliminates pileup events where, *e.g.*, two decays of ^{207}Bi occur close in time.

3) In any channel, the RMS noise computed with the 275 waveform samples used for either baseline averaging or energy measurement has a 5 sigma or greater excursion from the usual distribution of values. This cut eliminates noise bursts and events where the preamplifier or amplifier are at its maximum value.

In total these checks remove ∼10% of the total triggers included in the data set analyzed here.

In addition, the following cuts were applied to the spectra of both data and simulation to avoid pathological events either due to the small dimensions of the TPC or which are not properly modeled by the simulation:

- The event is required to hit exactly one X channel and exactly one Y channel, to ensure that its location is well known in the TPC. This eliminates events happening in the xenon volumes above the ceramic boards, which are not described by the simulation. This also removes events which have a light signal but no charge collected, which occurs when events interact with the LXe behind the cathode. This should roughly occur ∼50% of the time since the source emits radiation in all directions.

- The X and Y channels are required to be single-strip readout channels. The single-strip channels have the best energy resolution since they have the smallest area and correspondingly smallest capacitance. As a result of the chosen wiring map, this requirement constrains the event to be near the center of the detector, far from the ceramic boards were the drift field is less uniform and energy can be lost to charge depositing on the ceramic instead of the tile (see Figure 8.1). This cut, along with the number of signals cut, remove ∼88% of the remaining events which passed the data quality cuts.

- The time delay between the PMT trigger and the arrival of the charge event is required to be >9.5 µs. This cut rejects events closer than 19 mm from the

anode, where the signal height can vary by more than 1% due to electrostatic effects described in Section 8.3.2 and Appendix A (see Figure 8.18). This avoided introducing additional errors from imperfect modeling of these effects.

- The time delay between the PMT trigger and the arrival of the event is required to be $\leq$16.1 µs. Since the maximum expected drift time is 16.6 µs, this cut rejects energy deposits within 1 mm of the cathode, where the electric field is distorted, an effect attributable to the 3 mm spacing of the cathode mesh. This cut also eliminates pileup events (see Figure 8.6).

The drift time cuts remove an additional ∼56% of the total events that passed the other cuts. In total ∼5% of the total data included in the full data set is left for the final analysis.

8.4.3 Calibration

In order to achieve the optimal energy resolution the individual response of each charge channel is estimated to remove any variations. This is done using the γ peak at 570 keV in the ^{207}Bi energy spectrum. This peak is used since the range of a 570 keV γ is ∼1 cm, allowing for the possibility of a single channel to collect the entire energy deposit. In addition, as shown in Table 8.1 this peak has the highest intensity at ∼98%.

The calibration is done using a single-channel cut to ensure that there is no missing energy on the channel being calibrated, at least up to deposits buried beneath the RMS noise of neighboring channels. This selection produces the most prominent peak in each channel's energy spectrum. The 570 keV peak is than fit with a Gaussian function added with an exponentially-decaying background to determine the peak location. Figure 8.12 shows the result of this procedure for 4 channels in the detector and the overall calibration coefficients for each channel are shown in Figure 8.13.

In addition, the decay times of the preamplifiers are measured by fitting the pulse to an exponentially decaying function. The distribution of decay times are than fit to a landau distribution as shown in Figure 8.14 to extract the mean decay time for each channel. The resulting decay constant for each channel are shown in Figure 8.15.

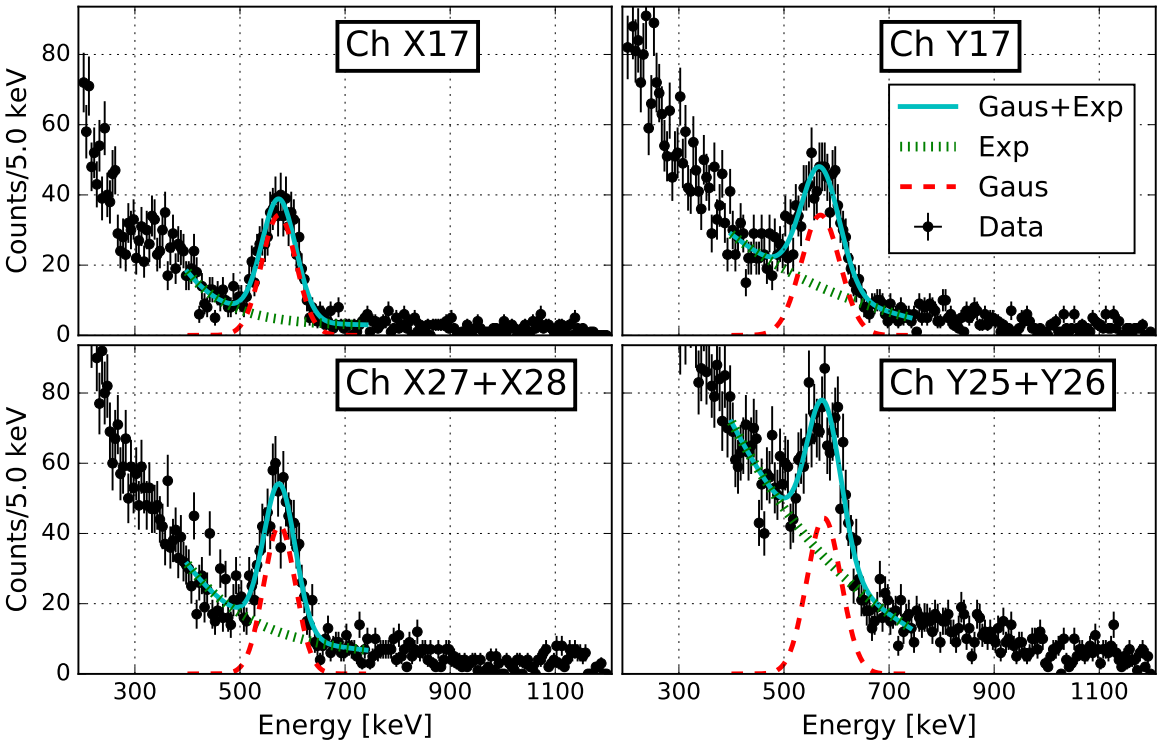

Figure 8.12: Example energy calibration fit to the 570 keV peak for single strip channels X17(top-left) and Y17(top-right) and double strip channels X27+X28 (lower-left) and Y25+Y26 (lower-right). A drift time cut and single-channel cut is included to select events from data. The full fit (solid blue) to the 570 keV for each channel is included as well as the individual Gaussian (dashed red) and the Exponential (hashed green) components which make up the full fit.

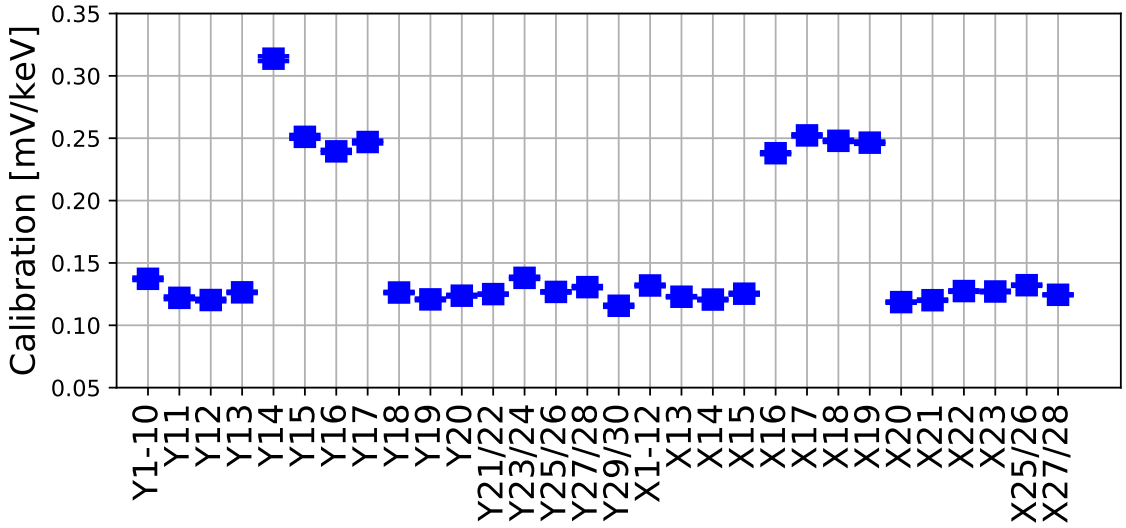

Figure 8.13: Calibration values in units of mV/keV for each charge channel. The different populations correspond to the 2 separate batches of pre-amplifiers which explain the large variations in response between them.

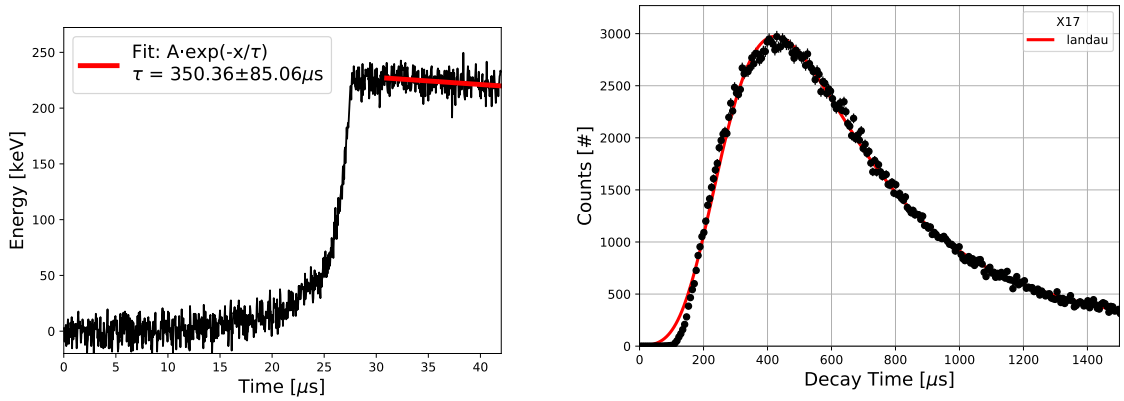

Figure 8.14: (left) Example fit to the decay time of a charge signal and (right) the resulting distribution of decay times over many events and the fit to a landau distribution used to extract the nominal constant.

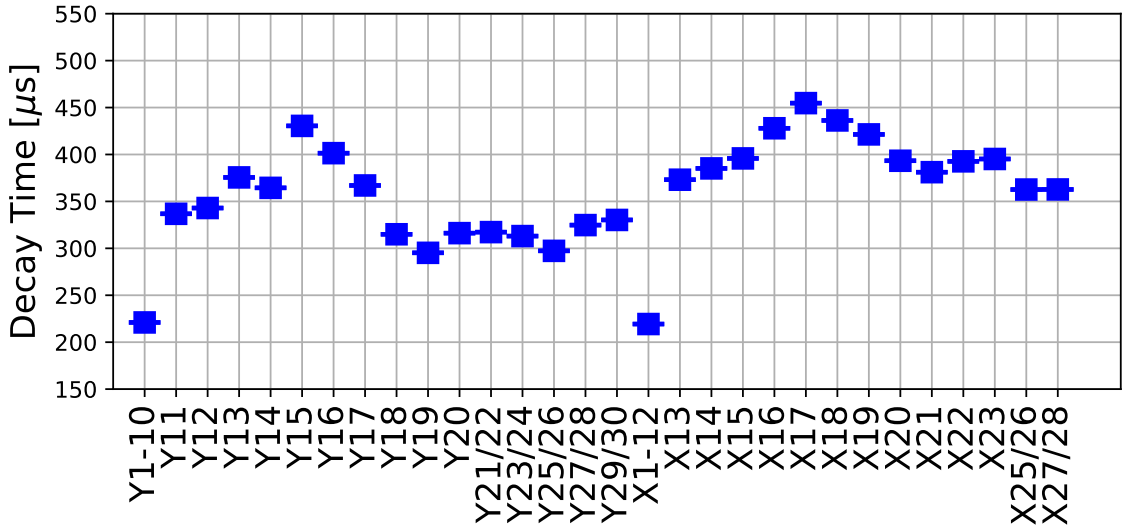

Figure 8.15: Measured decay times for each charge channel used to correct each channels signal.

8.4.4 Resolution

To characterize the charge only performance of the tile the energy resolution at the 570 keV peak was determined. by fitting the peak in data to a Gaussian with an exponential background model, as shown in Figure 8.16. The resolution is then estimated by taking the width of the gaussian ($\sigma/$) and subtracting the contribution from electronic noise, although the noise contribution is found to be negligible relative to the intrinsic fluctuations from recombination. The resulting charge only energy resolution determined by the fit was $\sigma/E = 5.5\%$ This resolution is consistent with the intrinsic resolution of liquid xenon measured by other investigators at the same field. [67, 68].

8.4.5 Agreement with Simulation

A comparison between the energy spectra from simulation and data is shown in Figure 8.17. The spectra for both are normalized to have equal area in the energy range 200 keV to 1200 keV. At energies above 200 keV, the simulation is in good agreement with data. Below 200 keV, the data spectrum has fewer counts than predicted by

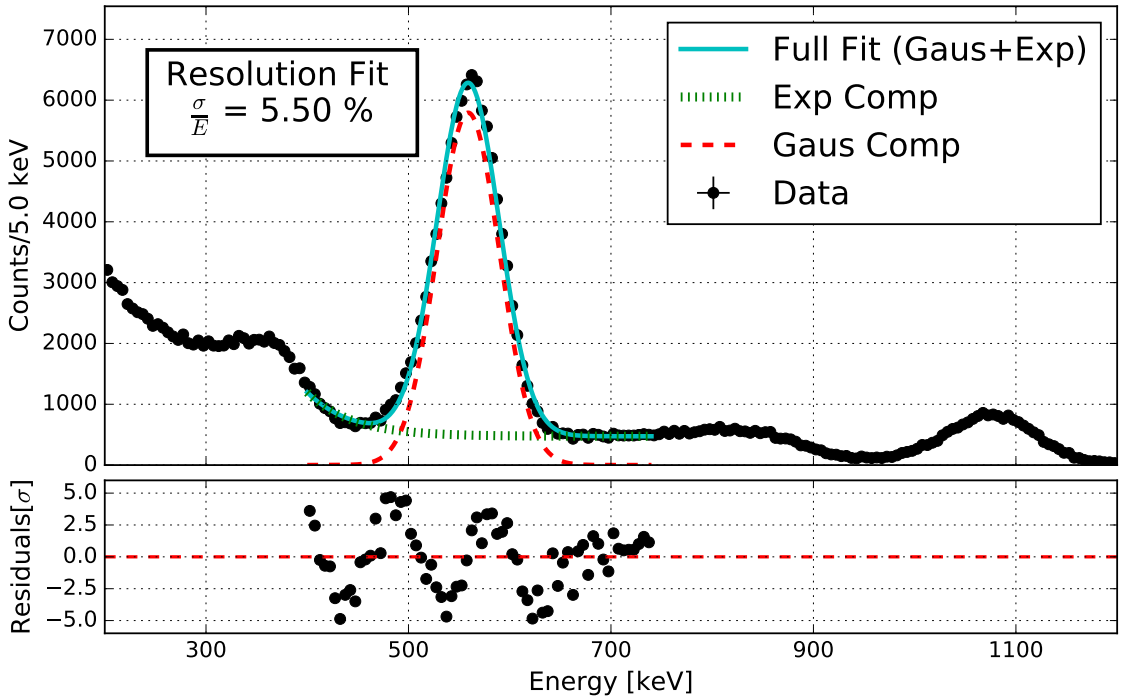

Figure 8.16: Fit to the 570 keV peak (solid blue) seen in data using the cuts described in the text. The fitting function is composed of a Gaussian (dashed red) summed with an exponential background (hashed green) which are also shown. The structure in the residuals is due to the simplifying assumptions of the fit model and doesn't affect the measured energy resolution.

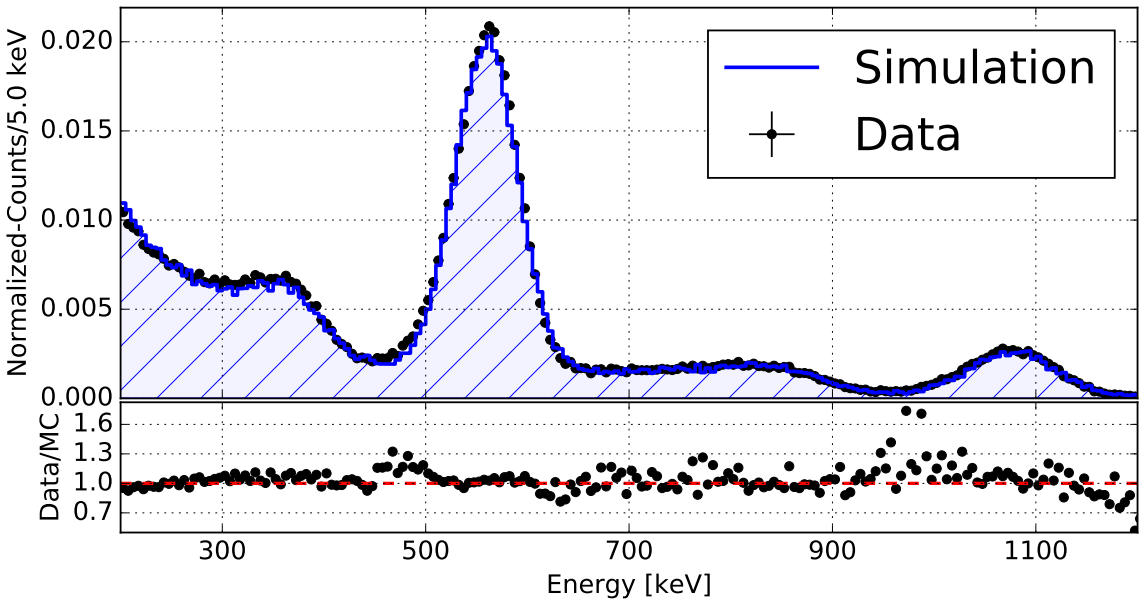

Figure 8.17: Energy spectra from data and simulation. Cuts are described in the text. The simulated spectrum was normalized to match the counts in the data spectrum. The lower panel shows the ratio of the data and MC spectra with a horizontal line in red to show unity.

simulations because of PMT trigger threshold effects. In addition Figure 8.18 shows the predicted and observed electrostatic effects in the current detector as a function of drift distance. The trend of decreasing reconstructed energy near the anode plane is observed in both supporting the model used in simulation.

8.5 Summary

The performance of a novel, segmented, grid-less ionization charge collection detector developed for the nEXO 5-tonne liquid xenon TPC for neutrinoless double beta decay was reported. The charge-only energy resolution measured in LXe is in line with the intrinsic value measured for LXe by numerous investigators. The data from the prototype "tile" shed light on non-conventional electrostatic effects arising from the absence of a shielding Frisch grid in front of the charge collection electrode. A study

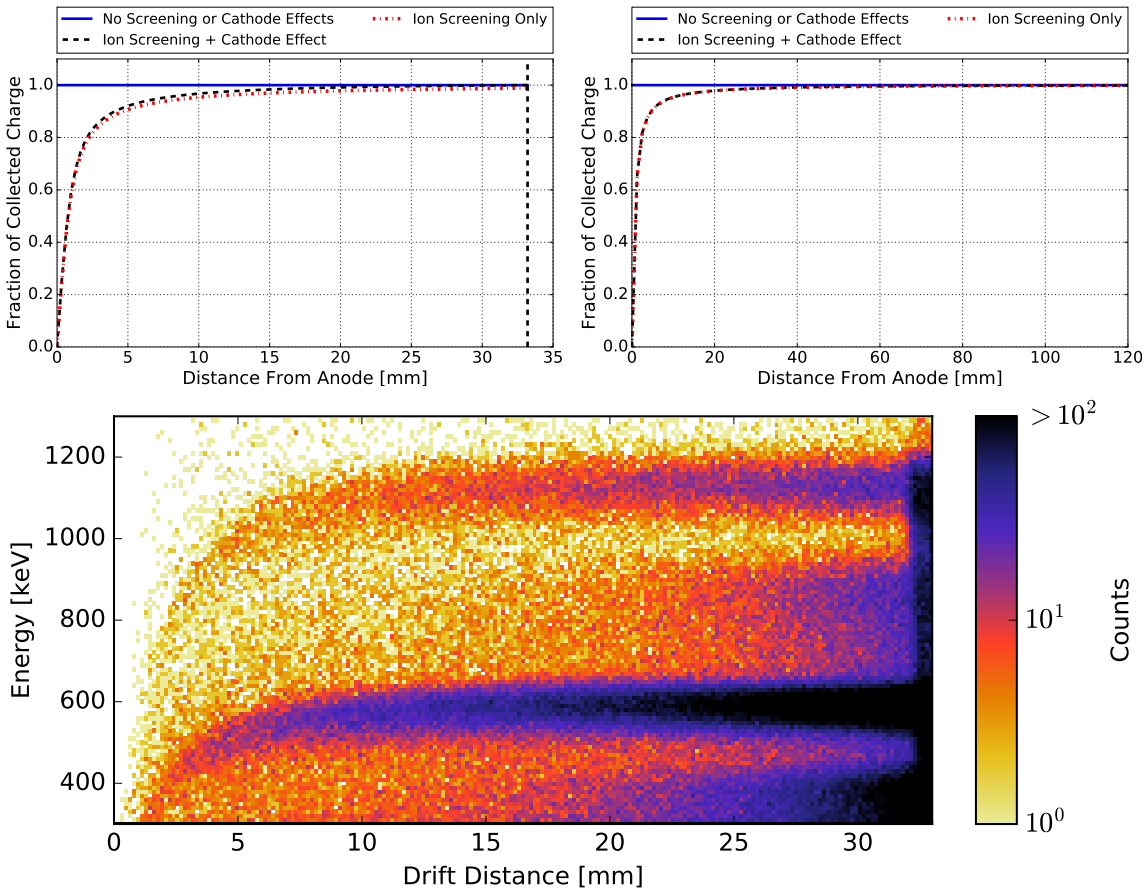

Figure 8.18: Summary of the electrostatic effects (a) in full volume of this test setup (33.2 mm drift length) and (b) the first 120 mm of a ∼1.3 m long detector such as nEXO. The vertical axis shows the magnitude of the signal induced on one strip in response to an energy deposit directly below the strip. This is shown for the hypothetical case where the detector response does not vary with distance from the anode (solid blue) as well the case were only effects of ion screening (dashed black) and ion screening and cathode suppression (hashed red) are considered. Details of these effects are summarized in Appendix A. Also shown (c) is the reconstructed energy vs drift distance observed in data where the predicted decrease in reconstructed energy close to the anode is observed.

of these effects for a nEXO sized TPC (≈130 cm drift length) are shown in Figure 8.18 along with a relative comparison to the same effects in the currently studied detector. This includes the position dependence of the reconstructed charge, including the ion screening and cathode effects described in Appendix A.

Work is in progress to refine the design of the charge tiles. The strip pitch is being optimized for use in nEXO. Future prototype tiles will run with integrated readout circuits placed in LXe. Finally, improvements in the light collection efficiency are being implemented to allow energy measurements taking advantage of the anti-correlation between the ionization and scintillation signals.

Appendices

Appendix A

Induced Charge on Strip

A.1 Calculation for Single Pad

The induced charge on a single square pad comes from considering a point charge above a rectangular plane of grounded electrodes. The induced charge per unit area on a conducting plane by a point charge of charge Q_0 located a distance h above the pad at x=y=0 is given by method of images as

$$\sigma = \frac{-Q_0 h}{2\pi(r^2+h^2)^{3/2}} \tag{A.1}$$

To find the charge in a rectangle which extends from x_1 to x_2 along the x-axis and y_1 and y_2 along the y-axis we calculate

$$Q(h) = -\frac{Q_0 h}{2\pi}\int_{x_1}^{x_2}\int_{y_1}^{y_2}\frac{dx\,dy}{(x^2+y^2+h^2)^{3/2}} \tag{A.2}$$

which can be evaluated as

$$Q(x,y,h) = -\frac{Q_0}{2\pi}\left[f(x_2,y_2,h) - f(x_1,y_2,h) - f(x_2,y_1,h) + f(x_1,y_1,h)\right] \tag{A.3}$$

where Q is the charge induced on the pad, Q_0 is the magnitude of the drifting change, the set of x_i and y_i are the distances from the drifting charge to the four corners of the pad, and h is the height of the charge above the anode. The function $f(x,y,h)$

is defined below:

$$f(x,y,h) = \arctan\left[\frac{xy}{h\sqrt{x^2+y^2+h^2}}\right]. \tag{A.4}$$

The value of h decreases with time according to the drift velocity, 2 mm/µs at 936 V/cm.

A.2 Corrections

Two corrections are applied to account for electrostatic effects from the positive Xe ions produced in the ionization process as well as the charge induced on the cathode.

The correction for the positive ion is an added term, Equation A.3, with the opposite sign to account for the positive charge and a constant value of h since the ion is approximately stationary during the duration of the event. The correction for the cathode is a linear scaling to the induced charge based on the distance from the cathode.

$$Q_{Full}(x,y,h) = \left(\frac{D-h}{D}\right)Q(x,y,h) - \left(\frac{D-h_0}{D}\right)Q(x,y,h_0) \tag{A.5}$$

where D is the distance between the cathode and the anode, h_0 is the initial height of the event, and h is a function of time.

A third correction for a finite electron lifetime was considered but was ignored since no effects of purity degradation were observed. A plot showing the relative charge induced on the cathode for a charge at different z-positions in the detector is shown in Figure 8.18, for the current test setup (left panel) as well as a ∼1 m detector such as nEXO (right panel).

The full expression for charge induced on a strip is calculated by summing Equation A.5 over all 30 pads in a strip for each voxel of charge considered in the simulation.

Bibliography

[1] J. N. Bahcall, M. H. Pinsonneault and S. Basu, *Solar models: Current epoch and time dependences, neutrinos, and helioseismological properties*, *Astrophys. J.* **555** (2001) 990–1012, [astro-ph/0010346]. doi:10.1086/321493.

[2] B. T. Cleveland, T. Daily, R. Davis, Jr., J. R. Distel, K. Lande, C. K. Lee et al., *Measurement of the solar electron neutrino flux with the Homestake chlorine detector*, *Astrophys. J.* **496** (1998) 505–526. doi:10.1086/305343.

[3] Super-Kamiokande collaboration, Y. Fukuda et al., *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562–1567, [hep-ex/9807003]. doi:10.1103/PhysRevLett.81.1562.

[4] Kamiokande-II collaboration, K. Hirata et al., *Observation of a Neutrino Burst from the Supernova SN 1987a*, *Phys. Rev. Lett.* **58** (1987) 1490–1493. [,727(1987)]. doi:10.1103/PhysRevLett.58.1490.

[5] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse and A. D. McGuire, *Detection of the free neutrino: A Confirmation*, *Science* **124** (1956) 103–104. doi:10.1126/science.124.3212.103.

[6] T. Araki et al., *Experimental investigation of geologically produced antineutrinos with KamLAND*, *Nature* **436** (2005) 499–503. doi:10.1038/nature03980.

[7] S. Weinberg, *Universal Neutrino Degeneracy*, *Phys. Rev.* **128** (1962) 1457–1473. doi:10.1103/PhysRev.128.1457.

[8] S. M. Bilenky, C. Giunti, J. A. Grifols and E. Masso, *Absolute values of neutrino masses: Status and prospects*, *Phys. Rept.* **379** (2003) 69–148, [hep-ph/0211462]. doi:10.1016/S0370-1573(03)00102-9.

[9] M. Fukugita and T. Yanagida, *Baryogenesis Without Grand Unification*, *Phys. Lett.* **B174** (1986) 45–47. doi:10.1016/0370-2693(86)91126-3.

[10] R. Carr et al., *Neutrino-based tools for nuclear verification and diplomacy in North Korea*, 1811.04737. doi:10.1080/08929882.2019.1603007.

[11] W. Pauli, *Dear radioactive ladies and gentlemen*, *Phys. Today* **31N9** (1978) 27.

[12] F. Reines, C. L. Cowan, F. B. Harrison, A. D. McGuire and H. W. Kruse, *Detection of the free anti-neutrino*, *Phys. Rev.* **117** (1960) 159–173. doi:10.1103/PhysRev.117.159.

[13] SNO collaboration, Q. R. Ahmad et al., *Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **89** (2002) 011301, [nucl-ex/0204008]. doi:10.1103/PhysRevLett.89.011301.

[14] KamLAND collaboration, K. Eguchi et al., *First results from KamLAND: Evidence for reactor anti-neutrino disappearance*, *Phys. Rev. Lett.* **90** (2003) 021802, [hep-ex/0212021]. doi:10.1103/PhysRevLett.90.021802.

[15] K. S. Babu and V. S. Mathur, *Radiatively Induced Seesaw Mechanism for Neutrino Masses*, *Phys. Rev.* **D38** (1988) 3550. doi:10.1103/PhysRevD.38.3550.

[16] R. N. Mohapatra and G. Senjanovic, *Neutrino Mass and Spontaneous Parity Nonconservation*, *Phys. Rev. Lett.* **44** (1980) 912. [,231(1979)]. doi:10.1103/PhysRevLett.44.912.

[17] S. Bilenky, *Neutrino Masses from the Point of View of Economy and Simplicity*, 1907.01472.

[18] ALEPH, DELPHI, L3, OPAL, SLD, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group collaboration, S. Schael et al., *Precision electroweak measurements on the Z resonance*, *Phys. Rept.* **427** (2006) 257–454, [hep-ex/0509008]. doi:10.1016/j.physrep.2005.12.006.

[19] B. Pontecorvo, *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*, *Sov. Phys. JETP* **26** (1968) 984–988. [Zh. Eksp. Teor. Fiz.53,1717(1967)].

[20] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the unified model of elementary particles*, *Prog. Theor. Phys.* **28** (1962) 870–880. [,34(1962)]. doi:10.1143/PTP.28.870.

[21] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. 2007.

[22] D. V. Forero, M. Tortola and J. W. F. Valle, *Neutrino oscillations refitted*, *Phys. Rev.* **D90** (2014) 093006, [1405.7540]. doi:10.1103/PhysRevD.90.093006.

[23] KATRIN collaboration, M. Aker et al., *An improved upper limit on the neutrino mass from a direct kinematic method by KATRIN*, 1909.06048.

[24] Planck collaboration, P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13, [1502.01589]. doi:10.1051/0004-6361/201525830.

[25] A. Baha Balantekin and B. Kayser, *On the Properties of Neutrinos*, *Ann. Rev. Nucl. Part. Sci.* **68** (2018) 313–338, [1805.00922]. doi:10.1146/annurev-nucl-101916-123044.

[26] F. T. Avignone, III, S. R. Elliott and J. Engel, *Double Beta Decay, Majorana Neutrinos, and Neutrino Mass*, *Rev. Mod. Phys.* **80** (2008) 481–516, [0708.1033]. doi:10.1103/RevModPhys.80.481.

[27] C. F. V. Weizsacker, *Zur Theorie der Kernmassen*, *Z. Phys.* **96** (1935) 431–458. doi:10.1007/BF01337700.

[28] *Nuclear Phenomenology*, ch. 2, pp. 33–69. John Wiley & Sons, Ltd, 2006. https://onlinelibrary.wiley.com/doi/pdf/10.1002/0470035471.ch2. 10.1002/0470035471.ch2. doi:10.1002/0470035471.ch2.

[29] W. Huang, G. Audi, M. Wang, F. Kondev, S. Naimi and X. Xu, *The ame2016 atomic mass evaluation (i). evaluation of input data; and adjustment procedures*, *Chinese Physics C* **41** (03, 2017) 030002. doi:10.1088/1674-1137/41/3/030002.

[30] N. data Services. https://www-nds.iaea.org/amdc/.

[31] M. Redshaw, E. Wingfield, J. McDaniel and E. G. Myers, *Mass and double-beta-decay Q value of Xe-136*, *Phys. Rev. Lett.* **98** (2007) 053003. doi:10.1103/PhysRevLett.98.053003.

[32] M. Goeppert-Mayer, *Double beta-disintegration*, *Phys. Rev.* **48** (Sep, 1935) 512–516. doi:10.1103/PhysRev.48.512.

[33] M. G. Inghram and J. H. Reynolds, *Double beta-decay of* te^{130}, *Phys. Rev.* **78** (Jun, 1950) 822–823. doi:10.1103/PhysRev.78.822.2.

[34] S. R. Elliott, A. A. Hahn and M. K. Moe, *Direct Evidence for Two Neutrino Double Beta Decay in* ^{82}*Se*, *Phys. Rev. Lett.* **59** (1987) 2020–2023. [,989(1987)]. doi:10.1103/PhysRevLett.59.2020.

[35] A. S. Barabash, *Average and recommended half-life values for two-neutrino double beta decay: upgrade-2019*, *AIP Conf. Proc.* **2165** (2019) 020002, [1907.06887]. doi:10.1063/1.5130963.

[36] XENON collaboration, E. Aprile et al., *Observation of two-neutrino double electron capture in* ^{124}*Xe with XENON1T*, *Nature* **568** (2019) 532–535, [1904.11002]. doi:10.1038/s41586-019-1124-4.

[37] S. S. Ratkevich, A. M. Gangapshev, Yu. M. Gavrilyuk, F. F. Karpeshin, V. V. Kazalov, V. V. Kuzminov et al., *Comparative study of the double K-shell-vacancy production in single- and double-electron capture decay*, *Phys. Rev.* **C96** (2017) 065502, [1707.07171]. doi:10.1103/PhysRevC.96.065502.

[38] A. P. Meshik, C. M. Hohenberg, O. V. Pravdivtseva and Ya. S. Kapusta, *Weak decay of Ba-130 and Ba-132: Geochemical measurements*, *Phys. Rev.* **C64** (2001) 035205. doi:10.1103/PhysRevC.64.035205.

[39] J. Schechter and J. W. F. Valle, *Neutrinoless Double beta Decay in SU(2) x U(1) Theories*, *Phys. Rev.* **D25** (1982) 2951. [,289(1981)]. doi:10.1103/PhysRevD.25.2951.

[40] J. Kotila and F. Iachello, *Phase space factors for double-β decay*, *Phys. Rev.* **C85** (2012) 034316, [1209.5722]. doi:10.1103/PhysRevC.85.034316.

[41] N. López Vaquero, T. R. Rodríguez and J. L. Egido, *Shape and pairing fluctuations effects on neutrinoless double beta decay nuclear matrix elements*, *Phys. Rev. Lett.* **111** (2013) 142501, [1401.0650]. doi:10.1103/PhysRevLett.111.142501.

[42] J. Menendez, A. Poves, E. Caurier and F. Nowacki, *Disassembling the Nuclear Matrix Elements of the Neutrinoless beta beta Decay*, *Nucl. Phys.* **A818** (2009) 139–151, [0801.3760]. doi:10.1016/j.nuclphysa.2008.12.005.

[43] J. Barea, J. Kotila and F. Iachello, *$0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements in the interacting boson model with isospin restoration*, *Phys. Rev.* **C91** (2015) 034304, [1506.08530]. doi:10.1103/PhysRevC.91.034304.

[44] J. Engel, F. Simkovic and P. Vogel, *Chiral Two-Body Currents and Neutrinoless Double-Beta Decay in the QRPA*, *Phys. Rev.* **C89** (2014) 064308, [1403.7860]. doi:10.1103/PhysRevC.89.064308.

[45] M. T. Mustonen and J. Engel, *Large-scale calculations of the double-β decay of ^{76}Ge, ^{130}Te, ^{136}Xe, and ^{150}Nd in the deformed self-consistent Skyrme*

quasiparticle random-phase approximation, *Phys. Rev.* **C87** (2013) 064302, [1301.6997]. doi:10.1103/PhysRevC.87.064302.

[46] J. Engel and J. Menéndez, *Status and Future of Nuclear Matrix Elements for Neutrinoless Double-Beta Decay: A Review*, *Rept. Prog. Phys.* **80** (2017) 046301, [1610.06548]. doi:10.1088/1361-6633/aa5bc5.

[47] F. F. Deppisch and J. Suhonen, *Statistical analysis of β decays and the effective value of g_A in the proton-neutron quasiparticle random-phase approximation framework*, *Phys. Rev.* **C94** (2016) 055501, [1606.02908]. doi:10.1103/PhysRevC.94.055501.

[48] J. T. Suhonen, *Value of the Axial-Vector Coupling Strength in β and $\beta\beta$ Decays: A Review*, *Front.in Phys.* **5** (2017) 55, [1712.01565]. doi:10.3389/fphy.2017.00055.

[49] Particle Data Group collaboration, M. Tanabashi et al., *Review of Particle Physics*, *Phys. Rev.* **D98** (2018) 030001. doi:10.1103/PhysRevD.98.030001.

[50] KamLAND-Zen collaboration, A. Gando et al., *Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen*, *Phys. Rev. Lett.* **117** (2016) 082503, [1605.02889]. [Addendum: Phys. Rev. Lett.117,no.10,109903(2016)]. doi:10.1103/PhysRevLett.117.109903,10.1103/PhysRevLett.117.082503.

[51] S. Dell'Oro, S. Marcocci, M. Viel and F. Vissani, *Neutrinoless double beta decay: 2015 review*, *Adv. High Energy Phys.* **2016** (2016) 2162659, [1601.07512]. doi:10.1155/2016/2162659.

[52] O. Cremonesi and M. Pavan, *Challenges in Double Beta Decay*, *Adv. High Energy Phys.* **2014** (2014) 951432, [1310.4692]. doi:10.1155/2014/951432.

[53] M. K. Moe, *New approach to the detection of neutrinoless double beta decay*, *Phys. Rev.* **C44** (1991) 931–934. [,1019(1991)]. doi:10.1103/PhysRevC.44.R931.

[54] B. J. Mount, M. Redshaw and E. G. Myers, *Double-beta-decay Q values of Se-74 and Ge-76*, *Phys. Rev.* **C81** (2010) 032501. doi:10.1103/PhysRevC.81.032501.

[55] GERDA collaboration, M. Agostini et al., *Probing Majorana neutrinos with double-β decay*, *Science* **365** (2019) 1445, [1909.02726]. doi:10.1126/science.aav8613.

[56] MAJORANA collaboration, S. I. Alvis et al., *A Search for Neutrinoless Double-Beta Decay in ^{76}Ge with 26 kg-yr of Exposure from the MAJORANA DEMONSTRATOR*, *Phys. Rev.* **C100** (2019) 025501, [1902.02299]. doi:10.1103/PhysRevC.100.025501.

[57] M. Redshaw, B. J. Mount, E. G. Myers and F. T. Avignone, III, *Masses of Te-130 and Xe-130 and Double-beta-Decay Q Value of Te-130*, *Phys. Rev. Lett.* **102** (2009) 212502, [0902.2139]. doi:10.1103/PhysRevLett.102.212502.

[58] CUORE collaboration, S. Pozzi et al., *New results from the CUORE experiment*, in *16th International Conference on Topics in Astroparticle and Underground Physics (TAUP 2019) Toyama, Japan, Sept. 9-13, 2019*, 2019.

[59] EXO-200 collaboration, G. Anton et al., *Search for Neutrinoless Double-Beta Decay with the Complete EXO-200 Dataset*, *Phys. Rev. Lett.* **123** (2019) 161802, [1906.02723]. doi:10.1103/PhysRevLett.123.161802.

[60] H. V. Klapdor-Kleingrothaus and I. V. Krivosheina, *The evidence for the observation of 0nu beta beta decay: The identification of 0nu beta beta events from the full spectra*, *Mod. Phys. Lett.* **A21** (2006) 1547–1566. doi:10.1142/S0217732306020937.

[61] U.S. Department of Energy, "Waste Isolation Pilot Plant." www.wipp.energy.gov, 2016. Accessed January 3, 2017.

[62] U.S. Department of Energy, "What happened at WIPP in February 2014." www.wipp.energy.gov, 2015. Accessed July 17, 2019.

[63] EXO-200 collaboration, M. Auger et al., *The EXO-200 detector, part I: Detector design and construction*, *JINST* **7** (2012) P05010, [1202.2192]. doi:10.1088/1748-0221/7/05/P05010.

[64] EXO-200 collaboration, J. B. Albert et al., *Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector*, *Phys. Rev.* **C89** (2014) 015502, [1306.6106]. doi:10.1103/PhysRevC.89.015502.

[65] EXO-200 collaboration, J. B. Albert et al., *Search for Majorana neutrinos with the first two years of EXO-200 data*, *Nature* **510** (2014) 229, [1402.6956]. doi:10.1038/nature13432.

[66] EXO collaboration, J. B. Albert et al., *Search for Neutrinoless Double-Beta Decay with the Upgraded EXO-200 Detector*, *Phys. Rev. Lett.* **120** (2018) 072701, [1707.08707]. doi:10.1103/PhysRevLett.120.072701.

[67] E. Aprile and T. Doke, *Liquid xenon detectors for particle physics and astrophysics*, *Rev. Mod. Phys.* **82** (Jul, 2010) 2053–2097. doi:10.1103/RevModPhys.82.2053.

[68] E. Conti et al., *Correlated fluctuations between luminescence and ionization in liquid xenon*, *Phys. Rev.* **B68** (2003) 054201, [hep-ex/0303008]. doi:10.1103/PhysRevB.68.054201.

[69] T. Doke, *Fundamental properties of liquid argon, krypton and xenon as radiation detector media*, *Portugaliae Physica* **12** (1981) 9–48.

[70] "3M^{TM} Novec^{TM} 700 Engineered Fluid." http://www.3m.com/.

[71] D. S. Leonard et al., *Systematic study of trace radioactive impurities in candidate construction materials for EXO-200*, *Nucl. Instrum. Meth.* **A591** (2008) 490–509, [0709.4524]. doi:10.1016/j.nima.2008.03.001.

[72] EXO-200 collaboration, J. B. Albert et al., *Cosmogenic Backgrounds to $0\nu\beta\beta$ in EXO-200*, *JCAP* **1604** (2016) 029, [1512.06835]. doi:10.1088/1475-7516/2016/04/029.

[73] EXO-200 collaboration, R. Neilson et al., *Characterization of large area APDs for the EXO-200 detector*, *Nucl. Instrum. Meth.* **A608** (2009) 68, [0906.2499]. doi:10.1016/j.nima.2009.06.029.

[74] COMSOL. http://www.comsol.com/.

[75] K. O'Sullivan, *The Search for $0\nu\beta\beta$ in $^{1}36Xe$ with EXO-200*. Dissertation, Stanford University, 2012. [http://grattalab3.stanford.edu/neutrino/Theses/Kevin-Osullivan-thesis-augmented.pdf].

[76] J. Allison et al., *Geant4 developments and applications*, *IEEE Trans. Nucl. Sci.* **53** (2006) 270. doi:10.1109/TNS.2006.869826.

[77] W. Shockley, *Currents to conductors induced by a moving point charge*, *J. Appl. Phys.* **9** (1938) 635–636. doi:10.1063/1.1710367.

[78] S. Ramo, *Currents induced by electron motion*, *Proc. Ire.* **27** (1939) 584–585. doi:10.1109/JRPROC.1939.228757.

[79] EXO-200 collaboration, J. B. Albert et al., *Measurement of the Drift Velocity and Transverse Diffusion of Electrons in Liquid Xenon with the EXO-200 Detector*, *Phys. Rev.* **C95** (2017) 025502, [1609.04467]. doi:10.1103/PhysRevC.95.025502.

[80] EXO Collaboration collaboration, M. Auger, D. J. Auty, P. S. Barbeau, E. Beauchamp, V. Belov, C. Benitez-Medina et al., *Search for neutrinoless double-beta decay in ^{136}Xe with exo-200*, *Phys. Rev. Lett.* **109** (Jul, 2012) 032505. doi:10.1103/PhysRevLett.109.032505.

[81] EXO-200 collaboration, C. G. Davis et al., *An Optimal Energy Estimator to Reduce Correlated Noise for the EXO-200 Light Readout*, *JINST* **11** (2016) P07015, [1605.06552]. doi:10.1088/1748-0221/11/07/P07015.

[82] M. Berger, J. Coursey, M. Zucker and J. Chang, "Stopping-power and range tables for electrons, protons, and helium ions." https://dx.doi.org/10.18434/T4NC7P.

[83] nEXO collaboration, S. A. Kharusi et al., *nEXO Pre-Conceptual Design Report*, 1805.11142.

[84] B. J. Mount et al., *LUX-ZEPLIN (LZ) Technical Design Report*, 1703.09144.

[85] T. Doke, *Recent Development of Liquid Xenon Detectors.*, *Nucl. Instrum. Meth.* **196** (1982) 87–96. doi:10.1016/0029-554X(82)90621-8.

[86] P. Sorensen, *Anisotropic diffusion of electrons in liquid xenon with application to improving the sensitivity of direct dark matter searches*, *Nucl. Instrum. Meth.* **A635** (2011) 41–43, [1102.2865]. doi:10.1016/j.nima.2011.01.089.

[87] L. G. H. Huxley and R. W. Crompton, *The diffusion and drift of electrons in gases*. Wiley New York, 1974.

[88] E. Shibamura et al., *Ratio of diffusion coefficient to mobility for electrons in liquid argon*, *Phys. Rev* **A20** (Dec, 1979) 2547–2554. doi:10.1103/PhysRevA.20.2547.

[89] EXO-200 collaboration, J. B. Albert et al., *Measurements of the ion fraction and mobility of α- and β-decay products in liquid xenon using the EXO-200 detector*, *Phys. Rev.* **C92** (2015) 045504, [1506.00317]. doi:10.1103/PhysRevC.92.045504.

[90] L. S. Miller et al., *Charge transport in solid and liquid Ar, Kr, and Xe*, *Phys. Rev.* **166** (Feb, 1968) 871–878. doi:10.1103/PhysRev.166.871.

[91] E. M. Gushchin et al., *Electron dynamics in condensed argon and xenon*, *Zh. Eksp. Teor. Fiz.* **82** (1981) 1114–1125.

[92] P. F. Sorensen, *A Position-Sensitive Liquid Xenon Time-Projection Chamber for Direct Detection of Dark Matter: The XENON10 Experiment*. PhD thesis, Brown University, 2008.

[93] XENON100 collaboration, E. Aprile et al., *Analysis of the XENON100 Dark Matter Search Data*, *Astropart. Phys.* **54** (2014) 11–24, [1207.3458]. doi:10.1016/j.astropartphys.2013.10.002.

[94] LUX collaboration, D. S. Akerib et al., *First results from the lux dark matter experiment at the sanford underground research facility*, *Phys. Rev. Lett.* **112** (Mar, 2014) 091303, [1310.8214]. doi:10.1103/PhysRevLett.112.091303.

[95] P. Benetti et al., *A Simple and effective purifier for liquid xenon*, *Nucl. Instrum. Meth.* **A329** (1993) 361–364. doi:10.1016/0168-9002(93)90954-G.

[96] K. Yoshino et al., *Effect of molecular solutes on the electron drift velocity in liquid Ar, Kr, and Xe*, *Phys. Rev.* **A14** (1976) 438–444. doi:10.1103/PhysRevA.14.438.

[97] W. E. Spear et al., *A possible explanation of the observed electron drift velocity saturation in solid ar, kr, and xe*, *Phys. Rev.* **178** (Feb, 1969) 1454–1456. doi:10.1103/PhysRev.178.1454.

[98] EXO-200 collaboration, G. Anton et al., *Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment*, 1908.04128.

[99] M. Szydagis et al., *Enhancement of NEST Capabilities for Simulating Low-Energy Recoils in Liquid Xenon*, *JINST* **8** (2013) C10003, [1307.6601]. doi:10.1088/1748-0221/8/10/C10003.

[100] E. Shibamura et al., *Ratio of the diffusion coefficient to the mobility for electrons in liquid argon, liquid argon-ethylene, and liquid xenon*, in *Proc. Int. Seminar on Swarm Experiments in Atomic Collision Research*, (Tokyo), p. 47, 1979.

[101] S. Kobayashi et al., *Ratio of transverse diffusion coefficient to mobility of electrons in high-pressure xenon*, *Jpn. J. Appl. Phys.* **43** (2004) 5568.

[102] A. Hoecker, P. Speckmayer, J. Stelzer, J. Therhaag, E. von Toerne, H. Voss et al., *Tmva - toolkit for multivariate data analysis*, 2007.

[103] EXO-200 collaboration, J. B. Albert et al., *Search for $2\nu\beta\beta$ decay of ^{136}Xe to the 0_1^+ excited state of ^{136}Ba with EXO-200*, *Phys. Rev.* **C93** (2016) 035501, [1511.04770]. doi:10.1103/PhysRevC.93.035501.

[104] B. P. Roe, H.-J. Yang, J. Zhu, Y. Liu, I. Stancu and G. McGregor, *Boosted decision trees, an alternative to artificial neural networks*, *Nucl. Instrum. Meth.* **A543** (2005) 577–584, [physics/0408124]. doi:10.1016/j.nima.2004.12.018.

[105] T. Fawcett, *An introduction to roc analysis*, *Pattern Recognition Letters* **27** (2006) 861 – 874. ROC Analysis in Pattern Recognition. doi:https://doi.org/10.1016/j.patrec.2005.10.010.

[106] EXO collaboration, S. Delaquis et al., *Deep Neural Networks for Energy and Position Reconstruction in EXO-200*, *JINST* **13** (2018) P08023, [1804.09641]. doi:10.1088/1748-0221/13/08/P08023.

[107] C. Szegedy et al., *Going deeper with convolutions*, *CoRR* **abs/1409.4842** (2014) , [1409.4842].

[108] F. Chollet et al., "Keras." https://keras.io, 2015.

[109] M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, 2015. Software available from tensorflow.org.

[110] "Working Document for 2019 $0\nu\beta\beta$ Search." https://confluence.slac.stanford.edu/display/exo/2019+0v+analysis.

[111] W. Verkerke and D. Kirkby, *The roofit toolkit for data modeling*, 2003.

[112] E. Browne and J. K. Tuli, *Nuclear Data Sheets for $A = 137$*, *Nucl. Data Sheets* **108** (2007) 2173–2318. doi:10.1016/j.nds.2007.09.002.

[113] D. Fudenberg, *Improved discrimination for neutrinoless double beta decay searches with EXO-200 and nEXO*. Dissertation, Stanford University, 2018. [purl.stanford.edu/vb237gq6176].

[114] nEXO collaboration, J. B. Albert et al., *Sensitivity and Discovery Potential of nEXO to Neutrinoless Double Beta Decay*, *Phys. Rev.* **C97** (2018) 065503, [1710.05075]. doi:10.1103/PhysRevC.97.065503.

[115] D. Sinclair, "Charge collection for nEXO." nEXO Internal Note, 2012.

[116] nEXO collaboration, M. Jewell et al., *Characterization of an Ionization Readout Tile for nEXO*, *JINST* **13** (2018) P01006, [1710.05109]. doi:10.1088/1748-0221/13/01/P01006.

[117] "LabVIEW - National Instruments." http://www.ni.com/en-us/shop/labview.html.

[118] "MKS 121A Baratron." https://www.mksinst.com/product/product.aspx?ProductID=1191.

[119] "SAES MonoTorr PF3C3R1." http://www.saespuregas.com/.

[120] "MKS Instruments." https://www.mksinst.com/.

[121] L. Fabris et al., *A fast, compact solution for low noise charge preamplifiers*, *Nucl. Instrum. Meth.* **A424** (1999) 545 – 551. doi:https://doi.org/10.1016/S0168-9002(98)01371-0.

[122] "Struck Innovative Systeme." http://www.struck.de/sis3316.html.

[123] "NuDat2." https://www.nndc.bnl.gov/nudat2/.

[124] S. Agostinelli et al., *Geant4 a simulation toolkit*, *Nucl. Instrum. Meth.* **A506** (2003) 250 – 303. doi:http://doi.org/10.1016/S0168-9002(03)01368-8.