Neutrinoless Double Beta Decay in Dual-phase Xenon Time Projection Chambers: LUX, LZ and Prospects for G3

Robert Taylor

Imperial College London
Department of Physics

A thesis submitted for the degree of
Doctor of Philosophy

July 2020

Abstract

The groundbreaking observation of neutrino oscillations indicates that the elusive neutrino is massive, providing tantalising evidence for physics beyond the standard model. However, given that neutrinos cannot obtain mass through the same mechanism as the charged leptons, unsolved problems still remain. Neutrinoless double beta decay ($0\nu\beta\beta$) is a hypothetical process that, if discovered, may shine light on the origins of neutrino mass. It would also confirm that Majorana fermions exist in nature, and that the neutrino is its own anti-particle, demonstrating that lepton number is not conserved through a fundamental symmetry of the universe.

In this thesis, the capability of past, present and future dual-phase xenon time projection chambers to search for ^{136}Xe $0\nu\beta\beta$ is demonstrated. A search for this process is performed in a short 19 kg·yr exposure of the LUX experiment, with the goal of illustrating the reconstruction of $\mathcal{O}$(MeV) events in such detectors, and demonstrating the techniques required for future experiments. Next, the sensitivity of the LUX-ZEPLIN (LZ) experiment to the $0\nu\beta\beta$ half-life $T_{1/2}$ is studied, revealing that after 1,000 days of exposure, LZ can exclude $T_{1/2} < 1.06 \times 10^{26}$ yr at 90% CL. A similar sensitivity projection is also performed for a hypothetical, large scale future detector that utilises $\mathcal{O}$(100 tonnes) of xenon. It is shown that at this scale it is possible to exclude $T_{1/2} < 1.0 \times 10^{28}$ yr, and probe almost the entire inverted neutrino mass hierarchy at 90% CL. Finally, a machine learning method for fast simulation of the detector response using a generative adversarial network is studied. It is demonstrated that this technique can accurately generate the digitised photomultiplier signals resulting from interactions in a dual-phase detector. This technique may provide a computationally inexpensive method for fast simulation in future detectors.

Declaration

I hereby declare that the work presented in this thesis is my own, and any work that is not my own is referenced appropriately. My specific contributions in each chapter are listed below.

Chapter 1 gives an overview of the phenomenon of neutrino oscillations and the requirement for neutrino masses. Neutrinoless double beta decay is introduced, followed by a survey of the main experimental efforts that are focused on the discovery of this process.

Chapter 2 describes the operating principles of dual-phase xenon time projection chambers, together with the underlying scintillation mechanisms in liquid xenon. The detector systems of the LUX and LUX-ZEPLIN experiments are introduced, which builds upon the previous work of both collaborations.

Chapter 3 is focused on the search for ^{136}Xe neutrinoless double decay in a data set obtained in the LUX experiment. For this I developed a new data processing framework, inspired by the official LUX framework, that improves the reconstruction of high energy interactions. I performed all of the subsequent data analysis, but the corrections and calibrations build heavily on the previous work performed by the LUX collaboration.

Chapter 4 introduces the projected sensitivity of the LUX-ZEPLIN experiment to ^{136}Xe neutrinoless double beta decay. Part of this work was performed in collaboration with Paulo Bras at the University of Coimbra, where we both developed the background model. My individual contribution was the implementation of the DECAY0 signal generator into the LUX-ZEPLIN simulation software, and the development of the statistical framework used to determine the final sensitivity. The results of this work are published in [1].

Chapter 5 studies the sensitivity of a 3rd generation dual-phase xenon experiment. The analysis and mathematical modelling is all my own work,

and the final sensitivity results utilise the statistical tools developed in the previous chapter.

Chapter 6 presents the results of using a generative adversarial network for fast simulation of the signal response in dual-phase detectors. This further develops and improves upon the initial work performed closely with two previous Imperial MSci students, Aman Ghataura and Trenton Delahaye. My contribution is the improvement of their network architecture, with the subsequent demonstration of an accurate reproduction of data from the LUX detector.

Robert Taylor
July 2020

Acknowledgements

The past three and a half years have been an intense but overall worthwhile and enjoyable experience. I am grateful to so many people throughout my PhD journey.

Firstly, thank you to the Imperial HEP group and STFC for allowing me to do this PhD. I am immensely grateful for everything I have learnt, the new skills I have developed and the many opportunities I have had to travel around the world as well as 4,850ft underground!

I would like to thank my supervisor Antonin for giving me the confidence to be an effective researcher. Throughout my PhD you provided critical guidance, and gave me the freedom to explore my own ideas and work on a variety of different projects, half of which aren't included in this thesis! I have learnt so much more than I hoped because of this. I am also grateful to Henrique Araujo and the rest of the LZ IC group for all your feedback and for providing an intellectually stimulating environment.

I cannot thank enough all my friends at Imperial. We had so many enjoyable times that made me always look forward to going into the office. Somehow through all the lengthy lunch and coffee breaks we managed to get through this. Thank you so much!

Thank you to Mum and Dad. Without your constant support throughout the years I would not be where I am today.

And finally, Jenny. Your patience, care and love over the past 8 years has been invaluable. This PhD would not have been possible without you.

Contents

List of figures

List of tables

Introduction

The neutrino has long been a source of intrigue in particle physics. It's existence was originally postulated by Wolfgang Pauli in 1930 [2], who supposed that an undetectable particle must be emitted during β-decays of unstable nuclei to ensure conservation of energy and momentum. Almost 30 years later, in 1956 the neutrino[1] was finally discovered through the reaction $\bar{\nu}_e + p^+ \rightarrow n + e^+$. In the mid 1960s, a significant deficit in the solar neutrino flux resulting from nuclear reactions in the sun was observed on earth [3]. This longstanding solar neutrino problem was not resolved until 2001, when the occurrence of neutrino oscillations was confirmed. This effect gives neutrinos the ability to oscillate between three different quantum mechanical states, and also hints at physics beyond the standard model.

It is understood that neutrinos are fundamental particles, and are included in the standard model as spin–$\frac{1}{2}$ fermions with three generations, or flavours: the electron (ν_e), muon (ν_μ) and tau (ν_τ) neutrino. However, neutrinos are distinctly unlike any other fundamental leptons in nature. Their only known interaction or production mechanism is via the weak interaction, and as a result only neutrinos with left-handed chirality (ν_L) are known to exist. Furthermore, they are the only fermion not to carry electric charge. This opens up the possibility for neutrinos to be Majorana fermions, where there is no distinction between particle and anti-particle. In fact, it is more theoretically appealing for them to be Majorana particles due to the fewer degrees of freedom required to describe the particle. Finally, they are known to have an astonishingly small mass, being at least one million times lighter than the electron. This mass cannot be acquired through the spontaneous symmetry breaking of the Higgs field that gives mass to all other fundamental particles. These facts demonstrate that there is much more to the origins of the neutrino than is currently understood, and that they are key to understanding physics beyond the standard model.

[1] more correctly the electron anti-neutrino

One process that may shed light on the nature of the neutrino is *neutrinoless double beta decay* ($0\nu\beta\beta$). Interestingly, the null observation of neutrinos would have profound implications for particle physics. It would confirm that the neutrino is a Majorana particle, demonstrate that lepton number is not conserved and can potentially be used to indirectly measure the neutrino mass. Searching for this process is therefore a direct method of probing physics beyond the standard model. However, with lower limits on the half-life being greater 10^{19} years, neutrinoless double beta decay, if exists, would be one of the rarest processes in nature. As a result, a successful observation requires incredibly sensitive detector technology with ultra-low background environments.

It is expected that $0\nu\beta\beta$ will occur in ^{136}Xe, an isotope of the noble gas xenon that has a natural abundance of 8.9%. It is therefore present in significant quantities in experiments that utilise xenon as a particle detection medium. In particular, dual-phase xenon time projection chambers (TPCs), traditionally utilised by experiments searching for dark matter, are a potential technology that can push the boundaries in the search for $0\nu\beta\beta$. This thesis aims to demonstrate the potential reach of such detectors, providing an analysis of the past (LUX), present (LZ) and future (G3) detectors. It is organised as follows:

Chapter 1 introduces the theoretical motivation for $0\nu\beta\beta$, beginning with neutrino oscillations and the requirement for massive neutrinos. The seesaw mechanism is provided as an example of how the standard model can be extended to include neutrino masses, and how this results in Majorana neutrinos. The light neutrino exchange mechanism is discussed, demonstrating the link between $0\nu\beta\beta$ and neutrino masses. Finally, the different search strategies are explored, providing an overview of the main experimental efforts currently searching for this process.

Chapter 2 details the operating principle of dual-phase xenon TPCs in addition to the detector systems of the LUX and LZ experiments. The use of xenon as a particle detection medium is motivated, and the physical processes resulting in observable scintillation signals are discussed.

Chapter 3 presents the results of a search for $0\nu\beta\beta$ in the LUX Run-3 data set, with the aim to demonstrate the steps required to successfully search for $0\nu\beta\beta$. A new software framework for the reconstruction of high energy interactions in LUX is introduced, which is used to establish an significant improvement in the energy resolution and background rejection for $\mathcal{O}$(MeV) events in LUX. The corrections and calibrations of the liquid xenon response are developed, leading to the calculation of a lower limit on the $0\nu\beta\beta$ half-life.

Chapter 4 assesses the sensitivity of the LZ experiment to $0\nu\beta\beta$. The expected observable signal is studied through simulation of the LXe response. A detailed discussion of the major backgrounds in LZ is presented, and background rates estimated from Monte Carlo simulations are given, resulting in a detailed background model for LZ that utilises both the energy and position distributions of the backgrounds. The statistical methodology used to project the half-life sensitivity is introduced, and sensitivity estimates for various scenarios of detector performance are produced.

Chapter 5 studies the sensitivity of a future third-generation (G3) detector, and compares the results to other planned experiments. A semi-analytical model of γ attenuation in a cylindrical detector is developed, and combined with simulations of the LZ backgrounds to predict the background rate arising from external backgrounds in a detector using $\mathcal{O}(100$ tonne) of xenon. The achievable sensitivity is estimated according to different background scenarios and levels of ^{136}Xe enrichment.

Chapter 6 introduces the use of a generative adversarial network (GAN) as a method for fast simulation of the response of a dual-phase xenon TPC. The goal is to demonstrate that this machine learning algorithm can produce accurate simulated data at a fraction of the computational cost and time of traditional Monte Carlo simulations. The method of training a GAN, and the Wasserstein GAN (WGAN) variant is explained. LUX data is used to train an architecture that can successfully reproduce the complex photomultiplier response. Several applications of machine learning algorithms were also explored, but the work presented here demonstrates the most promising and novel technique.

Chapter 7 summarises the key outcomes of this thesis, and provides an outlook for future work.

Chapter 1

Massive Neutrinos & Double Beta Decay

1.1 Neutrino Oscillations

The existence of neutrino oscillations, whereby neutrinos are observed to transform between three flavour eigenstates, not only implies that neutrinos are massive, but also provides a hint of physics beyond the standard model. Neutrino oscillations can be understood through the decomposition of three flavour eigenstates $|\nu_\alpha\rangle$ into a linear superposition of three definite mass eigenstates $|\nu_i\rangle$

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}|\nu_i\rangle, \tag{1.1}$$

where $U_{\alpha i}$ is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) unitary transformation matrix. This has the common parameterisation,

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.2}$$

where $s_{ij}(c_{ij}) = \sin\theta_{ij}(\cos\theta_{ij})$, θ_{ij} are the three independent mixing angles and δ is the charge-parity violating phase. To see how the existence of neutrino mass gives rise to oscillations, consider that neutrinos must produced in states of definite flavour through the weak interaction. However if each flavour state is a linear combination of mass eigenstates with definite mass m_i, after travelling a distance x with momentum p, the

state evolves in time according to

$$|\nu_i(t)\rangle = e^{-i\sqrt{|p|^2+m_i^2}t+ip\cdot x}|\nu_i(0)\rangle. \tag{1.3}$$

$$= e^{-i\frac{m_i^2}{2E}x}|\nu_i(0)\rangle \tag{1.4}$$

where the last equality is true in the ultra-relativistic limit. If the masses m_i are non-degenerate, each phase advances at a different rate and the resulting state at time t is a different mixture of mass eigenstates than at $t = 0$. However, as (2.1) implies, this corresponds to a different superposition of flavour eigenstates. An experiment may therefore observe an appearance or disappearance of neutrino flavours. The probability to transition from flavour α to flavour β is then [4],

$$P(\nu_\alpha \to \nu_\beta) = |\langle\nu_\beta(t)|\nu_\alpha\rangle|^2 \tag{1.5}$$

$$= \sum_i \sum_j U^*_{\beta i} U_{\alpha i} U^*_{\alpha j} U_{\beta j} e^{-i\frac{(m_i^2-m_j^2)}{2E}x}. \tag{1.6}$$

For $\alpha \neq \beta$ there is a non zero transition probability due to the constructive interference resulting from the squared mass splittings $\Delta m^2_{ij} = m_i^2 - m_j^2$. Observation of neutrino oscillations therefore implies that neutrinos have mass.

The phenomenon of neutrino oscillation explains the observed deficits in solar, and subsequent observations in reactor and long-baseline neutrino experiments have provided precise measurements of the neutrino mixing angles, together with constraints on the CP violating phase. A global analysis of all data gives values of $\Delta m^2_{21} = 7.4 \times 10^{-5}$ eV^2 and $|\Delta m^2_{31}| = 2.5 \times 10^{-3}$ eV^2 [5] for the squared mass splittings. Thus, neutrinos have mass but are extremely light, being $\mathcal{O}(10^6)$ times lighter than the electron. However, although neutrino oscillations have successfully confirmed the existence of neutrino mass, a number of open questions still remain:

- *Ordering*: the relative ordering of the neutrino masses is still unknown. This produces two possible scenarios, or hierarchies of the neutrino masses:

$$m_1 < m_2 < m_3 \quad \text{Normal hierarchy (NH)} \tag{1.7}$$

$$m_3 < m_1 < m_2 \quad \text{Inverted hierarchy (IH)} \tag{1.8}$$

- *Scale*: as oscillations are only sensitive to Δm^2_{ij} the absolute scale of neutrino mass cannot be inferred from oscillation experiments. Instead, the sum of the neutrino

masses $\sum_i m_i$ is accessible through cosmological measurements, with data from Planck observations producing an upper limit of $\sum_i m_i < 0.152$ eV [6]. In addition, precision measurements of β-decay end-points are sensitive to the electron neutrino mass $m_{\nu_e} = \sqrt{\sum_i |U_{ei}|^2 m_i^2}$, which is bounded by the $m_{\nu_e} < 1.1$ eV measurement from the KATRIN experiment [7]. As discussed later, neutrinoless double beta decay experiments are also sensitive to the neutrino mass scale.

- *Mechanism*: unlike other leptons, the standard model Higgs mechanism cannot be responsible for generating neutrino masses. In the standard model, only the left-handed neutrino ν_L is present. This has quantum numbers $(\frac{1}{2}, -1)$ under the electroweak symmetry group $SU(2)_L \times U(1)_Y$ and it is therefore not possible to generate a mass term for the neutrino without explicitly breaking gauge invariance. The mass generation mechanism must therefore require new physics beyond the standard model.

1.1.1 The Seesaw Mechanism

Perhaps the simplest approach in extending the standard model to include neutrino masses is the seesaw mechanism [8, 9, 10]. This postulates the existence of a heavy right-handed neutrino N_R field with quantum numbers $(1, 0)$ under $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, i.e it is sterile and does not take part in the electroweak interactions. The existence of such a particle is well motivated, for example it arises naturally in grand unified theories (GUTs) with a SO(10) symmetry group. In the three-generation scenario ($\alpha = e, \mu, \tau$) the most general neutrino mass term of a Lagrangian involving N_R^α and ν_L^α Weyl spinors is given by

$$L_M = Y_{\alpha\beta} L_i^\alpha \phi^{\dagger i} \bar{N}_R^\beta + \frac{M_L}{2} (\bar{\nu}_L^\alpha)^c \nu_L^\alpha + \frac{M_R}{2} (\bar{N}_R^\alpha)^c N_R^\alpha + \mathrm{h.c} \tag{1.9}$$

where $L_i^\alpha = (\nu_L^\alpha, e_L^\alpha)$ is the electroweak doublet, ϕ^i is a scalar Higgs doublet, $Y_{\alpha\beta}$ is a 3×3 Yukawa coupling matrix, M_L is a Majorana mass for the left-handed neutrinos. M_R is a Majorana mass for the right-handed neutrino and ψ^c denotes the charge-conjugate of the field ψ. Under the standard model symmetry group, the second term is forbidden as it is not gauge invariant, and therefore $M_L = 0$. On the other hand, as N_R transforms trivially under the gauge group, the corresponding Majorana mass M_R is allowed. In addition, this term changes lepton number (L) by two units ($\Delta L = 2$), violating the global symmetry that results in conservation of lepton number. After spontaneous symmetry breaking in

the Higgs sector, ϕ obtains a non-zero vacuum expectation value v, i.e $\phi \to (0, v)$, and therefore the mass terms become,

$$L_M = vY^{\alpha\beta}\nu_L^\alpha \bar{N}_R^\beta + \frac{M_R}{2}(\bar{N_R}^\alpha)^c N_R^\alpha + \text{h.c} \tag{1.10}$$

The first term is a standard Dirac mass term, where $M_D^{\alpha\beta} = vY^{\alpha\beta}$ can be interpreted as the 3×3 Dirac mass matrix. In terms of four-component spinors $\nu = (\nu_L^\alpha, (N_R^\alpha)^c)$ this becomes

$$L_M = \frac{1}{2}\left(\bar{\nu}_L^\alpha, (\bar{N}_R^\alpha)^c\right)\begin{pmatrix} 0 & M_D \\ M_D^T & M_R \end{pmatrix}\begin{pmatrix} (\nu_L^\alpha)^c \\ N_R^\alpha \end{pmatrix} + \text{h.c} \tag{1.11}$$

$$= \frac{1}{2}\bar{\nu}^c \mathcal{M}\nu^T + \text{h.c} \tag{1.12}$$

where the 9×9 non-diagonal $\mathcal{M}$ symmetric mass matrix implies that in this basis these fields are not mass eigenstates. In the limit $M_R \gg m_D$, diagonalising the mass matrix gives the eigenvalues [11]

$$m_L = -M_D^T M_R^{-1} M_D \tag{1.13}$$

$$m_R = M_R \tag{1.14}$$

with the corresponding eigenvectors

$$\nu_1^\alpha = \nu_L^\alpha + (\nu_L^\alpha)^c - \frac{m_D}{M_R^2}(N_R + N_R^c) \approx \nu_L^\alpha + (\nu_L^\alpha)^c \tag{1.15}$$

$$\nu_2^\alpha = N_R^\alpha + (N_R^\alpha)^c + \frac{m_D}{M_R^2}(\nu_L^\alpha + (\nu_L^\alpha)^c) \approx N_R^\alpha + (N_R^\alpha)^c \tag{1.16}$$

The mass eigenstates therefore approximately correspond to the left-handed light neutrinos and the heavy right-handed sterile neutrinos, where their mixing is heavily suppressed by $1/M_R^2$. Importantly, the mass eigenstates satisfy the Majorana condition $\nu = \nu^c$, thus the neutrinos are Majorana fermions. This mechanism by which the neutrinos obtain mass is termed the Type-I seesaw mechanism, as equation 1.13 implies that increasing the mass of the right-handed neutrinos reduces the mass of the light, left-handed neutrinos. A key feature of this mechanism is an explanation for the lightness of the neutrino masses; neutrinos masses of $\mathcal{O}(0.1$ eV) can be achieved with M_R of $\mathcal{O}(10^{14}$ GeV), roughly corresponding to the energy scale at which the standard model couplings unify in GUTs.

The light neutrino mass matrix m_L can diagonalised by $m_L = U m_\nu U^T$ where U is a matrix that can be identified with the PMNS matrix and m_ν is the diagonal matrix of light neutrino masses. The requirement of mass generation through the seesaw mechanism is that neutrinos must be Majorana fermions. This introduces two more physical parameters of the PMNS matrix known as the Majorana phases. In the Dirac case, these can be absorbed into a re-definition of the fields. However, for Majorana neutrinos this is not possible due to the condition $\nu = \nu^c$. A major consequence of Majorana neutrinos is the possibility of the existence of lepton number violating processes, such as neutrinoless double beta decay.

1.2 Double Beta Decay

Maria Goeppert-Mayer was the first to propose [12] that an unstable nucleus may decay via two simultaneous β-decays, i.e

$$(A, Z) \rightarrow (A, Z+2) + 2e^- + 2\bar{\nu_e} \tag{1.17}$$

This is a nuclear process involving the decay of an atom (A, Z) and the emission of two electrons and two neutrinos, known as double beta decay ($2\nu\beta\beta$). In addition, the complementary processes of double positron decay and double electron capture are also possible. As a second-order weak decay it is one of the rarest known processes; measured half-lives of $2\nu\beta\beta$ isotopes are in the range $10^{19} - 10^{21}$ yr, which is more than ten times longer than the age of the universe. A key property of this process is that it can only occur in a limited number of natural isotopes where the single β-decay is forbidden as a consequence of energy conservation. This can arise in nuclei that have both an even atomic number (A) and even proton number (Z), as illustrated in figure 1.1. Due to the positive pairing term in the nuclear binding energy E_b for A, Z odd atoms, the $Z-1$ isobar can have a higher nuclear binding energy than it's parent, and therefore this process cannot occur. However, the $Z-2$ isobar has a lower binding energy, and therefore the decay to this daughter can proceed via the emission of two simultaneous β decays, i.e $2\nu\beta\beta$. The total energy released in the decay is given by the Q-value $Q_{\beta\beta}$,

$$Q_{\beta\beta} = E_b(A, Z) - E_b(A, Z+2) \tag{1.18}$$

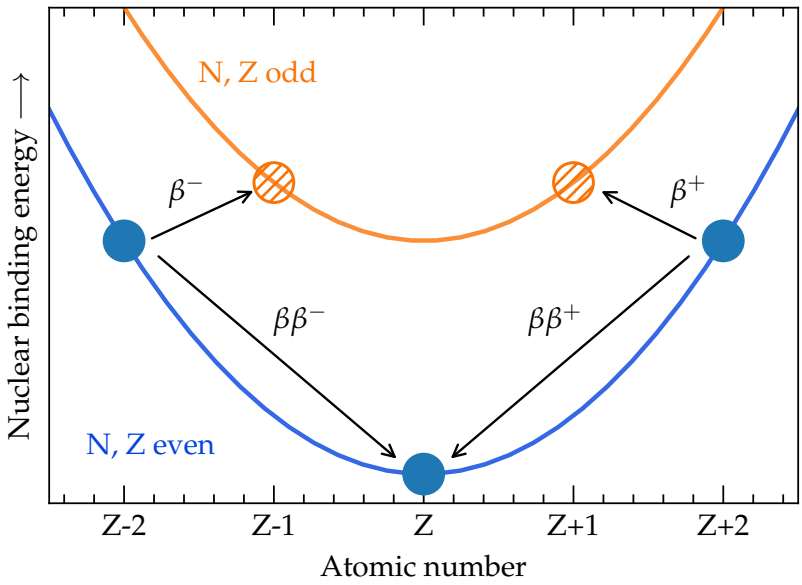

Figure 1.1: Schematic of nuclear binding energy and allowed $\beta^{\pm}$ and $\beta\beta^{\pm}$ decays for even A nuclei. Coloured lines show nuclear energy for isotopes of constant A. Filled markers indicate allowed nuclear transition and hatched markers shows forbidden transitions to isotopes that violate energy conservation.

Due to the large mass of the recoiling nucleus compared to the mass of the outgoing leptons, practically all the released energy is shared between the electrons and neutrinos.

Not long after Ettore Majorana developed the theory of Majorana fermions, Wendell Furry suggested [13] that if neutrinos are massive Majorana particles then double beta decay can occur without the emission of neutrinos, i.e

$$(A, Z) \rightarrow (A, Z + 2) + 2e^- \tag{1.19}$$

This alternate process is known as neutrinoless double beta decay ($0\nu\beta\beta$), and if discovered would imply that neutrinos are Majorana particles, i.e they are their own antiparticle. If discovered, it would therefore provide evidence that the seesaw mechanism may be responsible for the generation of neutrino mass. With no electron anti-neutrinos produced, another outcome is the violation of lepton number by two units ($\Delta L = 2$), which would break the accidental symmetry of the standard model. Many theories beyond the standard model such as grand unified theories predict the existence of lepton violating number processes such as $0\nu\beta\beta$, [11] making the search for this decay theoretically motivated and one of the most exciting avenues to look for new physics. In addition, depending on

the underlying mechanism responsible for $0\nu\beta\beta$, measurement of this process can also shed light on the absolute neutrino mass scale. The $0\nu\beta\beta$ isotopes commonly studied in experiments are given in table 1.1, together with the measured lower limits on their half-lives.

Table 1.1: Commonly studied $0\nu\beta\beta$ isotopes and the measured lower limits on their half-lives. Data from [14] and references therein.

Isotope	$T_{1/2}$ [$\times 10^{25}$ yr]	$Q_{\beta\beta}$ [keV]
^{76}Ge	> 8.0	2,039
^{82}Se	$> 3.6 \times 10^{-2}$	2,998
^{96}Zr	$> 9.2 \times 10^{-4}$	3,356
^{100}Mo	$> 1.1 \times 10^{-1}$	3,034
^{116}Cd	$> 1.0 \times 10^{-2}$	2,813
^{130}Te	> 1.5	2,527
^{136}Xe	> 10.7	2,458
^{150}Nd	$> 2.0 \times 10^{-3}$	3,370

The simplest mechanism giving rise to $0\nu\beta\beta$ is that of light neutrino exchange, mediated by a virtual light neutrino as shown in in figure 1.2. This does not introduce any new particles or physical processes beyond requiring that the neutrino is a Majorana fermion. The necessity for the neutrino to be a Majorana fermion for this process to occur can be seen by considering the chiral nature of the weak vertex, where only a right handed anti-neutrino or left-handed neutrino can be connected to the outgoing electrons. The virtual right-handed anti-neutrino emitted at the bottom vertex can only propagate towards the top vertex if it flips chirality and becomes a left handed neutrino.

It can be shown that the matrix element corresponding to to the leptonic part of the interaction is given by [11],

$$-i\mathcal{M} \propto \bar{e}\gamma_\mu(1-\gamma_5)\sum_i U_{ei}^2 \frac{\not{q}+m_i}{q^2-m_i^2} C\bar{e}\gamma_\nu(1-\gamma_5) \tag{1.20}$$

where C is the charge conjugation operator and q is the momentum of the exchanged neutrino with mass m. The introduction of the PMNS matrix element U_{ei} is necessary to project out the electron neutrino from the neutrino mass eigenstate that enters into

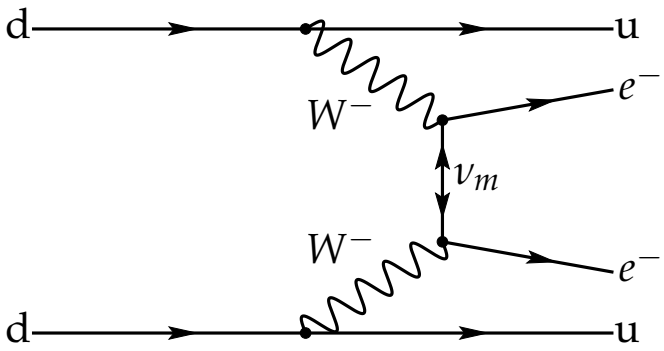

Figure 1.2: Feynman diagram for $0\nu\beta\beta$ mediated by the exchange of a light Majorana neutrino.

the propagator. It is seen that the decay rate is therefore proportional to term

$$|m_{\beta\beta}| = |\sum_{i=1}^{3} U_{ei}^2 m_i| \tag{1.21}$$

which is the effective mass of the Majorana neutrino that takes part in the interaction. In the case that neutrinos are massless this term vanishes, rendering $0\nu\beta\beta$ forbidden. The full expression for the decay rate is given by,

$$\Gamma = |m_{\beta\beta}|^2 |\mathcal{M}_{0\nu}|^2 G_{0\nu}(Q_{\beta\beta}) \tag{1.22}$$

where $|\mathcal{M}_{0\nu}|^2$ is the nuclear matrix element (NME) and $G_{0\nu}(Q_{\beta\beta})$ is the phase-space factor. The NME captures the physics of the nuclear transition, and cannot be determined analytically. The phase space factor can be calculated providing a suitable approximation of the Coloumb interaction between the outgoing electrons and the nuclear potential, and are in the range $10^{-15} - 10^{-14}$ yr^{-1} [15]. The light neutrino exchange mechanism establishes the connection between $0\nu\beta\beta$ half-life and the absolute neutrino mass scale. Figure 1.3 illustrates the dependence on the the effective mass with the neutrino mass hierarchy. There is a distinct dependence on the mass hierarchy, with the inverted hierarchy bounded below by $|m_{\beta\beta}| > 10$ meV. In the normal hierarchy scenario, the rate of $0\nu\beta\beta$ is at least an order of magnitude slower, and there is even a possibility

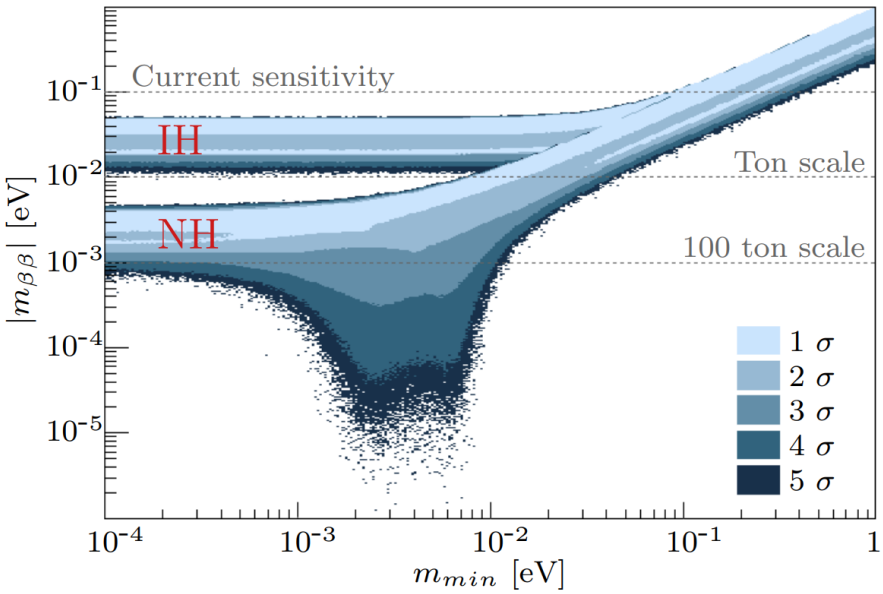

Figure 1.3: Effective Majorana neutrino mass $|m_{\beta\beta}|$ versus the mass of the lightest neutrino m_{min}. The shaded regions represent the statistical uncertainty in the mixing parameters, assuming a uniform distribution over the Majorana phases. Taken from [16].

around $m_{min} \sim 4$ meV where $|m_{\beta\beta}|$ vanishes entirely. It is important to note that there are other mechanism apart from light neutrino exchange that can result in $0\nu\beta\beta$. For example, in left-right symmetric models that propose an additional $\mathrm{SU}(2)_R$ symmetry to the standard model, the W_R boson can produce right-handed currents in the interaction resulting in a decay rate that does not depend on neutrino mixing matrix parameters [17], and is suppressed by the mass scale of the new right-handed gauge boson. Furthermore, if instead the mechanism is mediated by a heavy neutrino, the rate is again suppressed by the respective mass scale [18]. As a result, an observation of $0\nu\beta\beta$ cannot be translated into a measurement of the neutrino mass without first understanding the underlying mechanism. Aside from the $0\nu\beta\beta$ mechanism, the translation of the half-life to $|m_{\beta\beta}|$ is also complicated by the NME. These are of $\mathcal{O}(1)$ and have a large theoretical uncertainty due to the different methods of approximating the complex nuclear wave functions, as shown in figure 1.4. The result is that a measurement of the half-life translates into a large range of compatible $|m_{\beta\beta}|$. However, the Schechter–Valle theorem [19] necessitates that any mechanism responsible for $0\nu\beta\beta$ generates a non-zero Majorana neutrino mass, and therefore searching for this process is still key to the understanding of neutrino masses.

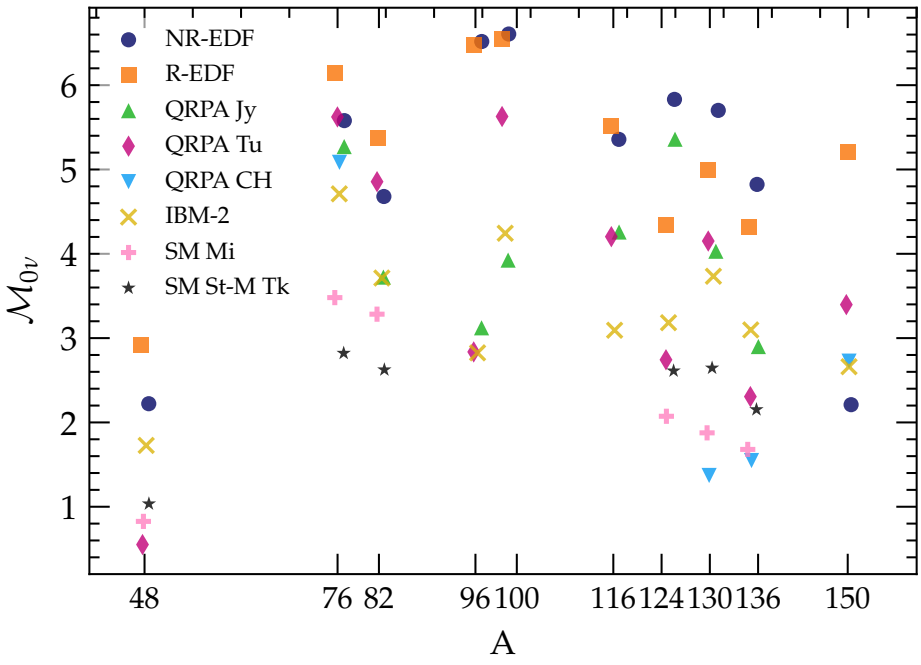

Figure 1.4: The range of calculated nuclear matrix elements for different $0\nu\beta\beta$ isotopes. Adapted from [20]

1.3 Experimental Strategies

The experimental signature of $0\nu\beta\beta$ is the detection of two electrons with a combined energy of $Q_{\beta\beta}$, observed as a mono-energetic peak at the tail end of the $2\nu\beta\beta$ electron energy spectrum. Given an observation of N_s $0\nu\beta\beta$ signal events, the corresponding half-life $T_{\frac{1}{2}}$ is given by

$$T_{\frac{1}{2}} = \frac{\varepsilon N_{\beta\beta} T}{N_s} \ln 2 \tag{1.23}$$

where ε is the signal detection efficiency, $N_{\beta\beta}$ is the number of $\beta\beta$ isotopes under observation and T is the measurement time of the experiment. In practice, this simple picture is complicated by the presence of background processes that imitate and obscure the signal, in addition to the finite detector resolution which imparts the reconstructed energy with an irreducible uncertainty. In this case, an experiment will measure a total of N events, of which N_b are background and N_s could be signal. In the presence of this background, there needs to be sufficient statistical significance in the number of signal events in order to claim discovery of $0\nu\beta\beta$ and measure the half-life.

The *sensitivity* can be used to characterise the ability of an experiment to exclude a hypothesised half-life, and provides a figure-of-merit through which comparisons of different experimental approaches can be made. Assuming N is Poisson distributed, this can be defined as the largest half-life that would result in a number of signal events that would be hidden by an n_σ upwards fluctuation of the background,

$$N_s \geq n_\sigma \sqrt{N_b}. \tag{1.24}$$

In a region of interest (ROI) around $Q_{\beta\beta}$ in which the background distribution is approximately flat, and if this background scales linearly with time and the mass of the $0\nu\beta\beta$ source M, the number of background events is given by,

$$N_b = b \cdot M \cdot \Delta E \cdot T \tag{1.25}$$

where b is the background rate in units of events per unit mass, time and energy, and ΔE is the width of the ROI, which is directly proportional to the energy resolution. Given a source of molar mass A with a $\beta\beta$ isotopic abundance α, the number of $\beta\beta$ atoms can be expressed as $N_{\beta\beta} = \alpha M N_A / A$ where N_A is Avogadro's constant. With this, the half-life sensitivity is given by

$$T_{\frac{1}{2}} \leq \frac{\varepsilon \alpha N_A}{n_\sigma A} \sqrt{\frac{MT}{b\Delta E}} \ln 2. \tag{1.26}$$

It is common to quote the sensitivity at 90% confidence level (CL), corresponding to $n_\sigma = 1.64$. The interpretation is then the largest half-life where the signal plus background hypothesis is excluded in 90% of repeated experiments. This expression highlights the four key parameters that can be tuned in an experiment:

- *Detection efficiency*: the ability to successfully observe $0\nu\beta\beta$ events requires a detection medium with a high efficiency to produce a response to the two electrons. Moreover, data analysis algorithms that reject backgrounds must minimise false negatives.

- *Exposure*: perhaps the most straightforward way of increasing sensitivity is by increasing the source mass and/or running the experiment for longer. If feasible, the isotopic enrichment can also be used to increase the number $\beta\beta$ isotopes, resulting in a significant increase in the sensitivity.

- *Background rate*: reducing contamination of the $\beta\beta$ source from natural radioisotopes, in addition to constructing the detector from radiopure materials and surrounding with shielding deceases the number of background interactions occurring in the ROI. Operation in deep underground laboratories is also required to limit the cosmic ray backgrounds.
- *Energy resolution*: utilising technology for improved energy resolution reduces the background in the ROI, for example by allowing the mono-energetic signal peak to be resolved from neighbouring γ-ray background peaks. In an experiment where only the energy of the electrons is available, this is also the only parameter that can be used to discriminate between $2\nu\beta\beta$ and $0\nu\beta\beta$ events.

Aside from these parameters, the choice of $\beta\beta$ isotope used in the experiment also plays a part. Ideally, it must have the the ability to be obtained in sufficiently large quantities, e.g $\mathcal{O}$(10-1000 kg), and be free from intrinsic radioactive backgrounds. Isotopes with large $Q_{\beta\beta}$ ($\geq$ 2 MeV) such as ^{130}Te will have an ROI situated far from the majority of naturally occurring radiation, and therefore the background rate has limited impact. The isotopes for which the active detector medium is also the source medium, such as ^{136}Xe, have the benefit of high detection efficiency. In addition there is a distinct advantage in high Z materials where the strong γ attenuation provides active self shielding against external backgrounds.

Aside from this, there are other methods of increasing sensitivity. For example, experiments with precise position reconstruction may observe the individual electron tracks, and can thereby exploit track topology and angular correlations to discriminate between signal and background events. The remainder of this chapter discusses the main detector technology currently used to search for $0\nu\beta\beta$.

1.3.1 Semiconductor Detectors

Semiconductor detectors measure ionising radiation through the production of electrons and holes in the solid-state phase. Under the application of an external field, the charge carriers drift through the conduction and valence bands to the electrodes where a current signal can be measured. High Purity germanium (HPGe) detectors utilising ^{76}Ge are a core technology in the search for $0\nu\beta\beta$. This is mainly due to the unparalleled energy resolution achievable with semiconductor diodes, where resolutions of $\sim$ 0.1% at $Q_{\beta\beta}$

have been achieved [21]. Moreover, the timing characteristics of these detectors allow for strong background rejection using pulse shape discrimination methods [22].

Notable ^{76}Ge experiments are the Germanium Detector Array (GERDA) [23], located at the underground Gran Sasso laboratory, and the MAJORANA DEMONSTRATOR (MJD) [24] at Sanford Underground Research Laboratory. With a 46.7 kg yr exposure and background rate of background rate of 1.0×10^{-3} events/kg/year/keV, GERDA placed an upper limit on the half life of $T_{1/2} \geq 8.0 \times 10^{25}$ yr [25]. The aim of MJD is to determine whether a future 1 tonne experiment can achieve a background rate of 1.0×10^{-4} events/kg/year/keV in the ROI.

1.3.2 Cryogenic Calorimeters

The minute change in temperature arising from an interaction with a crystal lattice is the operation principle behind cryogenic calorimeter, or bolometer, experiments. These operate a crystal absorber at mK temperatures, and temperature changes of $\mathcal{O}(0.1$ mK) can be measured by a highly sensitive thermometer. Crystals consisting of $\beta\beta$ isotopes such as ^{130}Te, ^{82}Sn and ^{100}Mo can be grown with minor background contamination, and energy resolutions comparable to HPGe detectors are achievable due to the large phonon statistics.

The Cryogenic Underground Observatory for Rare Events (CUORE) experiment, operated at Gran Sasso, utilises crystals of ^{130}Te. After an exposure of 24.0 kg y, a limit of $T_{1/2} > 1.3 \times 10^{25}$ yr was set, with a sensitivity of 9×10^{25} yr after 5 years of data taking [26]. Although, low backgrounds and strong energy resolutions are achievable, this technology is clearly disadvantaged by the requirement to operate at ultra-cold temperatures, which constrains the scale and running time of the experiment.

1.3.3 Tracking Calorimeters

A different approach to searching for $0\nu\beta\beta$ is the through the separation of the detection medium and source materials, such as in the Neutrino Ettore Majorana Observatory (NEMO) experiment [27]. Here, thin foils of a $\beta\beta$ isotope are placed between gaseous layers used for tracking of the two electrons. A surrounding calorimeter layer then measures their energy. The ability to track electron paths provides a strong method of background rejection, as both electrons have to be emitted from the same vertex.

In addition, this is the only technology that can measure the angular distribution of the emitted electrons, which may provide insight into the underlying $0\nu\beta\beta$ mechanism. Another advantage is that possibility of using multiple $\beta\beta$ isotopes. Indeed, the NEMO experiment has produced limits for ^{100}Mo, ^{82}Se, ^{96}Zr, ^{116}Cd and ^{150}Cd. The planned SuperNEMO experiment [28] builds upon NEMO, with a larger design and reduced background rate. A disadvantage of this approach is the difficulty in scaling up to large masses, and of course the restriction to studying the $\beta\beta$ isotopes that can be formed in to a thin foil.

1.3.4 Scintillators

Detection of the light produced from organic or inorganic scintillators is the principle method used by KamLAND-Zen [29] and the future SNO+ [30] experiment. This technology is relatively easy to operate at scale, and position reconstruction can be achieved with sufficient photo-coverage of the scintillator. KamLAND-Zen uses an organic liquid scintillator loaded with 380kg xenon enriched to 90% in ^{136}Xe. This is contained within a balloon that is suspended inside an spherical tank, and surrounded by 3,200 tonnes of water shielding. As of 2016, KamLAND-Zen has set the most stringent limits on the ^{136}Xe half-life of $T_{1/2} \geq 1.07 \times 10^{27}$ yr, and has provided the strongest probe of the neutrino mass hierarchy down to $\langle m_{\beta\beta} \rangle = 61 - 165$ meV. SNO+ is a planned experiment utilising a linear akyl benzene scintillator doped with 1,300 kg of ^{130}Te, with a projected sensitivity of $T_{1/2} \geq 1.9 \times 10^{26}$ yr.

A drawback of scintillator detectors is the relatively poor energy resolution, driven by a limited photon detection efficiency. For example, KamLAND-Zen achieved an energy resolution of $\sim 5\%$ at $Q_{\beta\beta}$, which results in a larger ROI and therefore an increased background rate.

1.3.5 Time Projection Chambers

The time projection chamber (TPC) can measure both the energy and three dimensional position of ionising radiation interactions in a gas or liquid scintillating detection medium. This is achieved by measuring the emitted photons, and drifting liberated charge carriers in an electric field towards an electrode, and using the time of arrival difference between this and the light signal to infer the position of energy deposits. Moreover, the precision of

the position reconstruction allows for identification of multi-site events such as background γ interactions.

^{136}Xe is a versatile $\beta\beta$ isotope that is gaseous at room temperature and can be cooled to a liquid at cryogenic temperatures. Xenon is also an efficient scintillator, and being an inert noble gas it is therefore well placed to be used as the detection medium in TPCs. Searches for $0\nu\beta\beta$ in ^{136}Xe commonly use xenon either in the liquid phase or gas phase. As a liquid, the technology can be currently scaled to $\mathcal{O}$(1,000) kg. At this scale, the xenon becomes self shielding, and attenuates much of the external γ background. The Enriched Xenon Observatory (EXO) [31] experiment utilised a xenon TPC enriched to 90% in ^{136}Xe, and achieved a lower limit of 3.5×10^{25} yr with an energy resolution of 1.15% at $Q_{\beta\beta}$. nEXO [32] is a next-generation experiment that plans to use $\mathcal{O}(10,\,000)$ kg of enriched xenon, resulting in a half-life sensitivity of 1.0×10^{28} yr after 10 years, with the ability to probe the inverted neutrino mass hierarchy.

In the gas phase, the higher density limits the scale of the experiments but provides improved energy resolution. This is the approach taken by the NEXT [33] and PandaX-III [34] experiments, which aim to operate $\mathcal{O}$(100 kg) of enriched xenon in as gas TPC. Gas TPCs also have the ability to image the tracks of ionising radiation, resulting in the ability to identify $0\nu\beta\beta$ events and reject backgrounds through an analysis of the track topology [35].

Recent attention is turning to detector technology that operates with both liquid and gas xenon, or dual phase TPCs. The strength of this technology has been demonstrated by the sensitivity of direct detection dark matter (DM) experiments, and the stringent limits placed on DM-nucleon scattering cross sections [36]. Using both the gas and liquid phase results in improved position reconstruction and the ability to achieve an energy resolution of below 1.0% at $Q_{\beta\beta}$, as demonstrated by the XENON1T experiment [37].

The remainder of this thesis focuses on ^{136}Xe $0\nu\beta\beta$ searches in dual phase xenon TPCs, specifically analysing the capability of past, present and future experiments.

Chapter 2

Dual-phase TPCs & Liquid Xenon Physics

2.1 Dual-phase Time Projection Chambers

Xenon dual phase time projection chambers (TPCs) are radiation detectors that utilise a combination of both liquid and gaseous xenon, and have the ability to reconstruct the energy and three-dimensional position of particle interactions. Experiments utilising dual phase xenon have consistently produced world leading results in searches for weakly interacting massive particles (WIMPs) [38, 39, 40] due to the high sensitivity afforded by the detector.

A schematic of a general TPC design is presented in Fig 2.1. Generally, a dual phase TPC can be split into two distinct regions: a bulk liquid xenon (LXe) volume and a region above containing gaseous xenon (GXe), with height of $\mathcal{O}$(mm). The energy loss of an incident particle in xenon results in the emission of vacuum ultraviolet (VUV) scintillation light, denoted *S1*, and production of ion-electron pairs. The former is promptly detected by a set of photo-detectors that observe the liquid from above and below. An electric field of $\mathcal{O}(100\ \mathrm{Vcm}^{-1})$ is applied to the LXe region to drift electrons away from the interaction site and towards the GXe. Upon reaching the liquid-gas interface, electrons are extracted into the GXe by a stronger field of $\mathcal{O}(2-10\ \mathrm{kVcm}^{-1})$. This produces an amplified VUV scintillation signal through the process of electroluminescence, denoted S2.

Observation of both of the S1 and S2 signals can be used to reconstruct interactions. Due to the constant drift velocity v_d of electrons in LXe, the depth z of an interaction

in the TPC is directly proportional[1] to the time difference ΔT between the the arrival of the S1 and S2 signals, and can therefore be reconstructed from $z = v_d \Delta T$, hence *time projection*; the positional information of an interaction is projected onto the time dimension. In addition, the pattern of light observed in the top PMTs can be used to infer the x, y position of the interaction; due to the small width of the GXe region, most of the scintillation light produced in the electroluminescence region is highly localised in the PMTs closest to where the electrons were extracted. Finally, the observed S1 and S2 signals can be calibrated to provide a measure of the energy transferred in the interaction.

Operating a xenon detector in dual-phase mode has a number of advantages over single phase: amplification of the charge signal results in a high efficiency to detect single electron signals. This provides enhanced sensitivity to low energy events where only

[1]This assumes the drift field is uniform. In reality, the drift field may show radial and azimuthal components, in which case the depth can be recovered using an electric field map.

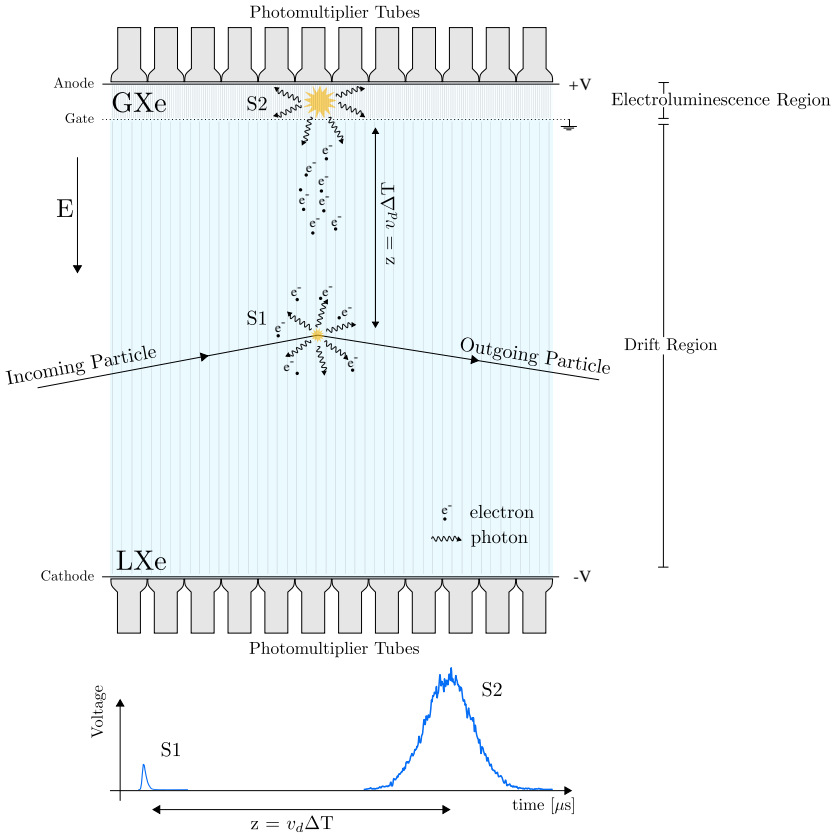

Figure 2.1: Schematic of an interaction in a dual phase xenon time projection chamber. The lower plot shows an example of the S1 and S2 signals observed with photomultiplier tubes.

a charge signal is detected, such as sub-GeV dark matter interactions [41]. Moreover, the x, y position reconstruction can have a resolution of $\mathcal{O}(mm)$ [42] which improves the ability to identify and reject background interactions that have a strong position dependence.

Current and previously operated dual phase xenon detectors that have utilised an active LXe mass of $\mathcal{O}(100$ kg), primarily for WIMP searches, are classified as first generation (G1) detectors. Planned, larger scale experiments using $\mathcal{O}(1$ tonne) of LXe form the class of second generation (G2) detectors. Although these detectors differ in scale, they all have a set of common characteristics:

Photo-detectors: observation of an interaction in a dual phase detector relies on the measurement of scintillation light. This is commonly achieved using photomultiplier tubes (PMTs) due to the existence of photocathodes that are sensitive to VUV light in addition to the low background construction and high efficiency to single scintillation photons. To increase the light collection of the detector two arrays of PMTs are commonly used: one at the bottom of the liquid region and another at the top of the gas region. Due to the refractive index of LXe, most of the S1 light is reflected at the liquid surface and subsequently detected in the bottom PMTs array. On the other hand, most of the S2 light is detected in the top array and is localised in the PMTs close to the extraction point of the electrons.

Electric fields: to create the electric fields necessary for electron drift and electroluminescence a set of three electrode wire grids are set at constant potentials: the cathode, which is positioned directly above the bottom PMTs; the gate which is situated just below the liquid level; and the anode which is found directly below the top PMTs. The potential difference between the cathode and gate defines the drift field and that between the gate and anode defines both the extraction field and electroluminescence field. Due to the difference in the dielectric constant between the LXe and GXe, the field between the gate and liquid surface is smaller than in the gas. The strength of this field must be large enough for drifted electrons to overcome the liquid-gas potential barrier and be extracted in to the gas. Correspondingly, the field in the gas must be large enough to initiate electroluminescence. To protect the PMTs from the high fields around the cathode and anode, a set of extra shield grids are placed between the cathode (anode) and bottom (top) PMTs to ensure the PMT photocathode is at ground. This creates a reverse field region below the cathode where the electrons drift away from the gate. Consequently, any interactions in this region do not produce an S2 signal.

Cooling: the triple point of xenon is 170 K at 1.6 bar and so cryogenic cooling systems are required to condense the xenon to the liquid state. To maintain this temperature, double-walled cryostats are used to contain the TPC and provide thermal insulation. Cooling can be provided by LN2, for example through the use of a thermosyphon system. This consists of a set of closed loop vertical tubes containing N_2 gas, where one end is enclosed in a LN2 bath above the detector and the other is attached to cooling elements of the inner cryostat vessel. LN2 flows down the tube to the cooling elements, extracts heat and boils off. The N_2 gas rises back to the LN2 bath where it condenses. This process continues until the temperatures reach an equilibrium.

Purification: Impurities in the xenon arising from radioactive progeny, leaks or outgassing can capture drifting electrons and diminish the S2 signal. To maintain a sufficient electron lifetime, the xenon must therefore be constantly purified. This is commonly achieved through a zirconium getter system that removes nearly all non-noble impurities. However, it is only capable of removing impurities from gas, so the LXe must also be continuously evaporated and condensed in a complex circulation system.

2.1.1 The Large Underground Xenon (LUX) Experiment

LUX [43] was a first generation (G1) dual-phase xenon experiment operated 4,850 ft underground in the Davis Cavern of Sanford Underground Research Facility (SURF), South Dakota. LUX operated between 2012 and 2016 with the primary goal of searching for weakly interacting massive particle (WIMPs), and placed world-leading limits on the WIMP-nucleon cross section [38]. There were two sustained periods of data taking for the WIMP analysis: Run-3 (WS2013) with a duration of $\sim$ 90 days, and Run-4 (WS2014-16) that lasted $\sim$ 600 days. The main parameters of the LUX detector are given in table 2.1.

TPC

The LUX TPC has a dodecagonal shape with a radius of 24 cm and drift region of 49 cm, containing 250 kg of active LXe. This is instrumented using a total of 122 PMTs, 61 above and 61 below. The cathode, gate and anode grids are made of thin steel wires and are kept at voltages of -10 kV, -1.5 kV and 3.5 kV, respectively. The gate grid is located approximately 5 mm below the liquid surface. This results in a drift field of 180 V/cm between the cathode and gate, a field of 2.9 kV/cm between the gate and liquid

Table 2.1: Parameters of the LUX experiment. The † denotes that these parameters are measured in Run-3.

Parameter	Value
TPC diameter	47 cm
TPC drift region	49 cm (322 μs)
PMTs	122 2" Hamamatsu R8778s
Active xenon mass	250 kg
Drift field	180 V/cm
Electroluminescence field	5.5 kV/cm
Electron extraction efficiency	49%†
Electron lifetime	650-950 μs
g_1	0.117† phe/photon
g_2	12.1† phe/e^-

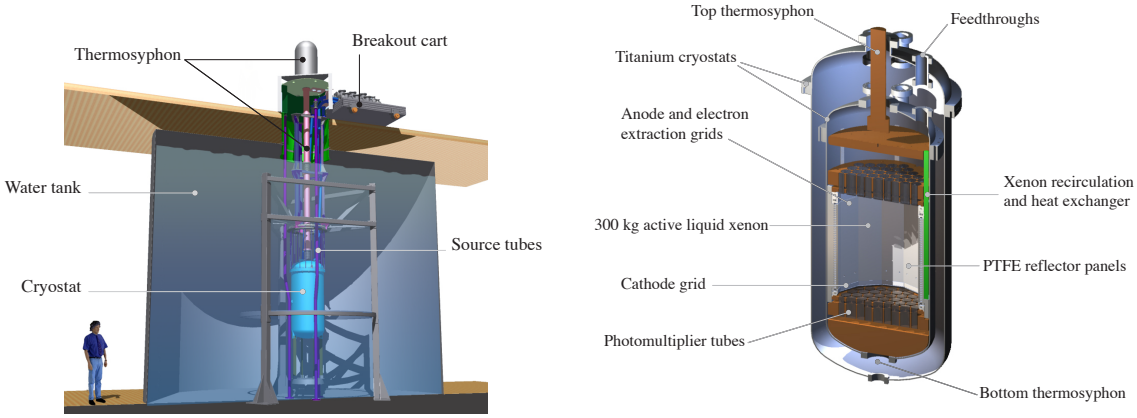

Figure 2.2: Left: Water tank containing the LUX cryostat. Right: The LUX TPC

surface for electron extraction, and a 5.5 kV/cm electroluminescence field between the surface and anode.

The walls of the TPC are made of polytetrafluoroethylene (PTFE) panels; an excellent reflector of VUV scintillation light that improves the light collection efficiency. The PTFE panels also house a series of copper field shaping rings, where the voltage is stepped from cathode to gate potentials through a resistor chain. This helps to maintain a uniform drift field in the active region.

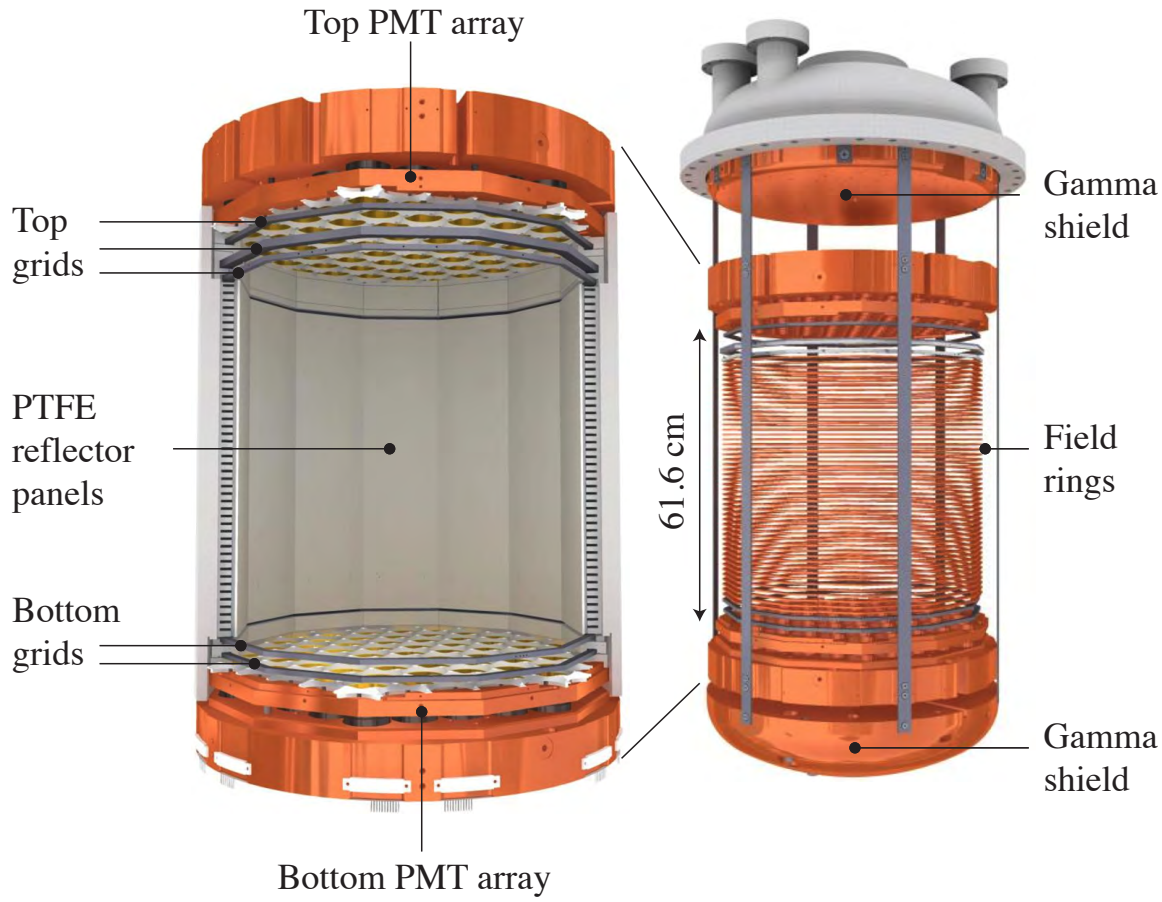

Figure 2.3: Left: Cut out view of the LUX TPC. Right: The TPC attached to the inner cryostat vessel flange.

The TPC is placed inside the inner vessel of a cryostat constructed of low background titanium, and is attached to the vessel's top flange through a copper disk that controls thermal gradients and also acts as shield for external γ-rays. Similarly, a 15 cm thick dome shaped copper block is located below the bottom PMT array to fill the empty inner vessel space and provide further shielding and temperature control. The inner vessel is situated within the outer vessel of the cryostat, where the space in between is held under constant vacuum to provide thermal insulation.

LUX uses 121 low background Hamamatsu R8778 2" PMTs, equally split between top and bottom arrays above the anode and below the cathode, respectively. The R8778 photocathode is sensitive to the 175 nm VUV scintillation light with a quantum efficiency (QE) of 33% at this wavelength. The current signal produced by the PMT is converted to a voltage signal by a 50 Ω load resistor, and with a nominal gain of 3.3×10^6 this results in a single photoelectron (SPE) pulse with an amplitude of ~ 4 mV.

Calibrations

LUX utilised a number of internal and external calibration sources to understand the response of LXe to electron and nuclear recoils over a range of energies. Self shielding makes calibration of the central active volume difficult, as most of the events occur close to the TPC walls. Instead, LUX pioneered the use of radioactive sources that are injected into the GXe during circulation. These sources eventually become uniformly distributed within the condensed LXe. Regular injections of $^{83}\mathrm{Kr}^{m}$ were used to determine the electron lifetime and to develop uniformity corrections for the S1 and S2 signals. This is a meta-stable isotope that decays via the emission of 32.1 keV and 9.4 keV conversion electrons with a half-life of 1.83 hours. The half-life is long enough for the source to become uniformly dispersed, but short enough that it does not result in a dominant background. During Run 3, an external ^{136}Cs source was used to study the high energy response via the detection of the 662 keV γ-ray. Further discussions of the electron recoil and nuclear recoil calibrations performed for the WIMP search can be found in [44].

Data acquisition

The LUX data acquisition (DAQ) system was designed to achieve a 95% detection efficiency of single SPE signals [45]; a requirement driven by the need to detect low energy WIMP interactions. Each analog PMT output is transmitted along coaxial cables

to detector breakouts where a set of pre-amplifiers provides impedance matching and amplifies the transmitted signals by a factor of 5. Post-amplifiers subsequently shape the signal and a provides a further 1.5× amplification. The post-amplifier output is split between an 14-bit analog to digital converter (ADC), discriminator and an FPGA trigger system [46]. The ADC boards are set at a sampling rate of 100 MHz (10 ns/sample) and operate in in pulse-only digitisation (POD) mode, a method of zero suppression that reduces the need to store the digitised baseline samples. In this mode, digitisation begins if the PMT signal rise above a threshold of 1.5 mV, and ends when the signal falls below 0.5 mV. An additional 24 samples before and 31 samples after the threshold cross are also digitised to ensure the full pulse is captured. The resulting digitised waveforms are referred to as PODs. The trigger system uses pulse shape and multiplicity information to select interesting events, and can detect detect S1 and S2 like pulses through two digital filters. Triggering can on occur on on either S1s, S2s, or both signals combined. For the S1 or the S2 mode, whenever the respective type of pulse is found a trigger signal is sent to the DAQ system. During Run 3 the trigger operated in S2 mode, and all detected pulses were recorded to disk. The application of the trigger information then occurred in software.

Water tank

The LUX cryostat is contained within a cylindrical water tank of height 6.1 m and radius 3.8 m. The large volume of water provides shielding against external γ-rays and neutrons produced from radioactive decays in the cavern rocks, in addition to neutrons produced during muon interactions. As a result, the external γ-flux is reduced by a factor of 8×10^{-9} by the time it reaches the TPC [43]. Further shielding is provided by a steel pyramid-shaped structure located directly beneath the detector in the cavern floor. The water is instrumented with 20, 10 inch PMTs that are capable of detecting Cherenkov radiation from incident muons. A coincident signal with an event in the TPC can then be used as a veto to remove background events. However, for Run 3 this veto system was not operational.

2.1.2 The LUX-ZEPLIN (LZ) Experiment

LUX-ZEPLIN (LZ) [47] is a second generation (G2) dual phase xenon physics experiment that will primarily search for WIMPs. The detector will be situated in the LUX water

tank, and, at the time of writing, detector construction is due to be completed in late 2020 with commissioning beginning shortly after.

LZ will utilise 7 tonnes of active LXe and has a design and operation that builds upon the experience gained from the LUX experiment. Background rejection is improved by the addition of an instrumented region between the TPC and inner cryostat, known as the xenon skin, together with a hermetic Gadolinium-loaded liquid scintillator (GdLS) outer detector system that surrounds the cryostat. These systems acts as anti-coincidence vetoes, providing the ability to reject background events that interact both in the TPC and these regions. The larger mass of LXe results in a two-fold improvement in the sensitivity to rare processes: there is a proportional increase in exposure, and a large increase in self shielding. Table 2.2 list the main parameters of the LZ TPC.

Table 2.2: Parameters of the LZ detector. The † denotes that these parameters are estimates from simulations.

Parameter	Value
TPC diameter	146 cm
TPC drift region	146 cm (890 μs)
PMTs	494 3" Hamamatsu R11410s
Active xenon mass	7,000kg
Drift field	310 V/cm
Electroluminescence field	10.2 kV/cm
Electron extraction efficiency	97.6%
Electron lifetime	$>$ 800 μs
g_2	$\sim 80^{\dagger}$ phe/e^-
g_1	$0.119^{\dagger}$ phe/photon

TPC

The LZ TPC is cylindrical with a diameter of 146 cm and drift region of 146 cm, containing 7,000 kg of active LXe. The cathode and gate grids are kept at potentials of -50 kV and -5.5 kV, respectively, creating a an average drift field of 310 V/cm. The anode voltage is $+5.5$ kV, and the gate grid is located approximately 8 mm below the liquid surface. This results in an extraction field of 5.2 kV/cm between the gate and liquid surface, and a 10.2 kV/cm electroluminescence field between the liquid surface and anode. The walls of

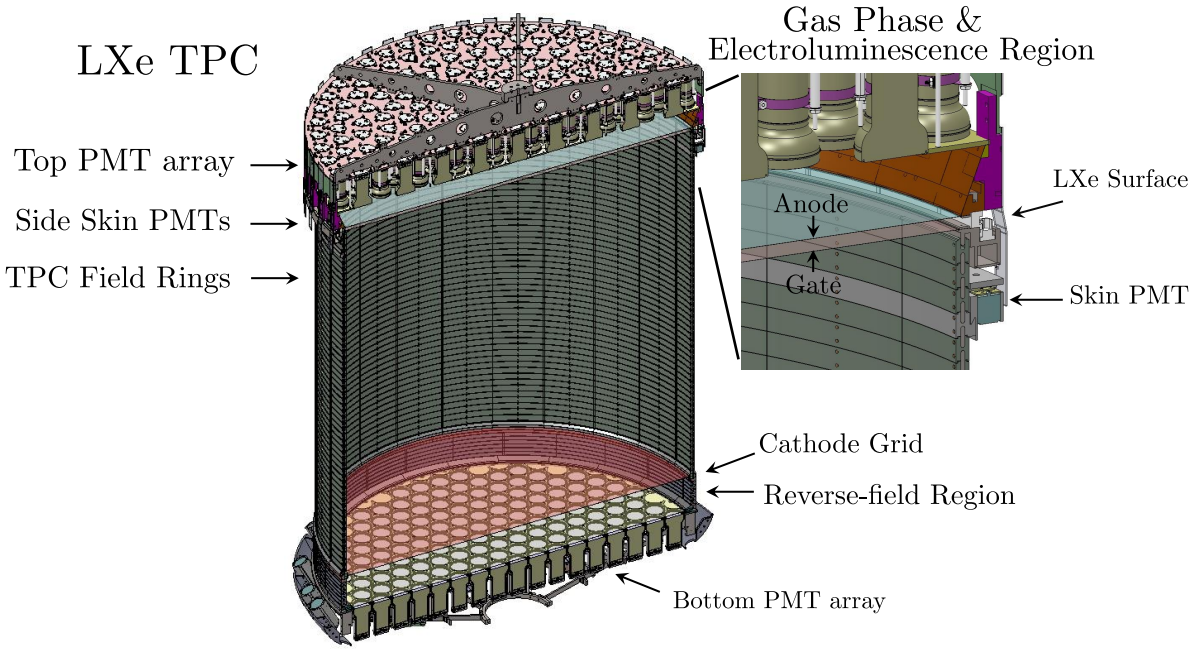

Figure 2.4: Schematic of the LZ TPC.

the TPC are made of reflective PTFE panels and surround the field cage. A total of 494 PMTs, 253 on the bottom array and 241 on the top array instrument the active region.

The TPC is placed inside the inner vessel of a cyrostat constructed of ultra radio-pure titanium [48]. Between the TPC and inner vessel is the xenon skin, a LXe region that is optically isolated from the main TPC. The inner vessel is situated within the outer vessel of the cryostat, where the space in between is held under constant vacuum to provide thermal insulation.

Photomultipliers

LZ will utilise 494 Hamamatsu R11410-20 3" PMTs to observe VUV scintillation light, with 253 on the top array and 241 on the bottom array. The R11410 photocathode is sensitive to the 175 nm VUV scintillation light with an average quantum efficiency (QE) of 32% [49]. The current signal produced by the PMT is converted to a voltage by a 50Ω load resistor, and with a nominal gain of 5.0×10^6 this results in a single photoelectron (SPE) pulse with an amplitude of ~ 7.3 mV.

To minimise the volume of inactive LXe, the skin system is instrumented with PMTs to detect scintillation produced from interactions in this region. 93 Hamamatsu R8250

1" PMTs are distributed in the top half of the skin, 20 Hamamatsu R8788 2" PMTs in the bottom half and 18 R8788 2" PMTs are used in the domed section of the skin below the bottom PMT array. Detection of scintillation light in this region that is coincident with an interaction in the main TPC can be used to veto multiple scatter backgrounds arising from γ-rays or neutrons.

Outer detector

Surrounding the LZ cryostat is the outer detector (OD) system. This is a set of acrylic tanks that contain $\sim 17,000$ kg of GdLS. As shown in figure 2.5, four large tanks enclose the cryostat from the sides and a further seven tanks enclose from the top and bottom. Interactions in the LS produce approximately 9,000 photons/MeV and are detected by 120 Hamamatsu R5912 8" PMTs located in the water tank. The purpose of the OD is for background rejection; background interactions from neutrons and γ-rays may interact in the active region of the TPC and leave the detector, however if the particle subsequently interacts in the OD it can be tagged and removed. Gadolinium has a high neutron capture cross section and during a capture event will emit a cascade of γ-rays. These are easily detected and provide improved efficiency for the detection of neutrons; a primary background for WIMP searches. In practice, any event in the TPC that has an OD signal in a pre-defined coincidence window and with an energy above a specific threshold will be vetoed. The threshold is determined by the scintillation yield and light detection efficiency of the LS together with the intrinsic OD background rate, which will result in accidental coincidences. For γ-ray backgrounds most events in the TPC and OD will be separated by less than 50 ns. In this small coincidence window, accidental coincidences from PMT dark rates and background events are negligible with a threshold as low as 100 keV [47].

Calibrations

LZ will have an extensive calibration campaign to study and calibrate the response of the detector over a range of energies. As with LUX, a number of dispersed internal sources, injected into the xenon gas during circulation, will be used to provide a uniform calibration of the active region. S1 and S2 uniformity corrections, together with electron lifetime determination, will be performed with regular injections of $^{131}\mathrm{Xe}^{m}$, a meta-stable isotope that decays via the emission of a 164 keV γ-ray. With a short-half life of 11.9 d,

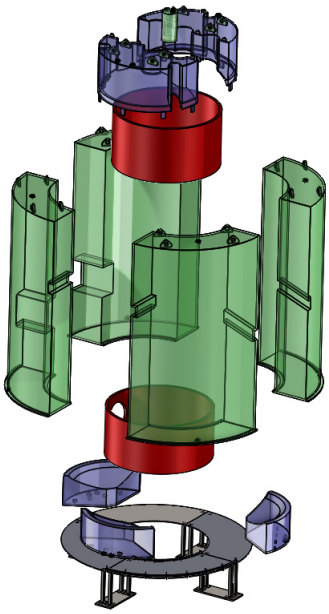

Figure 2.5: Exploded view of the LZ outer detector system, illustrating the scintillator tanks that surround the cryostat. Taken from [47].

this source decays away rapidly and does not contaminate the LXe for extended periods. With a volume of LXe that is much larger than LUX, uniform dispersion of the source occurs over longer timescales. As a result, $^{83}Kr^{m}$ calibrations will be less effective as many decays will occur before it becomes uniformly distributed in the LXe.

Several external calibration sources are deployed through source delivery pipes inserted between the inner and outer cryostat vessels. The timing calibration of the skin and OD detectors and synchronisation with the TPC is with carried out with a ^{22}N source that emits two back-to-back 511 keV γ-rays. The OD energy scale will be calibrated with a ^{228}Th source using the emission of a 2,614 keV γ-ray from the ^{208}Tl daughter. Moreover, this can also be used to calibrate the high energy response of the active xenon region.

Data acquisition

The LZ DAQ builds upon the lessons learned from LUX and retains many of the same features. A key difference is the use of a single amplifier that operates in dual gain mode and simultaneously produces low gain (LG) and high gain (HG) signals. This is motivated by the longer electroluminescence region in LZ that results in larger S2 signals which will more readily saturate the electronics. To minimise saturation but retain a reasonable SPE detection efficiency, the LG mode amplifies the signal by a factor of 4.

This provides an improved reconstruction of high energy events. On other hand, the HG mode amplifies by a factor of 40 and provides at least 90% detection efficiency for SPEs [47]. The Skin and OD PMTs use the same amplifier but only operate in HG mode to reduce the number of channels.

The amplified signals are digitised at 100 MHz (10 ns per sample) with a 14-bit resolution, using custom built FPGA-based DDC-32 digital signal processors. The digitisation occurs in POD-mode and has built-in digital S1 and S2 filters used for triggering. More detailed information about the signal processing and acquisition can be found in [50].

2.2 Liquid Xenon Physics

Xenon (Xe) is a naturally occurring noble gas that is an attractive medium for radiation detection. Experiments utilising Xe in either it's gas (GXe), liquid (LXe) or both phases (dual-phase) have consistently produced world leading results in searches for rare processes beyond the Standard Model, such as low energy dark matter interactions [38] [39] and neutrinoless double beta decay [51]. The success of these experiments relies on the key characteristics of Xe:

Scintillation and ionisation: energy loss from ionising particles in Xe results in the production of scintillation and ionisation with excellent efficiency; a 1 MeV electron will on average produce approximately 50,000 scintillation photons in LXe [52]. With a low ionisation W-value of 15.6 eV [53] charge is also readily produced and can be extracted by an electric field. The detection of these signals can then be used to reconstruct the energy and position of the interaction. Due to the high light yields, in addition to the anti-correlation with ionisation (see figure 2.6), an energy resolution of $\sigma/E < 1.5\%$ [54] (0.5%) [55] at 2.5 MeV is achievable in LXe (GXe), resulting in the ability to search for rare mono-energetic signals and constrain backgrounds through spectroscopy.

Intrinsic radioactivity: naturally occurring xenon has an intrinsically low background rate with only three unstable isotopes: ^{124}Xe, decaying via double electron capture, in addition to the $2\nu\beta\beta$ isotopes ^{134}Xe and ^{136}Xe. Each of these isotopes has a relatively low abundance, and with half-lives $> 10^{21}$ years they only become a significant background in experiments that utilise tonnes of Xe. Moreover, Xe can be readily purified to remove radioactive impurities arising from external contamination [57].

Particle discrimination: the ionisation and scintillation mechanisms are dependent on the energy and species of the ionising particle [59]. This results in different yields and scintillation decay constants and forms the basis of signal-background discrimination in dark matter searches. This can result in $> 99.5\%$ background rejection [60] for experiments utilising both GXe and LXe.

Self-shielding: With a large atomic mass and high density as a liquid, approximately $\sim 3 \mathrm{g/cm}^3$ at 170K and 1.8 bar, Xe is a strong attenuator of radiation. As a result, photons with energy below 1 MeV are unlikely to penetrate more than 1 mm before interacting, and most will be absorbed through the photoelectric effect as seen in figure 2.7. Higher energy photons are more likely to Compton scatter within 10 cm. Likewise, figure 2.8 demonstrates that electrons rapidly lose energy through scattering and consequently rarely travel further than 1 cm in LXe. A detector utilising a large mass of LXe will therefore have a central volume that is naturally shielded from external radiation.

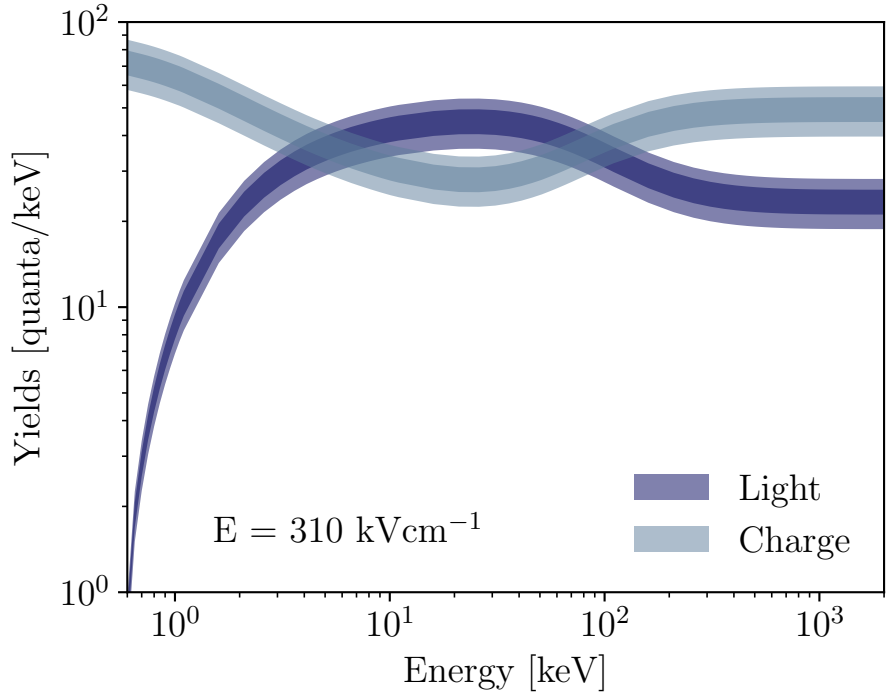

Figure 2.6: Predicted light and charge yields of e^- interactions in LXe at a field of 310 V/cm, according to the global fit obtained from NEST [56]. The shaded bands represent the 1σ and 2σ confidence intervals.

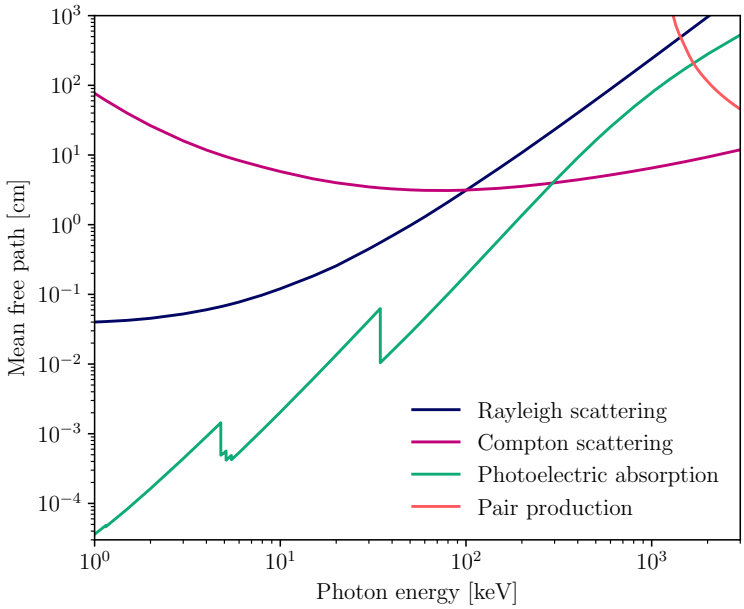

Figure 2.7: The mean free path of photons in LXe, assuming a liquid density of 3.0 g/cm^3, for different electromagnetic processes. Data obtained from [61].

Together, these properties result in a detection medium that is intrinsically low background with a sensitivity to energy transfers of $\mathcal{O}$(keV). However, efficiently detecting these interactions with an instrumented detector requires an understanding of the physical processes involved during a particle interaction in Xe.

Table 2.3: Abundance of naturally occurring xenon isotopes, in addition to their known decay mechanisms. Abundance data obtained from [58]. The † represents that only a lower limit on the half-life has been measured for this process.

Isotope	Abundance	Decay mechanism
^{124}Xe	0.095%	$2\nu\epsilon\epsilon$
^{126}Xe	0.089%	stable
^{128}Xe	1.91%	stable
^{129}Xe	26.4%	stable
^{130}Xe	4.07%	stable
^{131}Xe	21.2%	stable
^{132}Xe	26.9%	stable
^{134}Xe	10.4%	$2\nu\beta\beta^{\dagger}$
^{136}Xe	8.86%	$2\nu\beta\beta$

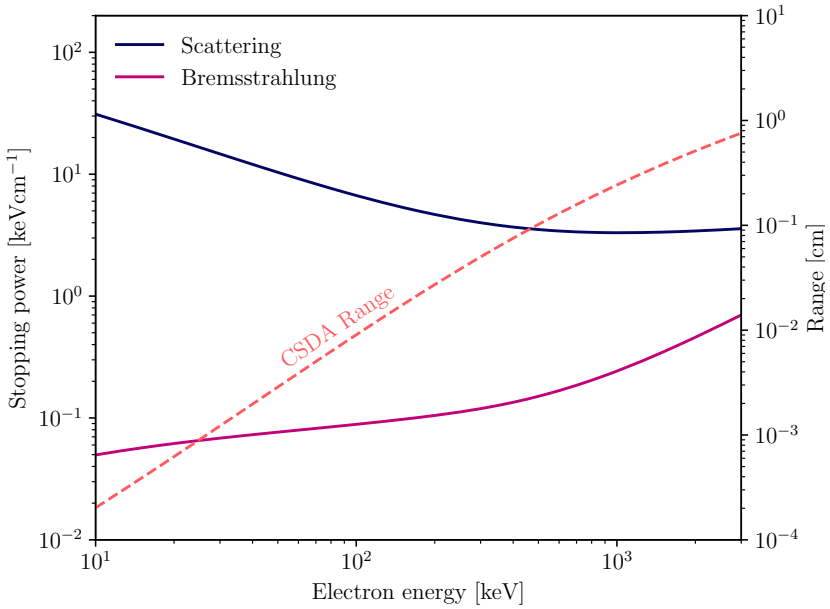

Figure 2.8: The stopping power of e^- in LXe for scattering and bremsstrahlung emission, assuming a liquid density of 3.0 g/cm^3. The right y-axis shows the corresponding range calculated using the continuously slowing down approximation (CSDA). Data obtained from [62].

2.2.1 Particle Interactions in Liquid Xenon

An ionising particle that passes through LXe can lose energy E_0 to the medium through two processes: *electron recoils*, where Coulomb interactions with electrons results in ionisation and excitation of atoms, and *nuclear recoils*, where the particle elastically scatters with a Xe atom and increases the thermal energy of the medium. The distribution of energy across these processes can be described by the Platzman equation [63]:

$$E_0 = N_{ex}E_{ex} + N_iE_i + \text{heat}. \tag{2.1}$$

Here, E_i and E_{ex} are the average energies required to ionise and excite an atom, N_i is the number of ion-electron pairs and N_{ex} is the number of excited atoms. Energy transferred to atomic motion or to electrons with an energy below the excitation or ionisation threshold contributes only to the thermal energy of the liquid. The N_{ex} excited atoms will eventually scintillate to produce N_{ph} photons, and in the presence of an external electric field N_e electrons can be extracted. This results in two observable signals: light

and charge. Predicting the yields of these signals is a difficult task due to the complex processes involved in an interaction:

Recombination: ion-pairs recombine to produce more excitation, resulting in a diminished charge signal. Under the assumption of complete recombination, the total number of photons is $N_{ph} = N_i + N_{ex}$. The probability of recombination is highly dependent on the ionisation density, track structure and electric field [64]. Generally speaking, particles with low $-dE/dx$ such as γ-rays produce sparse tracks where the ionised electrons must be within the Osnager radius to recombine. On the other hand, particles with large ionisation density such as α particles have much more rapid recombination. In the presence of an external field recombination is reduced, resulting in a higher charge yield. This results in a scintillation and charge yield that is heavily energy dependent, as seen in figure 2.6.

Excitation to ionisation ratio: the ratio N_{ex}/N_i directly determines the observed yields, and appears to have a dependence on the type of interaction. For electron recoils, values of 0.06-0.2 [52][53] are measured whereas for nuclear recoils the measurements tend to favour a value closer to unity [65]. It is unclear what physical mechanism may result in this difference.

Quenching: any energy transferred in elastic nuclear recoils results in lower charge and light yields. For electron recoils the light yield can be quenched in both the low and high ionisation density regime. For very low $-dE/dx$ it is possible for ionisation electrons to escape recombination with an ion, resulting in a lower light yield. For high $-dE/dx$, the light yield is decreased through the process of bi-excitonic quenching [66].

Observation of the light and charge signals can be used to infer the N_i and N_{ex} produced in an interaction, and therefore the corresponding energy transfer E_0. In a dual-phase detector the light yield is measured through the detection of scintillation photons arising from the initial excitation and recombination, or *primary scintillation*. The charge yield is measured from the detection of scintillation light produced during the extraction of drifted electrons in to the gas phase, or *secondary scintillation*.

Primary Scintillation

Primary scintillation light is produced during the relaxation of an excited diatomic xenon molecule Xe_2^*, and is emitted with a peak vacuum ultraviolet (VUV) wavelength of 175 nm. Xe_2^* production arises from two distinct processes [67]. The first involves the direct

excitation of a Xe atom by an ionising particle A,

$$A + \mathrm{Xe} \rightarrow \mathrm{Xe}^{*} + A \tag{2.2}$$

$$\mathrm{Xe}^{*} + \mathrm{Xe} \rightarrow \mathrm{Xe}_{2}^{*\nu} \tag{2.3}$$

The superscript * (ν) denotes that the excited diatomic molecule is produced in a excited electronic (vibrational) state. The relaxation to the vibrational ground state is non-radiative, and occurs through collisions with neighbouring Xe atoms, where as the electronic relaxation proceeds via the emission of a scintillation photon γ and subsequent dissociation of the molecule due to a repulsive ground state:

$$\mathrm{Xe}_{2}^{*\nu} + \mathrm{Xe} \rightarrow \mathrm{Xe}_{2}^{*} + \mathrm{Xe} \tag{2.4}$$

$$\mathrm{Xe}_{2}^{*} + \mathrm{Xe} \rightarrow \mathrm{Xe} + \mathrm{Xe} + \gamma, \tag{2.5}$$

The second process that results in scintillation light production comes from the recombination of electrons and Xe^{+} ions through the following steps:

$$\mathrm{Xe}^{+} + \mathrm{Xe} + \mathrm{Xe} \rightarrow \mathrm{Xe}_{2}^{+} + \mathrm{Xe} \tag{2.6}$$

$$e^{-} + \mathrm{Xe}_{2}^{+} \rightarrow \mathrm{Xe}^{*\nu} + \mathrm{Xe} \tag{2.7}$$

$$\mathrm{Xe}^{*\nu} + \mathrm{Xe} \rightarrow \mathrm{Xe}^{*} + \mathrm{Xe} \tag{2.8}$$

$$\mathrm{Xe}^{*} + \mathrm{Xe} + \mathrm{Xe} \rightarrow \mathrm{Xe}_{2}^{*} + \mathrm{Xe} \tag{2.9}$$

$$\mathrm{Xe}_{2}^{*} + \mathrm{Xe} \rightarrow \mathrm{Xe} + \mathrm{Xe} + \gamma \tag{2.10}$$

In both cases, the resulting scintillation photon arises from the radiative decay of the excited molecular $^{1}\Sigma_{u}^{+}$ singlet state or $^{3}\Sigma_{u}^{+}$ triplet state to the $^{1}\Sigma_{g}^{+}$ ground state, which have lifetimes of 2–4 ns and 21–28 ns [68] [69], respectively. Nevertheless, recombination is a slower process than direct excitation as it relies on the thermalisation time of the electrons and the surrounding ionisation density, resulting in an apparent third component of scintillation light with a lifetime of 30–45 ns [70]. However, in the presence of a strong electric field the recombination component is diminished and the distribution of the scintillation light is determined by the singlet and triplet components. The relative population of singlet and triplet states is dependent on the $-dE/dx$ in addition to the dynamics of recombination; high energy electrons are observed to populate mostly the triplet states, whereas larger $-dE/dx$ species such as α particles populate mostly singlet states. This produces an observable difference in the shape of the observed scintillation

pulse, and is the basis for the method of particle identification in Xe using pulse shape discrimination [60].

Combining the two processes of primary scintillation light production, the total number of photons N_{ph} produced in an interaction with a recombination probability of r is

$$N_{ph} = N_{ex} + rN_i \tag{2.11}$$

and the number of drifted electrons is therefore

$$N_e = (1 - r)N_i. \tag{2.12}$$

The efficiency to produce scintillation light is measured by the yield $\mathcal{Q}_l = N_{ph}/E_0$, or alternatively the W-value $W_{ph} = E_0/N_{ph}$. These are electric field, energy and particle-dependent quantities due to the differences in recombination. Using the fact that the charge and light yields are anti-correlated but the total number of quanta remains constant (i.e $N_{ph} + N_e = N_{ex} + N_i$), with zero electric field and assuming that all ion-electron pairs recombine, the maximum number of scintillation photons is $N_{ph} = N_{ex} + N_i$. Combing this with the ionisation W-value $W = \frac{E_0}{N_i}$ gives

$$W_{ph}^{min} = \frac{E_0}{(N_{ex} + N_i)} = \frac{W}{(1 + N_{ex}/N_i)}, \tag{2.13}$$

which can be interpreted as the minimum energy needed to produce a single scintillation photon. Assuming N_{ex}/N_i is independent of energy, this has a constant value of 13.7 $\pm$ 0.2 eV [64]. A simple rearrangement results in an expression for the energy transferred to ionisation and scintillation, or the *combined energy formula:*

$$E_0 = W_{ph}^{min}(N_{ex} + N_i) = W_{ph}^{min}(N_{ph} + N_e). \tag{2.14}$$

With the assumption that W_{ph}^{min} is energy independent, the measurement of both N_{ph} and N_e can be used to determine an estimate of the energy that is linear and independent of recombination. The anti-correlation between N_{ex} and N_i effectively removes the statistical recombination fluctuations and results in an energy scale with improved energy resolution, compared to using N_e or N_{ph} alone.

It should be stressed that this only measures the energy transferred to ionisation and excitation, and does not provide an estimate of the energy transfer to atomic motion

due to elastic scattering. In the case of nuclear recoils, the energy must be estimated with the addition of a quenching factor $\mathcal{L}$, defined as the fraction of energy transferred to excitation and ionisation from a recoiling nucleus:

$$E_0 = \frac{W_{ph}^{min}}{\mathcal{L}}(N_{ph} + N_e). \tag{2.15}$$

It has been shown that $\mathcal{L}$ can be estimated using the theory of Lindhard [65], and for $\mathcal{O}(10 - 100$ keV) nuclear recoils approximately 20% of the energy is transferred to excitation and ionisation, with the rest lost to heat.

The N_{ph} primary scintillation photons can ultimately be detected by VUV sensitive detectors such photo-multiplier tubes (PMT) that convert the photons into an observable current. The total measured signal, commonly denoted by $S1$ and given in units of the number of photons detected (phd), is then given by

$$S1 = g_1 N_{ph} \tag{2.16}$$

where g_1 is the photon detection efficiency of the detector. Due to the quantum efficiency of the light detectors and attenuation of VUV photons, either by absorption or Rayleigh scattering, typical values of g_1 for dual-phase detectors are in the range $0.05 - 0.15$ phd/photon. Due to the $\mathcal{O}(10$ ns) scintillation process, the S1 light is observed as a prompt signal with a rapid fall time corresponding to the lifetime of the excited states.

Secondary Scintillation

In a dual phase detector, the charge signal is measured by drifting the ionised electrons to a liquid-gas interface where they are extracted into the GXe region. The stronger electric field $\mathcal{O}(10$kV/cm) drifts electrons through the gas resulting in the production of further scintillation light through the process of electroluminescence (EL). This light is termed *secondary scintillation* to distinguish it from the light produced during the initial interaction.

Secondary scintillation light arises from the same mechanism as primary scintillation, where the excited Xe_2^* molecules are produced during collisions with the drifting electrons.

The mean EL light yield per path length for an electron drifting in GXe of pressure P under a field E can be estimated from,

$$\frac{dN_{EL}}{dx} = \alpha E - \beta P - \gamma \tag{2.17}$$

where $\alpha = 0.137$ V^{-1}, $\beta = 177$ bar cm^{-1} and $\gamma = 45.7$ cm^{-1} [71]. Therefore, at a typical operating pressure of 1.6 bar and an electric field of 10 kV a single electron can result in the emission of thousands of scintillation photons in a 1 cm GXe region. This charge amplification process results in sensitivity to single electrons and consequently very low energy interactions.

As with primary scintillation, the EL light can be detected resulting in a second signal, denoted S2 and again measured in units of photoelectrons. The total observed signal is given by

$$\mathrm{S2} = \varepsilon g_1^{gas} N_{EL} N_e = g_2 N_e \tag{2.18}$$

where ε is the efficiency to extract electrons into the gas gap and g_1^{gas} is the photon detection efficiency in the gas phase. The product of N_{EL} and g_1^{gas} represents the average pulse area of a single extracted electron $\langle SE \rangle$, therefore these factors are commonly grouped together to form the gain factor $g_2 = \varepsilon \langle SE \rangle$.

The observed shape of the S2 signal is dominated by diffusion. In the presence of an $\mathcal{O}(100$ V/cm) electric field, electrons will drift with a constant velocity of $v_d \sim 1$ mm/μs [72]. The spatial distribution of electrons can described as point source that evolves under the influence of longitudinal and transverse diffusion, and therefore the spread of an electron cloud after drifting a time Δt is characterised by

$$\sigma_{L,T} = \sqrt{2 D_L \Delta t} \tag{2.19}$$

Under the influence of an electric field, the diffusion is anisotropic with $D_T > D_L$ [73]. The longitudinal diffusion component results in extracted electrons entering the GXe with a spread of times, and consequently the shape of observed S2 pulse is roughly Gaussian with a width that is proportional to the drift time. This limits the precision of the z-coordinate reconstruction of an event and reduces the ability to identify events with multi site energy deposits, e.g multiple Compton scattering of γ-rays. The S2 signal shape can also complicated by delayed extraction of electrons into the gas phase in addition to photo-ionisation of the electric field grids. Due to the large potential barrier

between the liquid and gas phases, a population of electrons can remain trapped at the interface for extended time periods of $\mathcal{O}(10ms)$ [74]. Moreover, release of captured electrons from impurities can also result in an even slower component of the S2 light [75]. For high energy interactions this can result in a long trail of delayed S2 light after the main pulse that can have adverse affects on the energy resolution.

The total S2 signal can also be diminished by electron attenuation. Electronegative impurities present in the LXe, predominantly O_2, can capture electrons as they drift. Given an initial N_e^0 electrons, the the resulting charge that reaches the liquid-gas interface after drifting for a time Δt is given by

$$N_e = N_e^0 e^{-\lambda \Delta t} \tag{2.20}$$

The resulting electron lifetime is then $\tau = 1/\lambda$. Constant circulation and purification of the liquid xenon is required to remove impurities which would otherwise attenuate charge signal, a long with regular calibrations to monitor the electron lifetime. Typical electron lifetimes that have been achieved in liquid xenon detectors are $\mathcal{O}(500\mu s)$, approximately a factor of two longer than the maximum drift time in those detectors. For tonne-scale experiments with longer drift regions the effect of electron lifetime can severely affect the efficiency for the detection of small signals.

Chapter 3

Search for ^{136}Xe $0\nu\beta\beta$ in LUX

This chapter presents a search for ^{136}Xe $0\nu\beta\beta$ using the LUX 2013 WIMP search (Run-3) data set. As a first generation dual phase xenon experiment, the LUX detector and event reconstruction was optimised for searches of low energy, WIMP-induced nuclear recoils. In addition, the limited active mass of LXe contains only 22 kg of ^{136}Xe. With $T_{1/2} > 1.07 \times 10^{26}$, this corresponds to < 0.015 $0\nu\beta\beta$ events in the full 90 days of Run-3. An analysis of the LUX data is therefore not expected to discover this process, nor set an improved limit on the half-life. However, the aim of this chapter is to demonstrate the necessary event reconstruction, corrections and calibration methods required to perform a successful $0\nu\beta\beta$ search in a future dual phase xenon detector.

For this analysis an entirely new data processing framework was developed: LUXhep, or LUX High Energy Processing. The aim of LUXhep is to enhance the reconstruction of $\mathcal{O}$(MeV) interactions in LUX and to produce significant improvements in background rejection and energy resolution - both key parameters in $0\nu\beta\beta$ searches.

3.1 High Energy Event Reconstruction

All information about interactions that occur in the LXe are contained in the pulses recorded by the PMTs as they detect the S1 and S2 light. The identification and reconstruction of interactions in LUX can generally be decomposed into the following steps:

Pulse processing: following a trigger, individual PMT waveforms (PODs) that fall within a coincidence window are clustered together. Within each cluster, distinct pulses are identified using a pulse finding algorithm.

Pulse classification: each identified pulse is characterised by a set of quantities that describe the size and shape of the pulse, e.g the area, amplitude and timing information. These quantities are used to determine if the pulse corresponds to the production of S1 or S2 light, or whether it can be attributed to other processes such as the observation of a single electron or noise from the PMTs.

Pulse pairing: Identified S1 and S2 pulses are matched together, allowing the determination of the z position from the time difference between the S1 and S2 pulse, and the x, y position from the pattern of the PMT response observed during the S2 pulse. The S1 and S2 pulse areas correspond to the number of scintillation photons and electrons produced during the interaction, and the energy can therefore be calculated using equation 3.19.

For the WIMP, and other low-energy analyses performed in LUX, these tasks were accomplished using the data processing framework (DPF) described in [44]. This consists of a set of C++, Matlab and Python modules that were tuned for optimal reconstruction of $\mathcal{O}$(1-100 keV) events within the WIMP search region. However, the DPF suffers from two major pathologies that are detrimental to the accurate reconstruction of interactions with energy $>$ 1 MeV: the identification of S2 pulses, and the subsequent estimate of the pulse areas. These issues are clearly evident in figure 3.1, which shows an example of the summed PMT waveforms obtained from a 2.3 MeV multiple Compton scatter event.

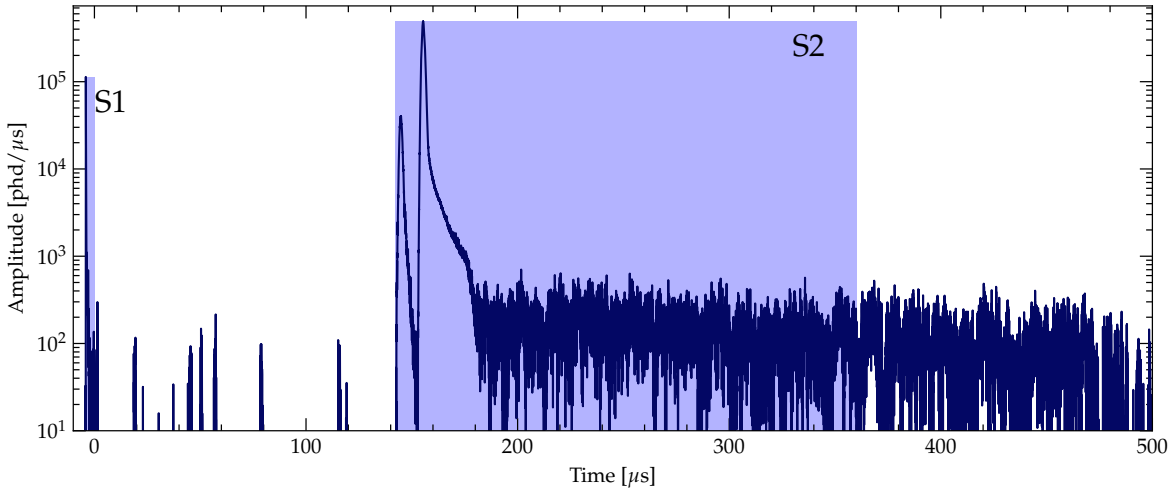

Figure 3.1: A 2.3 MeV γ-ray interaction in LUX. The shaded regions show the reconstructed S1 and S2 pulse boundaries obtained from the LUX data processing framework.

It is clear that that the pulse finder of the DPF fails to detect the presence of the two S2 pulses and instead they are clustered together. This event would therefore pass a

simple signal selection criteria that requires events with a single S2 pulse, resulting in a poor false positive rate. In addition, the x, y position of the interaction will also been inferred from the combination of both pulses, leading to a potentially erroneous position reconstruction. It can be seen that a proportion of the tail is attributed to the S2 pulse in an unsystematic manner. As discussed in section 2.2, at this energy the tail consists of a near-continuous stream of single electron signals arising from photo-ionisation of electrode grids and impurities in the LXe, together with the delayed emission of electrons trapped at the liquid surface. The inclusion of the tail therefore increases the variance in the estimate of the S2 pulse area and ultimately degrades the energy resolution.

Aside from the issues with the DPF, the LUX electronics chain suffers from non linear effects during large signals, as shown in figure 3.2 for the same 2.3 MeV event. These effects are predominantly observed in the top PMTs that that see the majority of the S2 light. The first effect arises from the signal exceeding the dynamic range of the digitiser, resulting in a truncated signal, and the second is a result of capacitor depletion in the PMT voltage divider circuit. Together, these cause a reduction in the observed S2 signal and a subsequent bias in the energy reconstruction that worsens with increasing signal size. From a signal processing perspective these issues are not trivial to correct. Instead, only the bottom PMTs are used for energy estimation as the observed S2 light is spread out across the entire array, resulting in significantly less saturation. The issues of pulse

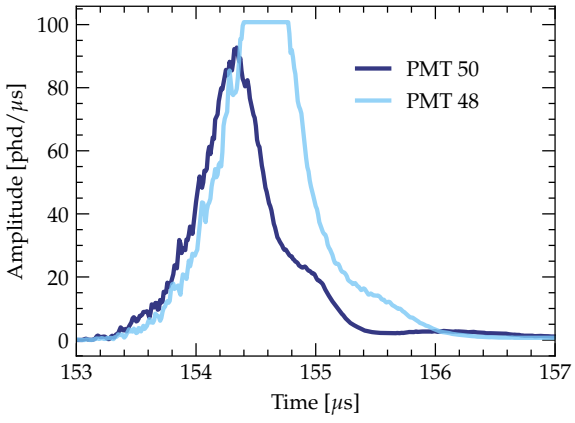

Figure 3.2: Saturation effects in two top PMTs observed during the S2 pulse. PMT 48 shows truncation due to the voltage exceeding the dynamic range of the digitiser in addition to capacitor depletion, identified by the rapid decay of the pulse. PMT 50 is within the digitiser range but still displays capacitor depletion.

classification and S2 pulse area estimation present in the LUX DPF necessitated the need to develop a new pulse finding algorithm that improves the ability to identify S2

pulses and provides a reduced variance estimate of the S2 pulse area. The results in both an increased efficiency to reject multiple scatter backgrounds and an improvement in the energy resolution. As a result, for the analysis of $\mathcal{O}$(MeV) interactions an entirely new data processing framework was developed, known as LUXhep.

LUXhep is a Python-based object orientated modular framework that performs the key processing and reconstruction of the signals from the LUX detector. The framework is built around a persistent *Event* object that stores all the required information from an event. It implements the following configurable modules:

- **EvtReader**: reads the digitised PMT waveforms stored in a binary format. Each trigger results in the acquisition of 500 μs of 122 PMT channels in zero-suppressed POD mode, sampled at 100 MHz. The data is stored as a sequence of individual PODs for each channel that also contains the start and end times and baseline.
- **PODCalibrator**: converts the digitised PMT voltage to photons detected (phd) per sample using pre-measured PMT gains [38]. This allows for measurement of the number of photons detected in a PMT, taking into account the probability for one photon to produce two photoelectrons [76].
- **PODClusterer**: groups together PODs that are separated in time by fewer than 15 ns.
- **PulseFinder**: identifies the start and end of distinct pulses within the clustered PODs that correspond to an underlying production process.
- **PulseParameteriser**: calculates key quantities of each pulse such as the top and bottom PMT areas, width and timing characteristics.
- **PulseClassifier**: uses the pulse characteristics to determine whether the pulse originates from S1 or S2 light, or whether it can can be attributed to other processes such as noise.
- **S2TailSeparator**: performs improved pulse finding on S2 pulses and removes a significant proportion of the S2 tail.
- **S2AreaIntegrator**: determines improved S2 pulse boundaries for area estimation by fitting a template function to the pulse.
- **PositionReconstructor**: infers the x, y position of an interaction using a maximum likelihood method that operates on the top PMT array pattern of an S2 pulse.

- **InteractionFinder**: pairs together S1 and S2 pulses within the 320 μs drift window and calculates the z position of the interaction from the difference in their arrival times.

After processing, the Event object is persisted on disk for further analysis through the following modules:

- **AwkwardWriter**: writes Events in a jagged-array format using Awkward-array [77]. This allows for fast, vectorised operations on the low-level Event attributes, such as individual pulses or PODs.
- **DataFrameWriter**: outputs attributes from Interactions as a columnar Pandas dataframe [78], with each row representing a single Interaction. Only the relevant quantities of an event are persisted, resulting in faster analysis and reduced disk space usage.

LUXhep is configured using a human-readable YAML file that specifies the sequence of modules to be used in the processing and their adjustable parameters. This makes it trivial to perform alternate processings for different datasets.

3.1.1 Pulse Finding

Pulse finding is at the core of any successful interaction reconstruction; identified pulses are used to determine the position, energy and multiplicity of an interaction, and these quantities are subsequently used to make inferences about the underlying physical process. LUXhep introduces a number of key improvements to the pulse finding compared to that used in the original LUX DPF that improves the identification of multiple S2 pulses and provides a reduced variance estimate of the S2 pulse area.

The LUXhep PulseFinder module uses a simple yet robust algorithm based on the clustering of local minima and maxima in the summed PMT waveform of a clustered POD. The algorithm proceeds via the following steps:

1. The summed waveform is smoothed using a Savitzky-Golay filter [79]. This fits a series of k-degree polynomials to the waveform samples in a sliding window of length L, resulting in a waveform with increased signal-to-noise ratio and reduction in local extrema due to noise.

2. The amplitude at each sample of the smoothed waveform is compared with that of it's n closest neighbours to determine the local maxima and minima. Only the maxima with an amplitude above a threshold T are retained in order to remove noise.

3. Each maximum is paired with it's closest minima that occurs before and after. This forms a proto-pulse P with a start t and end time t' determined by the time of respective minima.

4. Each proto-pulse P_i is clustered with the previous if it's start time is within d samples of the end time of P_{i-1}, extending the end time of P_{i-1} to t_i

This results in a set of identified pulses $\{P_0, ..., P_i\}$ for an event, as shown in figure 3.3. The set of parameters used for the high energy event processing are given in table 3.1, and were optimised based on the pulse finding performance assessed by a visual hand-scan of events.

Table 3.1: Parameters of the PulseFinder module used for first pass of pulse finding.

Parameter	Value	Description
k	2	degree of polynomial used in smoothing filter
L	11	length of smoothing window
n	5	number of neighbours for local extrema
T	0.05	local maxima threshold
d	50	clustering window for nearest pulse

A major difficulty with developing a pulse finding algorithm is that it must be able to separately identify S1 and S2 pulses, which both have very different time profiles, and must also perform efficiently over a wide range of pulse amplitudes. In addition, the identification of small spurious pulses following an S1 or S2 such as after-pulsing, SPE and SEs is necessary to reduce the variance in the pulse area. These signals are often orders of magnitude smaller than main S1 and S2 pulses. It was found that a "one-size fits all" approach to pulse finding was unable to efficiently identify multiple S2 pulses at the same time as identifying smaller pulses. Instead, a second round of pulse finding with alternate parameters is used on "S2-like" pulses, and is performed by the S2TailSeparator module. This is performed on pulses with an amplitude greater than 50 phd and a length greater than 800 samples, which translates to a loose selection of S2 pulses. The parameters used for this second round of pulse finding are given in table 3.2. The S2TailSeparator significantly enhances the S2 pulse finding, and is able to identify separate S2 pulses and can perform clustering of the tail.

Table 3.2: Parameters of the S2TailSeparator module used for the second pass of pulse finding of the S2.

Parameter	Value	Description
k	NA	degree of polynomial used in smoothing filter
L	NA	length of smoothing window
n	50	number of neighbours for local extrema
T	20	local maxima threshold
d	0	clustering window for nearest pulse

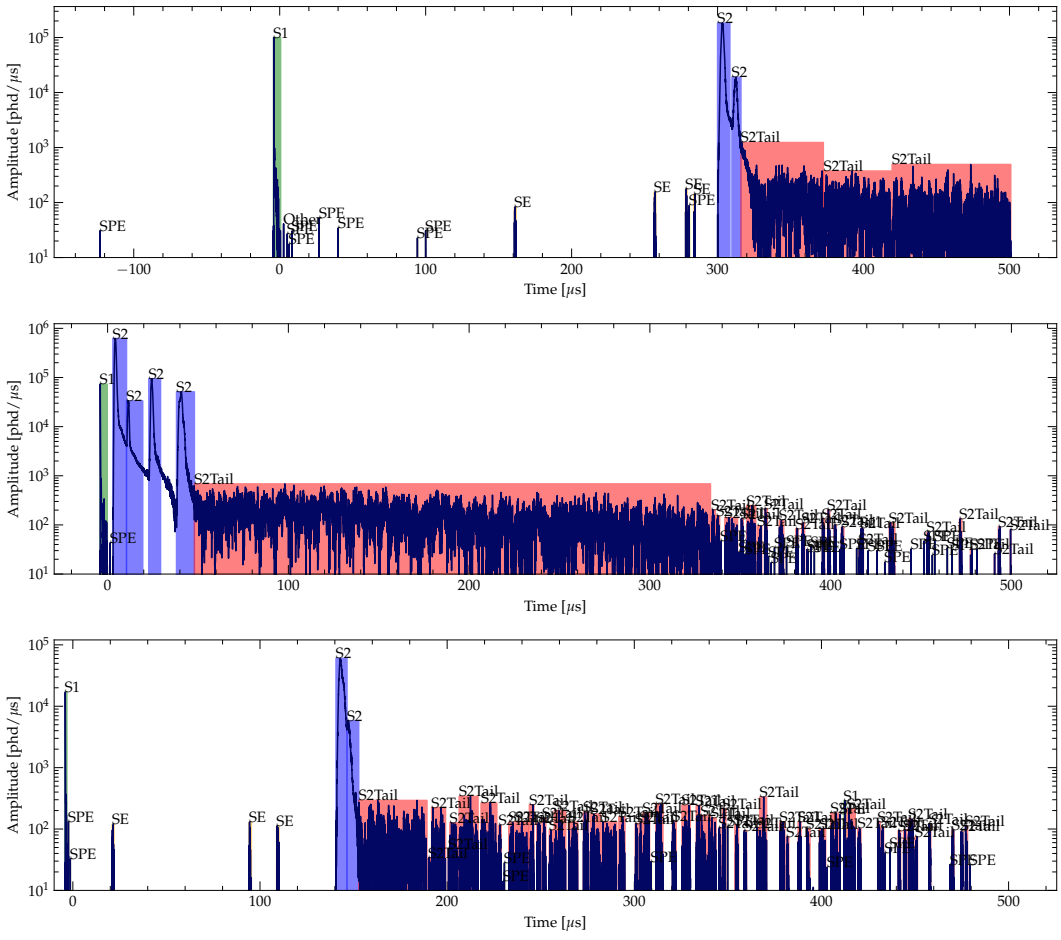

Figure 3.3: Examples of pulse finding and pulse classification performed by LUXhep. The coloured shaded regions identify the distinct pulses identified by the pulse finder.

This is a significant improvement over the LUX DPF, which as previously discussed performs poorly at identifying multiple S2 pulses for high energy interactions. Figure 3.4 shows the number of S2 pulses identified by the PulseFinder in a data set of events that the LUX DPF identified as consisting of only single scatters, i.e having only one S2 pulse. It is found that only 60% of these events are actually single scatters. Due to the

improved pulse finding, the multiple scatter background rate can therefore be reduced by $\sim 40\%$.

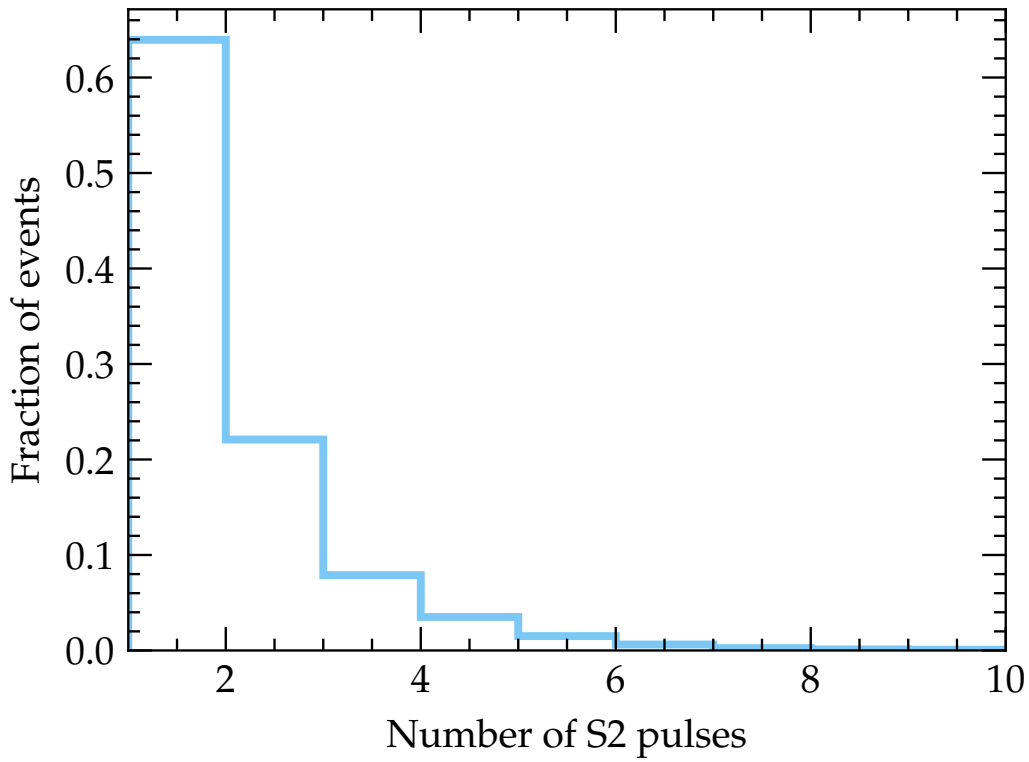

Figure 3.4: Fraction of events with n S2 pulses in a dataset of single scatter events (n=1) identified by the LUX DPF.

As discussed previously, the tail of the S2 pulse can be particularly troublesome. The majority of the tail consists of pile-up SE pulses that extend for $\mathcal{O}(100\ \mu s)$. Two possibilities exists for handling the tail: identifying individual pulses, or clustering the tail together into a single pulse. The first choice is better from an overall pulse finding efficiency perspective, but results in potentially hundreds of extra pulses to be processed. Moreover, many of these will classified as S2 pulses, making it non-trivial to identify events that have a single interaction. The alternative, and that which is used in LUXhep, is to cluster the tail together. The results is that it can easily be removed from the area estimation of the main S2 pulse, and it does not complicate the selection of single scatter interactions, resulting in improved signal efficiency. However, the second-pass of pulse finding on the S2-like pulses is also able to separate a large proportion of the tail, as seen in figure 3.3. As a result, any part of the S2-like pulse that falls outside of the newly identified pulse boundary is classified as an S2-tail pulse.

3.1.2 Pulse Classification

After pulse finding, a key part of reconstructing interactions is the subsequent classification of the identified pulses. Inferences of the underlying physical process rely on the accurate recognition of the S1 and S2 pulses, corresponding to the observation of scintillation light and ionisation, respectively. Fortunately, the S1 and S2 pulses are highly distinct; S1 light produces a prompt pulse with a sharp rise time of $\mathcal{O}(1$ ns) and a fast fall time of $\mathcal{O}(100$ ns), corresponding to the time constants involved in the emission of scintillation light. On the other hand, an S2 pulse is much broader and can extend over several μs due to the width of the GXe region. In addition, the pulse is further smeared due to diffusion of the drifting electrons. However, the picture is complicated by the existence of other pulses that originate either from single photoelectrons (SPE) from PMT dark current or S2 light arising from the extraction of a single electron (SE). Random coincidences of SPE can mimic small S1 signals, whereas coincidences of SE can be mis-classified as an S2 pulse. Moreover, as shown in figure 3.1, the tail of the S2 pulse consists of a continuous stream of SE pulses arising from photo ionisation of the gate and cathode grids.

Pulse classification relies on the calculation of several pulse parameters that characterise the size, shape and timing of the pulse. In LUXhep, these parameters are determined by the PulseParameteriser module. The key parameters are:

- *Pulse area*: integral of pulse waveform between start and end boundaries identified by the pulse finder. This is calculated separately for the bottom, top and all PMTs.
- *Coincidence*: number of PMT channels that observed a signal during the pulse.
- *Width*: standard deviation of the pulse waveform.
- *Top-bottom asymmetry*: difference between top and bottom area fractions.
- *Area fraction time*: time at which the pulse area reaches a specific fraction of the total. This is calculated for 0% to 100%, in intervals of 5%.

For pulses that are not identified as an S2-tail by the pulse finder, the PulseClassifier module in LUXhep uses these observables to classify pulses into four distinct categories: S1, S2, SPE and SE. Figure 3.5 illustrates the selections used to isolate separate populations corresponding to these pulses using a combination of the pulse width and the parameter $(t_{50\%} - t_{5\%})/t_{50\%}$, where $t_{50\%}$ ($t_{5\%}$) is the time at which the pulse area reaches 50% (5%)

of the total. In addition, the top-bottom asymmetry is used to further enhance the selection of S1 and SPE pulses.

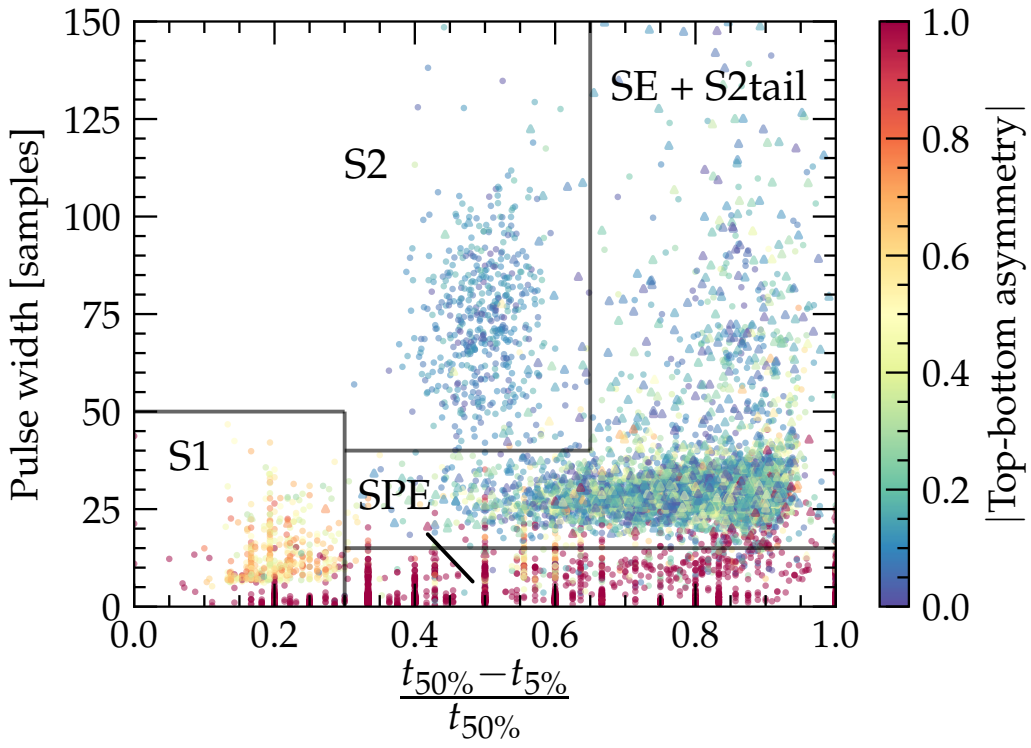

Figure 3.5: Selections used to define the pulse classifications. Each marker represents a pulse and the colour shows the absolute top-bottom asymmetry. Triangle markers show the pulses classified as S2tail.

3.1.3 S2 Area Estimation

The area of a pulse is related to the number of photons that were detected from an interaction, and can therefore be used to infer the number of scintillation photons N_{ph} and electrons N_e produced in the interaction through the relations

$$\mathrm{S1} = g_1 N_{ph} \tag{3.1}$$

$$\mathrm{S2} = g_2 N_e, \tag{3.2}$$

which are subsequently used for energy reconstruction. An accurate and precise estimation of the pulse area is therefore necessary to obtain the correct energy with an optimal energy resolution. The S1 and S2 areas are initially estimated by simple integration

of the pulse between the boundaries defined by the PusleFinder module. However, as discussed previously, the S2 pulse is impacted by photo-ionisation and delayed emission effects that lead to difficulties in obtaining a precise estimate of the area. The simple algorithm employed by the PulseFinder module is not robust to the variations induced by these processes, resulting in inconsistent pulse boundaries which introduces artificial noise into the estimate of the area. This becomes more pronounced at higher energies, where the S2 pulses are observed to have much longer tails. An improved estimate of the S2 area is therefore obtained using the S2AreaIntegrator module, which performs fits of analytical templates to the S2 pulse and the resulting tail.

The shape of an S2 pulse is approximately described by a Gaussian convoluted with an exponential. The Gaussian component arises from the longitudinal diffusion of the electron cloud, with the combination of singlet and triplet scintillation components resulting in a slight exponential tail [80]. Photo-ionisation of the gate grid results in a delayed component of scintillation light that occurs approximately 3 μs after the peak of primary S2 light. The quantum efficiency of the grid wires and the flux of light incident on them determines the area of the delayed pulse, and gives rise to a a predominantly exponential decay shape. The photo-ionisation of impurities in the LXe and the cathode wires also contributes to this tail. An approximation of the S2 pulse with area A is given by the sum of two exponentially modified Gaussians (EMG)

$$\mathrm{S2}(t) = A\sum_{i=1}^{2} f_i \frac{\lambda_i}{2} e^{\frac{\lambda_i}{2}(2(\mu_1+\Delta\mu_i)+\lambda_i\sigma^2-2t)}\mathrm{erfc}\left(\frac{\mu_1+\Delta\mu_i+\lambda_i\sigma^2-t}{\sqrt{2}\sigma}\right). \tag{3.3}$$

Here, each component is weighted by the area fraction f_i and shares the same width σ. The mean of each component is constrained by $\mu_i = \mu_1 + \Delta\mu_i$ to ensure the delayed components occur after the main pulse. This obviously requires $\Delta\mu_1 = 0$. The different exponential components are determined by independent λ_i. Figure 3.6 illustrates fits to representative S2 pulses where the two components are clearly visible.

Fitting this template function to the S2 pulse can be used to determine a reduced variance estimate of the pulse area. By integrating between $\mu_1 \pm n_\sigma$ of the first component, the contribution to the area arising from the photo ionisation tail is reduced. In addition, the χ^2 of the fit can be used to reject events with S2 pulses that do not confirm well to this shape, for example where the pulse finder failed to separate multiple merged S2 pulses. For the S2AreaIntegrator module, $n_\sigma = 4$ was found to provide a good trade-off between maximising the main S2 pulse area and minimising the tail contribution, but it was not

optimised rigorously. Figure 3.30 shows that this produces a significant improvement compared to the LUX DPF in the S2 resolution for 2,614 keV events.

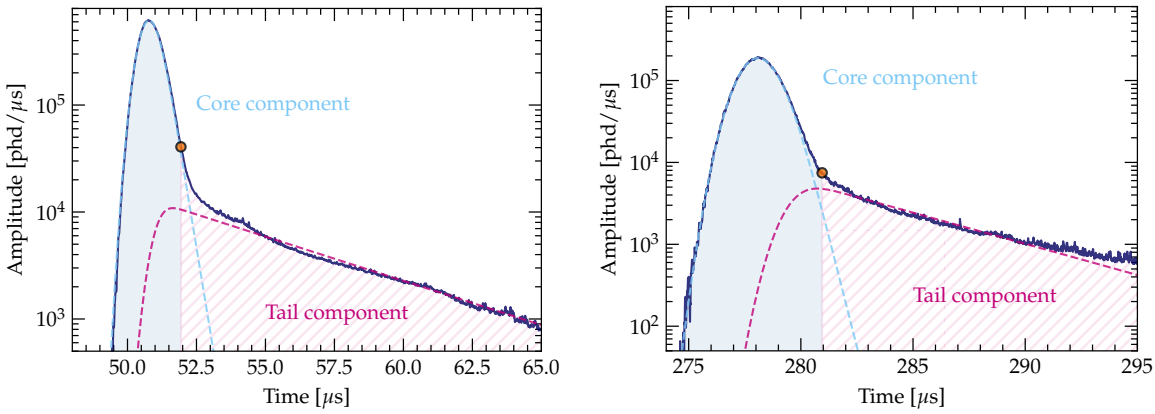

Figure 3.6: Fit of equation 3.3 to the S2 pulse of two single scatter events occurring near the top (left) and bottom (right) of the LXe. The orange marker indicates the updated pulse boundary. The solid shaded region determines the S2 pulse area, and the hatched region is the area attributed to the tail.

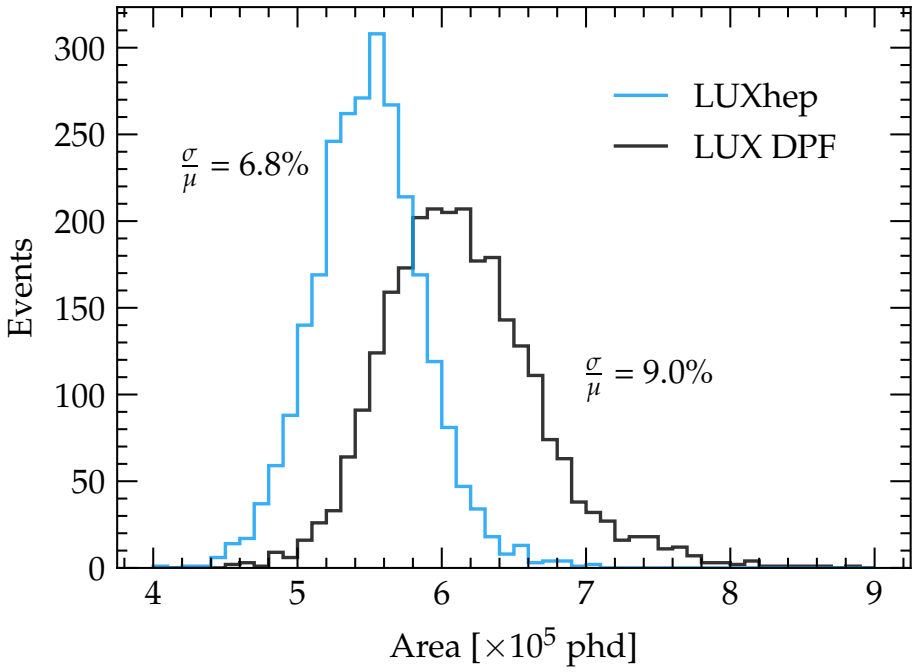

Figure 3.7: Distribution of the S2 area from ^{208}Tl 2,614 keV γ interactions for the LUXhep and LUX DPF processing.

3.1.4 Position Reconstruction

Accurate estimation of the position of interactions in the active region is vital to study the spatial response of the S1 and S2 signals, in addition to characterising background sources that have a high spatial dependence. Position reconstruction in the x, y plane in dual phase TPCs is performed using the pattern of light in the top PMT array, as this is strongly correlated with the position at which the drifting electron cloud is extracted into the GXe. As a result, the PMTs directly above the extraction point will collect the majority of S2 light.

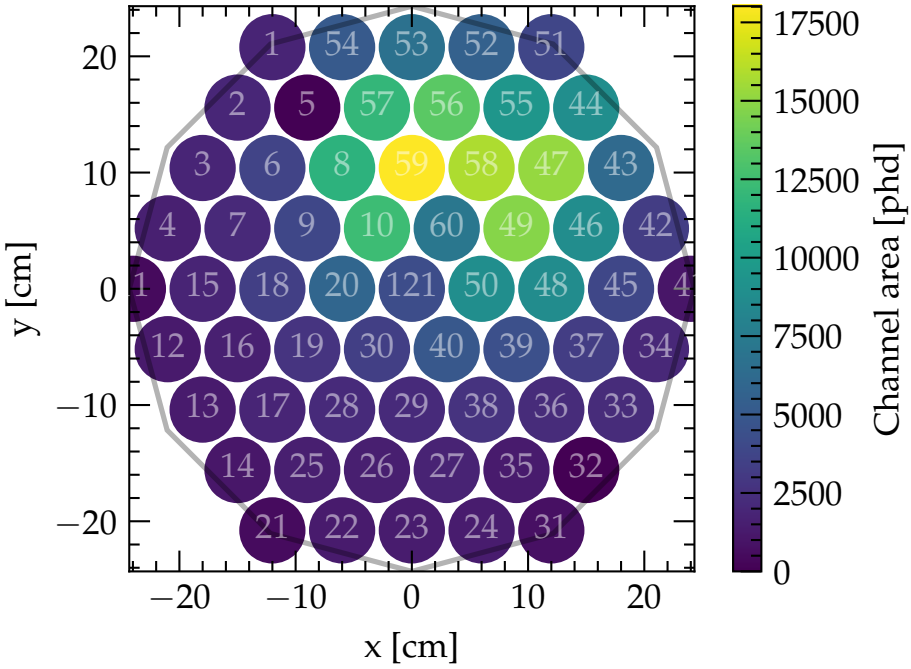

Figure 3.8: Distribution of S2 light in the top PMT array.

The PositionReconstructor module of LUXhep implements a maximum likelihood method known as Mercury [42] [81], that uses the per-PMT information in the S2 pulse to estimate the (x, y) position of an interaction. Given an S2 pulse with area N_{ph}, the number of photons detected in the i-th PMT N_i is given by,

$$N_i = N_{ph} f_i \tag{3.4}$$

where f_i is the fraction of light seen in the i-th PMT. For large signals with > 100 phd, N_i can be assumed to be distributed according to a normal distribution, i.e $N_i \sim \mathcal{N}(\mu_i(x, y), \sigma_i)$. Crucially, the mean $\mu_i(x, y)$ depends on the position x, y at which the S2

light was produced. This is a consequence of the simple fact that PMTs closer to the this position have a higher probability of collecting the light. The mean is determined by,

$$\mu_i(x,y) = \langle N \rangle \mathcal{L}_i(x,y) \tag{3.5}$$

where $\langle N \rangle$ is the expected number of photons observed and $\mathcal{L}_i$ is the light response function (LRF) of the i-th PMT. The LRF describes the probability that a photon emitted at (x,y) produces a response in the respective PMT. With this, the conditional probability distribution for the S2 response in the i-th PMT is given by

$$P(N_i \mid x,y) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{(N_{ph} f_i - \langle N_i \rangle \mathcal{L}_i)^2}{\sigma_i^2}} \tag{3.6}$$

The combined response from each of the 61 PMTs in the top array represents a set of measurements $\mathcal{N} = \{N_1, N_2, ..., N_{61}\}$ that determine the likelihood of an interaction at (x,y):

$$L(x,y \mid \mathcal{N}) = \prod_{i=1}^{61} \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{(N_{ph} f_i - \langle N \rangle \mathcal{L}_i(x,y))^2}{2\sigma_i^2}}. \tag{3.7}$$

The variance σ_i^2 of the PMT response is estimated to be a combination of a stochastic counting term and the intrinsic variance in single photoelectron (SPHE) response of the PMT,

$$\sigma_i^2 = N_i^2 (1 + (\sigma_i^{SPHE})^2) \tag{3.8}$$

Estimates of the position (x,y) and the expected number of photons $\langle N \rangle$ can then be determined through maximisation of the likelihood,

$$(\hat{x},\hat{y})_{MLE} = \underset{x,y}{\operatorname{argmax}}\, L(x,y \mid \mathcal{N}) \langle N \rangle_{MLE} = \underset{N}{\operatorname{argmax}}\, L(x,y \mid \mathcal{N}) \tag{3.9}$$

$\langle N \rangle_{MLE}$ can be determined analytically, and is simply equal to N_{ph}. For $(\hat{x},\hat{y})_{MLE}$, it is more convenient to minimise the negative log-likelihood,

$$-\ln L(x,y \mid \mathcal{N}) = \sum_{i=1}^{61} \frac{N_{ph}^2 (f_i - \mathcal{L}_i(x,y))^2}{2\sigma_i^2} + \ln 2\pi\sigma_i. \tag{3.10}$$

which is equivalent to a χ^2 minimisation between the observed response f_i and the expected $\mathcal{L}_i(x,y)$. The value of χ^2_{min} provides a measure of the goodness-of-fit. A large

value may indicate, for example, that the observed PMT response is not well described by a single position, and could indicate that multiple interactions in the detector that are displaced in (x, y).

In LUX, the light collection in PMTs is assumed to exhibit axial symmetry, and can therefore be described by a coordinate ρ, that depends only on distance between the PMT center and the (x, y) of the light emission. Moreover, due to the high reflectively of the PTFE walls, a polar term $\varepsilon(r, \rho)$ accounts for the light that is reflected back in to the PMT, and additionally depends on the radial coordinates r of the light emission, as light closer to the walls is reflected more. This effect is seen in figure 3.9. The LRFs used in LUXhep are identical to those used for previous LUX analyses. These are modelled by analytical functions informed by Monte Carlo simulations, with parameters that are fit to the data using an iterative approach described in [42].

The result of the position reconstruction algorithm performed on a uniformly distributed $^{83}\mathrm{Kr}^{m}$ calibration is shown in figure 3.10, demonstrating two key features. The first is a striped pattern, resulting from the placement of the grid wires; electrons drifting into the extraction region are focused by the electric field and funnelled in between the wires. There will therefore be a deficit of reconstructed events around the locations of the grid wires. This can be seen in figure 3.10 where the positions of the peaks roughly corresponds to the spacing of the grid wires. It has been shown that this ability to image the grid wires corresponds to a (x, y) reconstruction resolution of $\mathcal{O}(10\ \mathrm{mm})$ [42].

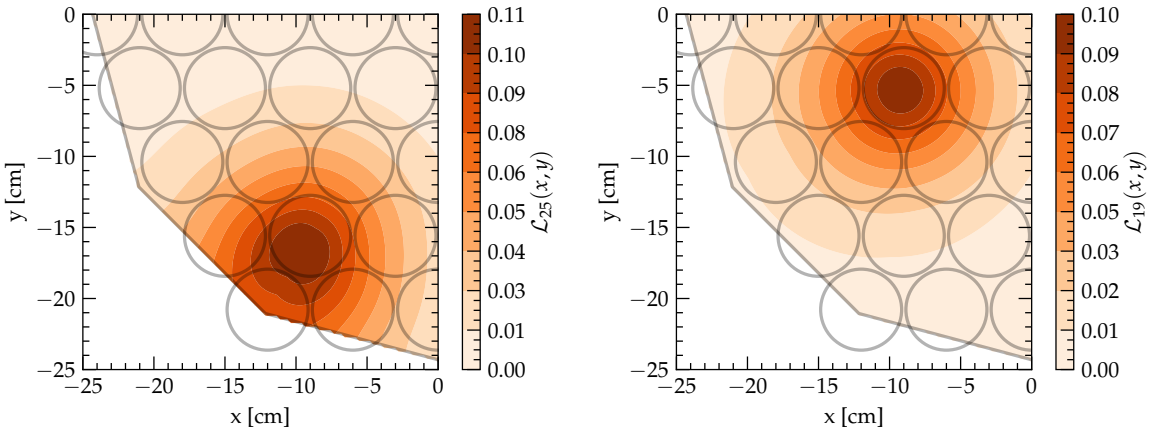

Figure 3.9: Light response functions for PMT 25 (left) and PMT 19 (right).

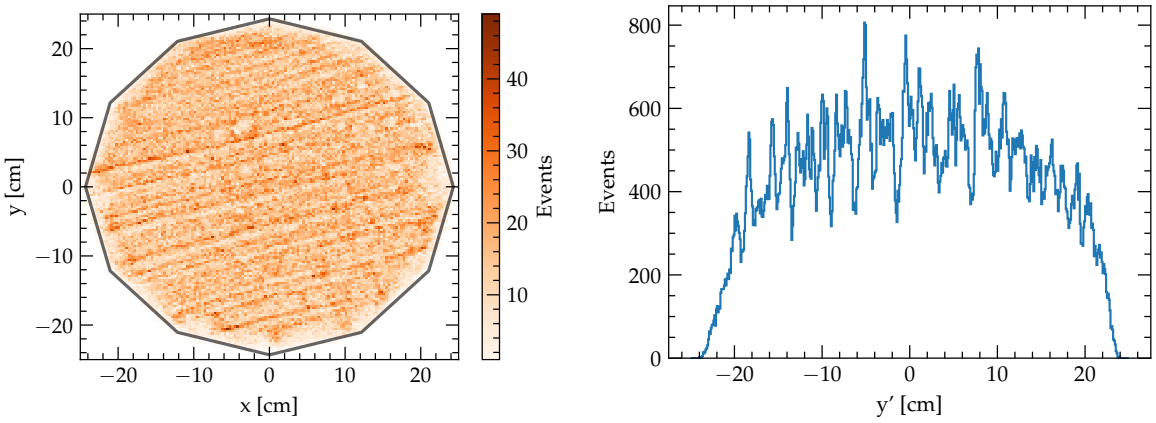

Figure 3.10: Left: distribution of reconstructed positions of $^{83}Kr^m$ in the active volume. Right: histogram of positions in a plane perpendicular to the grid wires.

3.2 Corrections & Calibrations

Dual phase TPCs are susceptible to a number of effects that result in a non-uniform light response across the detector, introducing a position dependence to the observed S1 and S2 signals. These can be caused by impurities, varying PMT gains, a non-uniform drift field, and variations in the liquid level or extraction field. Correcting for these effects is necessary when attempting to make inferences about the underlying processes, and especially when comparing to Monte Carlo simulations that do not include these effects. LUX utilised an internal $^{83}Kr^m$ source to study the position and time-dependence of the detector response. The 41.5 keV released during the decay produces S1 signals in the range 200-350 phd and S2 signals of 1000-20000 phd, both of which can easily be detected with near 100% efficiency. Weekly $^{83}Kr^m$ injections were performed over the course of Run-3, providing a method to study the stability of the detector and calculation of time-dependent corrections. Typical activities of ∼10 Bq provide enough statistics for the calibrations. Figure 3.11 shows how the population of $^{83}Kr^m$ is easily identifiable from the joint distribution of the S1 and S2 areas, with minor contamination from low energy background sources.

3.2.1 Electron Lifetime

Electronegative impurities in the LXe can capture drifting electrons, resulting in a diminished S2 signal. This effect worsens for interactions that occur towards the bottom

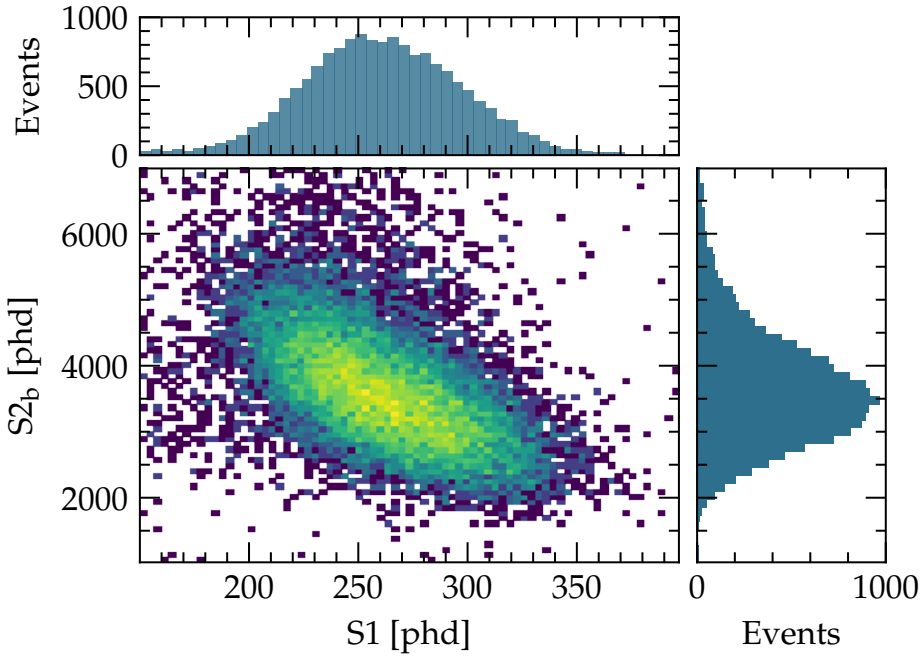

Figure 3.11: Distribution of $^{83}\mathrm{Kr}^{m}$ events in (S1, S2$_b$) space

of the detector as the electrons spend longer drifting and therefore a higher chance of being captured, leading to drift or z-dependence to the S2 signal. Given a population of drifting electrons n_0, and a probability of being captured by an impurity λ, the remaining number of electrons after drifting a time $t = v_d d$ is given by,

$$\mathrm{dn}(t) = -1/v_d \lambda t \ \mathrm{dt}. \tag{3.11}$$

Integrating over the full drift time t_d, the impact on the mean S2 signal is then,

$$S2 \rightarrow S2e^{-\frac{\lambda}{v_d}t_d}, \tag{3.12}$$

where $\tau = \lambda/v_d$ is the electron lifetime. By measuring τ, a correction can be determined that effectively scales the mean S2 response to that at top of the detector ($t_d = 0$)

$$S2_{\mathrm{z-corrected}} = S2e^{t_d/\tau}. \tag{3.13}$$

It is important to note that this only corrects the *mean* response. The fluctuations on the S2 pulse are unchanged; less variance, and therefore a better resolution, due to the reduced attenuation should still be expected at the top of the detector compared to the bottom.

The electron lifetime evolves during the running of the detector due to the dependence on the impurity levels in the LXe, Constant circulation of the xenon through the purification system results in an upwards trend and eventual plateau in electron lifetime. However, circulation failures, leaks and out-gassing will cause the electron lifetime to diminish. Consequently, time dependent corrections are necessary.

The regular injection of ^{83}Kr^m provides a convenient method to determine the electron lifetime. Figure 3.12 shows the mean S2 as a function of drift time for ^{83}Kr^m events at the start of WS2013. The mean S2 is determined by an unbinned maximum likelihood fit to

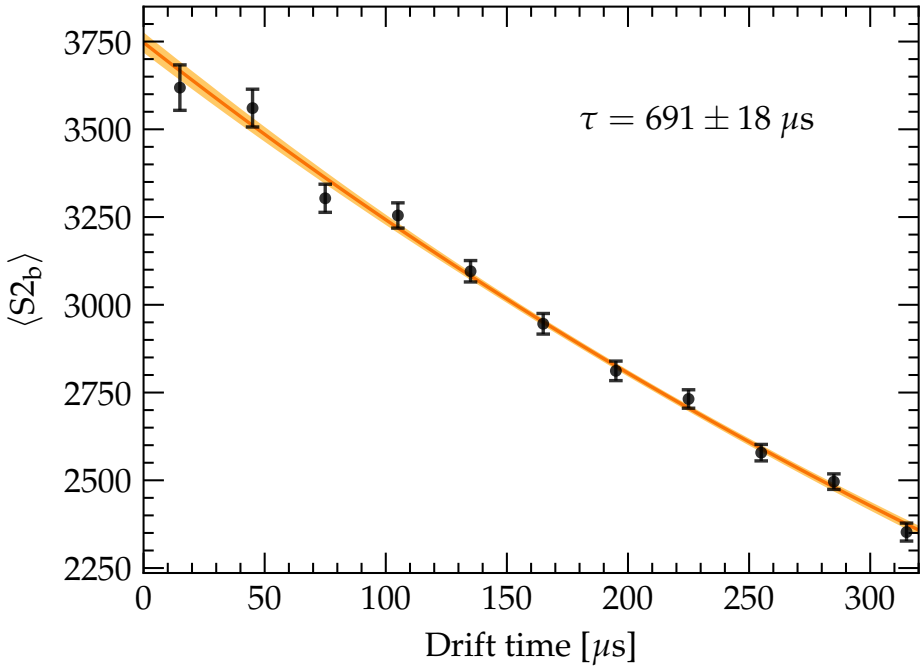

Figure 3.12: Mean S2 area as a function of the drift time in the LUX detector on 24th April 2013. The solid line shows the fit of equation 3.12 to extract the electron lifetime, with the shaded region indicating the corresponding 1σ uncertainty.

a normal distribution in 10 drift bins from 0 – 320 μs. A least-squares fit of equation 2.1 is then used to determine a measurement of τ. This is repeated for each ^{83}Kr^m injection in WS2013, with the results show in figure 3.13. The electron lifetime steadily evolves from 600 μs to a maximum value of 870 μs over the course of WS2013, resulting in a 30 – 40% reduction in the observed S2 signal that originates from an interaction at the bottom of the detector. In addition, there is an average uncertainty on the lifetime of 2.5 %, resulting in a non-negligible systematic uncertainty in the estimation of S2 areas. The lifetime at a specific time is approximated using a cubic spline interpolation, where the

uncertainty is determined by repeated Monte Carlo sampling of the measured lifetimes, assuming a Gaussian uncertainty.

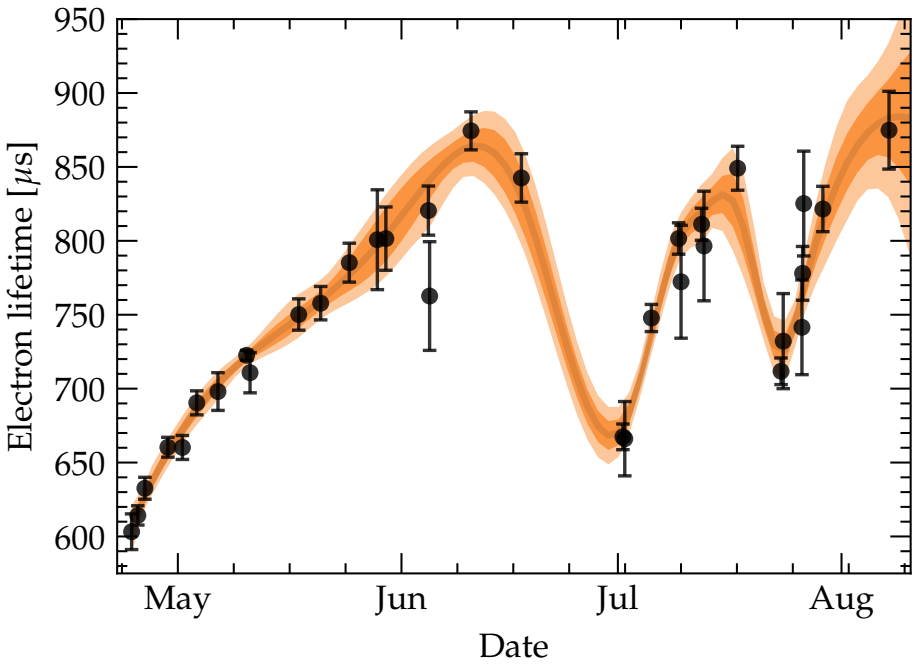

Figure 3.13: Measured electron lifetime over the duration of WS2013. The shaded regions represent the 1σ and 2σ uncertainties from a cubic spline fit.

3.2.2 Position Corrections

Light collection efficiency in dual-phase TPCs is far from spatially uniform; S1 light is affected by the detector geometry together with reflections off the liquid surface, resulting in most of the light being collected in the bottom PMT array. Variations in the electron extraction field, height of the liquid level and PMT gains introduce spatial dependencies into both the S2 light production and collection. As with the electro-lifetime correction, $^{83}\mathrm{Kr}^{m}$ can be used to develop corrections for both the S1 and S2 signals that result in a uniform response across the active region.

Variations in the mean S1 response with position are expected to be mostly due to geometric effects; reflections off the PTFE walls and liquid surface, together with attenuation result in a S1 light collection efficiency that predominantly varies with drift time, as shown in figure 3.14. The maximum light collection efficiency is seen towards

the lower center of the active region due to the proximity to the bottom PMT array and therefore reduced chance of attenuation.

A relative light collection efficiency map, $\mathcal{G}_1(x, y, \Delta t)$ is produced by binning $^{83}\mathrm{Kr}^m$ events in radius (r) and drift time (Δt) and measuring the ratio of the mean S1 response to that at the center of the detector through an unbinned maximum likelihood fit of a normal distribution. A bivariate cubic spline is then fit to this distribution to create a smooth estimate of the relative LCE throughout the detector. The S1 response of an interaction is then corrected by normalising to the center of the active region ($x = 0, y = 0, \Delta t = 165$ μs):

$$\mathrm{S1}_{\mathrm{xyz-corrected}} = \frac{\mathrm{S1}}{\mathcal{G}_1(x, y, \Delta t)} \tag{3.14}$$

S2 light production in the GXe is predominantly affected by non-uniformities in the

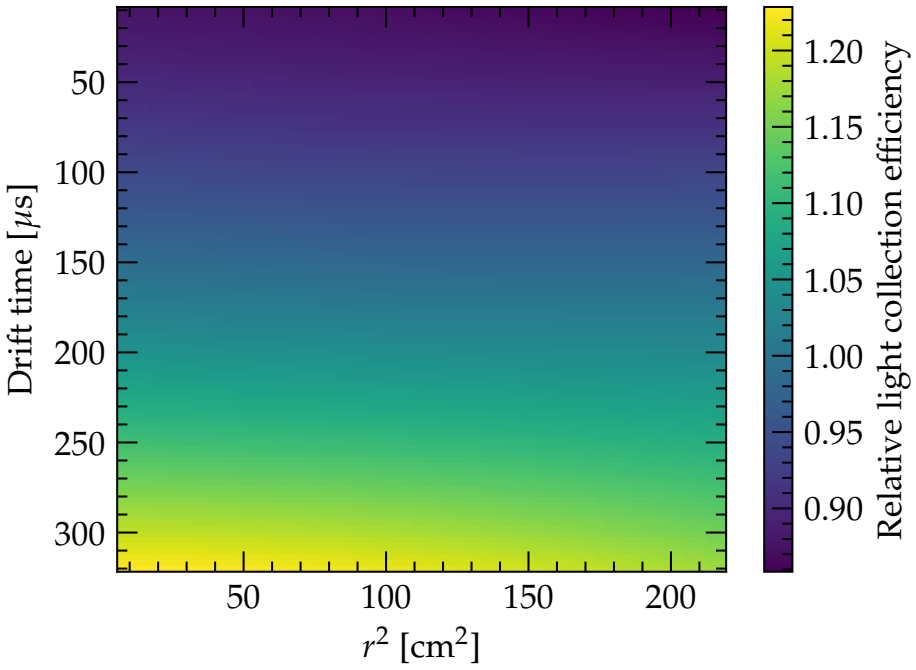

Figure 3.14: S1 light collection efficiency relative to the center of the active region.

extraction field strength and, or the height of the liquid level. Differences between PMT quantum efficiencies and gains can then lead to variations in the light collection efficiency (more so than compared to S1 as most of the S2 light is concentrated in only a few PMTs close to where the electrons are extracted into the GXe). These effects result in an S2

response that changes with (x, y). This is illustrated in figure 3.15, which shows the mean S2 response as a function of (x, y) in both the top and bottom PMT arrays.

As with the S1 response, a spatial correction map $\mathcal{G}_2(x, y)$ can be determined from the $^{83}\mathrm{Kr}^m$ calibration data. Events are binned in (x, y) and the mean $\mathrm{S2}_{\mathrm{z-corrected}}$ response is extracted from an unbinned maximum likelihood fit with a normal distribution. As with the S1 map, a cubic spline is used to create a smooth interpolation. The (x, y) correction to the S2 response is then given by normalising to the center of the detector:

$$\mathrm{S2}_{\mathrm{xy-corrected}} = \frac{\mathrm{S2}}{\mathcal{G}_2(x, y)} \tag{3.15}$$

The (x, y) and electron lifetime corrections to the S2 response are combined to produce

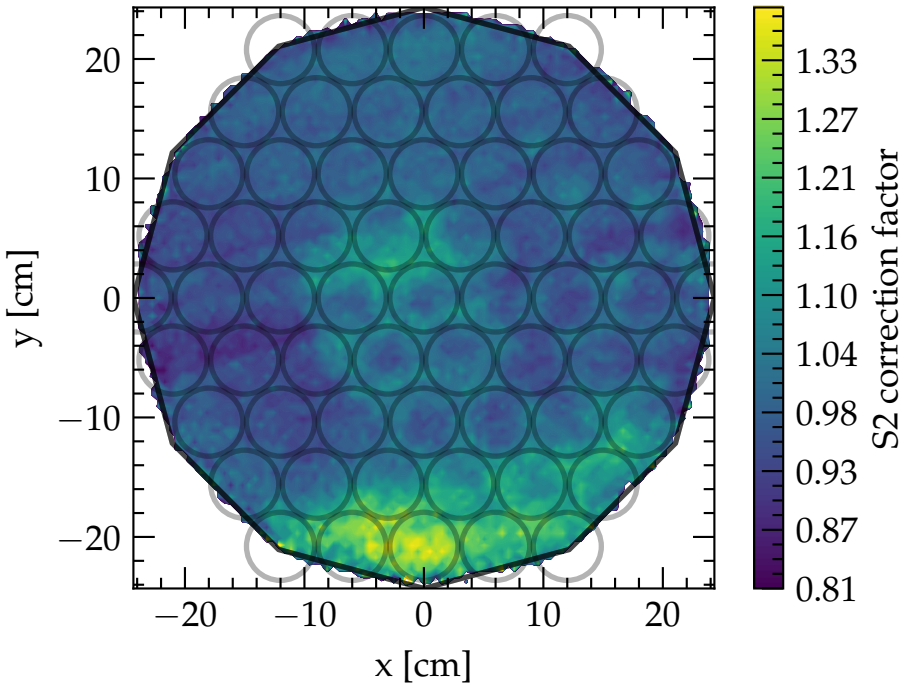

Figure 3.15: Mean $\mathrm{S2}_b$ response relative to the center of the detector ($x = 0, y = 0$).

an (x, y, z) correction:

$$\mathrm{S2}_{\mathrm{xyz-corrected}} = \frac{\mathrm{S2}}{\mathcal{G}_2(x, y)} e^{\Delta t/\tau} \tag{3.16}$$

Previous studies [82] found that a single correction map for both the S1 and S2 signals was sufficient for the entirety of WS2013, as the mean corrected S1 and S2 responses had variations within 2% during this period. However, figure 3.16 shows the mean $\mathrm{S2}_{\mathrm{xyz-corrected}}$ and $\mathrm{S1}_{\mathrm{xyz-corrected}}$ for $^{83}\mathrm{Kr}^{m}$ over the course of WS2013, illustrating that their is a time-dependent residual deviation, most notably in the S2 response. This may be due to changing detector conditions, such as in the drift and extraction fields, pressure, liquid level and PMT quantum efficiencies. To correct for this, a 3rd-order polynomial is fit to the mean S1 and S2 responses over time, producing the correction factors $\mathcal{T}_1(t)$, $\mathcal{T}_2(t)$ that normalise the respective response to that measured at the beginning of WS2013, producing the final corrected S1 and S2 signals:

$$\mathrm{S1}_{\mathrm{corrected}} = \mathrm{S1}_{\mathrm{xyz-corrected}} \frac{\mathcal{T}_1(t)}{\mathcal{T}_1(t_0)} \tag{3.17}$$

$$\mathrm{S2}_{\mathrm{corrected}} = \mathrm{S2}_{\mathrm{xyz-corrected}} \frac{\mathcal{T}_2(t)}{\mathcal{T}_2(t_0)} \tag{3.18}$$

The effect of the time dependent corrections can be seen in figure 3.17, resulting in a reduction in the variation to within 0.2% and 0.5% for S1 and S2 signals, respectively. The result of the time and position dependent corrections removes detector dependent effects from the signals, allowing for calibration of the S1 and S2 response to the physical quanta produced during an interaction, and therefore a calculation of the energy. Figure 3.18 shows that the resolution of each response is also improved, decreasing from $\sigma/\mu = 19\%$ to $\sigma/\mu = 16\%$ for S2 and $\sigma/\mu = 12\%$ to $\sigma/\mu = 8\%$ for S1.

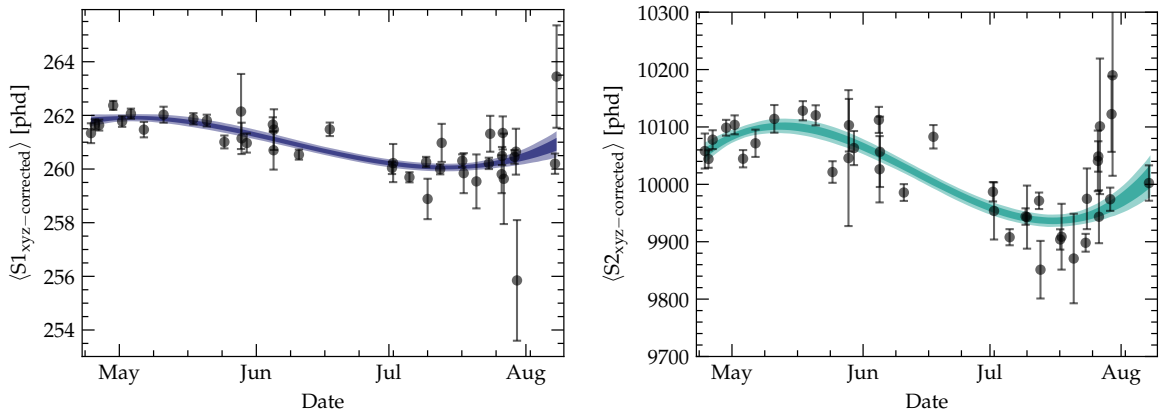

Figure 3.16: Mean S1 and S2 response over the course of WS2013.

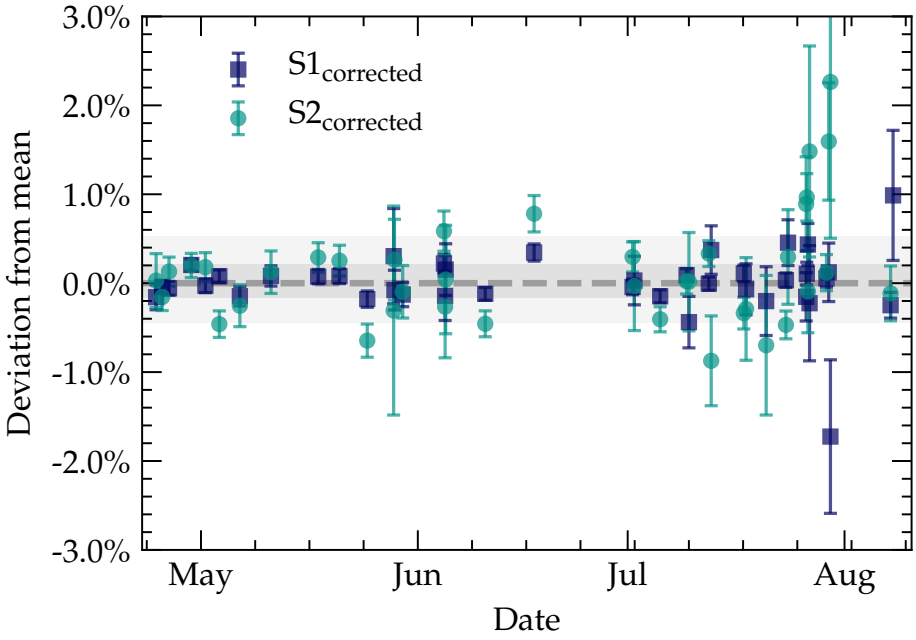

Figure 3.17: Percentage deviation form the mean over the course of Run-3 for the corrected S1 and S2 responses.

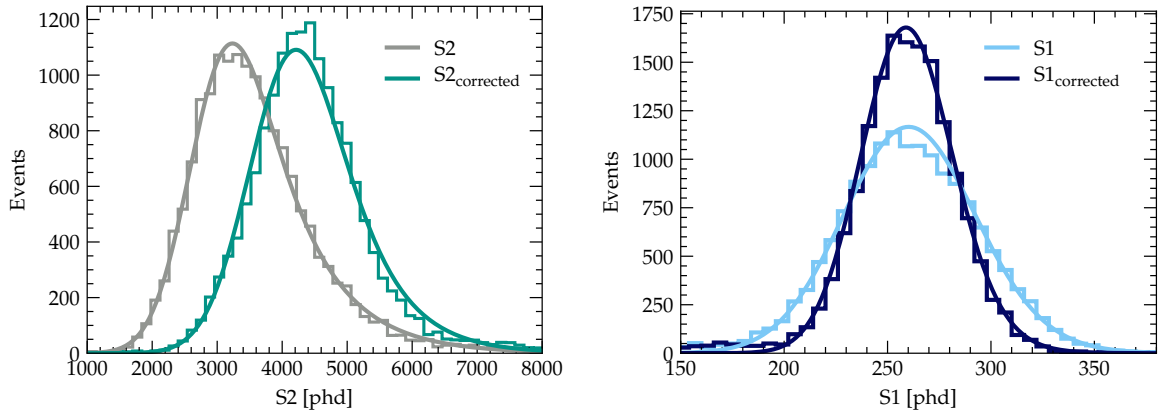

Figure 3.18: Left: $^{83}Kr^{m}$ S1 distribution before and after corrections. Right: $^{83}Kr^{m}$ S2 distribution before and after corrections.

3.2.3 Energy Calibration

The observables of an interaction in the active region are the S1 and S2 light, both of which are estimated from their identified and corrected pulse areas S1c and S2c. The

energy E of the interaction can the be determined using the combined energy formula:

$$E = W(\frac{\text{S1c}}{g_1} + \frac{\text{S2c}}{g_2}). \tag{3.19}$$

This determines a linear energy scale that is calibrated with the detector dependent factors g_1 and g_2. These can both be measured simultaneously using the Doke method [67], which utilises mono-energetic calibration or background sources. Equation 3.19 can be rearranged into the form

$$\frac{\text{S1c}}{E} = -\frac{g_1}{g_2}\frac{\text{S2c}}{E} + \frac{g_1}{W}, \tag{3.20}$$

which is the equation of a straight line, i.e $y = mx + c$ with,

$$m = -\frac{g_1}{g_2} \tag{3.21}$$

$$c = \frac{g_1}{W}. \tag{3.22}$$

A fit can then be performed to multiple measurements of different $(S2/E, S1/E)$ to extract the estimates of g_1 and g_2. In each of these measurements E is known a priori, either from a calibration source or identification of a background with a mono-energetic energy, e.g a γ-ray peak.

The sources used in this analysis are given in table 3.3. Apart from ^{83}Kr^m, all the sources are internal backgrounds originating from neutron activation or external backgrounds arising from the detector components. Although the latter have a high spatial dependence, with most of the events occurring close to the TPC boundaries, they are key to constraining the high energy response of the detector. Higher energy sources could be used, such as the 1764 keV or 2614 keV γ peaks from ^{214}Bi and ^{208}Tl, respectively. However, at these energies the saturation of the bottom array PMTs is observed to result in a non-linear relationship between S2 and E, and the application of 3.19 is therefore no longer valid.

The energy calibration is performed in nine equally spaced drift regions of the active region, from 30 μs to 300 μs. This is to account for imperfections in the position corrections, together with the fact that saturation in the bottom array is observed to be more severe for interactions occurring closer to the top of the active region. Each source is identified and selected in (S1, S2) space using a loose quadrilateral selection,

Table 3.3: Mono-energetic sources used for the determination of the energy calibration factors g_1 and g_2

Source	Energy	Type
$^{83}Kr^m$	41.5 keV	Calibration
$^{131}Xe^m$	163.9 keV	Background
^{127}Xe	236 keV	Background
^{208}Tl	583 keV	Background
^{214}Bi	609.3 keV	Background
^{60}Co	1173 keV	Background
^{60}Co	1337 keV	Background
^{40}K	1468 keV	Background

as shown in figures 3.19–3.21. In each selection, a bi-variate normal distribution is fit using an unbinned maximum likelihood. The mean S1 and S2 response is then estimated from data points within 3σ of the fitted mean of this distribution, in order to remove contamination from other sources. The error on the mean is the combination of the statistical error from the fit and the systematic errors from the S1 and S2 corrections, estimated to be 1% and 3%, respectively.

In each drift bin, a Bayesian linear regression is performed using a Markov chain Monte Carlo (MCMC) method to the set of (S2/E, S1/E) points, and the estimates for g_1 and g_2 obtained from the fitted gradient and intercept. Due to the degeneracy between g_1 and g_2 in equation 3.20, and the fact that both terms appear in the intercept, the MCMC method captures the correlations and uncertainties of the estimates more reliably than a simple least squares fit. The fits for each drift bin, known as 'Doke plots', are shown in figures 3.22–3.24. To highlight the effect of saturation, the points for the ^{214}Bi (1,764 keV) and ^{208}Tl (2,614 keV) γ peaks are also shown. At the top of the detector these points fall below the straight line, indicating that there is a reduction in the observed S2 light and the energy is therefore reconstructed below the true energy. This issue improves as drift time increases as the S2 signal is smaller due to capture of electrons on impurities, and also broader due to diffusion, resulting in a smaller current in the PMT.

The energy spectrum with the calibrated scale is shown in figure 3.25 for each drift time bin, together with the expected positions for each of the calibration sources. Above 1,500 keV, the saturation is seen to produce a significant bias in the energy reconstruction, with the 1,764 keV and 2,164 keV peaks under estimated by approximately $1\% - 6\%$. This is observed in all drift time bins but worsens near the top of the detector. In order

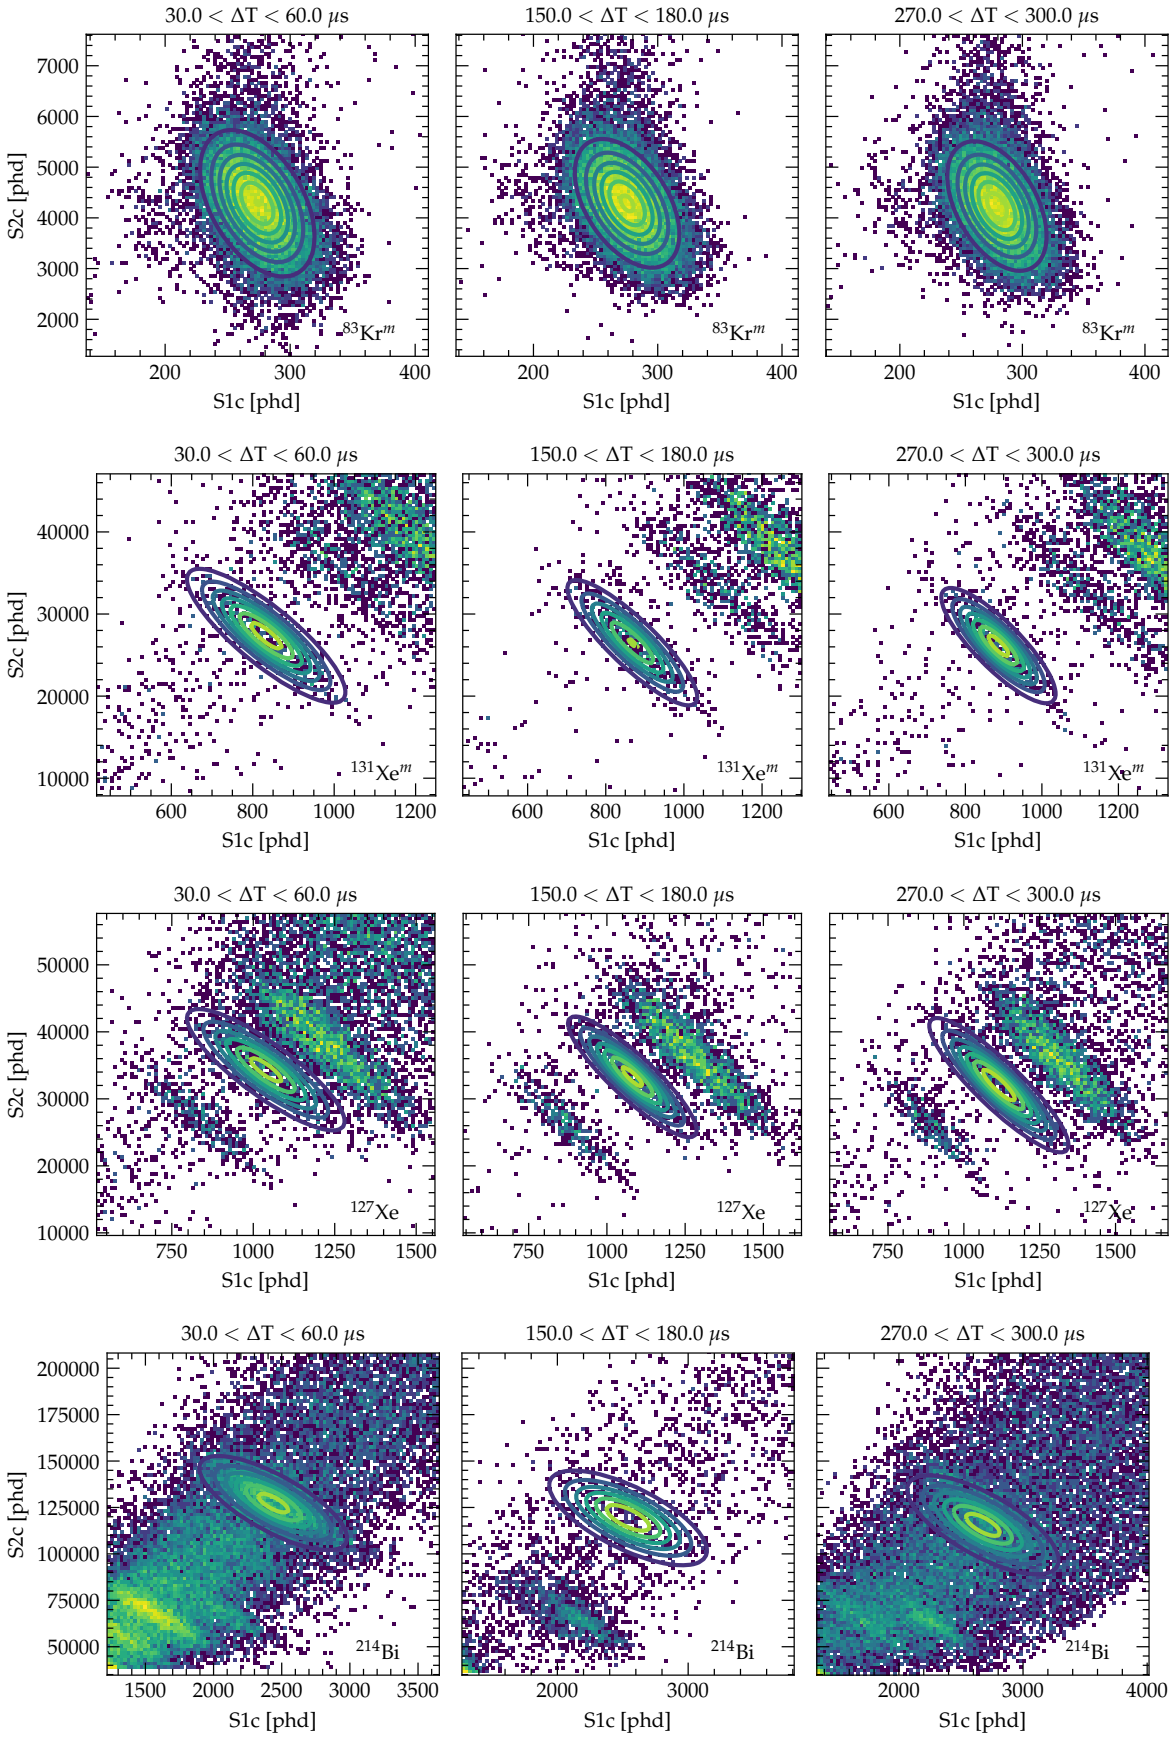

Figure 3.19: Distributions of ^{83}Kr^m, ^{131}Xe^m, ^{127}Xe and ^{214}Bi (609 keV) calibration sources in (S1c, S2c) space at the bottom, middle and top of LXe. The contours illustrate the fitted bi-variate normal distribution.

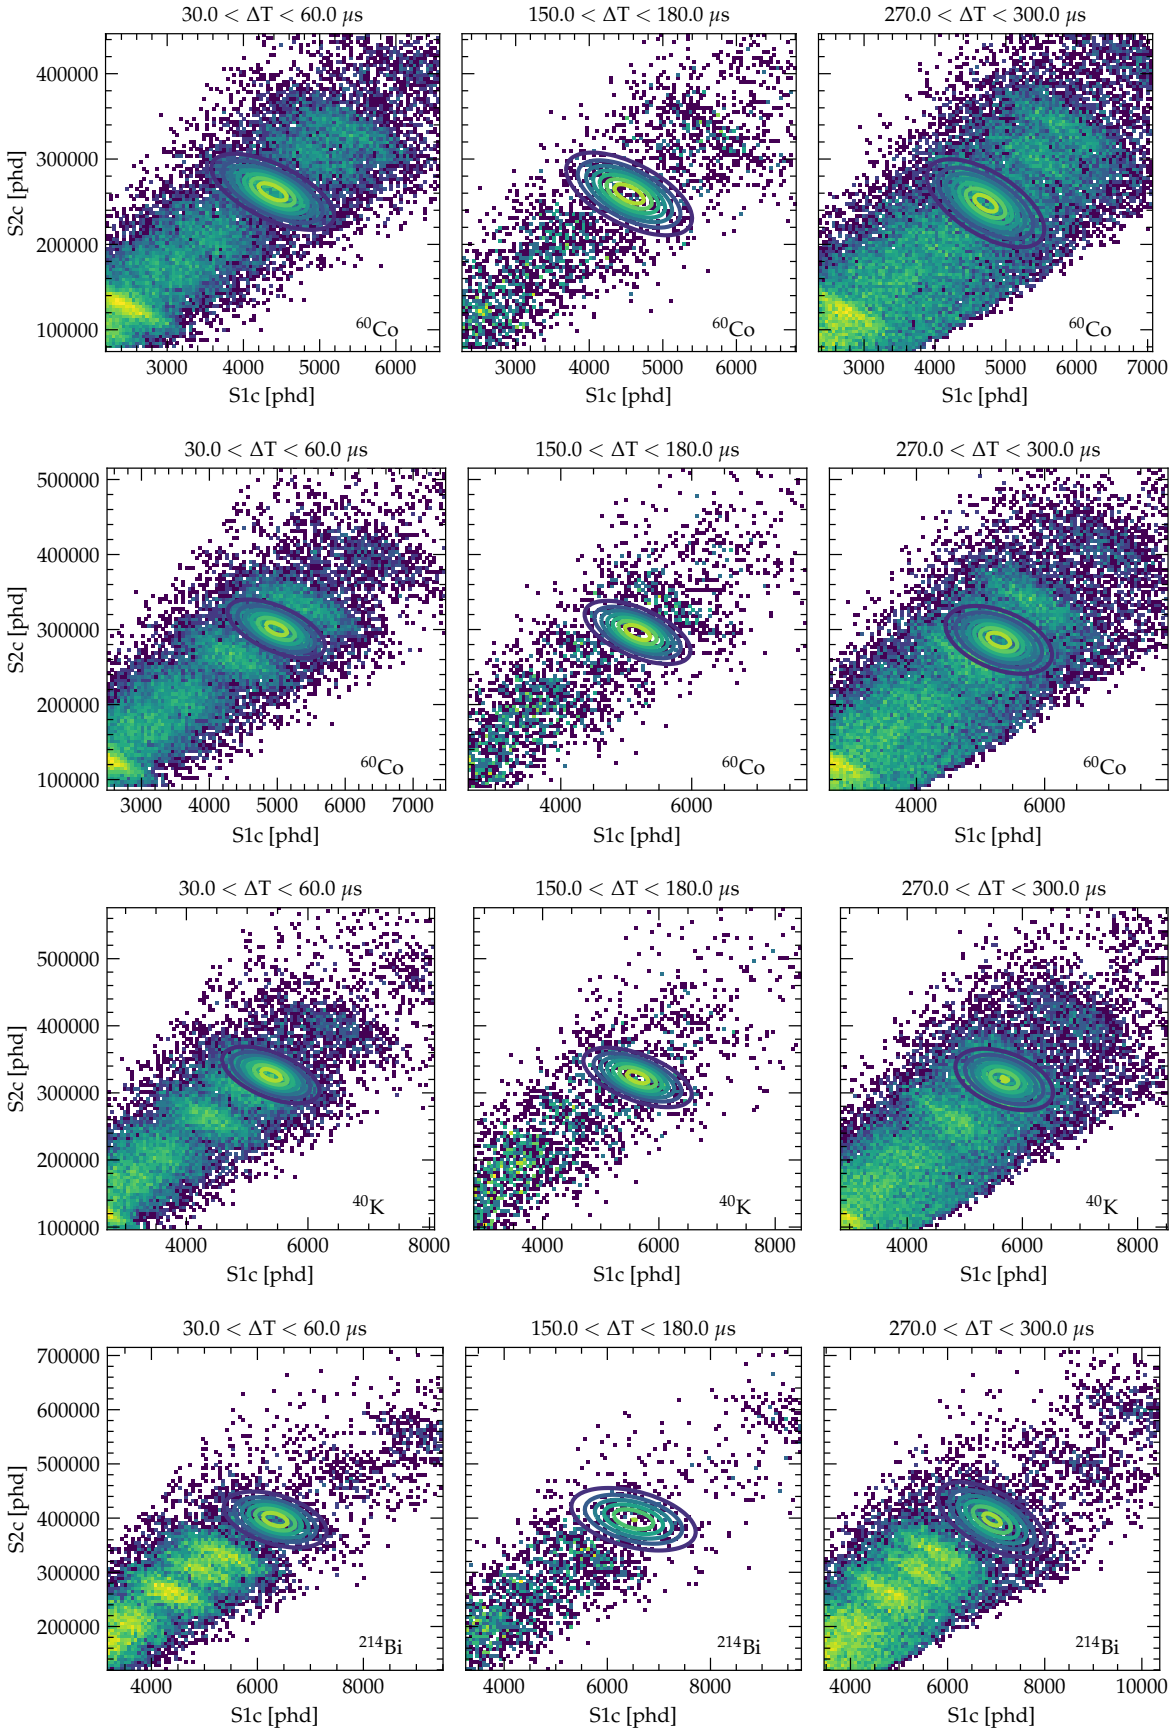

Figure 3.20: Distributions of ^{60}Co (1,337 keV), ^{60}Co (1,773 keV), ^{40}K (1,458 keV) and ^{214}Bi (1,764 keV) calibration sources in (S1c, S2c) space at the bottom, middle and top of the LXe. The contours illustrate the fitted bi-variate normal distribution.

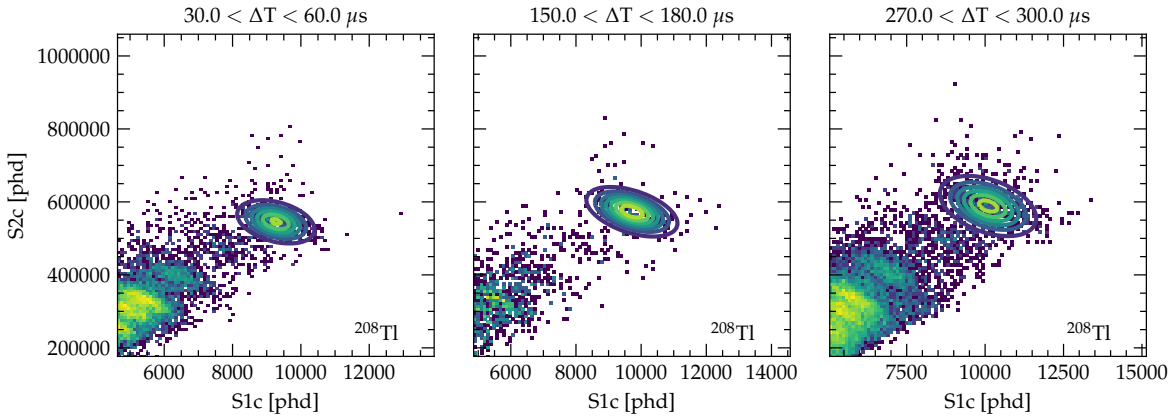

Figure 3.21: Distribution of ^{208}Tl calibration source in (S1c, S2c) space at the bottom, middle and top of LXe. The contours illustrate the fitted bi-variate normal distribution.

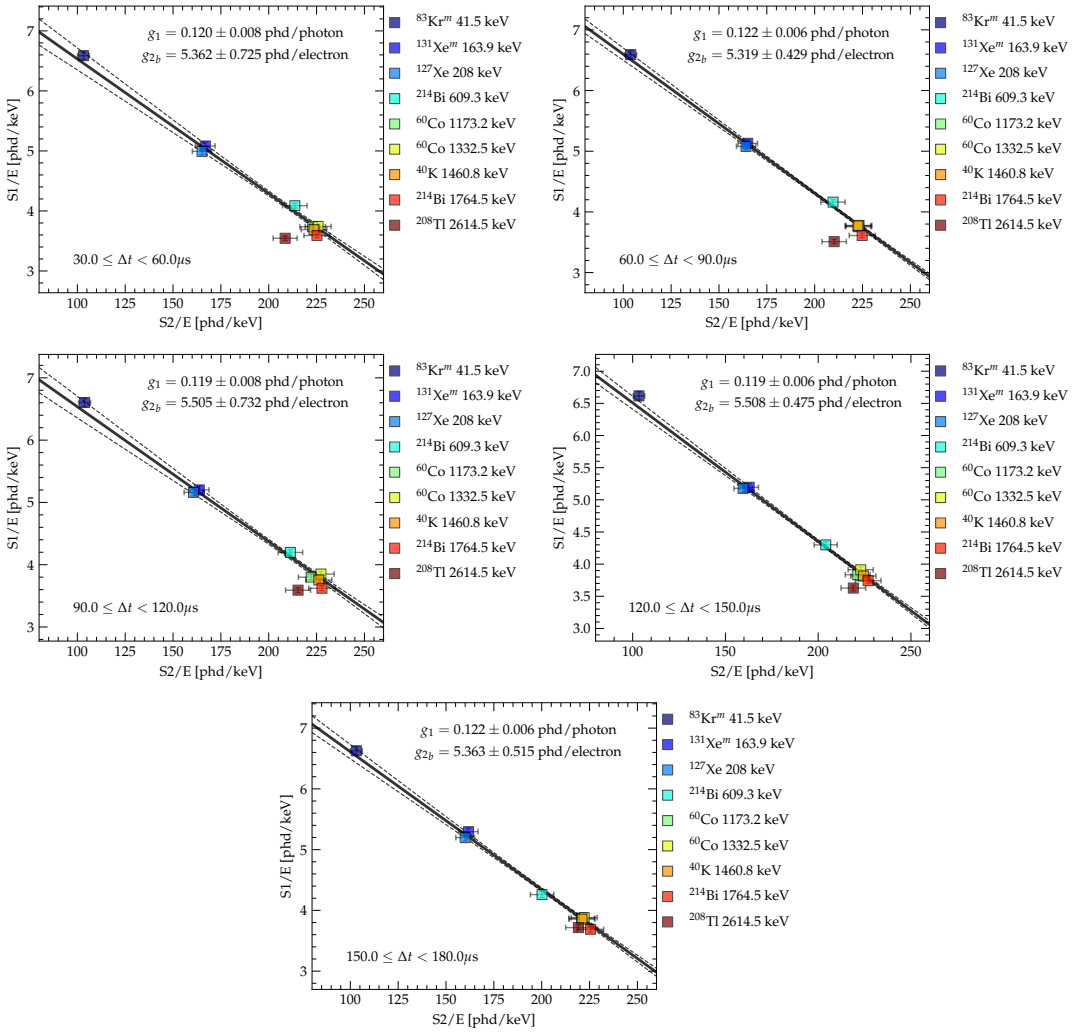

Figure 3.22: Doke plots in different drift bins ($30 \leq \Delta t < 180$ μs) used to measure g_1 and g_2. The fit of equation 3.20 and corresponding 1σ uncertainty is shown by the solid and dashed lines, respectively.

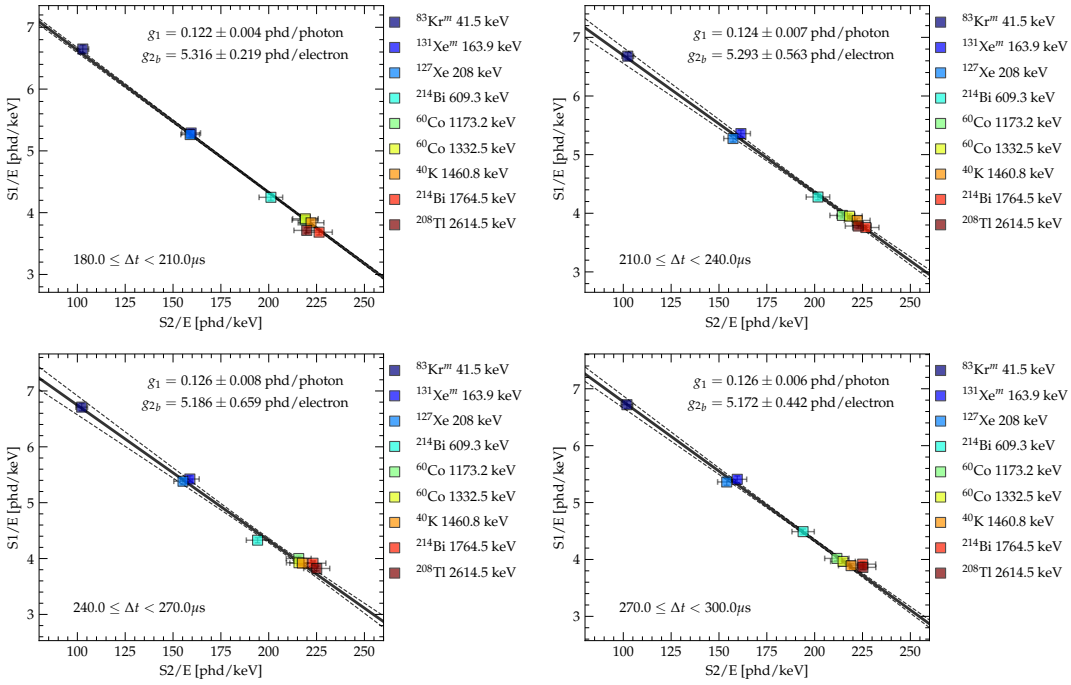

Figure 3.23: Doke plots in different drift bins ($180 \leq \Delta t < 300$ μs) used to measure g_1 and g_2. The fit of equation 3.20 and corresponding 1σ uncertainty is shown by the solid and dashed lines, respectively.

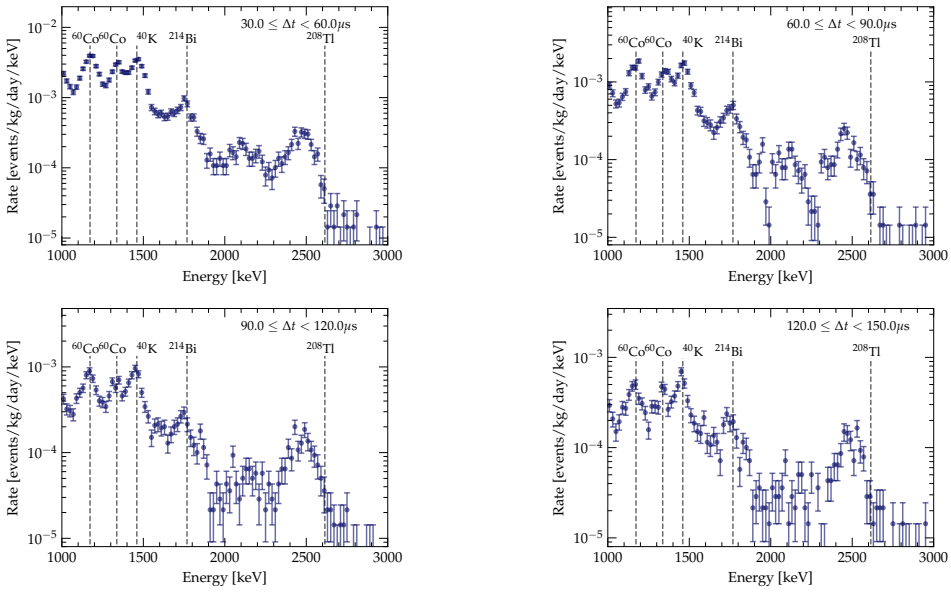

Figure 3.24: Reconstructed energy spectrum across different drift bins ($30 \leq \Delta t < 150$ μs). The dashed lines indicate the expected position of each calibration source.

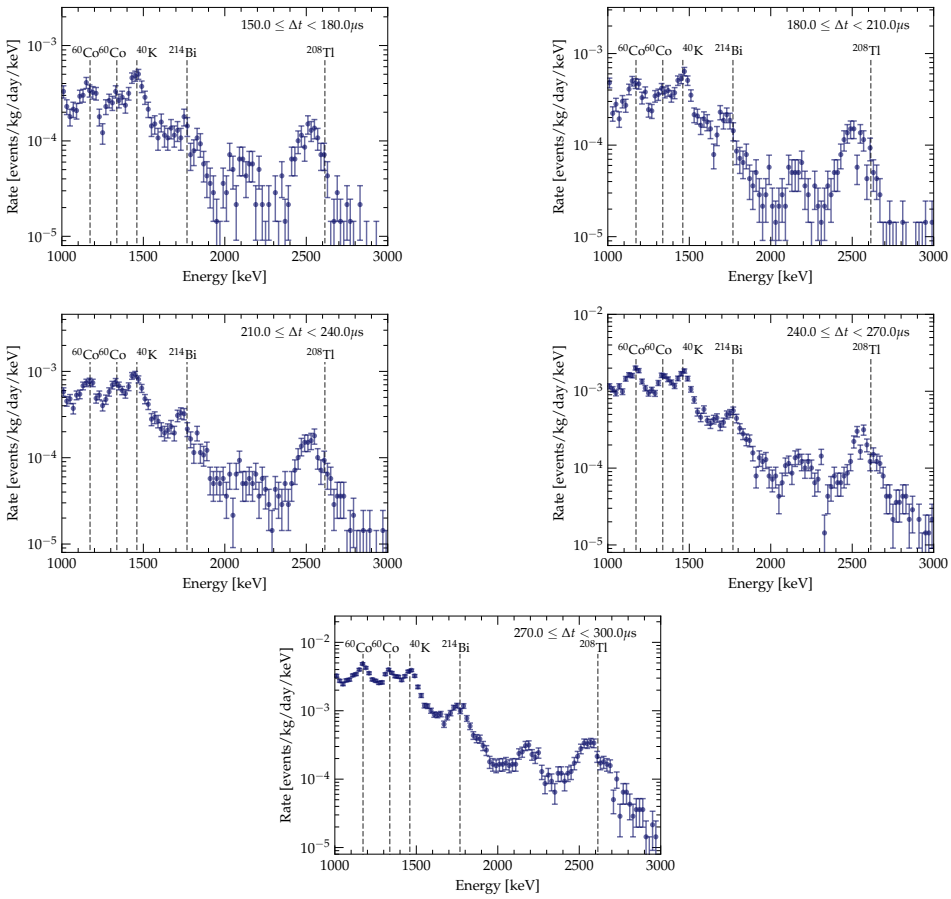

Figure 3.25: Reconstructed energy spectrum across different drift bins ($150 \leq \Delta t < 300$ μs). The dashed lines indicate the expected position of each calibration source.

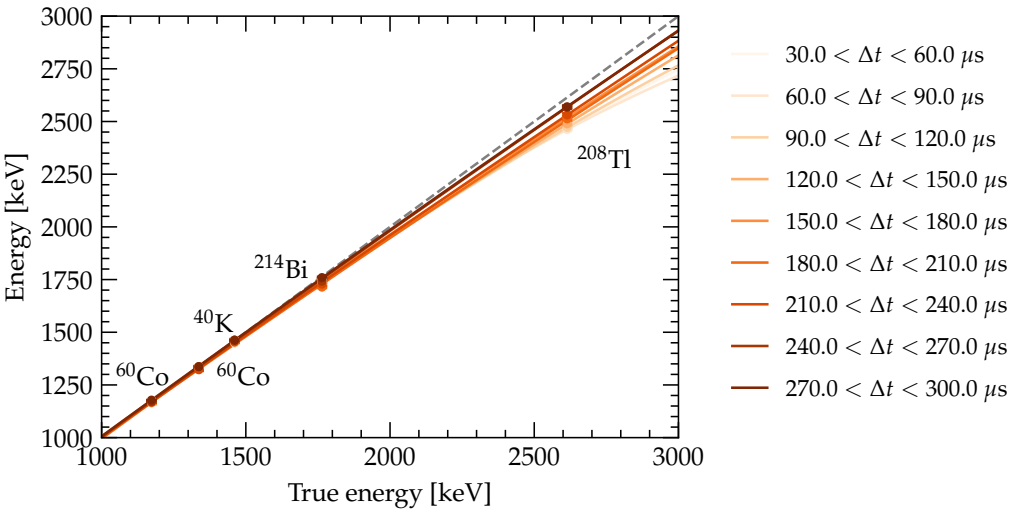

Figure 3.26: Comparison of reconstructed energy vs true energy for each calibration source. The solid lines show fits of a third-order polynomial used to correct the non-linearity.

to accurately quantify the background around $Q_{\beta\beta}$, this residual non-linearity needs to be corrected.

The corrections are determined by measuring the shift of the peak position for each calibration source, and comparing it to the expected position. To simplify this process, each calibration peak above 1,000 keV is modelled with a Gaussian and a single decaying exponential is used to model the residual Compton background. The 2,214 keV peak from ^{214}Bi is also included to fit the region below the 2,614 keV peak, but it is not used to determine the correction factors due to the the limited statistics in this region. The combined model is fit separately in each drift bin to the data in the region $1,100 - 3,000$ keV. The fitted peak positions compared to the expected are shown in figure 3.26. A linearity correction is then determined by fitting a second order polynomial to the ratio of reconstructed energy to expected energy. The application of the correction factors shifts the reconstructed energy towards the true energy, resulting in a linear energy scale up to 3,000 keV. The accuracy of the energy reconstruction near $Q_{\beta\beta}$ is verified by figure 3.27, which shows the reconstructed energy of the 2,614 keV peak is within 0.5% of the expected value across all drift bins.

The final ingredient in the analysis of the energy calibration is an estimate of the energy resolution $\sigma(E)/E$ of the detector. For each of the calibration sources, the mean and variance of the reconstructed energy is measured by fitting a Gaussian together with a suitable approximation of the background. An accurate estimate of the energy

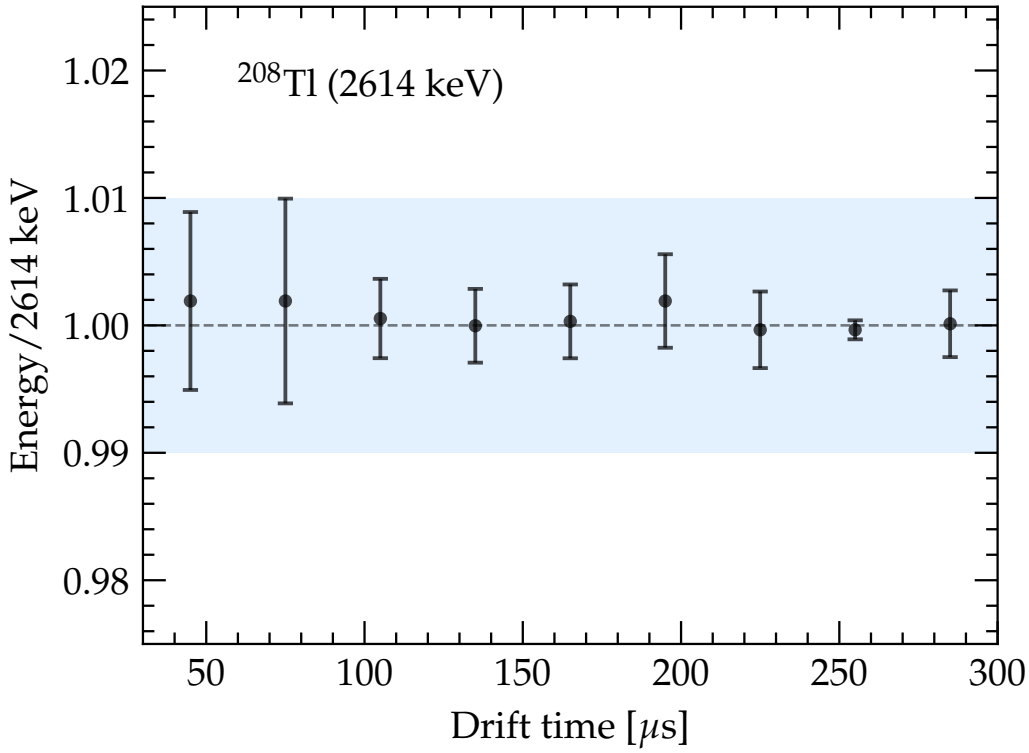

Figure 3.27: Reconstructed ^{208}Tl 2,614 keV energy as a function of drift time.

resolution requires careful modelling of the background and therefore for some of the calibration sources the neighbouring background peaks are also included in the fit. The fit components for each of the calibration sources are outlined in table 3.4. The ^{60}Co 1,337 keV and ^{40}K 1,458 keV peaks are fit simultaneously due to their proximity.

Table 3.4: Mono-energetic sources and background functions used for the determination of the energy resolution.

Source	Energy	Background
^{83}Kr^{m}	41.5 keV	Linear
^{131}Xe^{m}	163.9 keV	Exponential
^{127}Xe	208 keV	Linear
^{214}Bi	609.3 keV	Linear + ^{208}Tl 583 keV
^{60}Co	1,173 keV	Linear + ^{214}Bi 1,120 keV
^{60}Co, ^{40}K	1,337 keV, 1,458 keV	Linear
^{214}Bi	1,768 keV	Exponential
^{208}Tl	2,614 keV	Exponential

The fits for each of the calibration sources are shown in figure 3.28, and good agreement between the model and data is observed in all but the ^{83}Kr^{m} source, where a bias is

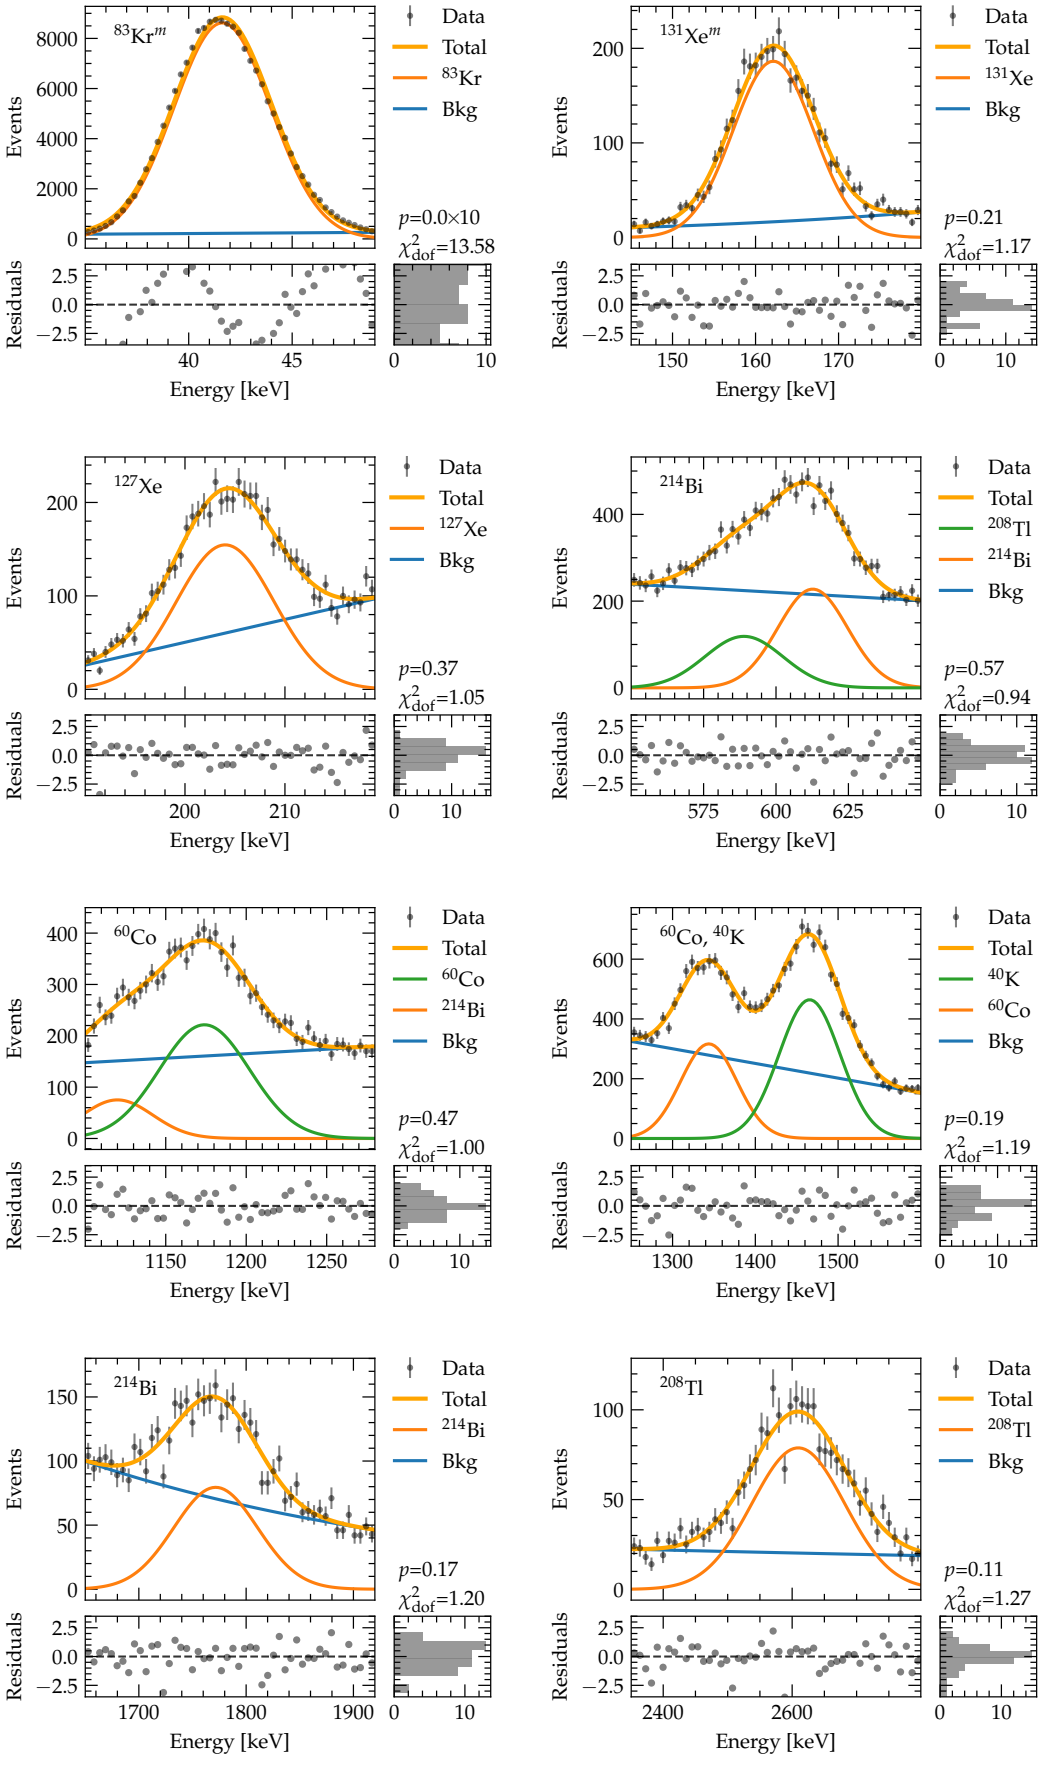

Figure 3.28: Fits to each calibration source used to determine the energy resolution.

observed in the reconstructed mean resulting in a poor χ^2. The cause of this bias is unknown, however, as this is a low-energy source, it does not affect the high energy reconstruction nor the estimate of the energy resolution. The mean and variance from each Gaussian fit is used to the obtain fractional energy resolution σ/E as function of energy, shown in in figure 3.29. It is seen that a single model cannot fit the resolution over the entire energy range. Below $1,000$ keV, the resolution is well described by the function with a stochastic and constant term, i.e

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} + b, \tag{3.23}$$

with $a = 0.35$ and $b = 0.014$. However, above $1,000$ keV there is a systematic shift in the resolution, which may be indicative of the onset of significant saturation in the bottom PMT array. In this region, the resolution is dominated by the constant term, resulting in a flat energy resolution of 2.4%. Fitting this function across the entire energy range results in a predicted resolution of 2.2% at $Q_{\beta\beta}$. The final resolution is taken as the mean of these two measurements, giving 2.3% at $Q_{\beta\beta}$. The difference between the two models then results in a 5% systematic uncertainty.

This is a significant improvement compared to the original LUX result of $\sim 3\%$ [82]. This enhancement predominantly arises from a decrease in the S2 resolution, as illustrated in figure 3.30, which compares the S2 resolution between LUXhep and the LUX DPF for each of the calibration sources. For all energies there is an absolute decrease in the S2 resolution by least 1%, and this difference is seen to improve at higher energies as a consequence of the improvements in the estimation of the S2 area.

3.3 Results

With the corrections presented in the previous section resulting in a uniform S1, S2 response that can be combined to produce a linear energy scale, a search for ^{136}Xe $0\nu\beta\beta$ can be performed. However, due to the unavailability of an accurate Monte Carlo simulation of LUX at high energies, an estimate of the expected background rate cannot be reliably obtained. Nonetheless, as less than 0.015 signal events are expected in full Run-3 exposure, all events can be safely regarded as backgrounds and a lower limit can be set on the $0\nu\beta\beta$ half-life.

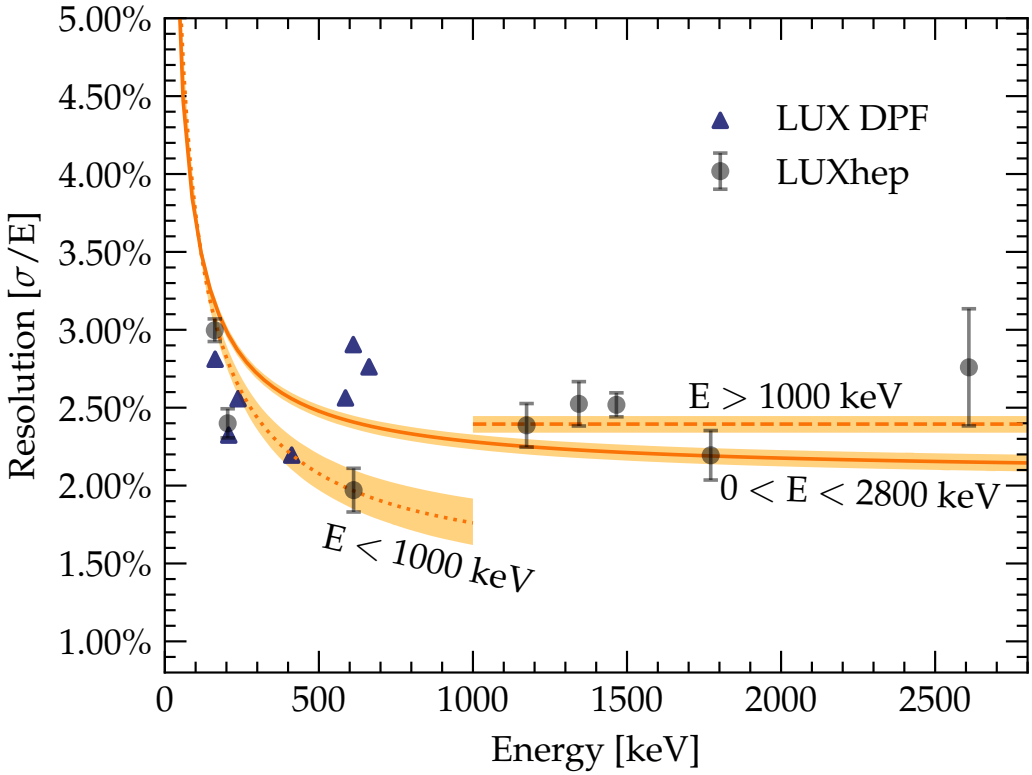

Figure 3.29: Measured energy resolution with LUXhep processing. Markers with error bars show measured resolution and uncertainty obtained from LUXhep processing. The solid lines and shaded regions demonstrate alternate fits to different energy regions. Blue triangles show the previously measured resolution using the LUX DPF.

To improve the sensitivity to $0\nu\beta\beta$ events, a number of selections are performed that reduce the background rate around $Q_{\beta\beta}$. These are:

- *Single scatter*: only events with a single S1 pulse and a single S2 pulse occurring within 320 μs are selected. In addition, events with a position reconstruction $\chi^2 > 200$ are removed. These conditions reject multiple scatter backgrounds arising from Compton scattered γ-rays.
- *Fiducial*: events within a central cylindrical volume defined by $r < 15$ cm and $30 < \Delta t < 300$ μs are selected. This removes a large proportion of the external γ background that interacts close to the PMTs and TPC walls, and corresponds to fiducial mass of $m_{\text{fid}} = 83$ kg of LXe.

At the obtained energy resolution, a 1σ ROI around $Q_{\beta\beta}$ is $2401 < E < 2514$ keV. Figure 3.31 demonstrates that the background in this range is approximately flat with a

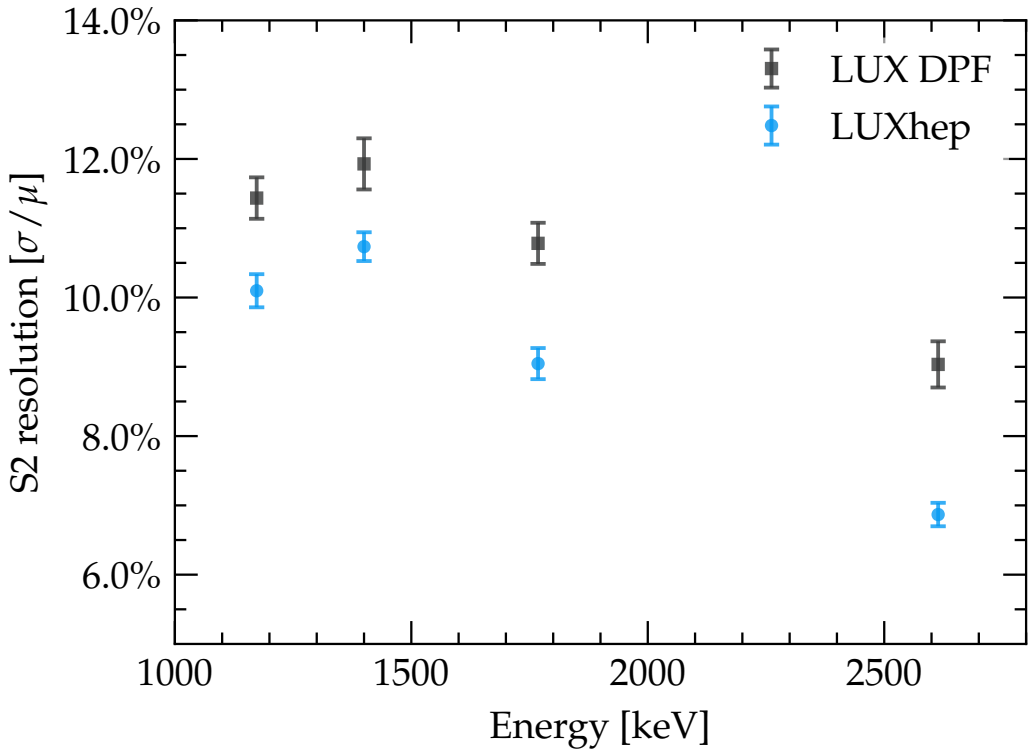

Figure 3.30: Comparison of S2 resolution obtained from original LUX DPF and LUXhep processing.

total of $N_b = 405$ events, corresponding to an average rate of 5×10^{-3} events/kg/day/keV. An upper limit on the expected number of $0\nu\beta\beta$ signal events μ_s can be determined assuming a Poisson distribution for the total number of events $N = N_s + N_b$,

$$P(N, \mu_s, \mu_b) = \frac{(\mu_s + \mu_b)^N e^{-((\mu_s+\mu_b))}}{N!} \tag{3.24}$$

where μ_b is the expected number of background events. A posteriori it is known that $\mu_s \approx 0$ and therefore the maximum likelihood estimate of background expectation is $\mu_b = N_b = 1,064$. The upper limit on μ_s at a confidence level (CL) of α is then that which satisfies the relation

$$1 - \alpha = \sum_{N=0}^{N_b} P(N, \mu_s, \mu_b). \tag{3.25}$$

This upper limit can be interpreted as the largest μ_s that has a probability of $1 - \alpha$ of under-fluctuating to produce the total observed number of events N or less. Solving this numerically gives $\mu_s \leq 42$ at 90% CL. This can then be translated to a lower limit on

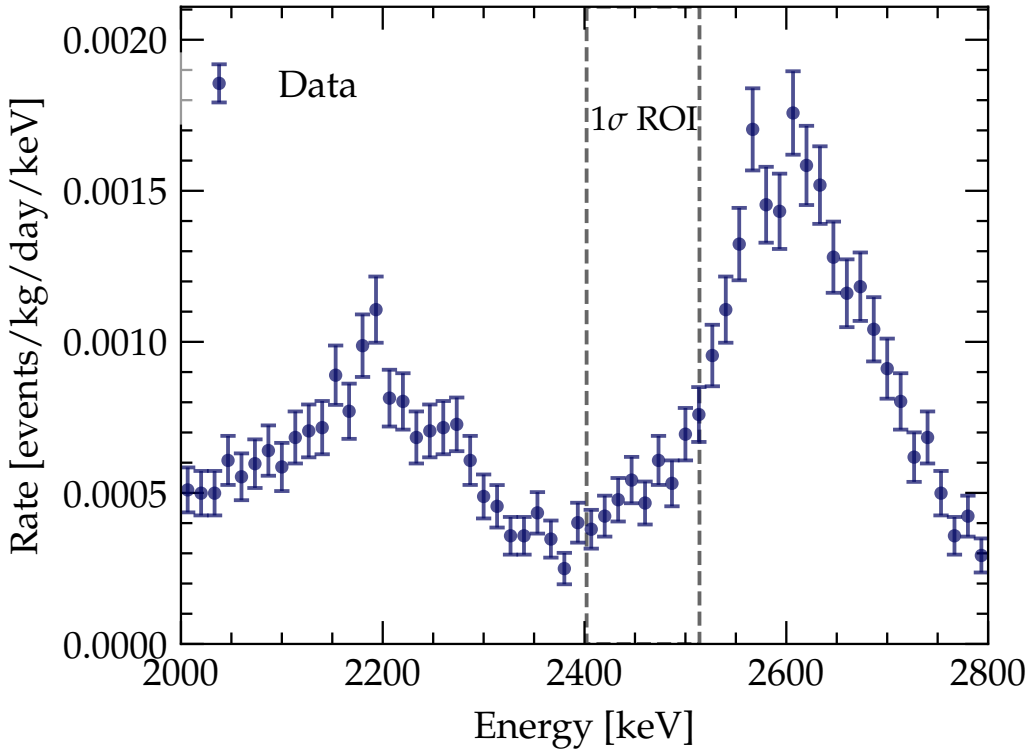

Figure 3.31: Rate of background events near $Q_{\beta\beta}$ after selections.

the half-life with

$$T_{1/2} \geq \varepsilon f_{\mathrm{ROI}} \frac{m_{\mathrm{fid}} A}{M_{\mathrm{Xe}} N_A} \frac{T}{\mu_s^{\mathrm{upper}}} \ln 2 \tag{3.26}$$

where f_{ROI} is the fraction of the signal covered by the ROI, ε is the signal detection efficiency, A is the abundance of ^{136}Xe, T is the total exposure time, M_{Xe} is the Xe molar mass and N_A is Avogadro's constant. The signal efficiency is estimated by simulating $0\nu\beta\beta$ events generated using the DECAY0 generator (introduced in chapter 4). Events are simulated using the LUXSim [83] software, which tracks energy deposits throughout the LUX detector and simulates PMT waveforms using the methods described in [80]. The events are processed with LUXhep to obtain a signal efficiency of 82% after the single scatter selection.

The corresponding lower limit on the half-life is then $T_{1/2} \geq 1.08 \times 10^{22}$ yr at 90% CL. As stated at the beginning of this chapter, the purpose of extracting this limit was not to obtain a competitive result, but to demonstrate the techniques required for $\mathcal{O}$(MeV) event reconstruction and mitigation of backgrounds in a dual-phase xenon TPC. This work also

establishes the dominant systematic uncertainties that would affect a dedicated search for $0\nu\beta\beta$. Due to the precision of the position reconstruction, $< 1\%$ uncertainty in the fiducial mass is achievable, and the exact ^{136}Xe abundance can be constrained through measurements of the ^{136}Xe $2\nu\beta\beta$ process. The uncertainties in the energy reconstruction, resulting from the uncertainty in g_1, g_2 and the linearity corrections leads to a significant systematic uncertainty on the background rate, which is estimated to be $\sim \pm 10\%$ in this work. Future experiments can reduce this uncertainty with dedicated high statistic calibrations, using sources such as ^{208}Tl which emits a $2,614$ keV γ and would there allow for a more precise energy calibration near $Q_{\beta\beta}$.

Chapter 4

Projected Sensitivity of LUX-ZEPLIN to ^{136}Xe $0\nu\beta\beta$

Although LZ is primarily designed to search for the $\mathcal{O}$(keV) scatters of weakly-interacting dark matter [84], the properties of the LXe detection medium and scale of the experiment affords it with a much wider physics reach. The active region of the TPC will contain more than 600 kg of ^{136}Xe, which is comparable to dedicated $0\nu\beta\beta$ experiments such as EXO and KamLAND-Zen, which both use enriched xenon. Coupled with the low background construction, significant self-shielding and expected run time of 1,000 days it is possible for LZ to competitively search for ^{136}Xe $0\nu\beta\beta$ and potentially improve on the current half-life limits.

This chapter presents the projected sensitivity of LZ to ^{136}Xe $0\nu\beta\beta$. Firstly, the analysis strategy is discussed, focusing on the methods of reconstructing high energy events and calibration of the detector. A brief description of the LZ simulation framework is given, and the major background contributions to this process are introduced. The simulations of the radioactivity in the LZ detector are used to inform the background model, and estimated background rates are obtained for each major source. Finally, the expected background in a 1,000 days is used to assess the half-life sensitivity using a likelihood based analysis. The effect of varying detector parameters such as the energy resolution is also considered, and their impact on the half-life sensitivity is assessed.

4.1 Search Strategy

The strategy to search for ^{136}Xe $0\nu\beta\beta$ in LZ is to look for an excess of events in a region of interest (ROI) around $Q_{\beta\beta} = 2,458$ keV. To be successful, this requires a precise estimate of the expected backgrounds together, a low variance energy reconstruction and techniques to reject dominant backgrounds. As will be seen, $0\nu\beta\beta$ events will predominantly result in single scatter interactions in the LXe, with one observable S1 pulse paired with a single S2 pulse. As ^{136}Xe will be uniformly dispersed within the LXe, the signal is expected to have a flat spatial distribution. This is in contrast to the dominant γ-ray background that can scatter multiple times throughout the LXe due to the high Compton scatter cross section. In addition, these background interactions will overwhelmingly occur close to the PMTs and TPC walls as a consequence of the strong γ attenuation of xenon. The multiplicity of interactions and the spatial distribution of events are therefore both strong observables that can be used to reject backgrounds.

The reconstruction of events in LZ is assumed to follow the general methods demonstrated in chapter 3. As discussed in section 2.2, interactions in the active region of the TPC result in the emission of scintillation photons and ionised electrons, which leads to the observation of S1 and S2 signals, respectively. The combination of these two signals can be used to reconstruct the interaction energy through

$$E = W(\frac{S1}{g_1} + \frac{S2}{g_2}). \tag{4.1}$$

Initial estimates of the detector dependent g_1 and g_2 factors, which relate the observed signals to the number of emitted photons and electrons, are are found to be $g_1 = 0.12$ phd/ph and $g_2 = 79.2$ phd/e$^-$, respectively [47]. Once the detector is operational, these quantities can be measured more accurately with the Doke calibration method used in section 3.2.3. It is assumed that the achieved energy resolution will be 1% at $Q_{\beta\beta}$, which is slightly more conservative than the 0.88% predicted by NEST [56]. In addition, the XENON1T experiment has demonstrated an energy resolution of 0.81% at $Q_{\beta\beta}$ [37]. This was measured with a drift field of 80 V/cm, which is much lower than the 310 V/cm field of LZ. This is significant as measurements from EXO [85] suggest that the the energy resolution is expected to improve with higher drift field. If similar techniques demonstrated in chapter 3 to estimate the S2 area are used, it is therefore not inconceivable that LZ can also reach a similar resolution. It is expected that any saturation of the PMTs and electronics can be corrected following similar methods outlined in [37]. In addition, the low-gain readout of the PMTs in LZ will

lessen the severity of saturated PMT signals, resulting in improved linearity in the energy reconstruction.

In addition to the energy, the two-dimensional X, Y position of interactions can be reconstructed using the maximum likelihood based approach used in section 3.1.4. Given that there are approximately four times as many PMTs in the top array of LZ compared to LUX, is it is assumed that a similar or improved X, Y position resolution of $\mathcal{O}(10$ mm) is achievable. The depth of the interaction in TPC can easily be determined from the difference between the time of arrival of the S1 and S2 pulses. It is expected that this resolution in the z-direction is $\mathcal{O}(1$ mm), which is limited by diffusion and the intrinsic width of the S2 pulse.

LZ will utilise the same injected $^{83}\mathrm{Kr}^{m}$ calibration source as used in LUX to study the response of the detector and develop spatial corrections for the response, as demonstrated in section 3.2. However, due to the larger size of LZ compared to LUX, $^{83}\mathrm{Kr}^{m}$ is likely to decay away before becoming uniformly dispersed within the LXe. As a result, a $^{131}\mathrm{Xe}^{m}$ source is also planned, which has a higher energy of 163.9 keV and a longer half-life of 12 days. In addition, further calibrations with internal ^{220}Rn and external ^{228}Th are planned. ^{208}Tl is a radioactive daughter of both isotopes emits a 2,614 keV γ-ray, which can be used to study the high energy response of the detector near $\mathrm{Q}_{\beta\beta}$.

4.2 Monte Carlo Simulations

Monte Carlo simulations of particle interactions in the LZ detector are necessary to characterise the response of the $0\nu\beta\beta$ signal and to determine the background rate. The generation of events and subsequent tracking of particles through the LZ detector is performed by BACCARAT [86], the LZ simulation framework built upon GEANT4 [87]. BACCARAT includes a detailed model of the LZ geometry that describes the shape and position of the detector components such as the inner and outer cyrostat vessels, together with small scale components such as the resistors in the TPC field cage. Each individual component can be linked with a GEANT4 or a custom event generator so that the distribution of radioactive sources and their intensities can be accurately modelled. The generators produce particles with realistic energy and momenta originating from the decay of radioactive isotopes present in each component. Particles produced by the generators are tracked through the LZ geometry, and the subsequent interactions can be propagated through two separate simulation chains:

- *Full chain*: energy deposits are clustered and NEST is used to generate the scintillation photons and ionised electrons. The ionised electrons are drifted to the gas phase where the secondary scintillation signal is simulated. The resulting scintillation photons are then propagated to the PMTs and generation of the S1 and S2 PMT signals is performed by the Detector Electronics Response (DER). This implements a realistic simulation of the PMT model and electronics chain. Although this simulation chain produces the most realistic response, the tracking of the scintillation photons and PMT signal generation is too computationally expensive to be used for studies requiring high statistics.

- *Fast chain*: only the energy deposits and their positions in different volumes of the detector are recorded. The resulting energy-based simulations can be used to determine event rates in different regions of the detector and produce high statistics templates of background distributions. As there is no generation of the S1 and S2 response nor reconstruction of events this chain is several orders of magnitude faster, but at the cost of reduced accuracy; detector resolution effects and reconstruction efficiencies have to be included through suitable approximations.

For the purposes of estimating the half-life sensitivity the fast chain is used to generate the simulations and obtain background predictions. For each simulated event the following observables are recorded:

- LXe energy: total energy deposited in the to the active region of the TPC.
- Veto energy: total energy deposited in the skin and outer detector systems.
- x, y, z position: average three-dimensional position of the interaction in the LXe, weighted by the energy.
- Δz : range of the energy deposits in the z-direction.

The affect of energy resolution is modelled by convolving the resulting energy distribution with the resolution function

$$\frac{\sigma}{E} = \frac{50\%}{\sqrt{E\ [\text{keV}]}}, \tag{4.2}$$

which ensures the resolution at $Q_{\beta\beta}$ is 1%. To determine whether an event is a single or multiple scatter, the range of the energy deposits in the z-direction (Δz) can be used. It is assumed that if multiple interactions are separated by $>$ 3mm in the z-direction, then each of the individual S2 pulses can be resolved. At a drift velocity of 1.51 mm/μs, this

translates to a separation of S2 pulses of > 2 μs, which is of the same order as the the maximum expected width of the S2 pulse due to diffusion [88]. With developments in the pulse-finding algorithms and utilising techniques such as pulse template fitting, it is expected that this is a realistic assumption.

4.3 Background Model

This section introduces the main sources of background that will contribute in a search ^{136}Xe $0\nu\beta\beta$ in LZ. For each background, the rate (events/kg/day/keV) is estimated in a signal ROI defined via the selections:

- Energy ROI: The energy window is defined as $2,433.3 < E < 2,482.4$ keV, corresponding to a $\pm 1\sigma$ energy region around $Q_{\beta\beta}$ with the assumed energy resolution (σ/E) of 1% at $Q_{\beta\beta}$.
- Inner volume: This inner volume provides the optimal signal-to-background ratio in the energy ROI and is defined as $26 < z < 96$ cm and $r < 39$ cm, containing approx. 967 kg of LXe.

In addition, each background is characterised by the distributions of two observables that are used to separately model each background component:

- Energy (2000, 2700 keV): The energy transferred to the LXe in the active region of the TPC.
- r^2 (0, 68.8^2 cm^2), z (2, 132.6 cm): the position of the event in the TPC represented by the squared radius and z-position. This volume corresponds contains approximately 5,613 kg mass of LXe.

4.3.1 Selections

The following selection criteria are applied to select signal-like events and reject dominant background sources:

- Single Scatter: As will be seen in section 4.4 ^{136}Xe $0\nu\beta\beta$ events in LXe are almost point-like and are therefore expected to result in single-scatter events, where as the dominant γ-ray background predominantly results in multiple scatter events. These

are rejected by requiring that multiple vertices are separated by less than 3 mm in the vertical direction, i.e., $\Delta z < 3$ mm.

- Veto: γ-ray backgrounds that produce a single scatter in the TPC but deposit more than 100 keV either in the outer detector (OD) or in the skin vetoes are rejected.
- Fiducial Volume: Backgrounds close to the TPC walls and/or grid wires are rejected by only keeping events with a fiducial volume defined as 4 cm from the TPC walls, 2 cm above the cathode grid and 13 cm below the gate grid. This defines a region containing approximately 5,613 kg of LXe.

4.3.2 From Monte Carlo to Background Rates

As will be seen, the majority of background in LZ arise from radioactive decays of the detector components. Accordingly, a comprehensive assay campaign has been carried out to determine the radioactivity of materials used within the construction of the detector systems. Using γ spectroscopy techniques, the activities of the naturally occurring radioactive isotopes ^{60}Co, ^{40}K, together with the ^{232}Th and ^{238}U decay chains have been measured in nearly all major components. These measurements are interfaced with the detailed detector simulations to predict the background rate.

For each radioactive component of the LZ detector an initial $N_{\text{simulated}}$ events are generated. The probability that an event ends up in the signal ROI after all selections is given by

$$\epsilon_s = \frac{N_{\text{survived}}}{N_{\text{simulated}}} \tag{4.3}$$

where N_{survived} is the number of events that pass all selections. To ensure the simulation results are not statistically limited $N_{\text{simulated}}$ is typically of $\mathcal{O}(10^8 - 10^9)$. For a detector component with measured activity A (Bq/kg) and mass m (kg), the resulting background rate is given by

$$R[\text{events/day}] = \epsilon_s \cdot A \cdot m \cdot 86,400 \tag{4.4}$$

where $86,400$ is the number of seconds in a day.

4.3.3 External Sources

Detector Components

The majority of backgrounds to ^{136}Xe $0\nu\beta\beta$ originate from radioactive decays of naturally occurring radioisotopes that contaminate both the surrounding underground environment and the materials that the detector components are constructed from. This contamination is unavoidable, but is mitigated in the LZ experiment by a comprehensive radio-assay campaign and rigorous cleanliness procedures [47].

Progeny of the ^{238}U and ^{232}Th decay chains undergo both α and β-decay, of which the later can result in the emission of de-excitation γ-rays from the daughter nucleus. With energies $\mathcal{O}(\text{keV} - \text{MeV})$ they are capable of traversing through the detector systems and transferring energy to the LXe. Those which deposit energy in the signal ROI in a single interaction will produce the same observable response as a signal event, i.e a single S1 and S2 pulse.

The decays important to consider are those of ^{60}Co, ^{214}Bi and ^{208}Tl; the latter of which are daughters of the ^{238}U and ^{232}Th decay chains, respectively. The measured activities of these isotopes in each of the major detector components is listed in Table 4.1.

^{60}Co is present in the majority of components constructed from steel and β-decays to the 4^+ excited state of ^{60}Ni with a branching ratio of 99.88% [89]. This promptly decays to the 0^+ state with the emission of two γ-rays of energy 1,173 keV and 1,332 keV. If both of these interact in coincidence and transfer all of their energy to the LXe, then the combined energy will fall within the energy ROI. However, these interactions are likely to occur in different regions in the detector and can therefore be rejected with the single scatter selection.

^{214}Bi has a multitude of β-decays which can result in the emission of several de-xcitation γ-rays. Many of these have energies outside of the energy ROI, however there is a branching ratio of 2.8% to decay to the 1^- state of ^{214}Po [89]. This subsequently decays to the 0^+ state with the emission of a 2, 447 keV γ-ray. As this is only 10 keV away from $Q_{\beta\beta}$, it cannot be completely resolved from the signal with the assumed energy resolution.

Finally, ^{208}Tl β-decays to the 0^+ state of ^{208}Pb, with the subsequent emission of a 2,614 keV γ-ray and together with at least one other γ-ray of energy no less than 583

Table 4.1: Masses and measured activities of ^{238}U, ^{232}Th and ^{60}Co for each the major components of the LZ detector systems. The subscript late refer only to the late chain activity, beyond ^{230}Th and ^{228}Th for ^{238}U and ^{232}Th, respectively, as several samples measured by LZ were found to be out of secular equilibrium [47]. † denotes that the an upper limit.

Component	Mass[kg]	$^{238}U_{late}$[mBq/kg]	$^{232}Th_{late}$[mBq/kg]	^{60}Co[mBq/kg]
TPC PMTs	91.92	3.22	1.61	2.99
TPC PMT bases	2.80	75.87	33.07	1.06
TPC PMT structures	165.75	1.60	1.06	0.20
TPC PMT cables	88.71	4.31	0.82	0.00
Skin PMTs and bases	8.59	45.98	14.91	23.6
PTFE walls	184.00	0.04	0.01	0.0
TPC sensors	5.02	5.82	1.88	1.32
Field grids and holders	89.11	2.63	1.46	10.8
Field-cage resistors	0.06	1347.49	2009.52	0
Field-cage rings	93.02	$0.35^{\dagger}$	$0.24^{\dagger}$	0.00
Ti cryostat vessel	2594.80	$0.08^{\dagger}$	$0.22^{\dagger}$	0.85
Cryostat insulation	13.78	$11.13^{\dagger}$	$7.79^{\dagger}$	1.11
Outer detector system	22 931.46	$4.71^{\dagger}$	$3.73^{\dagger}$	0.32
Other components	437.68	1.83	1.65	0.0

keV [89]. Although the total energy is outside of the ROI, these γ-rays can Compton scatter and deposit energy in the ROI.

The individual background contributions for each detector component and radioisotope are estimated using the simulations outlined in section 4.2. For the events that survive the selections, the resulting energy spectrum in the inner volume and spatial distribution of events in the energy ROI is shown in figure 4.1.

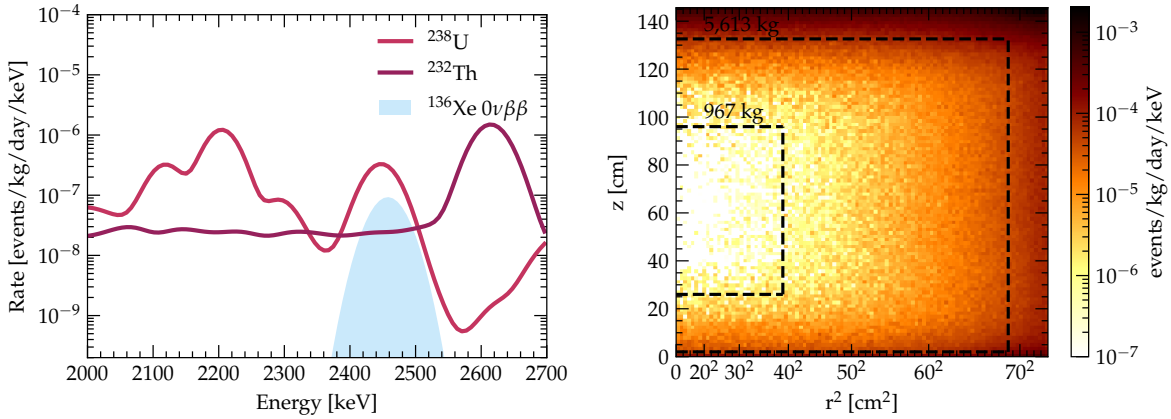

Figure 4.1: Left: energy spectrum of background events from the detector components for ^{238}U and ^{232}Th. The solid blue region shows the expected signal spectrum given a half-life corresponding to the current lower limit. Right: spatial distribution of this background in the 5,613 kg fiducial volume.

As expected, it is clear that the 2,447 keV γ-ray peak from ^{214}Bi (^{238}U chain) dominates in the ROI, where the majority of this background originates from the PMTs and field-cage resistors. There is also a significant component of the background arising from the Compton plateau of the 2,614 keV γ-ray from ^{208}Tl (^{232}Th chain). The bulk of this arises from the outer detector system which includes large acrylic tanks, PMTs and liquid scintillator. In the 967 kg inner volume, it is found that there is no background contribution resulting from the sum of the two ^{60}Co γ-rays, as these events are virtually all rejected by the single scatter selection. Inspection of the the spatial distribution demonstrates the strong attenuation of high energy γ-rays in LXe, with the majority of background interactions occurring near the outer edges of the TPC. The background rate is particularly high at the top of the TPC due to the proximity to the top PMTs, but this rapidly diminishes towards the center due to the strong attenuation. Conversely, the volume of LXe below the cathode in the reverse field region results in extra shielding that reduces the background rate near the bottom of the TPC. The consequence of this strong self-shielding is that the external γ-rays from the detector components rarely

penetrate into the inner volume, resulting in a very low background rate of only 3.0×10^{-7} events/kg/day/keV, or 15 events in 1000 days.

Cavern Rock

The surrounding rocks of the Davis Cavern have a formation of Rhyolite rock that has a naturally high ^{238}U and ^{232}Th content. As discussed, these decay chains produce backgrounds to $0\nu\beta\beta$ through the ^{214}Bi and ^{208}Tl isotopes. These activities have been measured to be 12.5 Bq/kg and 29.0 Bq/kg [90], respectively, resulting in a large flux of γ-rays directed towards the detector.

However, the water tank provides significant shielding against this γ flux. In addition, the outer detector provides even further mitigation that can be used to actively reject background events. For example, if the 2,614 keV γ-ray of ^{208}Tl deposits 135-185 keV in the outer detector and subsequently scatters in the ROI, the coincidence signal in the outer detector can be used to veto this interaction.

The backgrounds from the ^{238}U and ^{232}Th chains were estimated using simulations with γ-rays originating from the cavern rocks. However, the effects of the shielding provided by the water tank results in a statistically limited number of events reaching the inner volume. For this reason, the event biasing technique described in [86] was utilised to produce a spatial distribution with sufficient statistics. Unfortunately, due to the constraints of this technique, energy deposits in the outer detector were not recorded and the affect of the veto selection could not be taken into account. Instead, a separate simulation was used to determine the veto rejection efficiency as a function of position, shown in figure 4.2. This is then is then subsequently applied to the spatial distribution to model the effect of the veto selection. It is seen that the the veto selection reduces the γ-ray background by an extra $15\% - 20\%$.

For the events that survive the selections, the resulting energy spectrum in the inner volume and spatial distribution in the ROI is shown in figure 4.3. The ^{232}Th background dominates due to a combination of the the more penetrating 2, 614 keV γ-ray and higher rock activity. Overall, despite the shielding provided by outer detector and the water tank, the background from cavern rocks is significant with a rate in the ROI of 2.5×10^{-7} events/kg/day/keV, or 12 events in 1,000 days.

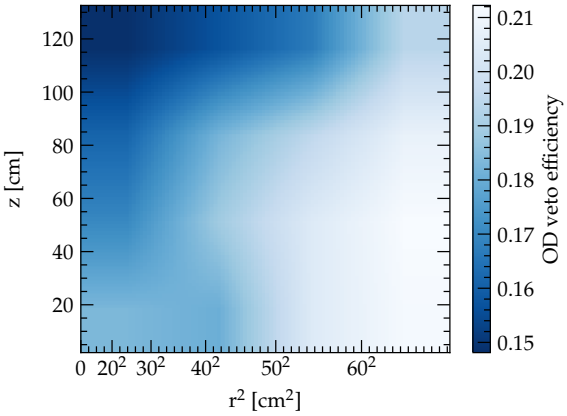

Figure 4.2: Spatial dependence of the outer detector veto efficiency.

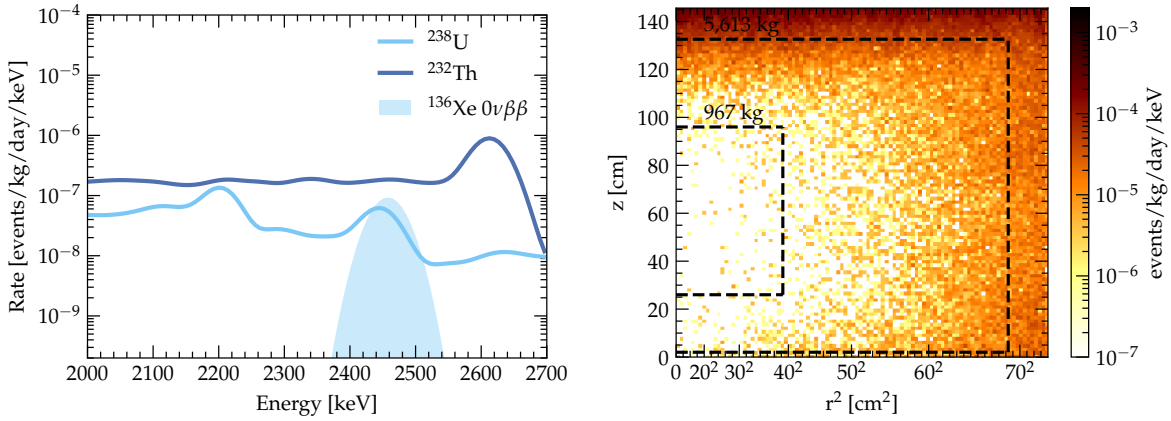

Figure 4.3: Left: energy spectrum of background events from the cavern rocks for both ^{238}U and ^{232}Th. The solid blue region shows the expected signal spectrum given a half-life corresponding to the current lower limit. Right: spatial distribution of the background in the 5,613 kg fiducial volume.

4.3.4 Internal and Astrophysical Sources

Radon

Radon gas can diffuse out of the detector components and TPC surfaces into the LXe, which is expected to result in the LXe being uniformly contaminated with ^{222}Rn and it's radioactive daughters. LZ has a requirement of < 2.0 μBq/kg of ^{222}Rn contamination, equivalent to 14 mBq in the active LXe.

^{222}Rn a is concern for $0\nu\beta\beta$ as one of it's daughters is ^{214}Bi. As discussed previously, this is a dominant background source due to the emission of the 2,447 keV γ-ray.

Fortunately, for decays that occur in the central volumes of the TPC, both the β particle and γ-ray will interact with the LXe, thereby resulting in a multiple scatter event that can be rejected.

A more problematic background from the ^{214}Bi is a naked-β decay, i.e., a β decay direct to the 0^+ state of ^{214}Po, and is therefore not accompanied by a de-excitation γ-ray. With a Q-value of 3,270 keV and branching ratio of 19.1% [89], 0.5% of naked ^{214}Bi decays can result in a single-scatter event in the ROI. However, the ^{214}Po promptly α-decays with a 164 μs half-life. This is is short enough that at least 99.99% of decays will occur within the same trigger window of the β-decay, resulting in a event that can be rejected due to the separation in time of the multiple scatters.

However, in LXe a fraction of the ^{222}Rn daughters are positively charged. This has been observed in EXO where the measured ion-fraction of ^{214}Bi$^+$ was found to be 76.4%.[91]. As positive ions these will drift towards the cathode and subsequently attach to it, where they will then eventually decay. In this case, it is possible that both the β-particle or α-particle are absorbed by the cathode wires and remain undetected, resulting in an irreducible background from the 2,447 keV γ-ray. The full extent of the radon background will not be known until LZ is operational, and so a conservative scenario is considered where the total ^{222}Rn background rate includes equal contributions from the uniformly distributed ^{214}Bi and the ^{214}Bi$^+$ that attaches to the cathode. The total background is estimated by simulating ^{222}Rn decays uniformly in the detector, together with ^{214}Bi$^+$ generated on the cathode, assuming and activity 2.0 μBq/kg for each. The Bi-Po coincidences are not modelled in the simulations and so a further rejection factor of 99.99% is applied to the background rate obtained from the simulations of the uniform ^{214}Bi.

For the events that survive the selections, the resulting energy spectrum in the inner volume and position distribution in the ROI is shown in figure 4.4. As expected, for the uniformly distributed ^{214}Bi only the naked β spectrum remains, and is reduced to an insignificant level due to the rejection from the α-coincidence. It is seen that the component of the background from ^{214}Bi$^+$ is highly localised towards the bottom of the LXe and does not penetrate far into the inner volume. The resulting background rate in the ROI is 6.4×10^{-9} events/kg/day/keV, corresponding to 0.3 events in 1,000 days.

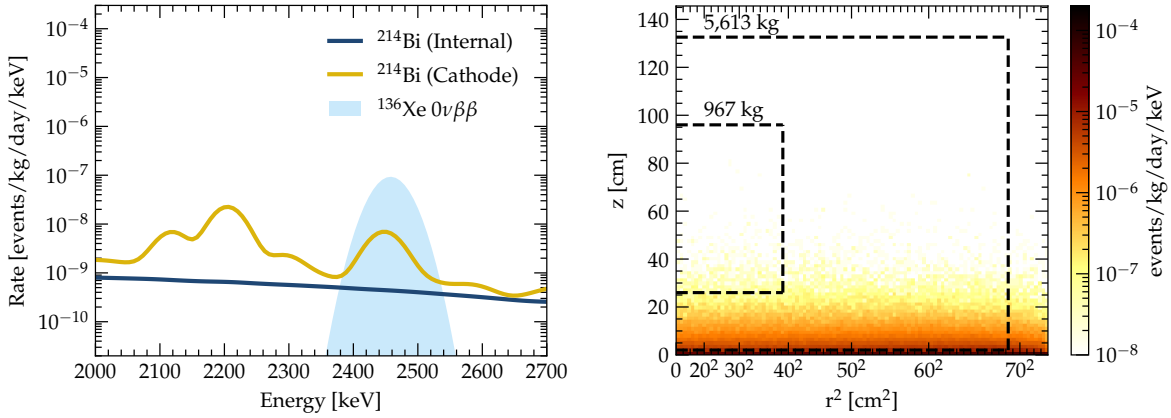

Figure 4.4: Left: energy spectrum of background events from radon contamination of the LXe for uniformly distributed ^{214}Bi and ^{214}Bi that is attached to the cathode. The solid blue region shows the expected signal spectrum given a half-life corresponding to the current lower limit. Right: spatial distribution of this background in the 5,613 kg fiducial volume

^{136}Xe

The $2\nu\beta\beta$ decay of ^{136}Xe to ^{136}Ba is known to occur with a measured half-life of 2.165×10^{21} years [92]. This process results in the emission of two neutrinos and two electrons with a total energy of $Q_{\beta\beta}$. The neutrinos will escape the detector undetected, whereas the electrons will transfer their share of the energy to the LXe. In the case that the combined electron energy falls within the ROI, such as an interaction would indistinguishable from a $0\nu\beta\beta$ event.

As the ^{136}Xe isotope is uniformly dispersed within the LXe this process will have a uniform position distribution. The energy distribution is modelled using the theoretical spectrum from [93], shown in figure 4.5. The $2\nu\beta\beta$ rate is significant outside of the ROI, with only a small leakage of these events into this region due to the energy resolution. With the measured half-life and 8.9% abundance, the background rate from $2\nu\beta\beta$ in the ROI is 3.2×10^{-11} events/kg/day/keV. In 1000 days, this corresponds to an insignificant 0.001 events.

Neutron Activated ^{137}Xe

Muon-induced and radiogenic neutrons can lead to the production of ^{137}Xe through ^{136}Xe(n,γ)^{137}Xe reactions. This isotope β-decays with a half-life of 3.8 minutes and Q-value of 4,173 keV. There is a 67% branching ratio to decay to the 0^+ state of ^{137}Ba

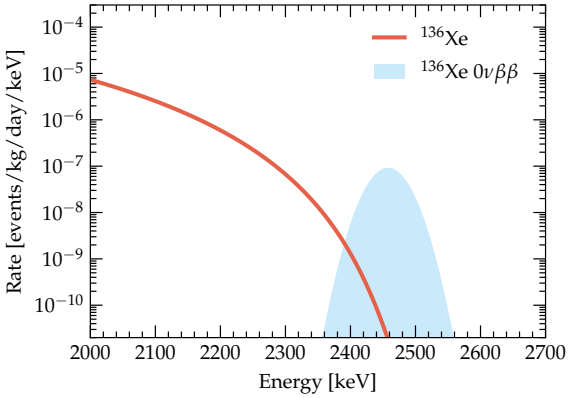

Figure 4.5: Energy spectrum of background events from ^{136}Xe $2\nu\beta\beta$. For comparison, also shown in solid blue is the expected signal spectrum given a half-life corresponding to the current lower limit.

[89] and subsequently no emission of dexcitation γ-rays. Such decays may therefore result in a background in the ROI.

The quantity of ^{137}Xe that will contaminate the LXe is difficult to estimate accurately as it depends on a number of factors such the muon flux and the abundance of thermal neutrons in the Davis Cavern. Measurements from EXO indicate approximately 70 events/(tonne ^{136}Xe)/year in the ROI [94]. However, LZ has a number of advantages that reduce the muon-induced ^{137}Xe background relative to EXO. The muon flux in the Davis cavern is estimated to be 6.2×10^{-9} cm^{-2} s^{-1} [86], nearly 100 times lower than observed in EXO. In addition, the veto systems and instrumented water tank can be used as an active muon veto. By looking back over several ^{137}Xe half-lives for these signatures, ^{137}Xe events can be rejected. These factors results in a factor of 100–1000 fewer events in the ROI per kg of ^{136}Xe compared to EXO. The muon-induced background is therefore considered negligible and is not modelled.

However, radiogenic neutrons may also cause ^{137}Xe production through thermal neutron capture. Though the thermal neutron flux is negligible within the LZ water tank [47], it is measured to be 1.7×10^{-6} cm^{-2} s^{-1} within the Davis cavern [95]. Xenon within the purification system is exposed to this thermal neutron flux, and given the 4000 kg/day xenon purification rate, 10 kg of xenon is delivered to LZ over the 3.8 minute half-life of ^{137}Xe, about half of which will go into the TPC. This background is simulated by generating ^{137}Xe decays uniformly within the LXe, and figure 4.6 shows the resulting energy spectrum. The resulting rate in the ROI is 3.5×10^{-9} events/kg/day/keV, or 0.24 events in 1,000 days.

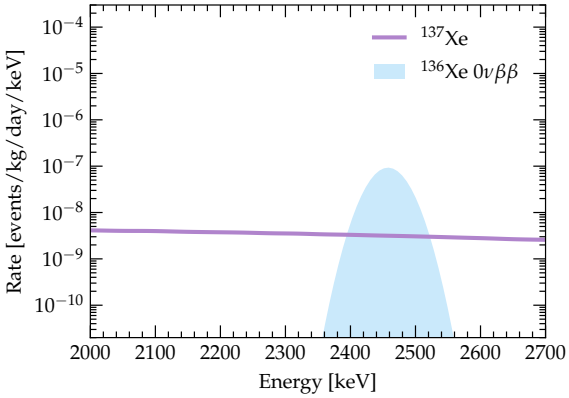

Figure 4.6: Energy spectrum of background events from the β-decay of ^{137}Xe. The solid blue region shows the expected signal spectrum given a half-life corresponding to the current lower limit.

Solar Neutrinos

^{8}B neutrinos produced in solar reactions can have energies exceeding 10 MeV, and therefore $\nu + e \rightarrow \nu + e$ scattering can result in a background in the ROI. These interactions will be uniformly distributed throughout the LXe, and the energy spectrum can calculated analytically using the differential cross section [96]

$$\frac{d\sigma\left(E_e, E_\nu\right)}{dE_e} = \frac{2G_F^2 m_e}{\pi}\left[g_L^2 + g_R^2\left(1 - \frac{E_e}{E_\nu}\right)^2 - g_L g_R \frac{m_e}{E_e E_\nu}\right], \tag{4.5}$$

where G_F is the Fermi constant, m_e is the electron mass, E_ν is the incoming neutrino energy, E_e is the outgoing electron energy and

$$g_L = \pm\frac{1}{2} + \sin^2\theta_W, \quad g_R = \sin^2\theta_W, \tag{4.6}$$

where θ_W is the Weinberg angle. The sign convention for g_L is positive v_e and negative for v_μ and v_τ. The differential electron energy spectrum is then calculated by convolution with the differential neutrino spectrum dN_ν/dE_ν,

$$\frac{dN_e\left(E_e\right)}{dE_e} = N_e \sum_i^{\text{flavours}} \int dE_\nu \frac{d\phi_i\left(E_\nu\right)}{dE_\nu}\frac{d\sigma_i\left(E_e, E_\nu\right)}{dE_e}. \tag{4.7}$$

Here N_e is the number of electrons in the LXe and the sum over flavours accounts for the possible neutrino oscillations between flavour states. Combining the differential neutrino

spectrum [97] with the measured cross section of 5.25×10^{-6} cm^{-2} s^{-1} [98] results in a rate in the signal ROI of 8.4×10^{-10} events/kg/day/keV, or 0.07 events in 1,000 days.

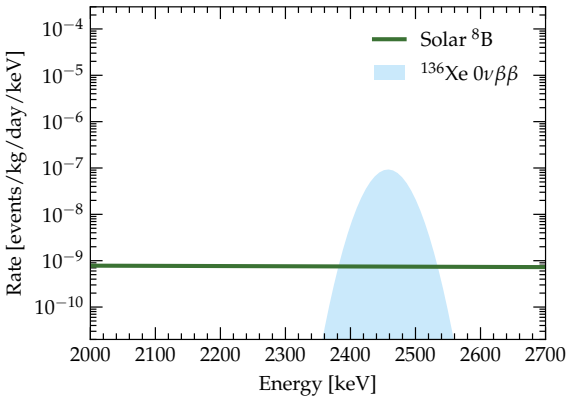

Figure 4.7: Energy spectrum of background events from ^{8}B solar neutrinos. The solid blue region shows the expected signal spectrum given a half-life corresponding to the current lower limit.

4.3.5 Summary

The total background rates in the signal ROI for each of the sources discussed in this section are given in table 4.2. The background in the ROI is dominated by the ^{238}U activity of the detector components and the ^{232}Th in the cavern rocks. The total background rate in the 967 kg fiducial volume, average over the ROI is 5.6×10^{-7} events/kg/day/keV. This corresponds to an expectation of 27 background events over 1000 days. Figure 4.8 shows the corresponding energy spectrum for each of these sources, and figure 4.9 demonstrates the impact of the selection criteria on the background rate. It is clear that the single scatter selection provides the strongest background rejection, resulting in more than a 10 times reduction in the background rate in the ROI. It is also seen that the skin and outer detector vetoes have a significant impact on reducing the background from the ^{208}Tl $2,614$ keV γ-ray.

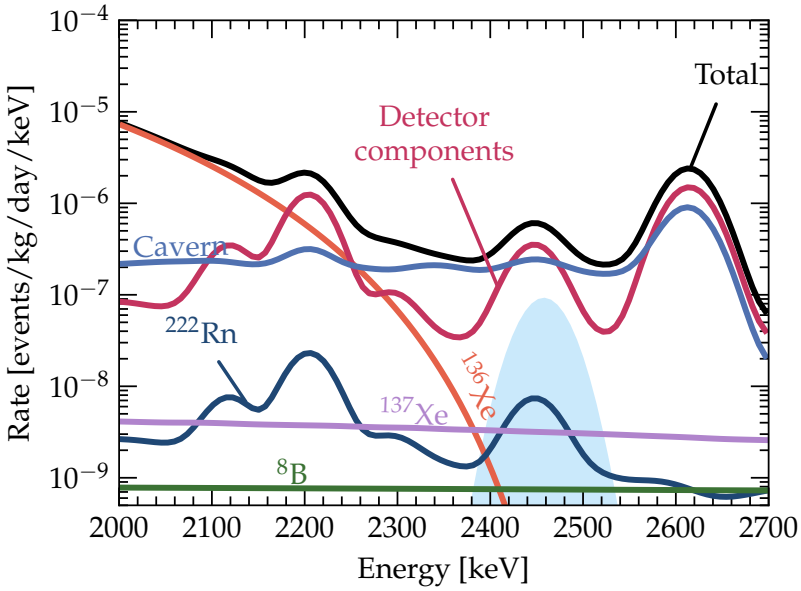

Figure 4.8: Energy spectrum in the 967 kg fiducial volume for each background source. The solid blue region shows the expected signal spectrum given a half-life corresponding to the current lower limit.

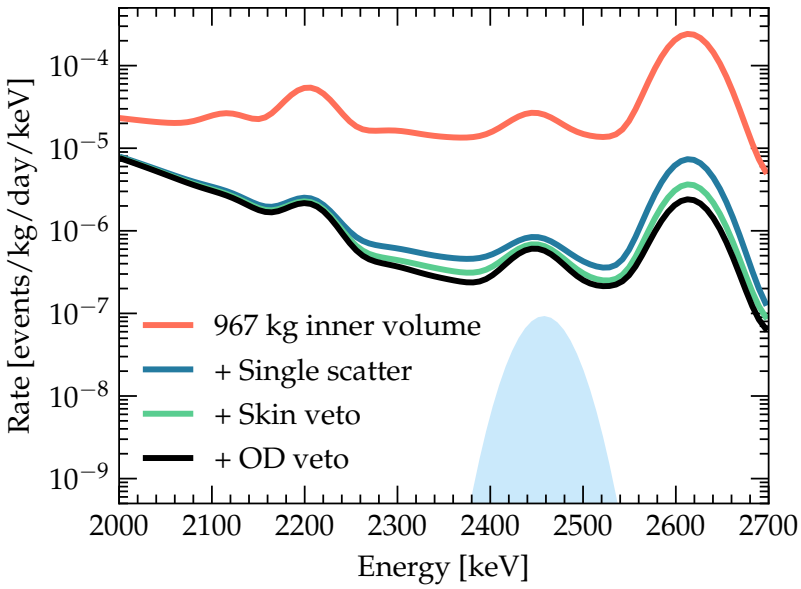

Figure 4.9: Effect of the successive selection criteria on the background energy spectrum in the 967 kg fiducial volume. The solid blue region shows the expected signal spectrum given a half-life corresponding to the current lower limit.

Table 4.2: Background rates, and expected events in 1,000 days, in a 967 kg fiducial volume and 1σ ROI around $Q_{\beta\beta}$ for each of each of the major background sources.

Background	Rate [events/kg/day/keV]	Events
Detector components	3.0×10^{-7}	15
Cavern rocks	2.5×10^{-7}	12
^{222}Rn	6.4×10^{-9}	0.3
^{137}Xe	3.5×10^{-9}	0.2
^{8}B	8.3×10^{-10}	0.04
^{136}Xe	3.2×10^{-11}	0.001
Total	$\mathbf{5.6 \times 10^{-7}}$	

4.4 Signal Response

The DECAY0 generator [99] was interfaced with BACCARAT in order to study the response of the detector to signal events and to estimate the reconstruction efficiency after selections. DECAY0 generates the initial kinematics of the two emitted electrons, which are then tracked through the LZ geometry by BACCARAT. Figure 4.10 shows the resulting single electron energy spectrum assuming the light neutrino exchange mechanism, demonstrating that on average each electron will share approximately 1,230 keV. Also shown is the corresponding distribution for the angle between the two emitted electrons. It is seen that they are preferentially emitted with an angle of 180°, however there is still significant probability for the emission across all other angles.

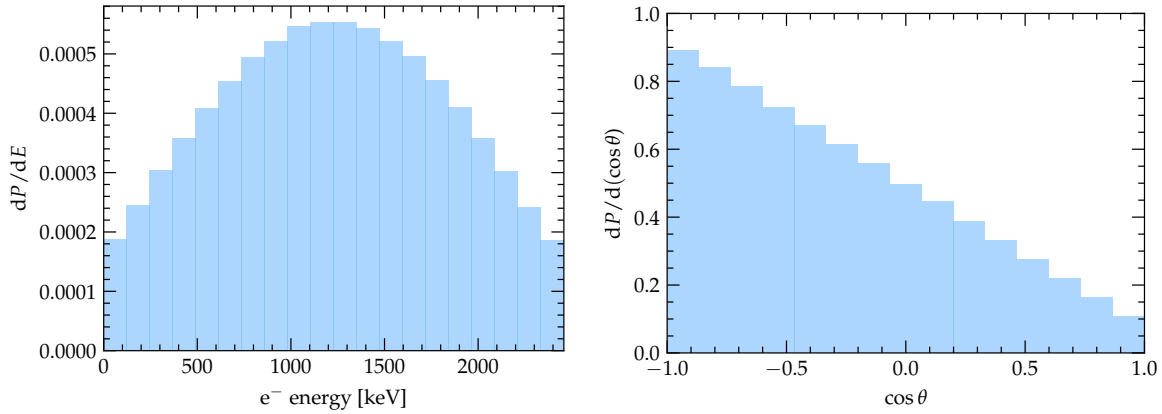

Figure 4.10: Distribution of the energy of a single electron emitted in ^{136}Xe $0\nu\beta\beta$ (left) and the angle between the two emitted electrons (right).

Figure 4.11 shows examples of the spatial distribution of energy deposits, or tracks, in the LXe resulting from the electrons interacting with the medium. Although it is most

probable that the two electrons are emitted back-to-back, the resulting paths that they take in LXe are from it; multiple scattering of the electrons results in tightly curled up tracks, with a high density of energy deposited at the end of the ends. The resulting distribution of the range of the tracks is strongly peaked at approximately 1 mm, as seen in figure 4.12. As this is below the expected position resolution of LZ, $0\nu\beta\beta$ events will therefore be reconstructed as single scatters with one S1 and one S2 pulse.

However, at the average energy of 1,230 keV energy losses via bremsstrahlung are probable. The emitted γ is likely to interact in the LXe at a different location than the electron emission, and if that location is displaced in the vertical direction, this can result in an event with multiple S2 pulses. This is responsible for the isolated energy deposits seen in figure 4.11. As the strongest method to reject backgrounds is through identification of multiple scatters, this results in a reduction in the signal efficiency, as shown in figure 4.12. After the selection that removes events with $\Delta z > 3$mm the resulting signal efficiency is 80%.

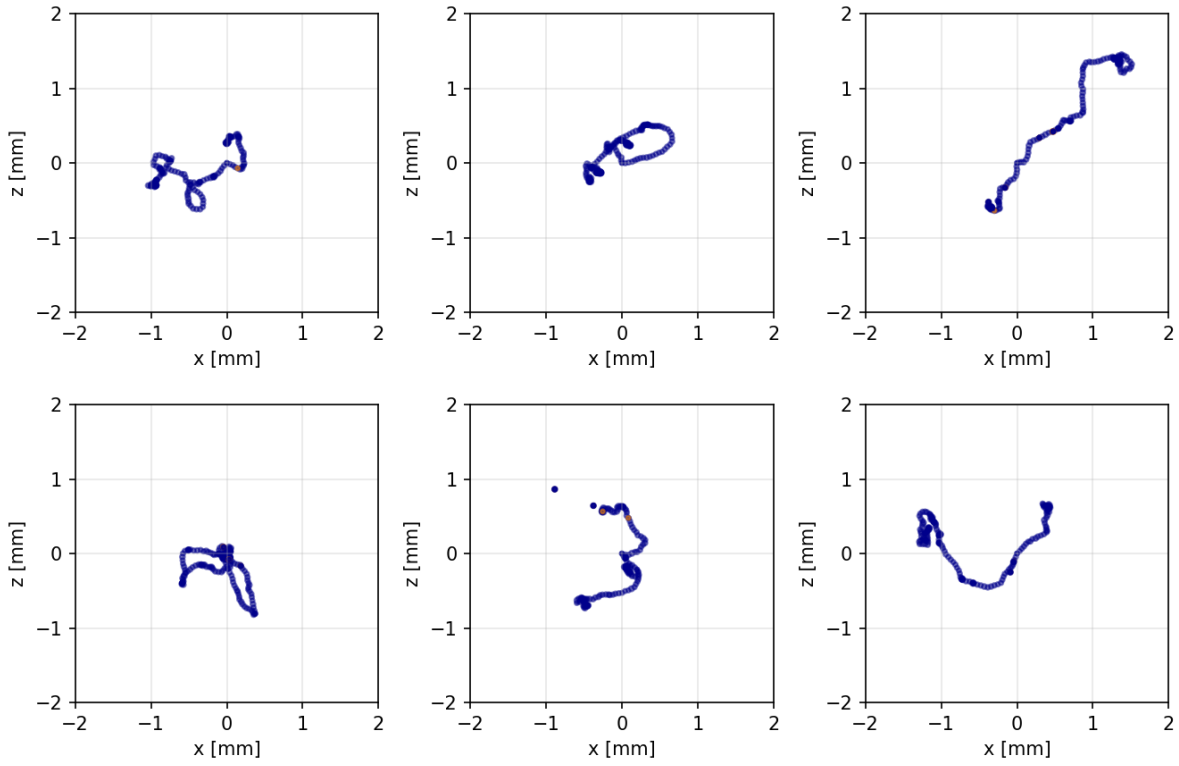

Figure 4.11: Simulated $0\nu\beta\beta$ electron tracks in LXe. The coloured dots represent the location of an energy deposit resulting from an electron (blue) or photon (orange).

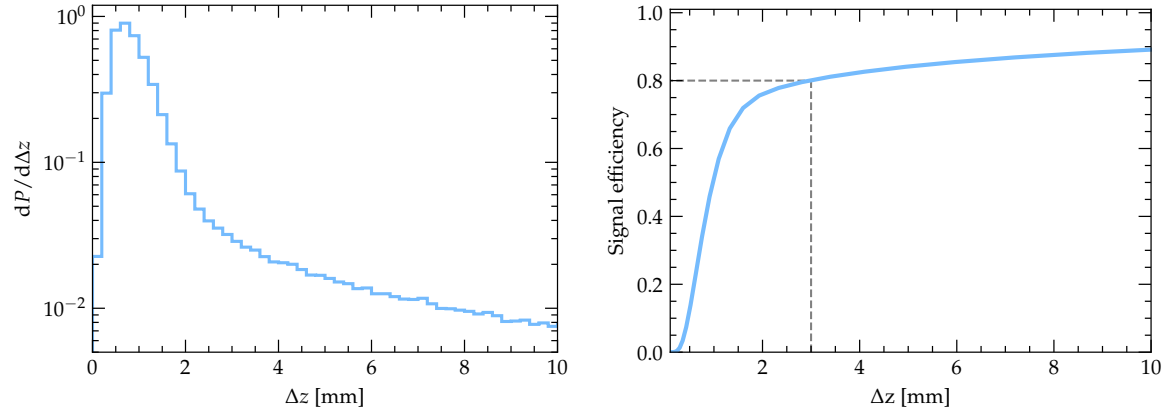

Figure 4.12: Distribution of Δz for simulated $0\nu\beta\beta$ events (left) and the resulting signal efficiency as a function of a Δz selection used to reject background events.

4.5 Sensitivity

The sensitivity to $0\nu\beta\beta$ provides a measure of the the half-lives that could be excluded in LZ. More formally, the it can be defined as the 90% confidence level (CL) upper limit on the number of signal events that would be obtained from the median background-only experiment. Given this upper limit, the lower limit on the half-life can then determined by

$$T_{1/2} \geq \frac{\epsilon_s N_{^{136}\mathrm{Xe}} T}{N_s^{\mathrm{upper}}} \ln 2 \tag{4.8}$$

where ϵ_s is the signal efficiency, $N_{^{136}\mathrm{Xe}}$ is the number of ^{136}Xe atoms, T is the number of live days of the experiment and N_s^{upper} is the upper limit on the number of signal events.

This section first introduces the statistical methods used to determine the sensitivity, and then presents the results using a likelihood based method that utilises the energy and spatial distributions of the signal and background events introduced in section 4.3. This method takes advantage of the strong position dependence of the backgrounds in LZ, resulting in improved background discrimination. The sensitivity is projected for a scenario assuming the natural 8.9% abundance of ^{136}Xe, in addition to a scenario where it is enriched to the 90% level. Finally, the effect of varying the energy resolution and background rate is studied.

4.5.1 Hypothesis Testing

Upper limits, and therefore sensitivity, can be determined using the framework of frequentist hypothesis testing. Hypothesis testing allows one to make statistical statements about the compatibility of a data set X with a null (H_0) and alternative hypothesis (H_1). Hypothesis testing relies on the use of a test statistic $T(X)$, which is a single-valued summary of the data set that should have a separable distribution under the two hypotheses. H_0 is rejected in favour of H_1 if

$$P(T \geq T(x)|H_0) \leq \alpha \tag{4.9}$$

where T is the distribution of the test statistic and α is the probability of rejecting the null hypothesis if it is true, or significance level. The probability of successfully rejecting the null hypothesis if it is not true is $1 - \alpha$, and is known as the confidence level (CL).

It is common in rare event searches to use the profile likelihood ratio (PLR) as the test statistic, defined as [100]

$$t_\mu = -2 \ln \frac{L(\mu, \hat{\hat{\nu}})}{L(\hat{\mu}, \hat{\nu})}. \tag{4.10}$$

Here μ is a hypothesised number of signal events, L is the likelihood and ν is a set of parameters that are not directly of interest, such as background rates from different background components, known as nuisance parameters. $\hat{\hat{\nu}}$ denotes the conditional maximum likelihood estimate (MLE) of the nuisance parameters under the tested hypothesis, and $\hat{\mu}$, $\hat{\nu}$ denote the global MLEs. This test statistic, ranging from 0 to ∞, measures the agreement of a hypothesis with a data set, where smaller values indicate the hypothesis is compatible and larger values indicate increasing incompatibility.

For the purposes of setting upper limits, the null hypothesis is that the data shows the presence of μ signal events plus a background (S+B), and the alternative hypothesis is the data contains only background events (B). It is also conventional to define the one-side PLR statistic,

$$q_\mu = \begin{cases} t_\mu & \hat{\mu} < \mu \\ 0 & \hat{\mu} \geq \mu. \end{cases} \tag{4.11}$$

If there is an under fluctuation of the background this definition prevents a signal smaller than the global MLE being incompatible with the data, which is consistent with the idea

of placing an upper limit. For a given data set with value of the test statistic q'_μ, the upper limit on μ at a CL of $1-\alpha$ is then the value of μ for which $p = 1-\alpha$, where

$$p = \int_{q'_\mu}^{\infty} f(q_\mu|S+B)\ \mathrm{d}q_\mu \tag{4.12}$$

is the p-value and $f(q_\mu|S+B)$ is the test-statistic distribution under the null hypothesis. The p-value is the probability that a value of q'_μ or greater would be obtained from the null hypothesis. In practice the upper limit is determined by scanning over a range of signal hypotheses, calculating the p-value and solving for μ numerically.

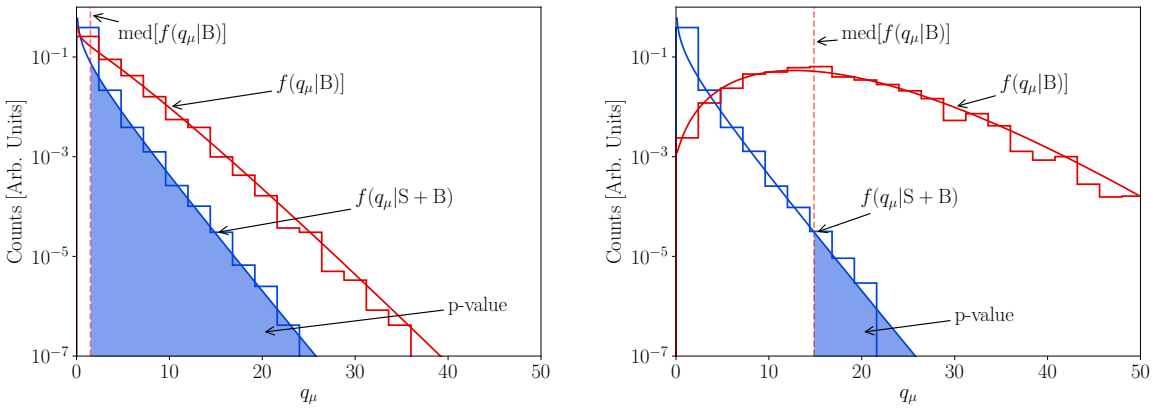

Figure 4.13: PLR test statistic distributions for signal plus background (S+B) and background only scenarios (B) where $\mu = 10$ (left) and $\mu = 60$ (right). The shaded region shows the p-value under the S+B hypothesis for the median B-only experiment.

The calculation of p-values requires the distribution of $f(q_\mu|S+B)$. This can be estimated from Monte Carlo simulations, and then the median upper limit can be obtained from a set of repeated background-only simulations. However, such simulations for complex models can be computationally expensive and therefore time consuming to run. Fortunately, in the large sample, or asymptotic limit, both the distribution of $f(q_\mu|S+B)$ and $f(q_\mu|B)$ can be determined analytically from Wilk's theorem. The median of $f(q_\mu|B)$ can then be used to represent the average background-only experiment.

Figure 4.13 shows examples of both the simulated and asymptotic distributions for a Poisson counting experiment with background B = 60 and where the values of $\mu = 10$ and $\mu = 30$ are tested. As μ increases, the background only distribution $f(q_\mu|B)$ shifts to higher values due the incompatibility of the data set with the signal hypothesis, which results in a lower p-value for the median experiment. If the p-value is smaller than the

chosen significance level then the signal hypothesis is rejected. It is clear that asymptotic approximation shows good agreement with simulations for both distributions.

4.5.2 Projections

The sensitivity analysis takes advantage of the precise multi-parameter reconstruction of events in the LZ TPC, namely the energy and 3-dimensional position, for enhanced discrimination of the backgrounds. Section 4.3 introduced the background model for $0\nu\beta\beta$, where each source is characterised by three observables: energy ($2000 < E < 2700$ keV), depth ($2 < z < 132.6$ cm) and radial position ($r < 68.8$ cm). In this fiducial volume there is approximately 5,613 kg of LXe, which results in 500 kg of ^{136}Xe assuming the 8.9% natural abundance. As demonstrated in Section 4.3, the self shielding LXe of LZ results in the majority of background events occurring close to the TPC walls and PMTs, producing a low background inner volume where the majority of signal sensitivity is expected. By using an extended fiducial volume, the backgrounds close to the TPC edges can constrain the rate in the inner volume. In addition, the shape of the spatial distribution can be used to discriminate between signal-like and background events. This results in improved sensitivity compared to that obtained from simple counting experiment.

Each background, along with the signal, is modelled with a probability density function (PDF) $P(E, r^2, z)$. For the detector components, cavern rocks, radon, and ^{137}Xe backgrounds, these PDFs are determined empirically from the energy deposit simulations discussed in section 4.2. However, due to computational constraints on the number of events that can be simulated, an empirical description of the full joint PDF is not possible without being statistically limited. As a result, it is instead approximated by decomposing into the factorised energy spectrum and spatial distributions $P(E, r^2, z) = P(E)P(r^2, z)$. For the backgrounds ^{8}B and ^{136}Xe backgrounds, which are expected to be uniformly distributed, this decomposition is exact and their analytical energy spectra is used with a uniform $P(r^2, z)$. Finally, the $0\nu\beta\beta$ signal is modelled with a flat spatial distribution and a Gaussian energy spectrum centered at $Q_{\beta\beta}$, with a width determined by the energy resolution.

The factorisation of the energy spectrum and spatial distributions is an approximation that assumes that the two observables are uncorrelated. In practice, this is not the case as low energy γ-rays are much more likely to be interact at larger r. However, in a suitable energy range this approximation is valid. To see this, consider that the joint

observed energy E and position x distribution $P(E, x)$ can be decomposed as

$$P(E, x) = P(E|x)P(x) \tag{4.13}$$

where $P(x)$ is the probability of an energy deposit at x and $P(E|x)$ is the probability of an energy deposit E conditioned on the position x. $P(E)$ depends on the distribution of the incoming external background energy E_s, and can be constructed by marginalising over that dimension

$$P(E|x) = \int \mathrm{d}E_s \; P(E, E_s|x) = \int \mathrm{d}E_s \; P(E|E_s, x)P(E_s|x) \; . \tag{4.14}$$

Combining this with Bayes' rule,

$$P(E_s|x) = \frac{P(x|E_s)P(E_s)}{P(x)}, \tag{4.15}$$

gives,

$$P(E, x) = \int \mathrm{d}E_s \; P(E|E_s, x)P(x|E_s)P(E_s) \; . \tag{4.16}$$

If the detector response is uniform, i.e the reconstructed energy has no intrinsic position dependence, then $P(E|E_s, x) = P(E|E_s)$ and,

$$P(E, x) = \int \mathrm{d}E_s \; P(E|E_s)P(x|E_s)P(E_s) \; . \tag{4.17}$$

So in general, $P(E, x) \neq P(E)P(x)$. $P(E|E_s)$ depends on cross section of the physical scattering process and the corresponding scattering energy distribution. Photoelectric absorption and Compton scattering are the dominant interactions processes for $E < 3$ MeV. To first-order, for photoelectric absorption

$$P(E|E_s) = \begin{cases} 1, & \text{if } E = E_s \\ 0, & \text{otherwise} \end{cases} \tag{4.18}$$

and for Compton scattering $P(E|E_s)$ depends on the scattering angle distribution $\mathrm{d}\sigma(E)_{KN}/\mathrm{d}\theta$ obtained from the Klein-Nishina equation. $P(x|E_s)$ depends on the incoming particle species, but for photons and neutrons it follows

$$P(x|E_s) \propto e^{-\mu(E_s)\Delta x} \; . \tag{4.19}$$

where Δx is the distance travelled through the LXe on the particle's path to point x, and $\mu(E_s)$ is the linear attenuation coefficient. This depends on the physics process and also determines the mean free path, $\lambda = 1/\mu$. Putting these pieces together gives,

$$P(E, x) \propto \int \mathrm{d}E_s \; P(E|E_s) e^{-\mu(E_s)\Delta x} \tag{4.20}$$

In a narrow energy ROI defined by $E_1 < E < E_2$ the total event rate is obtained by an integral over this range,

$$\text{Rate} \sim \int \mathrm{d}x \int_{E_1}^{E_2} P(E, x) \; \mathrm{d}E \; . \tag{4.21}$$

If only a limited subset of E_s contributes to the ROI, and in addition $\mu(E_s) = \mu$ is approximately constant for all relevant E in the integration, then

$$P(E, x) \propto e^{-\mu\Delta x} \int_{E_s} P(E|E_s) P(E_s) \; dE_s \tag{4.22}$$

$$= P(x)P(E), \tag{4.23}$$

i.e the energy spectrum is independent of the spatial distribution. For $0\nu\beta\beta$, the backgrounds that dominate result from the 2,447 keV γ-ray from ^{214}Bi and the 2,614 keV γ-ray from ^{208}Tl. In the limited ROI $\mu(E)$ is approximately constant, and therefore this approximation is considered to be valid with no significant impact on the sensitivity.

Given a data set of N events, the signal and background PDFs are combined to form the unbinned, extended likelihood function

$$L(\mu_s, \{\mu_b\}) = \frac{\mu^N e^{-\mu}}{N!} \prod_{i=1}^{N} \left[\left(\mu_s P_s(E_i, r_i^2, z_i) + \sum_{j=1}^{n_b} \mu_b^j P_b^j(E_i, r_i^2, z_i) \right) \prod_{k=1}^{n_b} \mathcal{N}(a_b^k, \sigma_b^k) \right], \tag{4.24}$$

where $\mu = (\mu_s + \sum_{i=1}^{N_b} \mu_b)$ is the total expected number of events with μ_s signal events and μ_b^i events for the i-th background source. The systematic uncertainties σ_b^k on the expected background rates a_b^k are included by assigning the nuisance parameters with the set of Gaussian constraint terms $\mathcal{N}(a_b^k, \sigma_b^k)$. Table 4.3 summarises each of the background sources included as parameters in the likelihood as well as the relative systematic uncertainties on their rate. The uncertainties on the detector component background rates are estimated from simulation and those on the cavern background are

Table 4.3: Table of the individual background sources and the relative uncertainties on their background rates used in the likelihood

Background	Source	Rel. uncertainty
Detector components	^{238}U	30%
	^{232}Th	30%
	^{60}Co	30%
Cavern rocks	^{238}U	50%
	^{232}Th	30%
Radon	^{214}Bi (internal)	50%
	^{214}Bi (cathode)	50%
Other	^{137}Xe	50%
	^{136}Xe	5%
	^{8}B	5%

from the uncertainty of the measured ^{232}Th and ^{238}U activities. The uncertainty for the ^{222}Rn component is driven by the range of the estimated contamination and those for ^{136}Xe and ^{8}B come from half-life and flux uncertainties, respectively. Finally, the internal ^{137}Xe background and ^{214}Bi cathode background are assigned a large uncertainty as their true rates will not be known until measured. However these are minor backgrounds and therefore do not significantly affect the sensitivity.

The 90% CL upper limit on the number of signal events is calculated using the profile likelihood ratio (PLR) method, utilising the asymptotic one-sided profile likelihood test-statistic introduced in Section 4.5.1 and the likelihood in 4.24. It has been verified that Wilk's theorem is valid and therefore that the asymptotic approximation of the test statistic distributions is is applicable.

Projection with Natural Abundance of ^{136}Xe

Figure 4.14 shows the median p-values obtained by scanning over the range of signal hypotheses for in addition to the 1σ and 2σ uncertainty, resulting in $\mu_{\text{upper}} = 32$ at 90% CL for a run of 1,000 days. This translates into a half-life sensitivity of $T_{1/2} \leq 1.06 \times 10^{26}$ yr at 90% CL. Although this is just below the current upper limit, it is significantly higher than the sensitivity of EXO and KamLAND-Zen, which are dedicated $0\nu\beta\beta$ experiments. Figure 4.15 demonstrates how the sensitivity increases with detector live time, showing that due to the low background environment of LZ, the sensitivity can surpass that of EXO within 1 year of running. After the full run of 1,000 days the sensitivity is

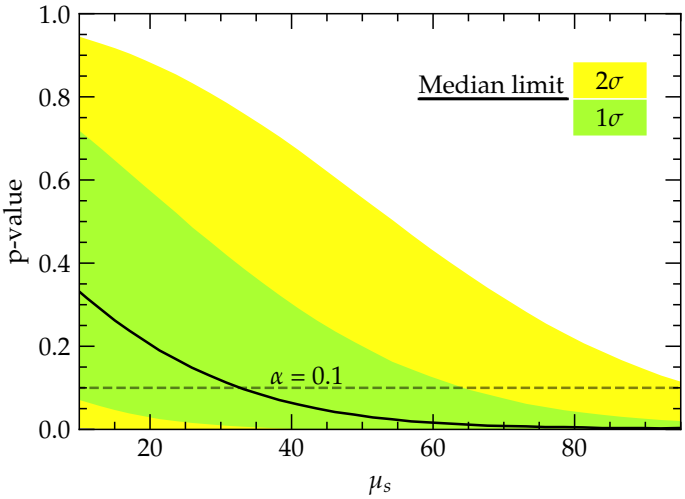

Figure 4.14: p-values of background-only experiments calculated assuming the null hypothesis of signal plus background. The solid line represents the median p-value and the intersection of this line with the dashed line indicates the 90% CL upper limit. The green (yellow) region indicates the $\pm 1\sigma$ ($\pm 2\sigma$) statistical uncertainty.

limited by the $1/\sqrt{B}$ behaviour of the background, resulting in minimal gains achieved by increasing the exposure. Under the assumption that light neutrino exchange is the driving mechanism for $0\nu\beta\beta$, the half-life sensitivity can be translated into the sensitivity to the effective Majorana neutrino mass $|m_{\beta\beta}|$ through the relation

$$\left(T_{1/2}^{0\nu}\right)^{-1} = \frac{|m_{\beta\beta}|^2}{m_e^2} G^{0\nu} |M^{0\nu}|^2. \tag{4.25}$$

Figure 4.16 shows that the expected sensitivity to $|m_{\beta\beta}|$ after 1,000 days is 53–164 meV, with the uncertainty driven by the range of estimates used for the nuclear matrix element [101, 102]. The phase space factor from [93] and an axial-vector coupling constant of $g_A = 1.27$ are used to calculate the effective neutrino mass.

Projection with 90% ^{136}Xe Enrichment

After completion of the full run of LZ, the exposure and therefore sensitivity to $0\nu\beta\beta$ could be increased by refilling the TPC with LXe that has been enriched to 90% ^{136}Xe. With more ^{136}Xe in the active volume, there will be a higher background rate from the $2\nu\beta\beta$ decay of ^{136}Xe and thermal-neutron induced ^{137}Xe. However, with 1.0% energy resolution, the additional background from $2\nu\beta\beta$ would still be insignificant in the

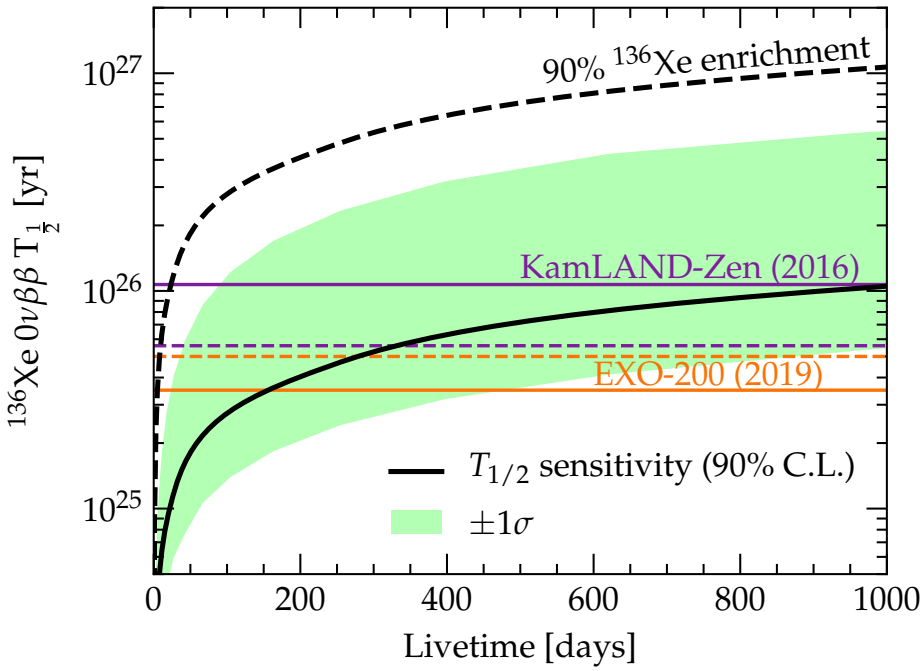

Figure 4.15: Half-life sensitivity as a function of detector live time. The green shaded region represents the $\pm 1\sigma$ statistical uncertainty on the sensitivity. The dashed black line shows the half-life sensitivity with 90% ^{136}Xe enrichment. The limits set by EXO (solid orange) and KamLAND-Zen (solid purple) are also shown, along with their respective sensitivities (dashed).

signal ROI and will therefore not significantly affect the sensitivity. Mitigation of the ^{137}Xe background would require the installation of a neutron shield around portions of the xenon purification system, but with sufficient shielding this background could be eliminated.

Figure 4.15 demonstrates that with 90% enriched xenon the sensitivity for a 1000 day run would reach $T_{1/2} \leq 1.06 \times 10^{27}$ yr, which would make LZ one of the most sensitive experiments searching for ^{136}Xe $0\nu\beta\beta$.

Detector Performance Studies

As the LZ experiment is not yet operational, a number of assumptions have been made regarding the performance of the detector. For this analysis, values for the the energy resolution at $Q_{\beta\beta}$ and rejection of multiple scatter backgrounds have been chosen based on reasonable assumptions. However, if these are not achieved then the sensitivity will clearly be reduced.

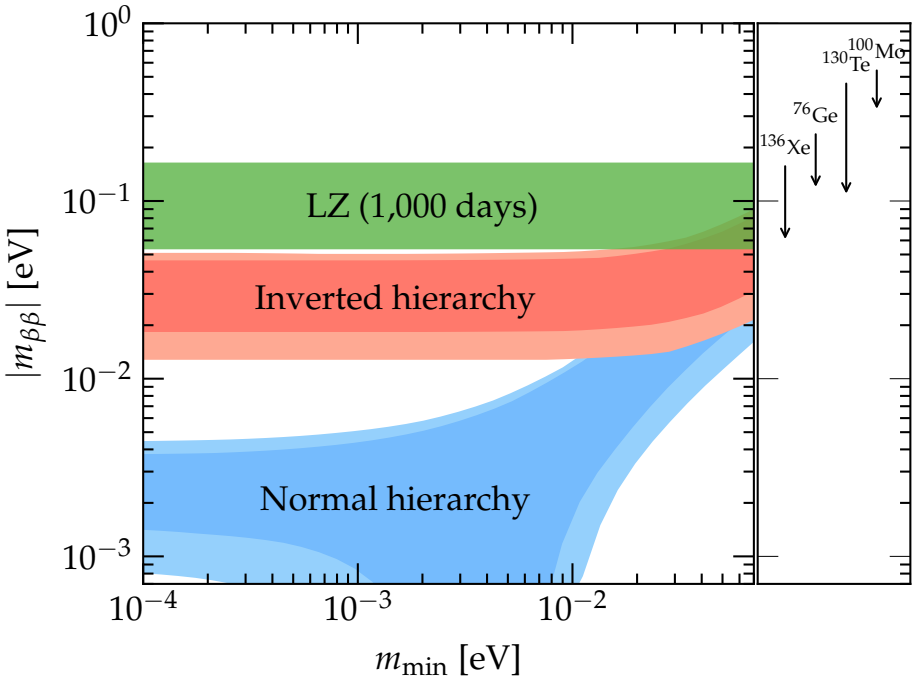

Figure 4.16: LZ projected sensitivity to $\langle m_{\beta\beta} \rangle$ and subsequently the neutrino mass hierarchy. The width of the green sensitivity band is due to the uncertainty in the nuclear matrix elements [101][102]. The red and blue contours show the allowed parameter space ($\pm$ 1 σ) for the inverted hierarchy and normal hierarchy neutrino mass scenarios, respectively. On the right are the current best limits and their uncertainties for different $2\nu\beta\beta$ isotopes, showing that ^{136}Xe provides the most stringent constraints on $|m_{\beta\beta}|$.

As the ability to distinguish signal events from the neighbouring ^{214}Bi and ^{208}Tl γ-ray peaks relies heavily on the energy resolution, the dependence of the sensitivity on the energy resolution at the ^{136}Xe Q-value is assessed and shown in figure 4.17. It is clear that an energy resolution slightly worse than the assumed 1.0% has a minor impact on the sensitivity. This is because the ^{214}Bi peak is already heavily overlapping with the signal ROI. However, if the energy resolution were 2.0% or larger, the backgrounds from the ^{208}Tl peak would begin to penetrate further into the ROI and reduce the sensitivity significantly. If LZ is capable of reaching a resolution of 0.8%, which has been achieved in the XENON1T experiment, then the projected sensitivity would be $T_{1/2} \leq 1.3 \times 10^{26}$ yr at 90% CL.

It is also assumed that multiple scatter events can be rejected with a depth-based vertex separation cut, as multiple energy deposits at different depths in the TPC will have multiple S2 pulses. For each of the background sources used in this analysis, a selection

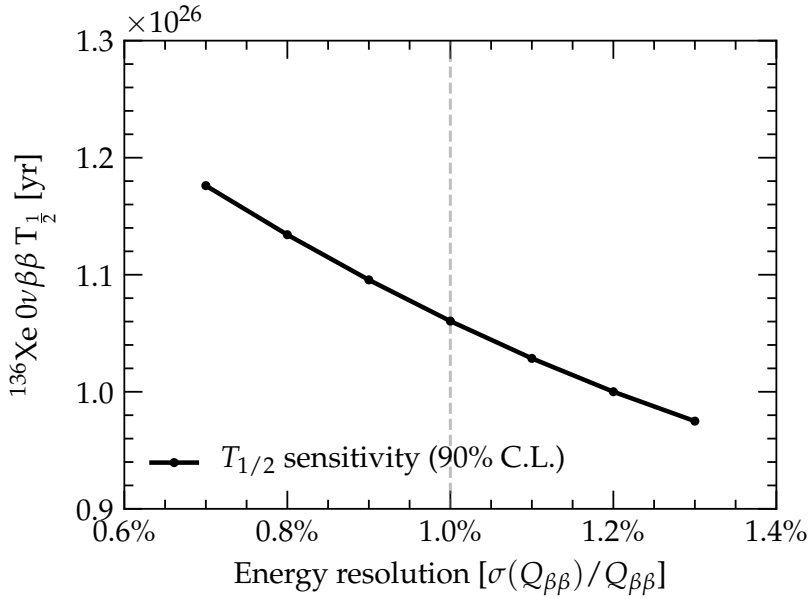

Figure 4.17: Half-life sensitivity after 1,000 live days as a function of energy resolution at $Q_{\beta\beta}$. The dashed line indicates the energy resolution assumed in this analysis.

of $\Delta Z < 3$ mm was used to reject multiple scatter events. However, in reality the ability to reject multiple scatters separated in depth depends on the performance of the pulse finder in the LZ reconstruction software. As a result, the impact on the sensitivity for different ΔZ selections is studied and shown in figure 4.18. Multiple scatter events are the dominant background source, and as the ability to separate multiple vertices worsens the background therefore increases significantly, resulting in a substantial decrease in sensitivity. As demonstrated in Section 4.4, removing events with $\Delta Z < 1.8$ mm also impacts the signal efficiency and therefore results in a decrease in sensitivity.

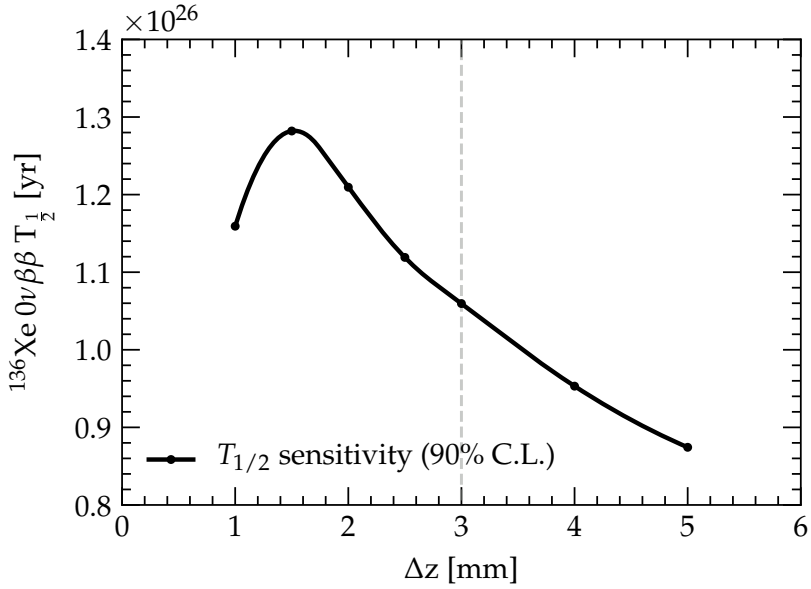

Figure 4.18: Half-life sensitivity after 1,000 live days as a function of the Δz selection. The dashed line indicates the Δz selection used in this analysis.

Chapter 5

Prospects for a Generation-3 Experiment

Although LZ is due to begin operations in late 2020, the research and development of future 3rd generation (G3) experiments is already underway. Scaling up the existing detector technology to an $\mathcal{O}$(10-100 tonne) LXe volume can provide orders of magnitude improvement in the ability to probe the parameter space of dark matter models in tandem with competitively searching for other processes such as $0\nu\beta\beta$. The DARWIN project [103] is one such R&D effort that proposes a detector with a dual phase TPC containing 40 tonnes of LXe. The significant increase in self-shielding results in a 5 tonne inner fiducial volume with an external background rate of < 1 events/year in the ^{136}Xe $0\nu\beta\beta$ ROI , resulting in a projected half-life sensitivity of 2.4×10^{27} yr after 10 years of exposure [104].

In addition to DARWIN, the UK dark matter community is currently developing proposals for a G3 experiment that would potentially be operated at the Boulby Underground Laboratory. This brief chapter presents the projected ^{136}Xe $0\nu\beta\beta$ half-life sensitivity for such an experiment. Predictions of the main background contributions in the ^{136}Xe $0\nu\beta\beta$ ROI are provided, and a semi-analytical model of the external γ attenuation is constructed to predict the background rate due to intrinsic radioactivity of the detector components.

5.1 A Generation 3 Rare Event Observatory

The goal of the UK G3 experiment (G3Xe) is to be a 'Rare Event Xenon Observatory', capable of performing world-leading science at both the low and high energy frontiers; specifically probing the parameter space of low mass dark matter, discovering coherent neutrino scattering in Xe, studying supernova neutrinos and searching for $0\nu\beta\beta$. Achieving sensitivity to a wide range of processes can arguably only be achieved with a multi-tonne scale detector, utilising an active LXe volume of $\mathcal{O}$(10t-100t). This is a consequence of the fact that sensitivity scales with exposure, and, coupled with the significant increase in self-shielding, the reduction in external backgrounds can result in orders of magnitude improvement in the ability to probe the remaining parameter spaces.

The G3Xe detector considered in this study is based on a scaled up version of LZ, with a cylindrical TPC of height 320 cm and radius 160 cm. This corresponds to an active mass of 75 tonnes of LXe, as depicted in figure 5.1. The LXe volume would be instrumented with two arrays of either PMTs or silicon-photomultipliers (SiPMs). Utilising SiPMs for light collection may offer the advantage of reduced radioactivity levels together with increased position resolution, resulting in improved ability to resolve multiple scatter and single scatter interactions [105][106]. As with LZ, the TPC would be placed inside a large double walled, low background titanium cryostat. The entire detector system would be operated deep underground, potentially at a dedicated cavern in the Boulby Underground Laboratory, and located inside a large instrumented water tank. This provides shielding from the cosmogenic muon flux and cavern γ radiation, and if the tank is instrumented it can act as Cerenkov detector to veto muon interactions.

5.2 Background Model

For this analysis, it is assumed that the energy resolution at $Q_{\beta\beta}$ is 1% and multiple scatter backgrounds with $\Delta z > 3$mm can be rejected. The backgrounds (described in more detail in section 4.3) are characterised by their rate in a $\pm 1\sigma$ ROI around $Q_{\beta\beta}$ (2433 keV – 2483 keV). The backgrounds considered are:

Materials: radioactivity of the detector materials produces external γ backgrounds arising predominantly from the decays of ^{214}Bi (2447 keV) and ^{208}Tl (2614 keV). It is assumed that the G3Xe detector is assembled from the same materials used in the construction of LZ, and therefore has the same contamination levels given in table 4.1.

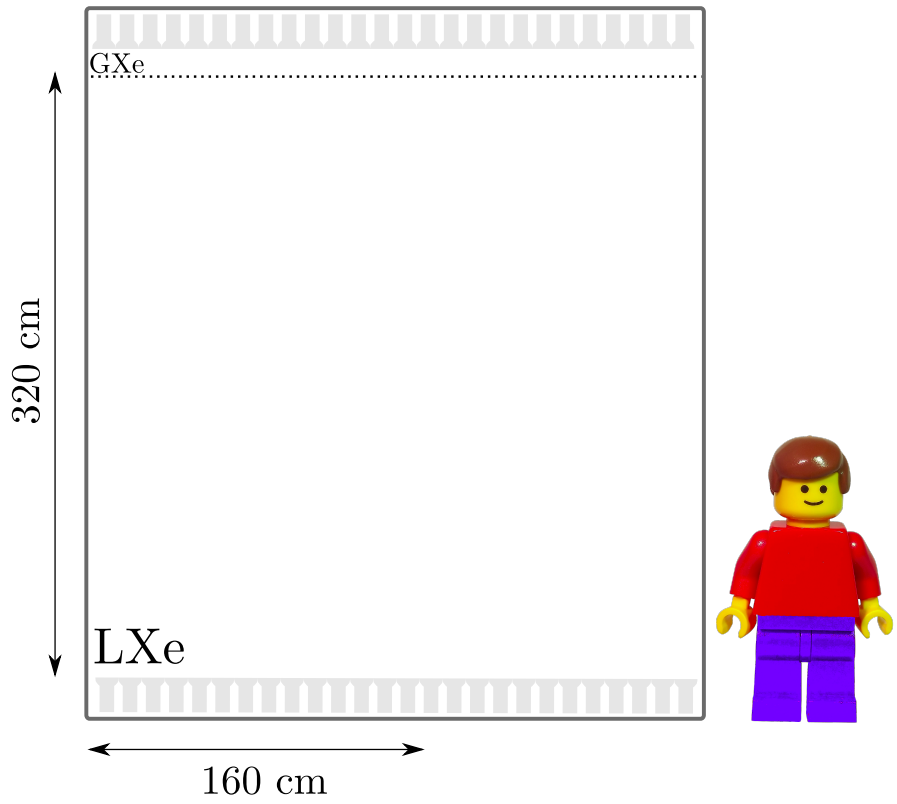

Figure 5.1: Cartoon demonstrating the scale of a G3 TPC, with dimensions corresponding to 75 tonnes of active LXe.

Due to the size of the G3Xe TPC, incident external γ radiation has to travel several attenuation lengths in order to reach the central volume of the TPC, and therefore the majority of this background occurs close to the PMTs and TPC walls. However, a precise prediction of this background from Monte Carlo simulations cannot be obtained at this stage, as there is no detailed design of the detector geometry. Instead, the spatial distribution of this background is estimated using a semi-analytical toy model of γ attenuation in a cylindrical volume of LXe, introduced in section 5.2.1.

^{137}Xe: muon-induced neutrons can capture on ^{136}Xe, resulting in the production of ^{137}Xe with a rate that depends on the muon flux at the location of G3Xe. Using the ^{137}Xe production rate measured by EXO-200 of approximately 3 atoms/tonne/year [94], and assuming operation at Boulby with a muon flux of $(4.09 \pm 0.15) \times 10^{-8}$ cm^{-2} s^{-1} [107], this equates to rate of 1.4×10^{-9} events/kg/day/keV in the ROI. This estimate is likely conservative; muon interactions in the water tank can be tagged, and a neutron capture event in the TPC is likely to be followed by the emission of de-excitation γs

that will multiple scatter in the TPC. Events in the ROI that occur in coincidence can therefore be vetoed.

^{222}Rn: A ^{222}Rn activity of 0.2 μBq/kg is assumed, which results in uniformly dispersed ^{214}Bi throughout the TPC. Assuming 99.99% of these decays can be rejected using the BiPo coincidence, the resulting rate in the ROI is 2.4×10^{-11} events/kg/day/keV. The capture of positive ions on the cathode is not considered, which is likely to reduce this background further.

^{8}B *solar* ν, ^{136}Xe $2\nu\beta\beta$: the irreducible $\nu - e$ elastic scattering and $2\nu\beta\beta$ results in a constant rate of background interactions that are uniform throughout the TPC, with rates in the ROI of 7.4×10^{-10} events/kg/day/keV and 2.9×10^{-11} events/kg/day/keV, respectively

5.2.1 Modelling External γ Backgrounds

The large LXe mass of G3Xe will significantly reduce the external γ background rate originating from the detector materials. The estimates for LZ obtained in chapter 4 used detailed simulations to predict this rate. However, instead of implementing a similar Monte Carlo simulation of G3Xe, this background rate is estimated through a simple model of γ attenuation.

Consider a γ source at a location $\overrightarrow{x_s}$ external to the LXe, and a infinitesimal volume element dVaround a position $\overrightarrow{x_d}$ in the LXe. The distance between the source and this volume is $r = |\overrightarrow{x_s} - \overrightarrow{x_d}|$, and the corresponding solid angle subtended by the volume element is $\mathrm{d}\Omega = \frac{\mathrm{d}V}{r^2\ \mathrm{d}r}$. For an isotropic source, the probability of a γ-ray being emitted in the direction of dV is then

$$\mathrm{d}P(\Omega) = \frac{1}{4\pi}\mathrm{d}\Omega. \tag{5.1}$$

Assuming a γ-ray of energy E travels along a straight path, and that outside of the LXe it is not attenuated by other components, the probability of reaching the position x_d without scattering or being absorbed is

$$P(x|E) = e^{-\mu(E)x} \tag{5.2}$$

where $\mu(E)$ is the linear attenuation coefficient of LXe and x is length of the straight line in the LXe that connects $\overrightarrow{x_s}$ and $\overrightarrow{x_d}$. The probability of an interaction between r

and $r + \mathrm{d}r$ is then

$$\mathrm{d}P(r|E) = \mu(E)e^{-\mu(E)x}\ \mathrm{d}r \tag{5.3}$$

Combining 6.1 and 6.3, the probability of a γ-ray emitted at x_s and interacting in $\mathrm{d}V$ is

$$\mathrm{d}P = \frac{\mu(E)}{4\pi}e^{-\mu(E)x}\mathrm{d}r\ \mathrm{d}\Omega = \frac{\mu(E)}{4\pi r^2}e^{-\mu(E)x}\ \mathrm{d}V, \tag{5.4}$$

which results in the probability density for an interaction throughout the detector volume,

$$\frac{\mathrm{d}P(r|E)}{\mathrm{d}V} = \frac{\mu(E)}{4\pi r^2}e^{-\mu(E)x}. \tag{5.5}$$

This expression relates only to a single γ of energy E emitted at a position $\overrightarrow{x_s}$. To determine the background rate for an arbitrary external background there are two missing ingredients: the spatial distribution of sources $\Theta(r_s, E)$ normalised to decays per unit time, and the probability that a γ-ray of energy E transfers an observable energy E_d to the LXe, $\mathcal{P}(E_d|E)$. As shown in figure 2.7, for $E < 3$ MeV interactions are dominated by Compton scattering and photoelectric absorption. For the latter, $\mathcal{P}(E_d|E)$ is calculated using the Klein-Nishina formula, and for the photoelectric effect all the energy is deposited and therefore providing that any secondary particles do not escape the detector, $\mathcal{P}(E_d|E) = 1$. The differential rate, in units of events/kg/day/keV, is then determined by,

$$\frac{\mathrm{d}R(r, E_d)}{\mathrm{d}M\mathrm{d}E_d} = \frac{1}{\rho}\int \mathrm{d}E \int \mathrm{d}r_s\ \mathcal{P}(E_d|E)\Theta(r_s, E)\frac{\mu(E)}{4\pi r^2}e^{-\mu(E)x} \tag{5.6}$$

where ρ is the density of LXe and $\mathrm{d}M = \rho\mathrm{d}V$ is the mass contained in an infinitesimal volume element. Unfortunately this expression is intractable as it requires integration over the unknown distribution of sources $\Theta(r_s, E)$. However, one approximation is to assume a finite number of sources N_s, with positions determined by delta functions

$$\Theta(r_s, E) = \sum_i^{N_s} \delta^3(r_s - r_s^i)\mathcal{S}_i(E) \tag{5.7}$$

where r_s^i is the position of the i-th source and $\mathcal{S}_i$ is the differential energy spectrum associated with that source. For γ-rays, this would be a discrete distribution of energies where the probability mass is determined by the branching ratio of the respective nuclear

transition. With this approximation, the differential event rate is

$$\frac{\mathrm{d}R(r,E_d)}{\mathrm{d}M\mathrm{d}E_d} = \frac{1}{\rho}\sum_i^{N_s}\int \mathrm{d}E \int \mathrm{d}r_s\ \mathcal{P}(E_d|E)\delta^3(r_s - r_s^i)\mathcal{S}_i(E)\frac{\mu(E)}{4\pi r^2}e^{-\mu(E)x} \tag{5.8}$$

$$= \frac{1}{\rho}\sum_i^{N_s}\int \mathrm{d}E\ \mathcal{P}(E_d|E)\mathcal{S}_i(E)\frac{\mu(E)}{4\pi r^2}e^{-\mu(E)x} \tag{5.9}$$

For a narrow ROI $E_1 < E_d < E_2$ in which $\mu(E)$ is approximately constant, this expression can be factorised into two separate components that independently describe the spatial distribution and the energy distribution of the rate. For ^{136}Xe $0\nu\beta\beta$ the majority of the external background rate is due to the decay of ^{214}Bi, which emits a γ-ray of $E = 2.447$ MeV, and in LXe has a corresponding linear attenuation coefficient of $\mu \approx 0.1$ cm^{-1}. Neglecting all other background sources, this becomes

$$\frac{\mathrm{d}R(r,E_d)}{\mathrm{d}M\mathrm{d}E_d} = \frac{1}{\rho}\sum_i^{N_s}\ \mathcal{P}(E_d|E)\mathcal{S}_i(E)\frac{\mu e^{-\mu x}}{4\pi r^2} \tag{5.10}$$

$$= N\frac{\sum_i \mathcal{S}_i(E)\frac{\mu e^{-\mu x}}{4\pi r^2}}{\int\ \mathrm{d}V \int\ \mathrm{d}E_d \sum_i \mathcal{S}_i(E)\frac{\mu e^{-\mu x}}{4\pi r^2}} \tag{5.11}$$

In the last step the normalisation has been made explicit, and N is the total integrated rate in the ROI. With this formulation it is now possible to approximate the differential rate using a discretised LXe volume, a simplified distribution of sources and an estimate of the total rate N.

The total rate is estimated using the LZ simulations described in section 4.2. The total rate from the LZ detector components in a 1σ ROI around $Q_{\beta\beta}$ is found to be 18 events/day. The corresponding rate in the G3Xe TPC is then obtained by scaling by the increase in TPC surface area, giving $N = 89$ events/day. The spatial distribution is approximated by a uniform placement of sources with $E = 2.447$ MeV on the surfaces of a cylinder of height 320 cm and radius 160 cm, approximating the ^{214}Bi backgrounds arising from the PMTs, PTFE walls and cryostats. It was found that using a ten-fold lower rate on the walls compared to the top and bottom surfaces produced a more accurate approximation.

The model is validated by comparing the integrated rate in a fiducial volume of mass M,

$$R_M(M) = \int_0^M dM \frac{\mathrm{d}R(r, E_d)}{\mathrm{d}M \mathrm{d}E_d}, \tag{5.12}$$

with that obtained from the LZ simulations. The comparison is shown in figure 5.2, and demonstrates that the model generally agrees well with Monte Carlo. However, for the inner volumes below 1 tonnes the rate is overpredicted by a factor of two. This disagreement is most likely due to the more accurate source placement and description of the detector geometry in the Monte Carlo, in addition to the modelling of second-order effects such as multiple scattering.

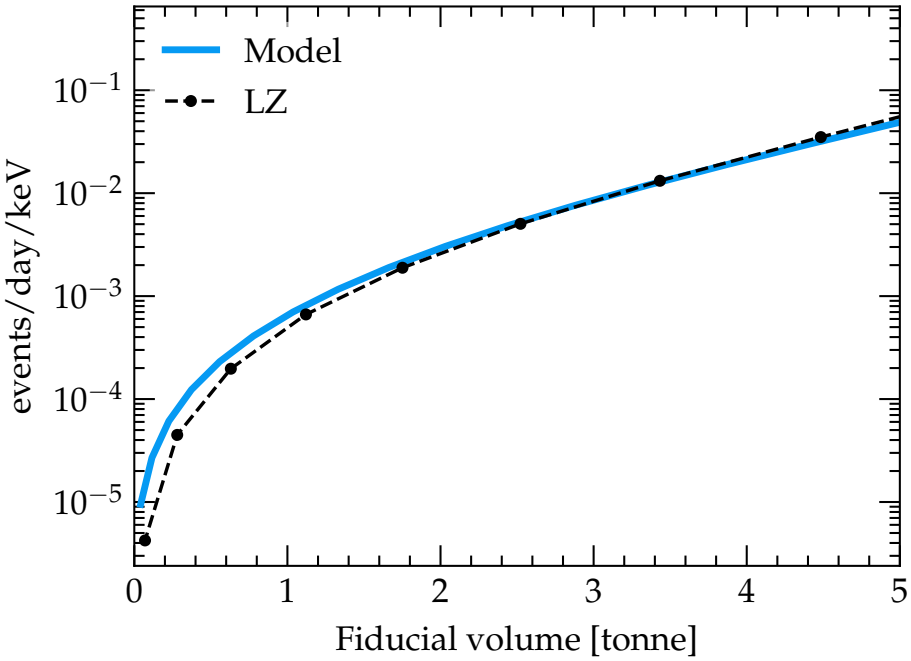

Figure 5.2: External background rate in LZ as a function of fiducial volume. The prediction from Monte Carlo simulations are shown by the dotted line, and the prediction from the model of γ attenuation is shown by the solid line.

The G3Xe differential rate in the ROI obtained from this model is shown in figure 5.3, and the comparison of the external rate with other background components in the ROI is shown in figure 5.4. The effect of self-shielding reduces the rate by over six orders of magnitude, resulting in an inner 10 tonne fiducial volume with a rate of 3.2×10^{-9} events/kg/day/keV. In this volume, the external rate approaches that of the total internal

background rate, and therefore it is approximately the optimal volume for minimising the external background and maximising exposure to the ^{136}Xe.

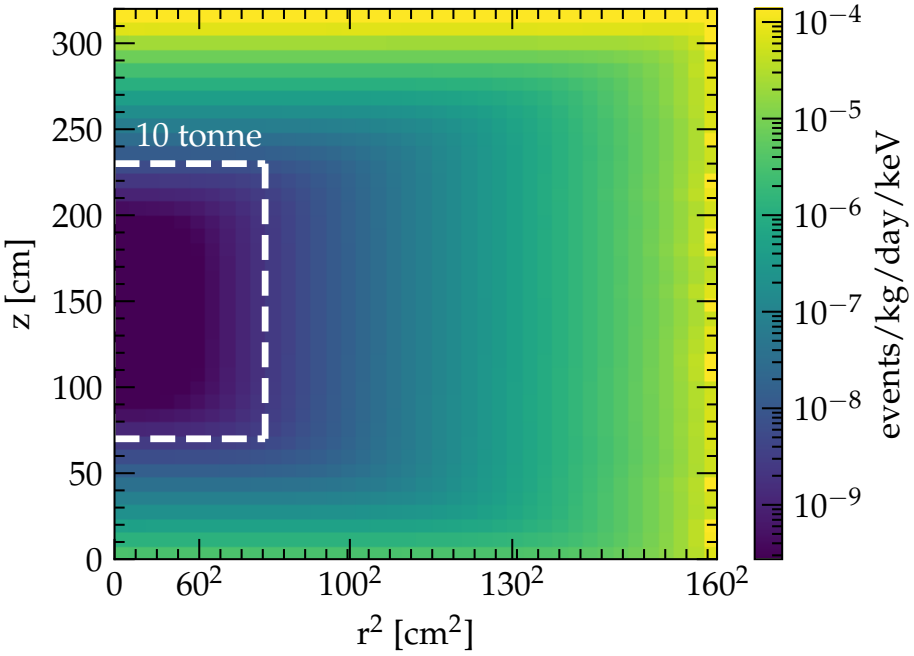

Figure 5.3: Predicted external background rate in the G3Xe TPC. The white dashed line indicates the 10 tonne fiducial volume.

5.3 Half-life Sensitivity

The ^{136}Xe $0\nu\beta\beta$ half-life sensitivity is estimated using an exposure in the 10 tonne fiducial volume, where the rates of each of the backgrounds in this volume are given in table 5.1. Two cases of ^{136}Xe enrichment are considered: a baseline scenario with natural abundance and a scenario with 90% enrichment. The purpose is to demonstrate the significant improvement in sensitivity that can be gained through enrichment, and allows for a direct comparison with dedicated multi-tonne scale enriched $0\nu\beta\beta$ experiments such as nEXO. In this scenario, the ^{137}Xe and ^{136}Xe $2\nu\beta\beta$ rates are increased by a factor of ten, and it is assumed no new backgrounds are introduced during the enrichment process. For both these cases, an extreme situation is considered where the external backgrounds are negligible and the rate is dominated by $2\nu\beta\beta$ and ^{8}B. Achieving this would require a significant reduction in both the muon flux and the radioactivity levels of detector components, or novel methods of distinguishing these interactions from $0\nu\beta\beta$ events.

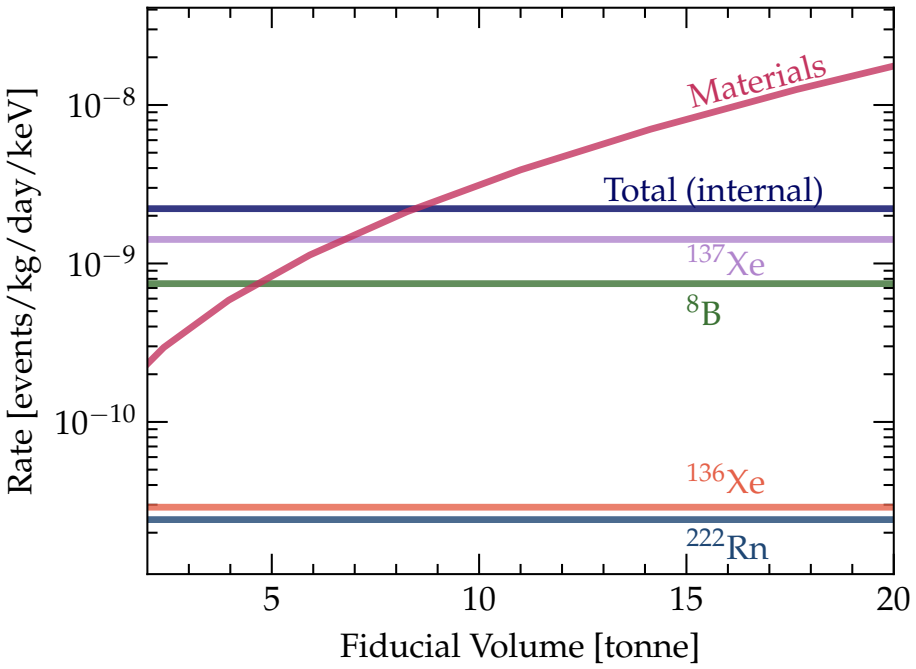

Figure 5.4: Rate in the ROI as a function of fiducial mass for each background component.

Table 5.1: Background rates in the 10 t fiducial volume, averaged over the 1σ ROI.

Background	Rate [events/yr]
^{238}U, ^{232}Th (Materials)	0.59
^{137}Xe	0.26
^{8}B	0.14
^{222}Rn	4×10^{-3}
^{136}Xe $2\nu\beta\beta$	5×10^{-3}
Total	**0.99**

The sensitivity is estimated using the frequentist profile likelihood method described in section 4.5.1. Each of the backgrounds is modelled by a probability density function $P(E)$ determined by the shape of the energy spectrum in the range 2000-2700 keV, shown in figure 5.5. The analysis uses the same $P(E)$ as that used in the LZ sensitivity projection (see section 4.5). Figure 5.6 demonstrates how the calculated half-life sensitivity evolves with the exposure time of the experiment. In the baseline scenario with natural abundance of ^{136}Xe, the projected half-life sensitivity at 90% CL is 3.2×10^{27} yr after 10 years of operation. For the 90% enrichment scenario, this increases to 1.8×10^{28} yr, and is competitive with the dedicated nEXO experiment. In the case that the background

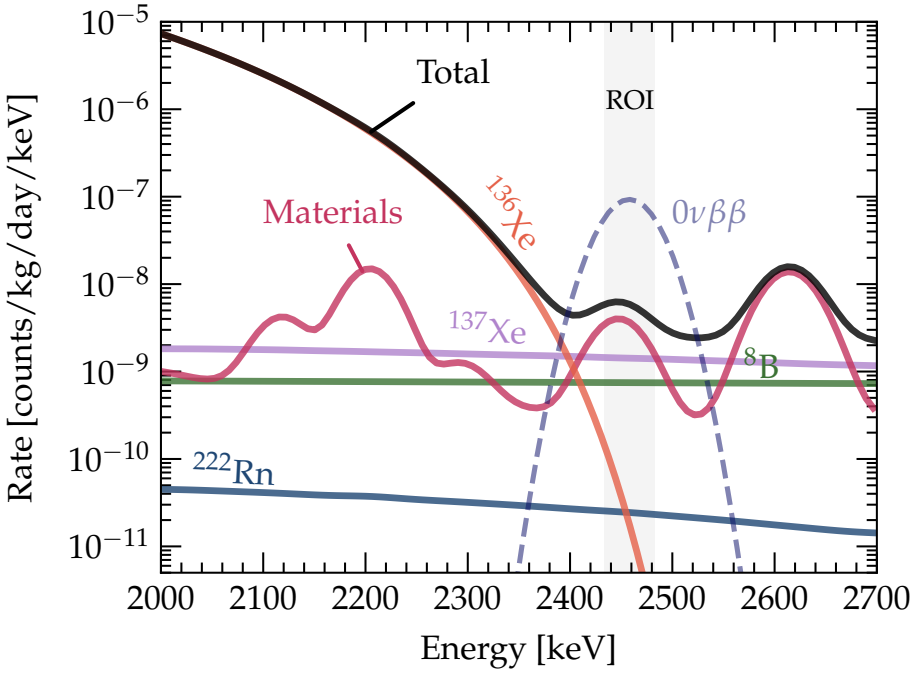

Figure 5.5: Background spectrum in the 10 t fiducial volume. The dashed signal spectrum corresponds to the expected rate at current upper limit on the half-life of 1.07×10^{26} yr.

is dominated by $2\nu\beta\beta$ and ^{8}B a significant a improvement in sensitivity is observed, reaching the nEXO sensitivity even with natural abundance. However, as shown in figure 5.8, this would require a factor 10-100 reduction in the background rate from the materials and ^{137}Xe. The sensitivity to the effective neutrino mass $\langle m_{\beta\beta} \rangle$ is shown in figure 5.7, where the range of values indicates the nuclear matrix element uncertainty. For the baseline scenario, it may be possible for G3Xe to cover all of the inverted hierarchy with a $\langle m_{\beta\beta} \rangle$ sensitivity of 30 meV to 10 meV, and in the enriched scenario, there is improvement in the ability to cover the normal hierarchy parameter space, with a $\langle m_{\beta\beta} \rangle$ sensitivity of 13 meV to 4 meV.

To conclude, a generation 3 dual-phase experiment utilising $\mathcal{O}(70$ t) of LXe could provide significant improvements in the search for ^{136}Xe $0\nu\beta\beta$, reaching sensitivity to $T_{1/2} > 10^{27}$ yr. Assuming the light neutrino exchange mechanism, such an experiment may therefore be able exclude much of the inverted hierarchy parameter space. Furthermore, by significantly reducing the radioactivity levels of detector components, and operating at a lab with low muon flux such that these backgrounds are negligible in a central fiducial volume, the sensitivity is competitive with the enriched nEXO experiment. Importantly,

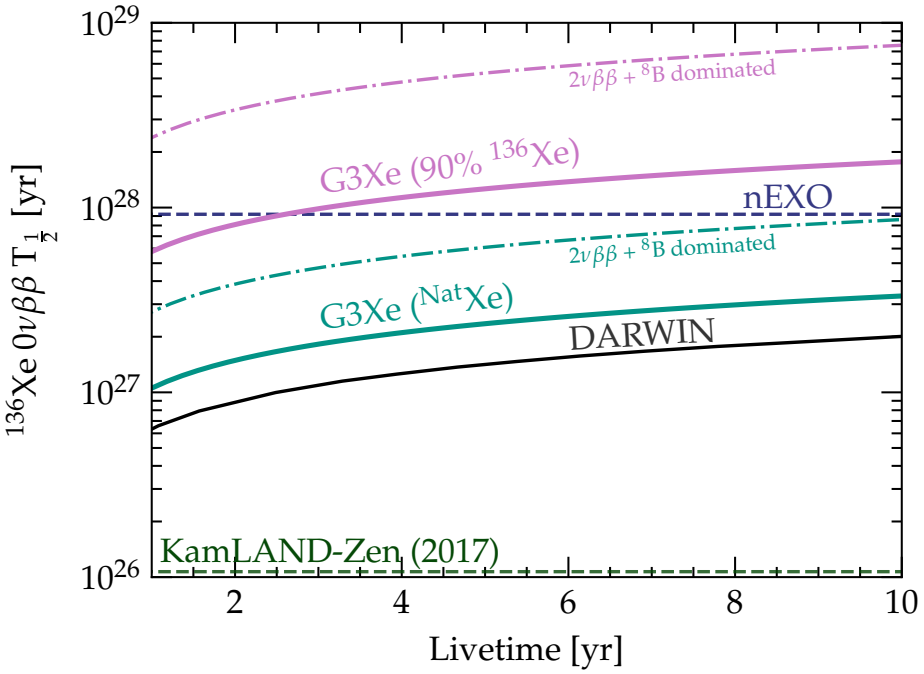

Figure 5.6: Projected ^{136}Xe $0\nu\beta\beta$ half-life sensitivity at 90% CL as a function of exposure time.

this sensitivity can be achieved even with natural abundance of ^{136}Xe, meaning that the sensitivity to other low energy searches is not diminished.

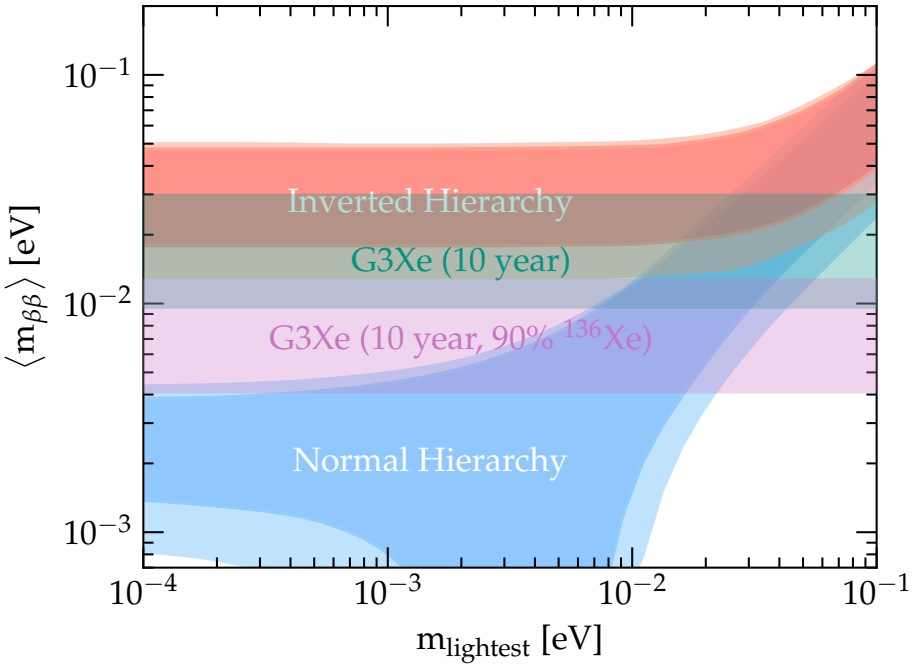

Figure 5.7: Effective electron neutrino mass sensitivity for a 10 year exposure in G3Xe, for both the baseline and 90% enrichment scenarios. The bands represent the uncertainty in the nuclear matrix elements.

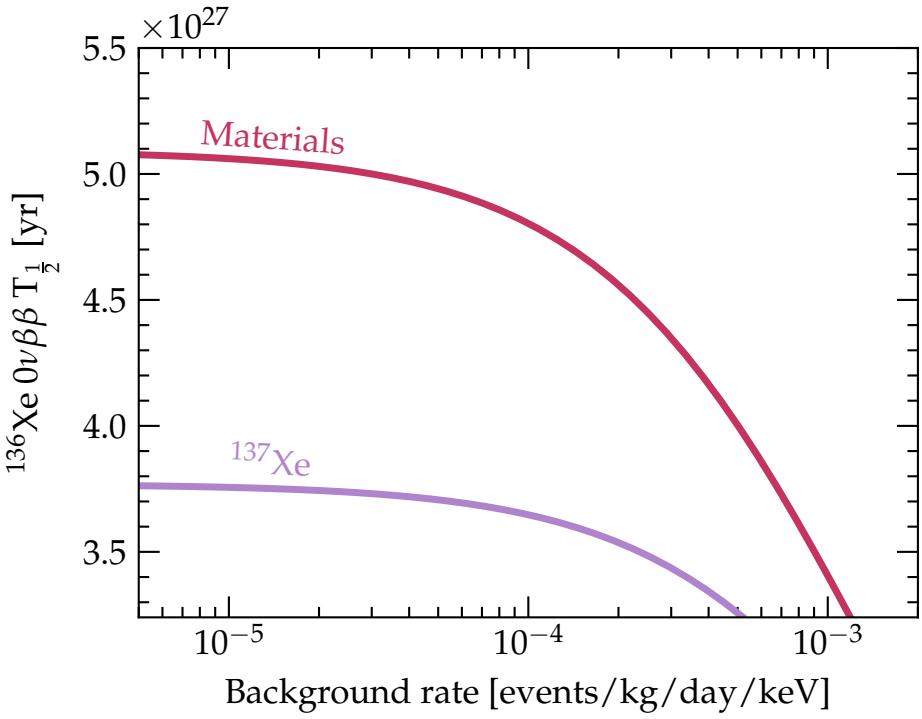

Figure 5.8: ^{136}Xe $0\nu\beta\beta$ half-life sensitivity versus background rate for the external materials and ^{137} backgrounds.

Chapter 6

Generative Models for Fast Simulation

Monte Carlo simulations are routinely used in high energy physics for a variety of purposes: e.g studying the response of a detector to radiation, generating precise background estimates, and calculating detection efficiencies. For large, complex detectors these simulations are both computationally and monetarily expensive. For example, the CMS experiment uses approximately 85% of it's total CPU capacity for simulation purposes at a predicted cost of more than 5 million US dollars per year [108]. With current generation dual phase Xe experiments, the large computational requirement of Monte Carlo simulations is already apparent, with many thousands of CPU hours required for large scale Monte Carlo productions. For G3 experiments this computing requirement will increase significantly. For example, for the detector considered in the previous chapter, the strong self-shielding means that 1 in every $10^5 - 10^6$ simulated events might reach the inner 10 t volume. Estimating an observable of the events in this volume with $\sim$ 1% statistical uncertainty would require simulating at least $\mathcal{O}(10^9 - 10^{10})$ initial events, which is several orders of magnitude more than required for LZ.

In LZ, the majority of processing power used for Monte Carlo simulations is spent on tracking the passage of particles through the detector system, and the subsequent digitisation of the detector readout. There are two main stages of the LZ simulation framework that perform these tasks:

BACCARAT: the Geant4 based framework that implements the LZ detector geometry, performs the generation of primary and secondary particles and tracks energy deposits along their paths through the detector. The Nobel Element Simulation Technique (NEST) [56] generates the initial S1 scintillation photons and ionised electrons at the sites of energy deposits in the LXe. The S1 photons are then ray-traced until they encounter a PMT photocathode. A ^{208}Tl 2.614 MeV γ-ray, for example, can produce

$\mathcal{O}(10^5)$ scintillation photons and a similar number of electrons. In the gas phase, each liberated electron can further produce thousands of scintillation photons, resulting in $\mathcal{O}(10^8)$ photons that must be tracked in order to model the full S2 signal. The simulation of S2 light is so intensive that the tracking of S2 photons is estimated to account for $> 95\%$ of the CPU time [86], and is therefore infeasible to use for large scale Monte Carlo productions. Instead, the spatio-temporal response to S2 light is approximated using a pre-generated lookup map, therefore bypassing the need to generate and track individual S2 scintillation photons, and resulting in a $\sim 20\times$ speed-up in simulation time. Delayed electron emission from the gate and cathode grids is also included to capture the structure of the S2 tail and model the backgrounds from spurious single electron signals. A subsequent module transforms the raw particle tracks and photon timing information into a Monte Carlo truth (MCTruth) event format used in the next stage of simulation.

Detector Electronics Response (DER): the simulation of the analog PMT signal generation and the subsequent signal processing performed by the amplifiers and digitisers of LZ. It reads in the individual MCTruth photon hit times, and passes the generated photoelectrons through an implementation of the electronics chain, resulting in simulated digitized waveforms for each PMT.

The processing time per event for each of these stages is given in figure 6.1 for different interaction energies. For nearly all energies the simulation of the electronics response is the most time consuming, requiring ∼50-100 seconds to generate the response for a single event. With expected event rates in LZ of $\mathcal{O}(10$ Hz), simulating even a small fraction of live time therefore requires significant computing resources. As a result, for the analyses in LZ that require high statistics simulations, such as sensitivity studies, the electronics simulation is not performed. Instead, only the energy deposit information from BACCARAT is used and passed to NEST to generate a simplified S1 and S2 response. Although this alternate simulation chain is significantly faster, it comes at the cost of reduced incompatibility with real data; PMT effects such as saturation are not modelled, and event reconstruction efficiencies cannot be taken into account.

There is a concerted effort within the HEP community [109] to upgrade existing solutions, and develop new technology to fulfill the computational requirements of future experiments. One component of this is the application of machine learning to simulation, specifically using generative adversarial networks (GANs) [110] for fast simulation. They have recently gained attention within the HEP community, with use cases including the simulation of jet formation [111], production of electromagnetic showers in multi-layer calorimeters [112], and generation of particle momenta [113]. This chapter studies the

effectiveness of using a GAN for the fast simulation of the digitised PMT response from an S2 signal in a dual phase LXe detector. A key benefit of GANs is their ability to be trained on real data, and therefore they can model the detector dependent effects that are not trivial to implement in simulation software. In addition, the underlying architecture automatically allows for parallelised data generation and can therefore result in significant speed-ups in processing time.

6.1 Generative Adversarial Networks

A GAN attempts to learn a generative model $q(x)$, of a distribution $p(x)$, from which it is possible to obtain samples $x \sim q(x)$ that are indistinguishable from $x \sim p(x)$. GANs are a form of latent variable model, i.e it models the distribution by integrating over an

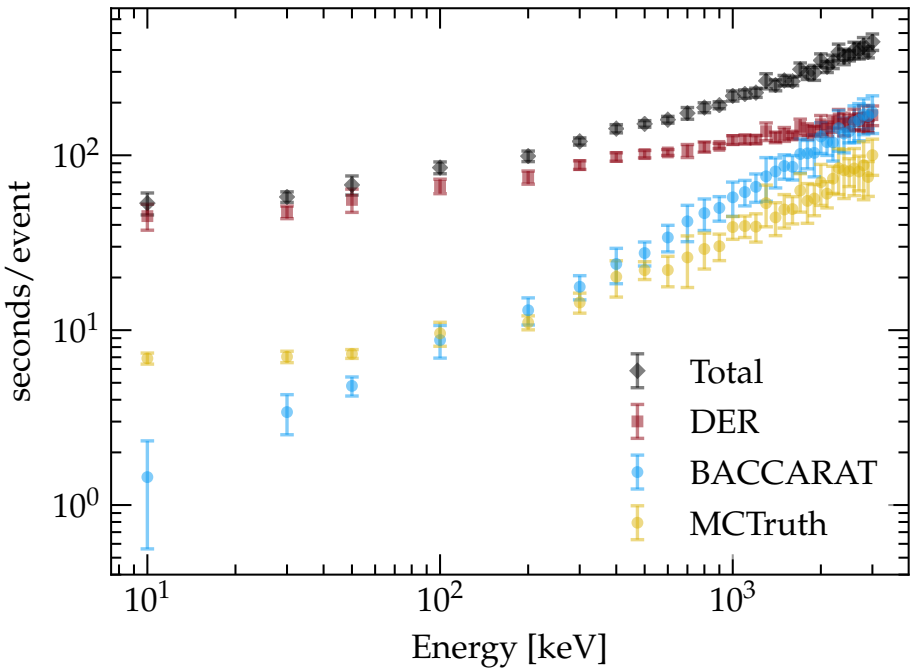

Figure 6.1: LZ Monte Carlo simulation processing time as a function of event energy, for both the Geant4 based framework BACCARAT and the detector electronics simulation (DER). Error bars represent the 1σ deviation around the mean obtained from 10 separate runs.

unobserved latent dimension z,

$$q(x) = \int \mathrm{d}z \; q(x,z) = \int \mathrm{d}z \; q_\theta(x|z)q(z) \tag{6.1}$$

where z has a prior distribution $p(z)$. It can be useful to think of z as an encoding of x in a lower dimensional space, where $q_\theta(\mathrm{x} \mid \mathrm{z})$ is a function that performs the decoding. In a GAN, this conditional distribution $q_\theta(x|z)$ is approximated by a neural network with parameters θ, called the generator $G(z)$. In theory, the optimal parameters $\hat{\theta}$ can then obtained by maximising the log-likelihood of generated samples

$$\hat{\theta} = \mathrm{argmax}_\theta \; \mathbb{E}_{x\sim p(x)} \int q(x|z)q(z) \; \mathrm{d}z \; . \tag{6.2}$$

However, in practice, this integral is intractable as it requires integrating over all possibilities of z . Instead, the parameters are optimised by introducing a second neural network $D(x)$ known as the discriminator, with an output that can be interpreted as the probability of a sample belonging to $p(x)$. The goal of the discriminator is to correctly classify samples as either belonging to the generator distribution $q(x)$ or the real distribution $p(x)$, thus the optimal discriminator D^* is given by

$$D^*_\theta = \mathrm{argmax}_\theta \; \mathbb{E}_{x\sim p(x)} \log\left[D_\theta(x)\right] + \mathbb{E}_{z\sim p(z)} \log\left[1 - D_\theta(G(z))\right]. \tag{6.3}$$

i.e the probability of the data being real and the probability of the generated data being fake is maximised. On the other hand, the objective of the generator is produce samples that fool the discriminator, thus the probability that they are fake must be minimised. The optimal generator is then given by

$$G^*_\theta = \mathrm{argmin}_\theta \; \mathbb{E}_{z\sim p(z)} \log\left[1 - D(G_\theta(z))\right]. \tag{6.4}$$

The generator and discriminator are then trained in an adversarial manner that can be seen as a the optimisation of a minimax game with the objective function

$$\mathrm{min}_G \mathrm{max}_D \; L(D,G) = \; \mathbb{E}_{x\sim p(x)} \log\left[D_\theta(x)\right] + \mathbb{E}_{z\sim p(z)} \log\left[1 - D_\theta(G(z))\right]. \tag{6.5}$$

It can be shown that optimising $L(D,G)$ is equivalent to optimising the Jenson-Shannon divergence between $q(x)$ and $p(x)$, with an optimum of $q(x) = p(x)$. In practice, the parameters of the generator and discriminator are optimised independently by a stochastic mini-batch gradient descent (SGD) algorithm, where each network is trained alternately

until convergence. The basic form of a SGD algorithm follows:

$$\theta_D \leftarrow \theta_D + \alpha \nabla_\theta \left[\frac{1}{m} \sum_{i=1}^{m} \log D_\theta(x_i) + \log(1 - D_\theta(G(z_i))) \right] \tag{6.6}$$

$$\theta_G \leftarrow \theta_G - \alpha \nabla_\theta \left[\frac{1}{m} \sum_{i=1}^{m} \log(1 - D_\theta(G(z_i))) \right] \tag{6.7}$$

where α is the learning rate (LR) and $\{x_1, x_2, ..., x_m\}$ is a randomly sampled batch of training data of size m.

This original GAN formulation suffers from a number of issues that result in an unstable optimisation problem that is unable to converge, and a generator that produces sub-optimal samples or suffers from mode collapse. This is where the generator does not capture the variance of $p(x)$ and learns to produce a single sample that always fools the discriminator. Many of the issues arising in the original GAN formulation can be understood to be affects caused by the specification of the objective function. Specifically, $L(D, G)$ suffers from vanishing gradients when $D(x)$ is too strong or if there is no overlap between $p(x)$ and $q(x)$. In this scenario, the parameters of the generator cannot be updated with a gradient descent algorithm and it therefore fails to produce realistic samples. A number of efforts have been made to improve GAN training [114] that solves these issues. Perhaps the most successful is the introduction of the Wasserstein GAN [115](WGAN) with gradient penalty (GP) [116]. The WGAN introduces an alternative objective function that uses the Wasserstein distance. This is provides a measure of the distance between two probability distributions $p(x)$ and $q(x)$ and is specified by

$$W(p, q) = \inf_{\gamma \sim \Pi(p,q)} \mathbb{E}_{(x,y)\sim\gamma}[\|x - y\|]. \tag{6.8}$$

It can be interpreted as the minimum work required to transform $q(x)$ to $p(x)$, where $\gamma(x, y)$ is a joint probability distribution that determines how much probability mass is moved from x to y, and $||x - y||$ is the distance travelled. In practice, the full set of distributions $\gamma(x, y)$ cannot be determined and the infimum cannot be found analytically. Instead the WGAN formulation uses a transformation of the the distance measure to form the objective function

$$W(p, q) = \max_{w \in W} \mathbb{E}_{x\sim p}[f_w(x)] - \mathbb{E}_{z\sim p(z)}[f_w(G_\theta(z))] \tag{6.9}$$

where f_w is a function with parameters w that is required to be K Lipschitz. This means it must satisfy,

$$|f_w(x) - f_w(x')| < K|x - x'| \tag{6.10}$$

With this formulation, $W(p, q)$ can be minimised to produce an approximation to $p(x)$. The aim of the discriminator is no longer to distinguish between real and generated samples, but to learn a function $f_w(x)$ that results in an accurate approximation of the Wasserstein distance for each G_θ. $f_w(x)$ is known as the critic, and can be interpreted as a function that provides a measure of how real a generated sample is. The objective function of the critic to minimise is,

$$L_c = \mathbb{E}_{z\sim p(z)}[f_w(G_\theta(z))] - \mathbb{E}_{x\sim p}[f_w(x)] \tag{6.11}$$

Given convergence of $f_w(x)$, the generator parameter updates then proceed via the minimisation of $W(p, q)$. The Lipschitz constraint on $f_w(x)$, which requires the normed gradients of $f_w(x)$ to be at most K everywhere, can be enforced using the gradient-penalty (GP) method. Here, a penalty term is added to the critic objective function to give

$$L_c = \mathbb{E}_{z\sim p(z)}[f_w(G_\theta(z))] - \mathbb{E}_{x\sim p}[f_w(x)] + \lambda\mathbb{E}_{\tilde{x}\sim\tilde{p}}[(||\nabla_{\tilde{x}} f_w(\tilde{x})|| - 1)^2], \tag{6.12}$$

where $\tilde{x}$ is an interpolated sample between x and $G(z)$. This term, regulated by the parameter λ, enforces the gradient of $f_w(x)$ to be close to unity between the space of real and generated samples. WGAN-GP has a number of benefits compared to the original GAN formulation:

Convergence: WGANs do not suffer from vanishing gradient problems, and are therefore much more stable during training. If the optimal critic function is obtained, they are also theoretically guaranteed to converge.

Evaluation metric: the Wasserstein loss can provide a measure of generated sample quality, with lower values indicating improved realism in the generated samples.

Sample quality: experiments suggest that WGANs produce more realistic samples and do not suffer from mode collapse.

6.2 PMT Waveform Generation with WGAN-GP

As discussed at the beginning of this chapter, the simulation of the S2 light production, followed by the subsequent generation of the PMT response, is the most computationally intensive component of Monte Carlo simulation in dual phase Xe detectors. Generative adversarial networks are an attractive alternative to the traditional MC methods currently used. Once trained, sampling from the generator is fast; a forward pass of the neural networks is just a sequence of matrix multiplications that can naturally be performed in parallel. In addition, if trained on real data, a GAN may be able to model subtle effects that are difficult to incorporate accurately in the traditional Monte Carlo, for example spatial non-uniformities in the extraction field or the distribution of light in the S2 tail. This results in a more accurate simulation and therefore the need to perform fewer corrections when comparing Monte Carlo to data.

This section introduces a GAN architecture that can produce a realistic, digitised PMT response to S2 light, demonstrating a fast simulation alternative to the traditional tracking of individual photons and their propagation into an observed PMT signal. This significantly reduces both the simulation time and the storage requirements for persisting the various stages of the MC simulation.

6.2.1 Architecture

A relatively small sample of 2800 ^{83}Kr LUX calibration events are used to train a WGAN-GP architecture. These events are uniformly distributed in the detector with a mono-energetic energy of 41.5 keV. For each event, the total S2 pulse summed across all PMTs is extracted and padded with baseline noise up to a maximum of 1000 samples (10μs). The architecture of the generator and discriminator networks is given in table 6.1, and follows that of the Deep Convolution GAN (DCGAN) [117] where convolutional layers are used throughout. The latent distribution $p(z)$ has 100 dimensions, where each value is sampled from a normal distribution of $\mu = 0$ and $\sigma = 1$. The Adam algorithm [118] with a learning rate of 0.005, $\beta_1 = 0.5$ and $\beta_2 = 0.999$ was used to optimise the WGAN-GP objective function, with $\lambda = 1$. The model was constructed using PyTorch [119], and trained on a single NVidia P100 GPU. Training was performed for a total of 5000 epochs with a batch size of 128. At each training iteration, the discriminator network was trained twice to improve the approximation of Wasserstein loss.

Table 6.1: Generator and discriminator network architectures for the generation of S2 pulses. A kernel size of 3 is used in all convolutional layers, with a stride of 1 in the generator and stride of 2 in the discriminator

Generator
Input: 100 dimensional noise $\sim \mathcal{N}(0, 1)$
Linear$(100, 250 \times 8)$
Upsample(2)
Conv1d(8, 32), BatchNorm1d, LeakyReLU(0.2)
Upsample(2)
Conv1d(32, 16), BatchNorm1d, LeakyReLU(0.2)
Conv1d(16, 1)
Discriminator
Input: 1000 dimensional vector
Conv1d(8, 32), Dropout(0.25), LeakyReLU(0.2)
Conv1d(16, 8), Dropout(0.25), LeakyReLU(0.2)
Linear(8000, 1)

6.2.2 Evaluation

The pulses generated by the GAN look realistic on visual inspection, possessing the approximate Gaussian shape with a noisy tail, as shown in figure 6.2. Also shown for each generated pulse is the nearest real pulse as measured by the Euclidean distance, demonstrating that the GAN has not simply memorised the training set and is capable of generating unseen samples. To further assess the quality of generated pulses, the distributions of different pulse parameters are compared with those of the real data. Figure 6.3 compares the pulse amplitude and pulse area distributions, and figure 6.4 compares the pulse timing, specifically the time at which the pulses reach 10%, 50% and 90% of the total area. Reasonable agreement is observed, with the GAN generated samples matching the non-trivial shapes of these distributions. The variance of the distributions is captured surprisingly well; quite remarkable considering the GAN has access to less than 3,000 individual training samples from which to learn the complex underlying distribution. It is likely that improved results would be obtained using a much larger training dataset.

The structure encoded in the generator latent space is investigated by interpolating along the line joining two random points z_1 and z_2. At various points on the line, generated pulses are produced and plotted in figure 6.5. It can be seen that the moving in the first direction results in pulses with increased width and reduced area, indicating that information about the interaction depth in the detector is being encoded. Moving a

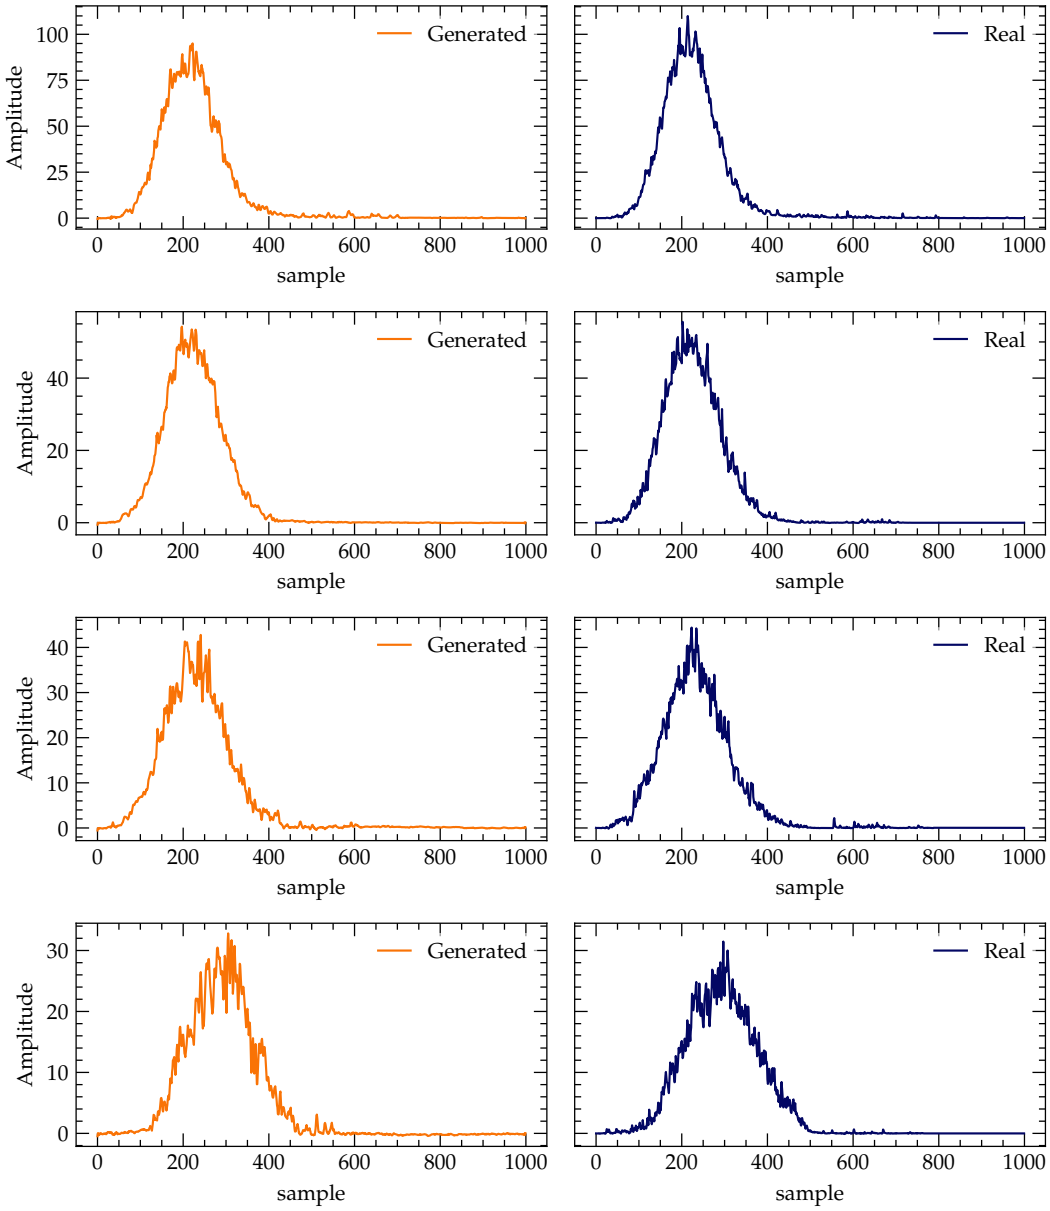

Figure 6.2: Examples of generated S2 pulses (left) shown together with the closest real S2 pulse in the training set (right).

long the second direction results in a clear decrease in pulse area, but a approximately constant width, suggesting that the latent space is capable of encoding information about the energy of the interaction.

Unlike the traditional MC simulation, GAN generation time is independent of energy. Once trained, only a forward pass of the sampled latent space through generator network is required. As this is nothing but a sequence of matrix operations, pulses can naturally

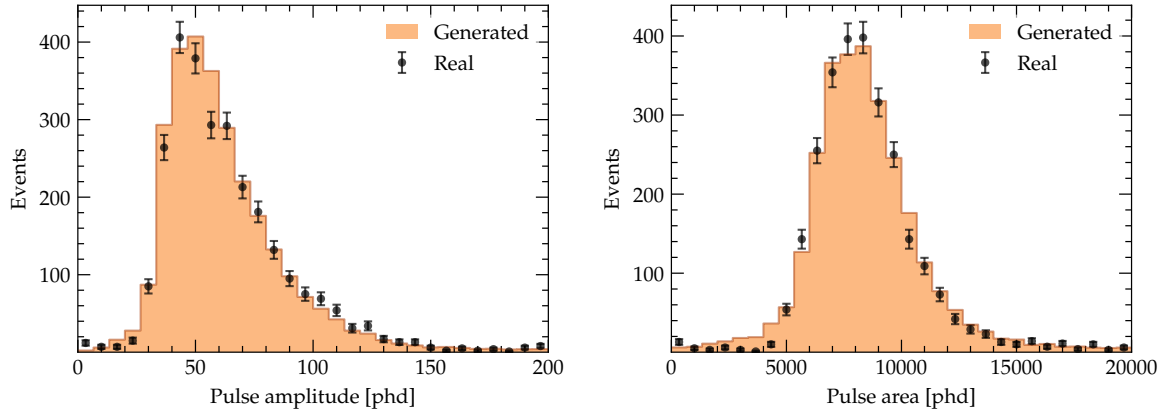

Figure 6.3: Comparison of S2 pulse amplitude (left) and pulse area (right) for the true distribution (markers) and the generated distribution (histogram).

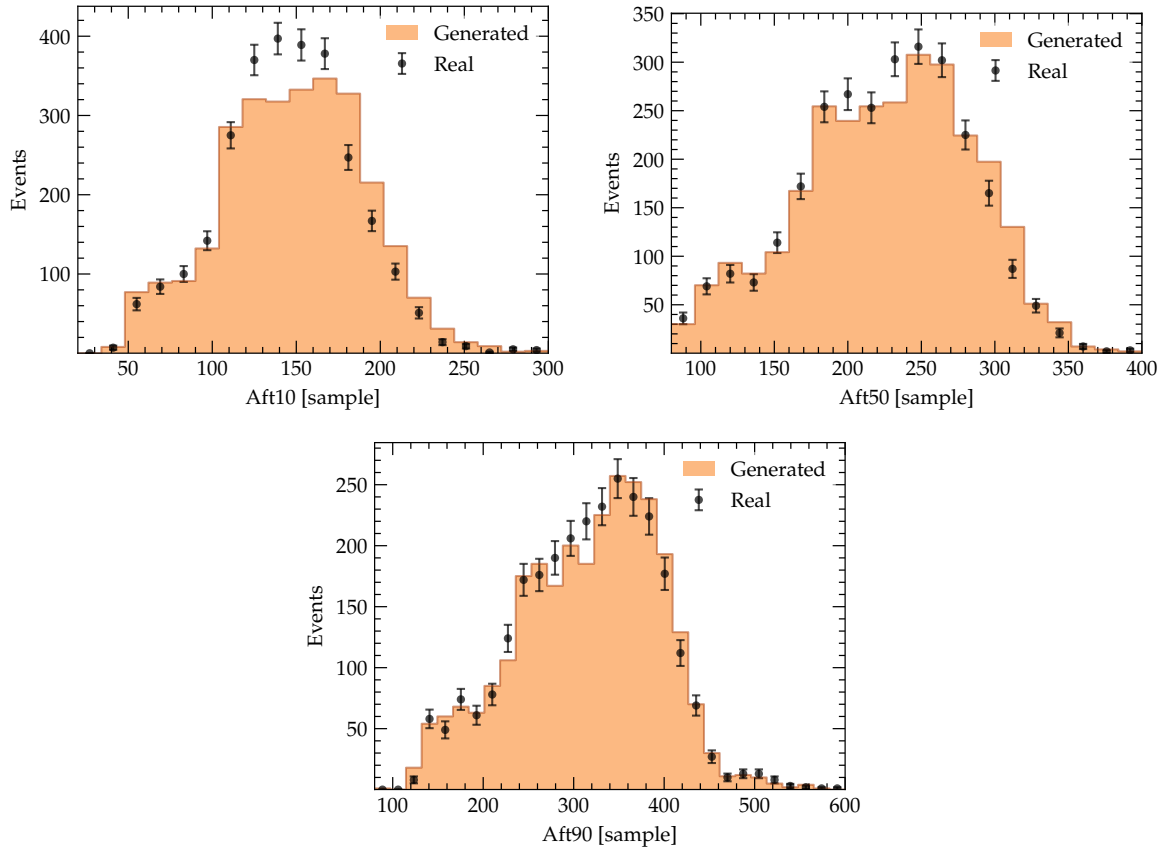

Figure 6.4: Comparison of S2 pulse shape parameters for the true distribution (markers) and the generated distribution (histogram). Aft10, Aft50 and Aft90 is the time at which 10%, 50% and 90% of the pulse area is reached, respectively.

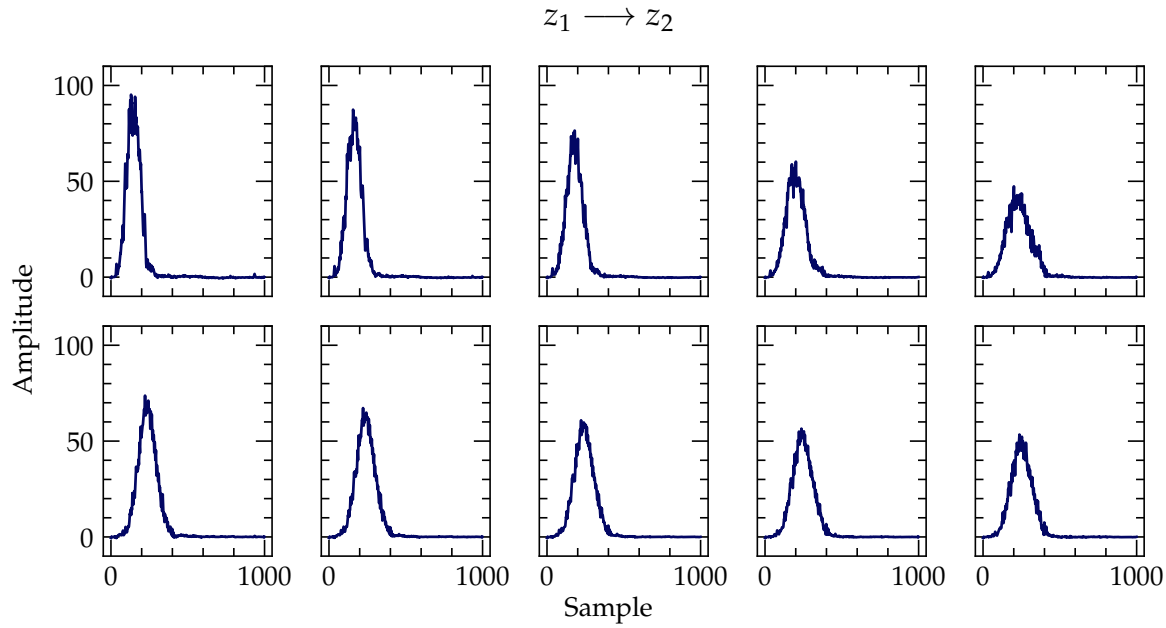

Figure 6.5: Examples of generated S2 pulses obtained by interpolating between two random points (z_1, z_2) in the generator latent space.

be generated in parallel, resulting in significant speed-ups. On a single NVidia P100 GPU, generating a batch of 1000 pulses in parallel takes approximately 5 ms. Assuming the same amount of time to generate the equivalent sized batch of S1 pulses, and including generation of a separate pulse in all 494 PMTs results in a generation time ∼5 ms/event. Comparing to the DER processing time in figure 6.1, this equates to a $\mathcal{O}(10^4)$ decrease in simulation time.

There are many avenues for future work in this area. Perhaps the most important for this technique to be a viable alternative to traditional MC methods is the ability to generate pulses with specific attributes. This may include, for example, a specific pulse area or position in the detector. This can be framed as modelling the conditional distribution $p(x|C)$, where $C = \{c_1, c_2, ..., c_i\}$ is a set of fixed conditions. This can naturally be performed within the GAN setting, with C being just a set of fixed latent variables. Both the conditional GAN (CGAN) [120] and auxiliary classifier GAN (ACGAN) [121] architectures can generate from the conditional distribution. Such a method could be paired with PMT light response functions (see section 3.1.4) to generate individual waveforms for each PMT in the array. An obvious limitation of the GAN architecture considered here is the requirement to generate waveforms of fixed length, which is of course not a realistic property of real S2 pulses. This may be overcome by

utilising a recurrent neural network (RNN) for both the generator and discriminator, which naturally accommodates data with varying lengths.

Chapter 7

Conclusions

The work presented in this thesis has focused on the sensitivity of ^{136}Xe $0\nu\beta\beta$ decay for three generations of dual-phase xenon time projection chambers (TPCs): LUX, LUX-ZEPLIN (LZ) and a future 3rd generation (G3) experiment. Chapter 3 introduced a new reconstruction framework for the LUX experiment. This demonstrated significantly improved reconstruction of $\mathcal{O}$(MeV) events, and the developed techniques resulted in an enhanced energy resolution of $\sim 2\%$ at 2,458 keV, together with a notable reduction in the background rate. Although LUX was designed to search for $\mathcal{O}$(keV) interactions of dark matter particles, this work has demonstrated that dual-phase xenon TPCs are not just limited to searching for dark matter, but are versatile detectors that can simultaneously search for a wide range of processes.

At the time of writing, the LZ experiment is nearing completion of construction and is planned to begin operation in late 2020. Although the primary purpose of the experiment is to search for weakly-interacting dark matter, with $>$ 600kg of ^{136}Xe in the active region, it will also be possible to search for $0\nu\beta\beta$. The background model introduced in chapter 4 shows that the estimate background rate near the $0\nu\beta\beta$ Q-value is 5.6×10^{-7} events/kg/day/keV in a $\sim 1,000$ kg fiducial volume. After the full 1,000 day run, LZ would have the sensitivity to exclude $T_{1/2} < 1.06 \times 10^{26}$ yr at 90% CL. This result puts the LZ experiment on par with dedicated $0\nu\beta\beta$ experiments that use enriched xenon. It is hoped that the techniques demonstrated in chapter 3 will inspire the future search for $0\nu\beta\beta$ in LZ.

In the near future, it is likely that research and development of a future G3 dual-phase xenon experiment will enter the next phase. Chapter 5 has studied the $0\nu\beta\beta$ sensitivity for a hypothetical experiment with $\mathcal{O}$(100 tonnes) of xenon. It is predicted that the scale of such a detector could have an inner 10 tonne fiducial volume with a background rate as

low as 6×10^{-9} events/kg/day/keV, resulting in the ability to exclude $T_{1/2} < 3.2 \times 10^{27}$ yr at 90% CL after 10 years of exposure. With this sensitivity, a G3 experiment could exclude almost all of the inverted hierarchy of neutrino masses. If significant steps are taken to reduce external backgrounds to negligible levels, then such an experiment could also begin to probe the normal hierarchy of neutrino masses, and finally illuminate the nature of the neutrino.

Finally, a brief study was performed to demonstrate the potential of generative adversarial networks (GAN) as a method to generate fast Monte Carlo simulations. Chapter 6 illustrated a GAN architecture that could generate part of the complex detector response of a dual-phase xenon TPC. It was shown that this method can be orders of magnitude faster compared to current simulation techniques, and it is hoped that the developed work will be utilised in future, large scale experiments. Given that the human impact on the climate is becoming ever more apparent, and that we now live in a world where scientific research budgets are tightening, it is paramount to research technology that can reduce both the energy intensive computing requirements and the cost of future high energy physics experiments.

Bibliography

[1] LUX-ZEPLIN (LZ) Collaboration Collaboration, "Projected sensitivity of the LUX-ZEPLIN experiment to the $0\nu\beta\beta$ decay of ^{136}Xe", *Phys. Rev. C* **102** (Jul, 2020) 014602, doi:10.1103/PhysRevC.102.014602.

[2] S. Bilenky, "Neutrino. History of a unique particle", *Eur. Phys. J. H* **38** (2013) 345–404, doi:10.1140/epjh/e2012-20068-9, arXiv:1210.3065.

[3] J. Davis, Raymond, D. S. Harmer, and K. C. Hoffman, "Search for neutrinos from the sun", *Phys. Rev. Lett.* **20** (1968) 1205–1209, doi:10.1103/PhysRevLett.20.1205.

[4] G. Bellini, L. Ludhova, G. Ranucci et al., "Neutrino oscillations", *Adv. High Energy Phys.* **2014** (2014) 191960, doi:10.1155/2014/191960, arXiv:1310.7858.

[5] I. Esteban, M. C. Gonzalez-Garcia, A. Hernandez-Cabezudo et al., "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of 23, CP, and the mass ordering", *Journal of High Energy Physics* **2019** (Jan, 2019) doi:10.1007/jhep01(2019)106.

[6] S. Roy Choudhury and S. Choubey, "Updated Bounds on Sum of Neutrino Masses in Various Cosmological Scenarios", *JCAP* **09** (2018) 017, doi:10.1088/1475-7516/2018/09/017, arXiv:1806.10832.

[7] KATRIN Collaboration, "Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN", *Phys. Rev. Lett.* **123** (2019), no. 22, 221802, doi:10.1103/PhysRevLett.123.221802, arXiv:1909.06048.

[8] R. N. Mohapatra and G. Senjanović, "Neutrino Mass and Spontaneous Parity Nonconservation", *Phys. Rev. Lett.* **44** (Apr, 1980) 912–915, doi:10.1103/PhysRevLett.44.912.

[9] M. Gell-Mann, P. Ramond, and R. Slansky, "Complex Spinors and Unified Theories", *Conf. Proc. C* **790927** (1979) 315–321, arXiv:1306.4669.

[10] T. Yanagida, "Horizontal Symmetry and Masses of Neutrinos", *Progress of Theoretical Physics* **64** (09, 1980) 1103–1105, doi:10.1143/PTP.64.1103.

[11] W. RODEJOHANN, "NEUTRINO-LESS DOUBLE BETA DECAY AND PARTICLE PHYSICS", *International Journal of Modern Physics E* **20** (Sep, 2011) 1833–1930, doi:10.1142/s0218301311020186.

[12] M. Goeppert-Mayer, "Double Beta-Disintegration", *Phys. Rev.* **48** (Sep, 1935) 512–516, doi:10.1103/PhysRev.48.512.

[13] W. H. Furry, "On Transition Probabilities in Double Beta-Disintegration", *Phys. Rev.* **56** (Dec, 1939) 1184–1193, doi:10.1103/PhysRev.56.1184.

[14] M. J. Dolinski, A. W. Poon, and W. Rodejohann, "Neutrinoless Double-Beta Decay: Status and Prospects", *Annual Review of Nuclear and Particle Science* **69** (Oct, 2019) 219–251, doi:10.1146/annurev-nucl-101918-023407.

[15] J. Kotila and F. Iachello, "Phase-space factors for double-β decay", *Phys. Rev. C* **85** (Mar, 2012) 034316, doi:10.1103/PhysRevC.85.034316.

[16] G. Benato, "Effective Majorana mass and neutrinoless double beta decay", *The European Physical Journal C* **75** (Nov, 2015) doi:10.1140/epjc/s10052-015-3802-1.

[17] F. ˇ Simkovic, D. s. ˇ Stefánik, and R. Dvornický, "The λ Mechanism of the $0\nu\beta\beta$-Decay", *Fron. in Phys.* **5** (2017) 57, doi:10.3389/fphy.2017.00057, arXiv:1804.04223.

[18] M. Blennow, E. Fernandez-Martinez, J. Lopez-Pavon et al., "Neutrinoless double beta decay in seesaw models", *Journal of High Energy Physics* **2010** (Jul, 2010) doi:10.1007/jhep07(2010)096.

[19] J. Schechter and J. W. F. Valle, "Neutrinoless double-β decay in SU(2)U(1) theories", *Phys. Rev. D* **25** (Jun, 1982) 2951–2954, doi:10.1103/PhysRevD.25.2951.

[20] J. Engel and J. Menéndez, "Status and future of nuclear matrix elements for neutrinoless double-beta decay: a review", *Reports on Progress in Physics* **80** (Mar, 2017) 046301, doi:10.1088/1361-6633/aa5bc5.

[21] G. Benato, V. D'Andrea, C. Cattadori et al., "Improvement of the GERDA Ge Detectors Energy Resolution by an Optimized Digital Signal Processing", *Physics Procedia* **61** (2015) 673 – 682, doi:https://doi.org/10.1016/j.phpro.2014.12.069. 13th International Conference on Topics in Astroparticle and Underground Physics, TAUP 2013.

[22] S. I. Alvis, I. J. Arnquist, F. T. Avignone et al., "Multisite event discrimination for the majorana demonstrator", *Physical Review C* **99** (Jun, 2019) doi:10.1103/physrevc.99.065501.

[23] "Background-free search for neutrinoless double- decay of 76Ge with GERDA", *Nature* **544** (Apr, 2017) 47–52, doi:10.1038/nature21717.

[24] S. Elliott, N. Abgrall, I. Arnquist et al., "Initial Results from the Majorana Demonstrator", *Journal of Physics: Conference Series* **888** (sep, 2017) 012035, doi:10.1088/1742-6596/888/1/012035.

[25] M. Agostini, A. Bakalyarov, M. Balata et al., "Searching for neutrinoless double beta decay with GERDA", *Journal of Physics: Conference Series* **1342** (jan, 2020) 012005, doi:10.1088/1742-6596/1342/1/012005.

[26] A. Caminata et al., "Results from the Cuore Experiment", *Universe* **5** (2019), no. 1, 10, doi:10.3390/universe5010010.

[27] R. Arnold, C. Augier, J. Baker et al., "First Results of the Search for Neutrinoless Double-Beta Decay with the NEMO 3 Detector", *Physical Review Letters* **95** (Oct, 2005) doi:10.1103/physrevlett.95.182302.

[28] R. Arnold, C. Augier, J. Baker et al., "Probing new physics models of neutrinoless double beta decay with SuperNEMO", *The European Physical Journal C* **70** (Nov, 2010) 927–943, doi:10.1140/epjc/s10052-010-1481-5.

[29] KamLAND-Zen Collaboration, "Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen", *Phys. Rev. Lett.* **117** (2016), no. 8, 082503, doi:10.1103/PhysRevLett.117.082503, arXiv:1605.02889. [Addendum: Phys.Rev.Lett. 117, 109903 (2016)].

[30] V. Lozza, "The SNO+ Experiment for Neutrinoless Double-Beta Decay", *Nuclear and Particle Physics Proceedings* **273-275** (2016) 1836 – 1841, doi:https://doi.org/10.1016/j.nuclphysbps.2015.09.296. 37th International Conference on High Energy Physics (ICHEP).

[31] J. Albert, G. Anton, I. Badhrees et al., "Search for Neutrinoless Double-Beta Decay with the Upgraded EXO-200 Detector", *Physical Review Letters* **120** (Feb, 2018) doi:10.1103/physrevlett.120.072701.

[32] J. B. Albert, G. Anton, I. J. Arnquist et al., "Sensitivity and discovery potential of the proposed nEXO experiment to neutrinoless double- decay", *Physical Review C* **97** (Jun, 2018) doi:10.1103/physrevc.97.065503.

[33] J. Gomez-Cadenas, "The NEXT experiment", *Nuclear and Particle Physics Proceedings* **273-275** (2016) 1732 – 1739, doi:https://doi.org/10.1016/j.nuclphysbps.2015.09.279. 37th International Conference on High Energy Physics (ICHEP).

[34] K. H. and, "PandaX-III: Searching for Neutrinoless Double Beta Decay with High Pressure Gaseous Time Projection Chambers", *Journal of Physics: Conference Series* **1342** (jan, 2020) 012095, doi:10.1088/1742-6596/1342/1/012095.

[35] NEXT Collaboration, "First proof of topological signature in the high pressure xenon gas TPC with electroluminescence amplification for the NEXT experiment", *JHEP* **01** (2016) 104, doi:10.1007/JHEP01(2016)104, arXiv:1507.05902.

[36] M. Schumann, "Dual-phase liquid xenon detectors for dark matter searches", *Journal of Instrumentation* **9** (Aug, 2014) C08004–C08004, doi:10.1088/1748-0221/9/08/c08004.

[37] XENON Collaboration, "Energy resolution and linearity in the keV to MeV range measured in XENON1T", arXiv:2003.03825.

[38] LUX Collaboration, "Results from a Search for Dark Matter in the Complete LUX Exposure", *Physical Review Letters* **118** (Jan, 2017) doi:10.1103/physrevlett.118.021303, arXiv:1608.07648.

[39] XENON Collaboration, "Dark Matter Search Results from a One Ton-Year Exposure of XENON1T", *Phys. Rev. Lett.* **121** (2018), no. 11, 111302, doi:10.1103/PhysRevLett.121.111302, arXiv:1805.12562.

[40] PandaX-II Collaboration, "Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment", *Phys. Rev. Lett.* **119** (2017), no. 18, 181302, doi:10.1103/PhysRevLett.119.181302, arXiv:1708.06917.

[41] XENON Collaboration, "Light Dark Matter Search with Ionization Signals in

XENON1T", *Phys. Rev. Lett.* **123** (2019), no. 25, 251801, doi:10.1103/PhysRevLett.123.251801, arXiv:1907.11485.

[42] LUX Collaboration, "Position Reconstruction in LUX", *JINST* **13** (2018), no. 02, P02001, doi:10.1088/1748-0221/13/02/P02001, arXiv:1710.02752.

[43] LUX Collaboration, "The Large Underground Xenon (LUX) Experiment", *Nucl. Instrum. Meth.* **A704** (2013) 111–126, doi:10.1016/j.nima.2012.11.135, arXiv:1211.3788.

[44] LUX Collaboration, "Calibration, event reconstruction, data analysis, and limit calculation for the LUX dark matter experiment", *Phys. Rev.* **D97** (2018), no. 10, 102008, doi:10.1103/PhysRevD.97.102008, arXiv:1712.05696.

[45] D. S. Akerib et al., "Data Acquisition and Readout System for the LUX Dark Matter Experiment", *Nucl. Instrum. Meth.* **A668** (2012) 1–8, doi:10.1016/j.nima.2011.11.063, arXiv:1108.1836.

[46] LUX Collaboration, "FPGA-based Trigger System for the LUX Dark Matter Experiment", *Nucl. Instrum. Meth.* **A818** (2016) 57–67, doi:10.1016/j.nima.2016.02.017, arXiv:1511.03541.

[47] B. J. Mount et al., "LUX-ZEPLIN (LZ) Technical Design Report", arXiv:1703.09144.

[48] LZ Collaboration, "Identification of Radiopure Titanium for the LZ Dark Matter Experiment and Future Rare Event Searches", *Astropart. Phys.* **96** (2017) 1–10, doi:10.1016/j.astropartphys.2017.09.002, arXiv:1702.02646.

[49] B. López Paredes, H. M. Araújo, F. Froborg et al., "Response of photomultiplier tubes to xenon scintillation light", *Astropart. Phys.* **102** (2018) 56–66, doi:10.1016/j.astropartphys.2018.04.006, arXiv:1801.01597.

[50] LZ Collaboration, "The Data Acquisition System for LZ", *JINST* **11** (2016), no. 02, C02072, doi:10.1088/1748-0221/11/02/C02072, arXiv:1511.08385.

[51] A. Gando, Y. Gando, T. Hachiya et al., "Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen", *Physical Review Letters* **117** (Aug, 2016) doi:10.1103/physrevlett.117.082503, arXiv:1605.02889.

[52] T. Doke, A. Hitachi, J. Kikuchi et al., "Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles", *Jap. J. Appl. Phys.* **41** (2002)

1538–1545, doi:10.1143/JJAP.41.1538.

[53] T. Takahashi, S. Konno, T. Hamada et al., "Average energy expended per ion pair in liquid xenon", *Phys. Rev.* **A12** (1975) 1771–1775, doi:10.1103/PhysRevA.12.1771.

[54] EXO-200 Collaboration, "An Optimal Energy Estimator to Reduce Correlated Noise for the EXO-200 Light Readout", *JINST* **11** (2016), no. 07, P07015, doi:10.1088/1748-0221/11/07/P07015, arXiv:1605.06552.

[55] NEXT Collaboration, "The NEXT high pressure xenon gas TPC for neutrinoless double beta decay searches", *JINST* **13** (2018), no. 01, C01048, doi:10.1088/1748-0221/13/01/C01048.

[56] M. Szydagis, J. Balajthy, J. Brodsky et al., "Noble Element Simulation Technique v2.0", July, 2018. doi:10.5281/zenodo.1314669.

[57] A. Dobi, D. S. Leonard, C. Hall et al., "Study of a zirconium getter for purification of xenon gas", *Nucl. Instrum. Meth.* **A620** (2010) 594–598, doi:10.1016/j.nima.2010.03.151, arXiv:1002.2791.

[58] M. Berglund and M. Wieser, "Isotopic compositions of the elements 2009", *Pure and Applied Chemistry* **83** (01, 2011) 397–410, doi:10.1351/PAC-REP-10-06-02.

[59] A. Hitachi, "Properties of liquid xenon scintillation for dark matter searches", *Astropart. Phys.* **24** (2005) 247–256, doi:10.1016/j.astropartphys.2005.07.002.

[60] LUX Collaboration, "Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector", *Phys. Rev. D* **97** (Jun, 2018) 112002, doi:10.1103/PhysRevD.97.112002, arXiv:1802.06162.

[61] M. Berger, J. Hubbell, S. Seltzer et al., "XCOM: Photon Cross Section Database (version 1.5)". http://physics.nist.gov/xcom, 2010.

[62] M. Berger, J. Coursey, M. Zucker et al., "ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions (version 1.2.3)". http://physics.nist.gov/Star, 2005.

[63] R. Platzman, "Total ionization in gases by high-energy particles: An appraisal of our understanding", *The International Journal of Applied Radiation and Isotopes* **10** (1961), no. 2, 116 – 127,

doi:https://doi.org/10.1016/0020-708X(61)90108-9.

[64] C. E. Dahl, "The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter". PhD thesis, Princeton University, 2009.

[65] P. Sorensen and C. E. Dahl, "Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter", *Phys. Rev.* **D83** (2011) 063501, doi:10.1103/PhysRevD.83.063501, arXiv:1101.6080.

[66] A. Hitachi, T. Doke, and A. Mozumder, "Luminescence quenching in liquid argon under charged-particle impact: Relative scintillation yield at different linear energy transfers", *Phys. Rev.* **B46** (1992), no. 18, 11463, doi:10.1103/PhysRevB.46.11463.

[67] T. Doke, K. Masuda, and E. Shibamura, "Estimation of absolute photon yields in liquid argon and xenon for relativistic (1 MeV) electrons", *Nucl. Instrum. Meth.* **A291** (1990) 617–620, doi:10.1016/0168-9002(90)90011-T.

[68] A. Hitachi, T. Takahashi, N. Funayama et al., "Effect of ionization density on the time dependence of luminescence from liquid argon and xenon", *Phys. Rev.* **B27** (1983) 5279–5285, doi:10.1103/PhysRevB.27.5279.

[69] S. Kubota, M. Hishida, M. Suzuki et al., "Dynamical behavior of free electrons in the recombination process in liquid argon, krypton, and xenon", *Phys. Rev.* **B20** (1979), no. 8, 3486, doi:10.1103/PhysRevB.20.3486.

[70] E. Hogenbirk, J. Aalbers, P. A. Breur et al., "Precision measurements of the scintillation pulse shape for low-energy recoils in liquid xenon", *JINST* **13** (2018), no. 05, P05016, doi:10.1088/1748-0221/13/05/P05016, arXiv:1803.07935.

[71] V. Chepel and H. Araujo, "Liquid noble gas detectors for low energy particle physics", *JINST* **8** (2013) R04001, doi:10.1088/1748-0221/8/04/R04001, arXiv:1207.2292.

[72] EXO-200 Collaboration, "Measurement of the Drift Velocity and Transverse Diffusion of Electrons in Liquid Xenon with the EXO-200 Detector", *Phys. Rev.* **C95** (2017), no. 2, 025502, doi:10.1103/PhysRevC.95.025502, arXiv:1609.04467.

[73] 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, doi:10.1016/j.nima.2011.01.089, arXiv:1102.2865.

[74] P. Sorensen and K. Kamdin, "Two distinct components of the delayed single electron noise in liquid xenon emission detectors", *JINST* **13** (2018), no. 02, P02032, doi:10.1088/1748-0221/13/02/P02032, arXiv:1711.07025.

[75] P. Sorensen, "Electron train backgrounds in liquid xenon dark matter search detectors are indeed due to thermalization and trapping", arXiv:1702.04805.

[76] C. H. Faham, V. M. Gehman, A. Currie et al., "Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes", *JINST* **10** (2015), no. 09, P09010, doi:10.1088/1748-0221/10/09/P09010, arXiv:1506.08748.

[77] J. Pivarski, C. Escott, M. Hedges et al., "scikit-hep/awkward-array: 0.12.0rc1", July, 2019. doi:10.5281/zenodo.3275017.

[78] J. Reback, W. McKinney, jbrockmendel et al., "pandas-dev/pandas: Pandas 1.0.0", January, 2020. doi:10.5281/zenodo.3630805.

[79] A. Savitzky and M. J. E. Golay, "Smoothing and Differentiation of Data by Simplified Least Squares Procedures.", *Analytical Chemistry* **36** (1964), no. 8, 1627–1639, doi:10.1021/ac60214a047.

[80] J. Mock, "A Search for Dark Matter with the Large Underground Xenon Detector ". PhD thesis, Case Western Reserve University, 2008.

[81] V. N. Solovov, V. A. Belov, D. Y. Akimov et al., "Position Reconstruction in a Dual Phase Xenon Scintillation Detector", *IEEE Transactions on Nuclear Science* **59** (Dec, 2012) 3286–3293, doi:10.1109/tns.2012.2221742.

[82] D. Akerib, S. Alsum, H. Araújo et al., "Signal yields, energy resolution, and recombination fluctuations in liquid xenon", *Physical Review D* **95** (Jan, 2017) doi:10.1103/physrevd.95.012008.

[83] D. Akerib, X. Bai, S. Bedikian et al., "LUXSim: A component-centric approach to low-background simulations", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **675** (May, 2012) 63–77, doi:10.1016/j.nima.2012.02.010.

[84] D. Akerib, C. Akerlof, S. Alsum et al., "Projected WIMP sensitivity of the

LUX-ZEPLIN dark matter experiment", *Physical Review D* **101** (Mar, 2020) doi:10.1103/physrevd.101.052002.

[85] EXO-200 Collaboration, "Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment", arXiv:1908.04128.

[86] D. S. Akerib, C. W. Akerlof, A. Alquahtani et al., "Simulations of Events for the LUX-ZEPLIN (LZ) Dark Matter Experiment", arXiv:2001.09363.

[87] S. Agostinelli, J. Allison, K. Amako et al., "Geant4—a simulation toolkit", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506** (2003), no. 3, 250 – 303, doi:https://doi.org/10.1016/S0168-9002(03)01368-8.

[88] A. Bailey, "Dark matter searches and study of electrode design in LUX and LZ ". PhD thesis, Imperial College London, 2016.

[89] M.-M. Bé, V. Chisté, C. Dulieu et al., "Table of Radionuclides", volume 8 of *Monographie BIPM-5*. Bureau International des Poids et Mesures, 2016.

[90] LUX-ZEPLIN (LZ) Collaboration, "Measurement of the Gamma Ray Background in the Davis Cavern at the Sanford Underground Research Facility", *Astropart. Phys.* **116** (2020) 102391, doi:10.1016/j.astropartphys.2019.102391, arXiv:1904.02112.

[91] J. B. Albert, D. J. Auty, P. S. Barbeau et al., "Measurements of the ion fraction and mobility ofand-decay products in liquid xenon using the EXO-200 detector", *Physical Review C* **92** (Oct, 2015) doi:10.1103/physrevc.92.045504.

[92] J. B. Albert, M. Auger, D. J. Auty et al., "Improved measurement of the2half-life of136Xe with the EXO-200 detector", *Physical Review C* **89** (Jan, 2014) doi:10.1103/physrevc.89.015502.

[93] J. Kotila and F. Iachello, "Phase-space factors for double-decay", *Physical Review C* **85** (Mar, 2012) doi:10.1103/physrevc.85.034316.

[94] J. Albert, D. Auty, P. Barbeau et al., "Cosmogenic backgrounds to 0 in EXO-200", *Journal of Cosmology and Astroparticle Physics* **2016** (Apr, 2016) 029–029, doi:10.1088/1475-7516/2016/04/029.

[95] A. Best, J. Görres, M. Junker et al., "Low energy neutron background in deep

underground laboratories", *Nucl. Instrum. Meth. A* **812** (2016) 1–6, doi:10.1016/j.nima.2015.12.034, arXiv:1509.00770.

[96] J. N. Bahcall, "Neutrino - Electron Scattering and Solar Neutrino Experiments", *Rev. Mod. Phys.* **59** (1987) 505, doi:10.1103/RevModPhys.59.505.

[97] J. N. Bahcall, E. Lisi, D. E. Alburger et al., "Standard neutrino spectrum fromB8decay", *Physical Review C* **54** (Jul, 1996) 411–422, doi:10.1103/physrevc.54.411.

[98] B. Aharmim, S. Ahmed, A. Anthony et al., "Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory", *Physical Review C* **88** (Aug, 2013) doi:10.1103/physrevc.88.025501.

[99] O. Ponkratenko, V. Tretyak, and Y. Zdesenko, "The Event generator DECAY4 for simulation of double beta processes and decay of radioactive nuclei", *Phys. Atom. Nucl.* **63** (2000) 1282–1287, doi:10.1134/1.855784, arXiv:nucl-ex/0104018.

[100] G. Cowan, K. Cranmer, E. Gross et al., "Asymptotic formulae for likelihood-based tests of new physics", *The European Physical Journal C* **71** (Feb, 2011) doi:10.1140/epjc/s10052-011-1554-0.

[101] 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. C* **87** (Jun, 2013) 064302, doi:10.1103/PhysRevC.87.064302.

[102] N. L. Vaquero, T. R. Rodríguez, and J. L. Egido, "Shape and Pairing Fluctuation Effects on Neutrinoless Double Beta Decay Nuclear Matrix Elements", *Phys. Rev. Lett.* **111** (Sep, 2013) 142501, doi:10.1103/PhysRevLett.111.142501.

[103] DARWIN Collaboration, "DARWIN: towards the ultimate dark matter detector", *JCAP* **11** (2016) 017, doi:10.1088/1475-7516/2016/11/017, arXiv:1606.07001.

[104] F. Agostini, S. E. M. A. Maouloud, L. Althueser et al., "Sensitivity of the DARWIN observatory to the neutrinoless double beta decay of ^{136}Xe", arXiv:2003.13407.

[105] A. Jamil, T. Ziegler, P. Hufschmidt et al., "VUV-Sensitive Silicon Photomultipliers for Xenon Scintillation Light Detection in nEXO", *IEEE Transactions on Nuclear*

Science **65** (Nov, 2018) 2823–2833, doi:10.1109/tns.2018.2875668.

[106] L. Baudis, M. Galloway, A. Kish et al., "Characterisation of Silicon Photomultipliers for liquid xenon detectors", *Journal of Instrumentation* **13** (Oct, 2018) P10022–P10022, doi:10.1088/1748-0221/13/10/p10022.

[107] M. Robinson, V. Kudryavtsev, R. Luscher et al., "Measurements of muon flux at 1070 meters vertical depth in the Boulby underground laboratory", *Nucl. Instrum. Meth. A* **511** (2003) 347–353, doi:10.1016/S0168-9002(03)01973-9, arXiv:hep-ex/0306014.

[108] V. Daniel Elvira, "Impact of detector simulation in particle physics collider experiments", *Physics Reports* **695** (Jun, 2017) 1–54, doi:10.1016/j.physrep.2017.06.002.

[109] J. Albrecht, A. A. Alves, G. Amadio et al., "A Roadmap for HEP Software and Computing RD for the 2020s", *Computing and Software for Big Science* **3** (Mar, 2019) doi:10.1007/s41781-018-0018-8.

[110] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza et al., "Generative Adversarial Networks", arXiv:1406.2661.

[111] M. Paganini, L. de Oliveira, and B. Nachman, "CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks", *Physical Review D* **97** (Jan, 2018) doi:10.1103/physrevd.97.014021.

[112] M. Erdmann, J. Glombitza, and T. Quast, "Precise Simulation of Electromagnetic Calorimeter Showers Using a Wasserstein Generative Adversarial Network", *Computing and Software for Big Science* **3** (Jan, 2019) doi:10.1007/s41781-018-0019-7.

[113] C. Ahdida, R. Albanese, A. Alexandrov et al., "Fast simulation of muons produced at the SHiP experiment using Generative Adversarial Networks", *Journal of Instrumentation* **14** (Nov, 2019) P11028–P11028, doi:10.1088/1748-0221/14/11/p11028.

[114] T. Salimans, I. Goodfellow, W. Zaremba et al., "Improved Techniques for Training GANs", arXiv:1606.03498.

[115] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN",

arXiv:1701.07875.

[116] I. Gulrajani, F. Ahmed, M. Arjovsky et al., "Improved Training of Wasserstein GANs", arXiv:1704.00028.

[117] A. Radford, L. Metz, and S. Chintala, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks", arXiv:1511.06434.

[118] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization", arXiv:1412.6980.

[119] A. Paszke, S. Gross, F. Massa et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library", arXiv:1912.01703.

[120] M. Mirza and S. Osindero, "Conditional Generative Adversarial Nets", arXiv:1411.1784.

[121] A. Odena, C. Olah, and J. Shlens, "Conditional Image Synthesis With Auxiliary Classifier GANs", arXiv:1610.09585.