BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

ENHANCING THE NEUTRINOLESS DOUBLE BETA DECAY SENSITIVITY OF SNO+ WITH DEEP LEARNING

by

DIANA M. GOODING

B.A., Transylvania University, 2014
M.S., Tufts University, 2016

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2025

Approved by

First Reader

Christopher P. Grant, PhD
Associate Professor of Physics

Second Reader

Zeynep Demiragli, PhD
Associate Professor of Physics

נו
גיי שוין לייענען

Nu?
Go read already!

Acknowledgments

Every few weeks since the start of my graduate career, I have promised different coworkers, friends, and family their own chapters of acknowledgments in my dissertation. Today, I am tasked with condensing several Webster-sized volumes of gratitude and memory into a single page.

First, the topic of this dissertation would not exist without the SNO+ collaboration, or without the immense brilliance and efforts of Aobo Li, who invented the neural network at the heart of the following analysis. I also want to extend my warmest thanks to So Young Jeon, Hasung Song, and Ömer Penek, for the long days, late nights, code snippets, and friendship that have sustained me throughout my last year of school.

Although the entire SNO+ collaboration is a wealth of knowledge and community, I find it very necessary to emphatically thank Sam Naugle, Ben Land, James Shen, Daniel Cookman, and Teal Pershing for teaching me everything they know about analysis and computing. Their patience and swift attention to my questions (or unawareness of Slacks mute button) were crucial to establishing my analysis framework and independence as a graduate student.

I also thank my closest friends and family—my parents Steve and Pam, my sister Nicole, my identical twin Barbara, my husband Charles, my in-laws and dear friends Phyllis and Earl, and my oldest friends, Alex, Jon, John, Aimee, Zak, and Matteo—for their concern, kindness, and high tolerance for repetitive conversations about school-related stress.

And Finally—Everything I have achieved during my doctoral program required teachers. I owe a great debt of gratitude to Prof. Chris Grant for his endless patience in seeing me through these seven years. I thank Prof. Lindley Winslow for graciously giving me my start in particle physics, and my undergraduate advisor Dr. George Kaufman for giving me my start in research.

ENHANCING THE NEUTRINOLESS DOUBLE BETA DECAY SENSITIVITY OF SNO+ WITH DEEP LEARNING

DIANA M. GOODING

Boston University, Graduate School of Arts and Sciences, 2025

Major Professor: Christopher P. Grant, PhD
Professor of Physics

ABSTRACT

Measurements of neutrino oscillations have confirmed that neutrinos are massive, but other questions remain: What are the absolute neutrino masses? Which neutrino is the lightest? Answers to these questions don't only provide all of the remaining unknown quantities in The Standard Model. They also have consequences for our understanding of how fundamental particles acquire mass. One way to investigate neutrino mass is to determine whether or not the neutrino is its own anti-particle, or Majorana in nature. This can be done experimentally by measuring neutrinoless double-beta decay. Once the SNO+ experiment is loaded with natural tellurium, containing 34% ^{130}Te, it will begin to look for this hypothetical process. This dissertation describes a log-likelihood approach to quantifying current background levels in the SNO+ detector, as well as a toy Monte-Carlo study to calculate the experiment's projected sensitivity to neutrinoless double beta decay after 5.25×10^{-2} ton-years of exposure. This work also describes the additional background rejection achieved by KamNet, A novel deep neural network developed by the KamLAND collaboration. At 90% confidence, the deployment of this tool increases the projected neutrinoless double beta decay half-life limit of SNO+ by 98.5%, from $T^{0\nu}_{1/2} > 3.41 \times 10^{25}$ years to $T^{0\nu}_{1/2} > 6.67 \times 10^{25}$ years. Additionally, Asimov datasets were used to project the decay half-life limit for different tellurium exposures. At a benchmark of three years of livetime, KamNet is projected to achieve a limit of $T^{0\nu}_{1/2} > 1.63 \times 10^{26}$ years. This limit represents a savings of two years for a traditional analysis of the same tellurium loading. And, although 1.5% tellurium loading is projected to achieve this sensitivity in just under two years, it does so at an extra cost of roughly $2.6M.

Contents

List of Tables

List of Figures

List of Abbreviations

$0\nu\beta\beta$	Neutrinoless Double Beta Decay
$2\nu\beta\beta$	Two Neutrino Double Beta Decay
CNN	Convolutional Neural Network
DAQ	Data Acquisition
IBD	Inverse β Decay
MC	Monte Carlo simulation
PMT	Photomultiplier Tube
QE	Quantum Efficiency
SM	Standard Model
SOFT	SO(3) Fourier Transform
TOF	Time of Flight

Chapter 1

Introduction

One of the most profound and unresolved questions in physics is how our universe came to exist in its present form, particularly the origin of its matter content. According to our current understanding, the universe began with the Big Bang, an event that marked the start of an extremely hot and dense state where energy was continuously transformed into particles of matter and antimatter, and vice versa. In this primordial soup, the laws of physics, as we know them, dictate that matter and antimatter should have been created in perfectly equal amounts. However, when we observe the universe today, we find it overwhelmingly dominated by matter, with antimatter conspicuously absent. This asymmetry poses a significant puzzle: what caused the imbalance that allowed matter to prevail?

This observation challenges the long-held assumption in particle physics that matter and antimatter are always produced symmetrically in particle interactions. Processes that break this symmetry, known as charge-parity (CP) violation, are key to understanding the matter-antimatter imbalance. Although CP violation has been observed in certain interactions involving quarks, the degree of violation in the Standard Model is far too small to account for the universe's pronounced asymmetry. This discrepancy strongly suggests the existence of new physical processes or particles that go beyond the Standard Model and amplify CP violation to the levels required to explain the matter-dominated universe.

An essential candidate for such mechanisms involves leptogenesis, a theoretical framework tied to neutrinos. Leptogenesis posits that the asymmetry originated from CP-violating decays of heavy right-handed neutrinos in the early universe. These decays would have created an excess of leptons over antileptons, which was then converted into a baryon asymmetry (i.e., an excess of matter over antimatter) through processes governed by the Standard Model. For this scenario to work, neutrinos must be Majorana particles (their own antiparticles), a hypothesis that can be directly tested by searching for neutrinoless double beta decay ($0\nu\beta\beta$). Observing this hypothetical nuclear decay would provide crucial evidence for the Majorana nature of neutrinos and bolster the case for leptogenesis as the mechanism behind the matter-antimatter imbalance.

The search for this decay starts 6,800 ft underground inside the Vale Creighton Nickle Mine in Sudbury, Ontario. This mine is home to SNOLAB, the world's cleanest and deepest underground laboratory. Within SNOLAB, the SNO+ experiment has been constructing a world-class neutrinoless double beta decay detector with ^{130}Te as the target isotope. SNO+ was recently filled with approximately 780 metric tons of ultra-pure liquid scintillator and plans to initially load 1.3 tons of ^{130}Te (in the form of 3.9 tons natural tellurium) into the liquid scintillator in late 2025. My thesis focuses on our recent data taken with pure liquid scintillator and provides one of the first comprehensive analyses of the scintillator backgrounds in the $0\nu\beta\beta$ region-of-interest of SNO+. I also explore the application of a new deep learning algorithm on the SNO+ scintillator data to improve the rejection of backgrounds and, in turn, boost SNO+ sensitivity to $0\nu\beta\beta$. To get us started, I briefly describe the fundamental physics of neutrinos and double beta decay in Chapters 2 and 3. Chapter 4 gives an overview of the SNO+ experiment and its various subsystems. Chapter 5 describes the makeup of the SNO+ liquid scintillator, its various chemical properties, and the planned tellurium loading. Chapter 6 describes how physics events are modeled and reconstructed in SNO+ and Chapter 7 describes the backgrounds for the neutrinoless double beta decay search.

The culmination of my thesis work on SNO+, the application of deep learning classification on the data and a comprehensive analysis of the $0\nu\beta\beta$ backgrounds, is presented in Chapters 8 and 9. Chapter 10 summarizes my $0\nu\beta\beta$ sensitivity predictions for SNO+ as the experiment approaches tellurium deployment.

Chapter 2

Neutrino theory

Eighty-five years before neutrinos became the subject of the 2015 Nobel Prize in Physics, they were one of three competing theories behind β decay. After Becquerel's discovery of radioactivity in 1896, Rutherford and Villard identified the decay products of uranium: α, β, and γ. In 1914, Chadwick discovered that, unlike uranium's discrete α and γ spectra, its β spectrum was continuous. Instead of always emerging with an energy $E_\beta = m_M - m_D = Q$ (for masses m_M and m_D of the mother and daughter nuclei), most of the β's appeared to emerge with less. Where was the missing energy?

One possibility, proposed by Meitner, suggested that the β's were initially emitted with $E_\beta = Q$. Then, before reaching the detector, they underwent secondary interactions that depleted their kinetic energy, like collisions with other atomic electrons, or the emission of an accompanying γ. Another possibility proposed by Bohr was that the electrons were emitted at exactly these deficient energies, up to and including Q. In this picture, energy is only statistically conserved.

Chadwick's original experiment used a Geiger counter and magnetic spectrometer to measure the energies of β's from ^{210}Bi, which has a Q-value of 1161.0 keV. To test Meitner's hypothesis, Chadwick and Charles Ellis independently performed calorimetry experiments on ^{210}Bi. A calorimeter measures the energy of an entire process, start to finish, including any secondary processes. If Meitner's hypothesis was correct, these experiments would always measure 1161.0 keV.

Ellis measured 344 ± 34.4 keV, and Meitner 337 ± 20.2 keV. These results seemed to support Bohr's theory of broken energy conservation, but there was yet one more problem neither hypothesis addressed: spin statistics. Parent and daughter nuclei in beta decay differ by one unit of spin, meaning there has to be one additional spin-1/2 particle accompanying the electron in the final state. To resolve both the missing energy and spin, Pauli did "something no theorist should ever do", and proposed a new, electrically neutral, and undetectable spin-1/2 particle to carry away the missing energy. Pauli's hypothesis led Fermi to develop a theory of β decay that successfully predicted the shape of ^{64}Cu's β spectrum – the first of several experimental results supporting the neutrino's existence.

Today, neutrinos are a well-established component of the Standard Model of Particle Physics, with known masses (albeit very small), three flavors, and a rich phenomenology that includes oscillations and potential connections to physics beyond the Standard Model. This chapter provides a theoretical overview of neutrinos and their properties within the framework of the Standard Model.

2.1 Neutrinos in the Standard Model

The Standard Model of Particle Physics is a foundational theory that describes all known fundamental particles and their interactions. In this framework, particles fall into two categories: spin-1/2 fermions (matter particles) and spin-1 bosons (force carriers). Fermions consist of six quarks and six leptons. Leptons are made up of three charged particles: e^-, μ^-, and τ^-. Each charged lepton comes with a corresponding neutrino: ν_e, ν_μ, and ν_τ. Measurements of Z^0 decays have given clear proof that these are the only active neutrino flavors [1]. All charged fermions have a corresponding antiparticle with opposite electric charge. Neutrinos and their corresponding antiparticles are electrically neutral.

Particles interact with each other according to a gauge theory that follows the symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$. Quantum chromodynamics governs the quark sector, and is described by the $SU(3)_C$ (*color*) symmetry group. The eight generators of this group correspond to the eight massless gluons. The $SU(2)_L \times U(1)_Y$ group describes electroweak interactions, and has four generators corresponding to four bosons, the massive $W^{\pm}$ and Z, and the massless γ. We note that leptons have no color charge, meaning that they do not interact with gluons. Hence, we only need to consider the $SU(2)_L \times U(1)_Y$ group when describing lepton interactions.

The symmetry group $SU(2)_L$ is referred to as *weak isospin* and has three corresponding operators $\boldsymbol{I} = (I_1, I_2, I_3)$. The symmetry group $U(1)_Y$ is referred to as *weak hypercharge* and has one corresponding operator Y_W. The Gell-Mann-Nishijima equation relates the third component of weak isospin I_3 and the weak hypercharge Y_W to the electric charge Q as follows:

$$Q = I_3 + \frac{Y_W}{2}. \tag{2.1}$$

To obtain the correct phenomenology for weak interactions we choose a representation where the left-handed chiral components of the fermion fields are grouped into weak isospin *doublets*:

$$L_e = \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix} \quad L_\mu = \begin{pmatrix} \nu_{\mu L} \\ \mu_L \end{pmatrix} \quad L_\tau = \begin{pmatrix} \nu_{\tau L} \\ \tau_L \end{pmatrix} \tag{2.2}$$

This choice fixes the generators of the $SU(2)_L$ to $I_a = \tau_a/2$, where the τ_a ($a = 1,2,3$) are the three Pauli matrices. The weak isospin and hypercharge operators act on the weak isospin doublets in the following way:

$$\boldsymbol{I} L_L = \frac{\boldsymbol{\tau}}{2} L_L \qquad \text{and} \qquad Y_W L_L = -L_L \tag{2.3}$$

Within this isospin space all leptons have total weak isospin $I = 1/2$, the charged leptons have third component of weak isospin $I_3 = +1/2$, and left-handed neutrinos have $I_3 = -1/2$. From Equation 2.3 we see that weak hypercharge $Y_W = -1$ for all left-handed leptons, such that when these eigenvalues are plugged into Equation 2.1 we end up with electrically neutral neutrinos and charged leptons with $Q = -1$, as required. To accommodate the observation of parity violation (discussed in Section 2.2) right-handed fermions are assigned *singlet* representations in $SU(2)_L \times U(1)_Y$, resulting in $\boldsymbol{I} = 0$ and $I_3 = 0$. This means that right-handed charged leptons must have $Y_W = -2$ to obtain the correct charge of $Q = -1$. On the other hand, electrically neutral right-handed neutrinos must have $Y_W = 0$ meaning that they do not participate in the weak interaction. The existence of right-handed neutrinos is currently an open question in particle physics.

2.2 The Weak Interaction

In a first attempt to describe nuclear β decay, the weak interaction was originally formulated as four fermions (neutron, proton, electron, neutrino) at a single vertex by Enrico Fermi in 1933 [2]. To explain earlier experimental observations of a continuous spectrum of electron energies, Fermis theory allowed the neutrino carried any remaining energy that was not carried by the emitted electron.

This initial formulation of weak interactions was followed by a landmark discovery in 1957 where Chien-Shiung Wu and collaborators first observed parity violation in the β-decay of polarized ^{60}Co, which proceeds as:

$$^{60}\mathrm{Co} \rightarrow ^{60}\mathrm{Ni}^{*} + \mathrm{e}^{-} + \overline{\nu}_{\mathrm{e}} \tag{2.4}$$

Cooled near absolute zero, the nuclei align with spin parallel to the magnetic field, and remain almost thermally still. Were parity conserved, β^- particles would emanate equally along both ends of the spin axis. Instead, Wu and collaborators measured an excess emission against the direction of the nuclear spin, indicating that β^- emission by polarized ^{60}Co is not equal to that of its mirror image.

Parity violation is a key component of the unified framework of electroweak interactions developed by Sheldon Glashow, Abdus Salam, and Steven Weinberg [3], which successfully describes a large collection of experimental results. We now know that weak interactions are mediated by three massive vector bosons, the W^+, W^-, and Z^0.

Weak interactions involving the W^+ or W^- bosons are referred to as *charged current* (CC) interactions and involve an exchange of electric charge between initial and final state leptons. To accommodate parity violation, the charged current Lagrangian is constructed as a mixture of vector (**V**) and axial vector (**A**) currents of the form (**V** – **A**) in the following manner:

$$\mathscr{L}^{CC} = -\frac{g}{2\sqrt{2}}\overline{\nu_\alpha}\gamma^\mu(1-\gamma^5)l_\alpha W^+_\mu + \text{h.c.} \tag{2.5}$$

where $\overline{\nu_\alpha}$ is the adjoint neutrino spinor of flavor α, l_α is the lepton spinor of flavor α, $\gamma^\mu = \left(\gamma^0, \gamma^1, \gamma^2, \gamma^3\right)$ are the Dirac matrices, $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$, and g is the weak coupling constant associated with the symmetry group $SU(2)$. The fields that create/annihilate $W^\pm$ bosons are defined as a linear combination of the first two massless vector gauge boson fields A^1_μ and A^2_μ associated with $SU(2)_L$ group as: $W^\pm_\mu = \left(A^1_\mu \mp iA^2_\mu\right)/\sqrt{2}$. The CC interactions for leptons can be represented by four vertices shown in Figure 2·1.

Weak interactions mediated by the Z^0 boson are referred to as *neutral current* (NC) interactions and leave the charges of the initial and final state leptons intact. The NC interactions of leptons are described by the following Lagrangian:

$$\mathscr{L}^{NC} = -\frac{g}{2\cos\theta_W}\left\{\overline{\nu_{\alpha L}}\gamma^\mu\nu_{\alpha L} - \left(1-2\sin^2\theta_W\right)\overline{e_L}\gamma^\mu e_L + 2\sin^2\theta_W\overline{e_R}\gamma^\mu e_R\right\}Z_\mu$$
$$+ g\sin\theta_W\overline{e}\gamma^\mu eA_\mu \tag{2.6}$$

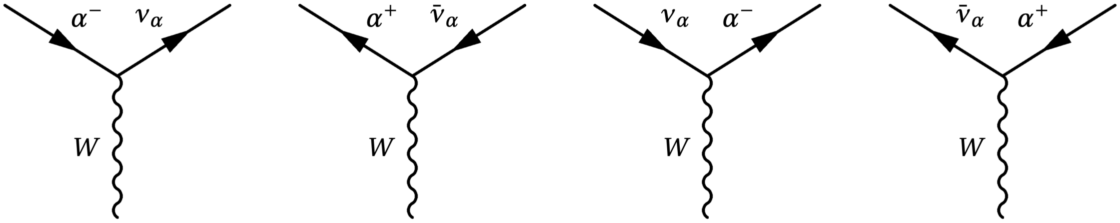

Figure 2·1: Charged current weak interaction vertices involving leptons.

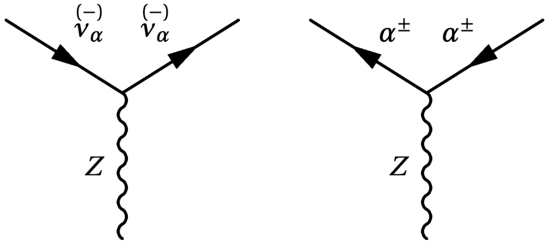

Figure 2·2: Neutral current weak interaction vertices involving leptons.

where Z_μ is the vector boson field representing Z^0 and A_μ is the electromagnetic field representing the photon. Both Z_μ and A_μ are expressed in terms of linear combinations of the third vector gauge boson field A^3_μ of the $SU(2)_L$ group and the vector gauge boson field B_μ of the $U(1)_Y$ group in the following way:

$$A^\mu = \sin\theta_W A^\mu_3 + \cos\theta_W B^\mu \tag{2.7}$$

$$Z^\mu = \cos\theta_W A^\mu_3 - \sin\theta_W B^\mu \tag{2.8}$$

The angle θ_W is called the *weak mixing angle*. The weak mixing angle is connected to the coupling constants g from $SU(2)_L$ and g' from $U(1)_Y$ such that $g\sin\theta_W = g'\cos\theta_W$. The NC interactions of leptons can be represented by two distinct vertices shown in Figure 2·2.

2.3 Neutrino Oscillations

One of the most important milestones in particle physics was the discovery of neutrino oscillations for which Art McDonald of the SNO collaboration and Takaaki Kajita of the Super-K collaboration were awarded the 2015 Nobel Prize in Physics [4]. The concept of flavor oscillations was first proposed by Bruno Pontecorvo in the late 1950s when looking for a lepton phenomenon analagous to $K^0 \leftrightarrow \overline{K}^0$ oscillations [5]. Neutrino oscillations arise due to the fact that neutrino flavor states are not mass eigenstates. During a weak interaction, a neutrino is generated in a definite flavor eigenstate along with its associated charged lepton. The flavor ($\alpha = e$, μ, or τ) eigenstates are related to the mass ($k = 1$, 2, and 3) eigenstates as follows:

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k} \, |\nu_k\rangle \tag{2.9}$$

which can also be written in matrix form as:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.10}$$

The weights $U_{\alpha k}$ are elements of the Pontecorvo-Maka-Nakagawa-Sakata (PMNS) matrix, characterized by three angles of rotation (θ_{12}, θ_{23}, and θ_{13}), a CP-violating phase δ_{CP}, and two Majorana phases ξ_1 and ξ_2. The most common parameterization of the PMNS matrix is given by:

$$\begin{aligned} U = &\begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta_{23} & \sin\theta_{23} \\ 0 & -\sin\theta_{23} & \cos\theta_{23} \end{pmatrix} \begin{pmatrix} \cos\theta_{13} & 0 & \sin\theta_{13} e^{-i\delta_{\mathrm{CP}}} \\ 0 & 1 & 0 \\ -\sin\theta_{13} e^{i\delta_{\mathrm{CP}}} & 0 & \cos\theta_{13} \end{pmatrix} \\ &\times \begin{pmatrix} \cos\theta_{12} & \sin\theta_{12} & 0 \\ -\sin\theta_{12} & \cos\theta_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\xi_1} & 0 \\ 0 & 0 & e^{i\xi_2} \end{pmatrix} \end{aligned} \tag{2.11}$$

The last 3×3 matrix with the Majorana phase component drops out when performing the inner product. It is not necessarily an important aspect of neutrino oscillation experiments, since the Majorana phase component drops out when performing the inner product to calculate oscillation probabilities. However, it could play an important role in future neutrinoless double beta decay experiments.

To provide some context as to how neutrino oscillation parameters are measured by experiments, a derivation of neutrino oscillation probabilities in a vacuum is presented here. Unlike the quarks, which are only ever observed in bound systems of two or more, neutrinos can propagate through space as free particles. One can generally assume the massive neutrino states (ν_k) evolve in time as plane waves making them solutions of the Schrödinger equation. The plane wave solution and corresponding energy eigenvalues E_k can be written as:

$$|\nu_k(t)\rangle = e^{(-iE_k t)}\,|\nu_k\rangle\,, \quad \text{and} \quad E_k = \sqrt{m_k^2 + \vec{p}^2} \tag{2.12}$$

where $|\nu_k\rangle$ is a mass eigenstate with eigenvalue m_k at rest, and momentum p_k. If we now consider the flavor state $|\nu_\alpha(t)\rangle$ which describes a neutrino created with a definite flavor α at time $t = 0$, the time evolution of this state is given by:

$$|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k}^* \, e^{-iE_k t}\,|\nu_k\rangle \tag{2.13}$$

We can now invoke the unitary relation:

$$U^\dagger U = \mathbb{1} \Longleftrightarrow \sum_\alpha U_{\alpha k}^* U_{\alpha j} = \delta_{jk} \tag{2.14}$$

This allows us to substitute Equation 2.9 into Equation 2.13 and write the time evolution of a flavor state in terms of another flavor state at a later time:

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_k U_{\alpha k}^* \, e^{-iE_k t}\, U_{\beta k} \right) |\nu_\beta\rangle \tag{2.15}$$

Using Equation 2.15, the probability that neutrinos generated in one flavor state ν_α will be detected in another ν_β, can be calculated as:

$$P(\nu_\alpha \to \nu_\beta) = |\,\langle \nu_\beta | \nu_\alpha(t) \rangle\,|^2 = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t} \tag{2.16}$$

Let us now focus on the phase in exponent in Equation 2.16. Since neutrinos are ultra-relativistic particles and the neutrino masses are much smaller than their momentum $(m_k \ll |\vec{p}|)$, an approximation can be used for the energy-momentum relation:

$$E_k \simeq p_k + \frac{m_k^2}{2p_k} \approx E + \frac{m_k^2}{2E} \tag{2.17}$$

The indice k disappears in this approximation since it is assumed that neutrinos with the same momentum propagate in the same direction ($p_k = p \approx E$). Our original phase can now be modified as follows:

$$E_k - E_j = \frac{\Delta m^2_{kj}}{2E}, \quad \text{where } \Delta m^2_{kj} \equiv m^2_k - m^2_j \tag{2.18}$$

For neutrino oscillation experiments the propagation time t is not a measured quantity. Instead it is the distance, L, between the source and the detector that is known, and the ultra-relativistic approximation $t \simeq L$ can be used. This allows us to write Equation 2.16 in terms of the total neutrino energy E and source-detector distance L:

$$P(\nu_\alpha \to \nu_\beta) = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} e^{-i(\Delta m^2_{kj})L/2E} \tag{2.19}$$

Another useful way to write the oscillation probability in Equation 2.19 is to separate the real and imaginary parts $U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}$:

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} - 4 \sum_{k>j} \mathfrak{Re}\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right] \sin^2\left(\frac{\Delta m^2_{kj} L}{4E}\right) \\ + 2 \sum_{k>j} \mathfrak{Im}\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right] \sin\left(\frac{\Delta m^2_{kj} L}{2E}\right) \end{aligned} \tag{2.20}$$

The PMNS matrix terms determine the amplitude of the oscillation, and the phase in the argument this expression is responsible for the oscillation itself and its frequency. The phase is currently written in the form of natural units ($\hbar = c = 1$), but is more useful for experimentalists when written, for example, in terms of L[km], E[GeV], and Δm^2_{jk}[eV^2]:

$$\frac{\Delta m^2_{jk} L}{2E} \approx 1.27 \frac{\Delta m^2_{jk}[\text{eV}^2]\, L[\text{km}]}{E[\text{GeV}]} \tag{2.21}$$

Experimentalists will build their neutrino detectors at a well defined distance from a source of neutrinos from which the neutrino energy is well known. We can see that an experiment can maximize its sensitivity to oscillation frequency if the argument in the sinusoidal terms in Equation 2.20 are close to unity. For example, if we desire sensitivity to oscillations with a frequency driven by $\Delta m^2_{jk} \sim 10^{-3}\,\mathrm{eV}^2$, then we would aim to build our detector at a distance of $L = 1{,}000\,\mathrm{km}$ from a source of roughly $E =$ 1 GeV neutrinos. On the other hand, if an experiment was to choose values such that $\Delta m^2_{jk} L/2E \gg 1$, then they would only be sensitive to the average transition probability, and hence, only the mixing angles of the PMNS terms. This is due to the frequency of the oscillations becoming so rapid that they are essentially incoherent. If an experiment chooses values such that $\Delta m^2_{jk} L/2E \ll 1$, then they wouldn't observe any oscillations at all due to the oscillation probability being close to zero.

It is important to note that we have not included a proper description of neutrino oscillations in matter, which becomes critically important for neutrinos that propagate through the most dense regions of the core of our Sun or neutrinos that travel a significant distance through the Earth. Matter oscillations, as they are called, are beyond the scope of this work but the general theory is described in detail by Reference [6].

Within the last 25 years, most of the neutrino oscillation parameters (θ_{12}, θ_{23}, θ_{13}, Δm^2_{21}, and Δm^2_{31}) have been measured to several percent uncertainty or better. This has been accomplished with a wide range of different neutrino detectors measuring oscillations from neutrino sources such as the Sun, the Earth's atmosphere, nuclear reactors, and particle accelerators. Only two big questions remain with respect to neutrino oscillations, mainly, the value of δ_{CP} and the ordering of the neutrino mass eigenstates. Since past and current experiments haven't been able to determine the sign of Δm^2_{31} (or Δm^2_{32}), there remain two possible neutrino mass orderings: normal ($m_3 > m_2 > m_1$) and inverted ($m_2 > m_1 > m_3$). The uncertainty around the mass ordering arose due to the inability of past and current experiments to determine the sign of Δm^2_{13}. Future experiments such as JUNO [7], Hyper-Kamiokande [8], and DUNE [9] are positioned to shed light on these remaining mysteries in the very near future. Table 2.1 is recreated from NuFIT version 6.0 in Reference [10] and provides a summary of all the best-fit data for neutrino oscillation parameters coming from a global fit that includes

Oscillation parameter	Normal	Inverted
$\Delta m^2_{21}\,[10^{-5}\,\text{eV}^2]$	$7.49^{+0.19}_{-0.19}$	$7.49^{+0.19}_{-0.19}$
$\Delta m^2_{3k}\,[10^{-3}\,\text{eV}^2]$	$+2.513^{+0.021}_{-0.019}$	$-2.484^{+0.020}_{-0.020}$
$\sin^2\theta_{12}$	$0.308^{+0.012}_{-0.011}$	$0.308^{+0.012}_{-0.011}$
$\sin^2\theta_{23}$	$0.470^{+0.17}_{-0.13}$	$0.550^{+0.012}_{-0.015}$
$\sin^2\theta_{13}$	$0.02215^{+0.00056}_{-0.00058}$	$0.02231^{+0.00056}_{-0.00056}$
$\delta_{CP}\,[°]$	212^{+26}_{-41}	274^{+22}_{-25}

Table 2.1: Best-fit values $\pm 1\sigma$ from a global analysis of neutrino oscillation parameters reproduced from NuFIT version 6.0 in Reference [10]. Note that $\Delta m^2_{3k} \equiv \Delta m^2_{31} > 0$ for normal ordering and $\Delta m^2_{3k} \equiv \Delta m^2_{32} < 0$ for inverted ordering.

atmospheric neutrino data from the Super-Kamiokande experiment. A graphical representation of the neutrino mass ordering and flavor composition of the mass eigenstates is shown in Figure 2·3, where Δm^2_{sol} refers to Δm^2_{21} and Δm^2_{atm} refers to Δm^2_{32} or Δm^2_{31}. This labeling is quite common in literature as solar neutrino (and some reactor neutrino) experiments were generally sensitive to Δm^2_{21} and atmospheric neutrino experiments were generally sensitive to $|\Delta m^2_{32}| \approx |\Delta m^2_{31}|$.

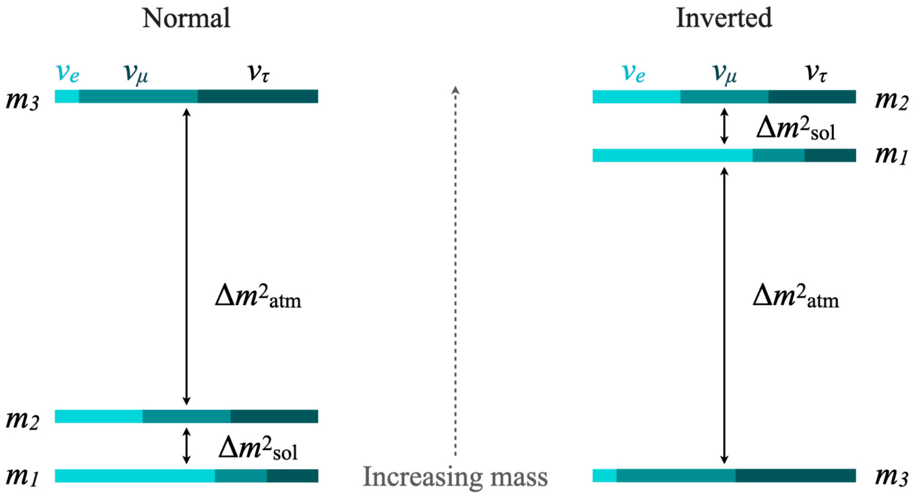

Figure 2·3: Diagram of the neutrino mass orderings and the flavor composition of each mass eigenstate.

As seen from Equation 2.20, neutrino oscillation experiments are only sensitive to the squared mass differences Δm^2_{jk} and have no sensitivity to the values of the individual mass eigenstates. Therefore, determining the absolute masses of neutrinos requires a different experimental approach. One such technique, employed by the KATRIN experiment, is to directly measure the amount of energy carried away by the neutrino during the β decay of tritium with a Q-value of 18.6 keV. In principle, a massive neutrino must carry away a very small part of the total energy released during the β decay. This means that the difference between the emitted β energy and the Q-value places a limit on the effective mass of the *electron* flavor neutrino, which exists as a combination of mass eigenstates:

$$m_{\nu_e} = \sqrt{\sum_i |U_{ei}|^2 m_i^2} \tag{2.22}$$

A successful measurement is extremely challenging and requires exquisite energy resolution of order single electron-volts or better. To date, the best limit on the effective electron neutrino mass set by KATRIN is $m_{\nu_e} < 0.8$ eV at 90% C.L. [11]. Another future experiment, called Project 8, aims to improve on this measurement using a technique called Cyclotron Radiation Emission Spectroscopy (CRES) [12].

An alternative, indirect approach for measuring neutrino masses comes from cosmological measurements, due to the fact that neutrinos play an important role in the evolution and shaping of large-scale structure that we observe in our universe. A combination of one or more measurements of the Cosmic Microwave Background (CMB), Baryon Acoustic Oscillations (BAO), and Redshift Space Distortions (RSD), are commonly used to set a limit on the sum of the mass eigenstates $\sum m_i = m_1 + m_2 + m_3$. At present, the most stringent cosmological bound on the sum of neutrino masses comes from Reference [13], where $\sum m_i < 0.09$ eV at 95% C.L.

There is yet another very interesting and extremely sensitive method for measuring the absolute mass scale of neutrinos, which involves a hypothetical nuclear process called *neutrinoless double beta decay* ($0\nu\beta\beta$). Before jumping into the topic of $0\nu\beta\beta$ decay, which is discussed in further detail in Chapter 3, we will first review some theoretical aspects of neutrino mass in the following Section.

2.4 Neutrino Mass

The discovery of neutrino oscillations provided experimental proof that at least two out of three neutrino mass states must have non-zero mass. This is in direct conflict with the original formulation of the Standard Model where it was assumed that neutrinos are massless. Hence, the discovery of massive neutrinos indicates the existence of new, undiscovered physics. What follows is a brief review of some necessary modifications of the Standard Model that allow for massive neutrinos.

2.4.1 Dirac Neutrino Masses

In the Standard Model, all fermions gain mass via the Higgs mechanism [14, 15, 16] which invokes *spontaneous symmetry breaking* of the $SU(2)_L \times U(1)_Y$ gauge group. To do this, we construct a mass term that is invariant under $SU(2)_Y \times U(1)$ by using the complex Higgs doublet:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} \tag{2.23}$$

which is written in the unitary gauge thereby setting the Goldstone boson components to zero. Here, v is the Higgs vacuum expectation value and h is the Higgs field. This doublet, and its conjugate ϕ_c, can be coupled to the fermion fields to form a *Dirac* mass term Lagrangian. This can be seen if we look at the following example for the $SU(2)_L$ doublet containing the electron:

$$\begin{aligned}
\mathscr{L}_e^D &= -\frac{y_e}{\sqrt{2}}\left[\left(\overline{\nu_e}\,,\,\overline{e}\right)_L \begin{pmatrix} 0 \\ v+h \end{pmatrix} e_R + \overline{e_R}\,(0\,,\; v+h)\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L \right] \\
&= -\frac{y_e}{\sqrt{2}} v \left(\overline{e_L} e_R + \overline{e_R} e_L\right) - \frac{y_e}{\sqrt{2}} h \left(\overline{e_L} e_R + \overline{e_R} e_L\right) \\
&= -\frac{y_e}{\sqrt{2}} v \bar{e} e - \frac{y_e}{\sqrt{2}} h \bar{e} e
\end{aligned} \tag{2.24}$$

where we've introduced the constant y_e representing the Yukawa coupling of the electron to the Higgs field. We can choose to write this coupling constant in a way that is consistent with the observed electron mass:

$$m_e = \frac{y_e v}{\sqrt{2}} \tag{2.25}$$

Upon inspection, we see that the first term in Equation 2.24 gives the electron mass by coupling it to the Higgs field through a non-zero vacuum expectation value. The second term gives rise to a coupling between the electron and the Higgs boson itself. To give mass to fermions in the upper part of the $SU(2)_L$ doublet, like the u, c, and t quarks, we can simply replicate the procedure for the electron but instead use the conjugate Higgs doublet (also written in the unitary gauge):

$$\phi_c = -i\tau_2\phi^* = -\frac{1}{\sqrt{2}}\begin{pmatrix} v+h \\ 0 \end{pmatrix} \tag{2.26}$$

If we step back for a moment, we'll notice that our approach for creating Dirac mass terms requires a coupling between the left-handed and right-handed Dirac spinors. However, we know that the neutrinos (anti-neutrinos) are left(right)-handed in nature [17]. By original design, right-handed neutrino singlets ν_R don't exist in the Standard Model. If the neutrino is a Dirac fermion, and also gains mass in the same way that the other fermions do, then we're required to add right-handed (sterile) states into

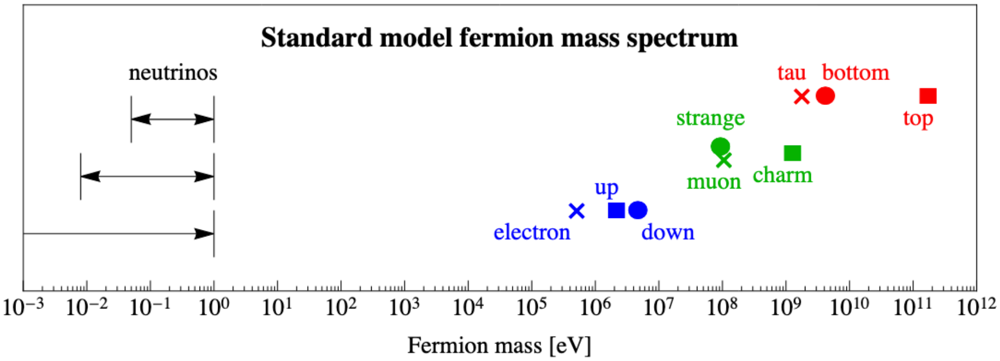

Figure 2·4: Diagram adapted from Reference [18] showing the relative sizes of fermion masses in the Standard Model.

the Standard Model. This is sometimes referred to as a ‘minimal’ extension.

Even if we go along with the idea that right-handed neutrinos exist in nature and neutrinos are Dirac fermions, we are left without an explanation for the relative magnitude of the neutrino mass. As far as we know, neutrinos are roughly six orders of magnitude less massive than electrons. The bizarre rift between masses of neutrinos and the other matter particles is illustrated in Figure 2·4. The idea that neutrinos have Dirac mass terms would require them to have a very different Higgs-Yukawa coupling than the other fermions. This is also very unsettling and motivates our search for other possible mass-generating mechanisms.

2.4.2 Majorana Neutrino Masses

An alternative theory of spin-1/2 particles, where particles and antiparticles enter symmetrically into the Dirac equation, was created by Ettore Majorana [19]. Majorana made a very crucial assumption, mainly:

$$\psi_R = C\overline{\psi_L^T}, \qquad C = i\gamma^2\gamma^0 \tag{2.27}$$

This statement allows the Dirac equation to be written only in terms of left-handed field ψ_L and also implies that we can write a *Majorana* field as:

$$\psi = \psi_L + \psi_R = \psi_L + C\overline{\psi_L^T} = \psi_L + \psi_L^c \tag{2.28}$$

where we've defined the charge-conjugated field as $\psi_L^c = C\overline{\psi^T}$. Notice that this implies the following must be true:

$$\psi^c = (\psi_L + \psi_L^c)^c = \psi_L^c + \psi_L = \psi \tag{2.29}$$

Equation 2.29 requires that ψ is its own antiparticle. Obviously, this can only be applied to fermions that are electrically neutral, and the neutrino is the only such particle. The neutrino being a Majorana fermion has profound consequences. For example, Dirac neutrinos are assigned a lepton number of $L = +1$ and Dirac antineutrinos are assigned $L = -1$, but particle processes involving Majorana neutrinos can't possibly conserve lepton number as there's no way to distinguish a neutrino from an antineutrino.

At this point we might be tempted to form a Majorana mass term for our Lagrangian analogous to the Dirac mass term without introducing right-handed neutrino states. We could write this as:

$$\mathscr{L}_\nu^M = -\frac{1}{2}M\overline{\nu_L^c}\nu_L + \text{h.c.} \tag{2.30}$$

where we have written the mass M in place of the Yukawa coupling constant ($M = y_\nu v/\sqrt{2}$). Unfortunately, such a term is does not end up being gauge invariant with the Higgs doublet ϕ that we introduced before. One would have to add something like a Higgs triplet with $I_3 = -1$ and $Q = 0$ into the mix. However, there's still no clear experimental evidence that a Higgs triplet exists in nature.

Another approach is to write down a mass term that only includes right-handed Majorana neutrino fields, which would preserve gauge invariance with a Higgs doublet. This can be done in a very general form with a combination of Dirac and Majorana terms including right-handed neutrino fields:

$$\mathscr{L}_{\nu}^{M+D} = -\frac{1}{2}\left(m_D\overline{\nu_L}\nu_R + m_D\overline{\nu_R^c}\nu_L^c + M\overline{\nu_R^c}\nu_R\right) + \text{h.c.} \tag{2.31}$$

where we note that $\overline{\nu_L}\nu_R = \overline{\nu_R^c}\nu_L^c$ has been used for the middle term. Equation 2.32 is often written in a matrix form:

$$\mathscr{L}_{\nu}^{M+D} = -\frac{1}{2}\left(\overline{\nu_L}\,,\,\overline{\nu_R^c}\right)\begin{pmatrix} 0 & m_D \\ m_D & M \end{pmatrix}\begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + \text{h.c.} \tag{2.32}$$

We can find the physical states of this system from the basis where the mass matrix is diagonalized the usual way when solving eigenvalue problems. Solving the characteristic equation yields:

$$m_{\pm} = \frac{M \pm M\sqrt{1+4m_D^2/M^2}}{2} \tag{2.33}$$

This gives us two physical solutions for the neutrino mass. Things get really interesting if we assume that the Majorana mass is much greater than the Dirac mass: $M \gg m_D$. We then find one solution representing a light neutrino state m_ν and another solution representing a very heavy state m_N:

$$m_\nu \approx \frac{m_D^2}{M} \quad \text{and} \quad m_N \approx M \tag{2.34}$$

This result is what is known as a *seesaw mechanism* (Type-I). The m_ν is the familiar mass of a very light left-handed Majorana neutrino, and the m_N is the mass of a very heavy right-handed (sterile) neutrino. Suppose the Dirac mass of the neutrino is close to the value of the top quark mass ($m_D \sim 170$ GeV). If the mass of the heavy right-handed partner is around 10^{15} GeV (comparable to the GUT scale), then the mass of the light neutrino will be $O(10)$ meV and consistent with our best estimates to date. This is a very interesting hypothesis to explain the small mass of neutrinos relative to the other matter particles without the need for an unusually small Higgs Yukawa coupling. It requires that the neutrino be a Majorana particle, and that there exists an extremely heavy partner to the neutrino with a mass too large to be easily created in a laboratory setting. The possibile existence of these heavy neutrinos are an important component in another intriguing idea called *leptogenesis* which is intimately linked to the question of the matter-antimatter asymmetry of the universe [20]. Leptogenesis proposes that as the universe cooled, extremely heavy Majorana neutrinos decayed into lighter left-handed neutrinos or right-handed antineutrinos, accompanied by Higgs bosons, which subsequently decayed into quarks. If the decay probabilities of these heavy neutrinos favored left-handed neutrinos over right-handed antineutrinos, an imbalance would emerge, leading to a surplus of quarks over antiquarks and resulting in a matter-dominated universe. The process is closely tied to the conservation of the quantum number $B-L$ (baryon number minus lepton number). A violation of lepton number L in the decays of heavy Majorana neutrinos would translate into a baryon number B violation, offering a possible explanation for the absence of antimatter. While CP violation in heavy neutrinos is not directly linked to CP violation in light neutrinos, the possibility of such effects continues to drive experimental efforts to measure CP violation in long-baseline neutrino experiments today [8, 9, 21].

Chapter 3

Double Beta Decay

As discussed in Chapter 2, the fact that neutrinos are massive provides strong theoretical arguments for new physics beyond the Standard Model. One of the most exciting possibilities is that the neutrino is a Majorana particle. This chapter describes the best experimental probe of the Majorana nature of neutrinos: *neutrinoless double beta decay*. A brief overview of two-neutrino double beta decay and neutrinoless double beta decay is provided, while highlighting important theoretical considerations in the search for neutrinoless double beta decay. A general guide for maximizing an experiment's sensitivity in the search for neutrinoless double beta decay is outlined near the end.

3.1 Two-neutrino Double Beta Decay ($2\nu\beta\beta$)

Most nuclei are unstable due to the weak interaction and decay via the emission of a single electron (single β^- decay) or single electron capture (EC). These decays produce final nuclei that more tightly bound and with the same number of nucleons. In β^- decay a neutron turns into a proton, while the opposite occurs in EC, so that electric charge is conserved. In addition, either neutrinos (in EC) or antineutrinos (in β^- decay) are emitted to conserve energy, momentum and lepton number. In some nuclei, β^+ decay can turn a proton into a neutron, but this is often disfavored relative to EC because it produces a positron and reduces the available energy: $Q_{\beta^+} = Q_{EC} - 2m_e$.

In 1935, Maria Goeppert-Mayer proposed and calculated the transition probability for the simultaneous emission of two electrons and two neutrinos ($2\nu\beta\beta$) [22]:

$$(A, Z) \rightarrow (A, Z+2) + 2e^{-} + 2\overline{\nu}_e \tag{3.1}$$

This lepton-number-conserving process occurs through two simultaneous β^- decays is allowed by the Standard Model regardless of the neutrino's Dirac or Majorana nature. In addition to the Goeppert-Maiers proposed $2\nu\beta\beta$ decay, there are three additional second-order transitions are allowed in the Standard Model that depend on the relative numbers of protons and neutrons in the nucleus:

$$\begin{aligned} \text{Double } \beta^+ \text{ decay}: &\quad (A, Z) \rightarrow (A, Z-2) + 2e^{+} + 2\nu_e \\ \text{Double EC decay}: &\quad (A, Z) + 2e^{-} \rightarrow (A, Z-2) + 2\nu_e \\ \text{EC} + \beta^+ \text{ decay}: &\quad (A, Z) + e^{-} \rightarrow (A, Z-2) + e^{+} + 2\nu_e \end{aligned} \tag{3.2}$$

The energy released in these three processes is smaller compared to the double β^- decay in Equation 3.1. In turn, they have lower probabilities due to the smaller available phase space and are much more challenging to experimentally observe. When discussing $2\nu\beta\beta$ decay, the double β^- process in Equation 3.1 will always be the standard reference unless otherwise noted.

The $2\nu\beta\beta$ decay process changes the nuclear charge Z by two units and leaves the atomic mass number A unchanged. This process can occur only if the conversion of two protons into two neutrons leads to a more tightly-bound nuclear system, with a positive $2\nu\beta\beta$ decay Q-value, $Q_{\beta\beta}$. The $Q_{\beta\beta}$ is defined to be the available kinetic energy for the two ejected electrons and can be calculated using the mass difference between the initial and final products:

$$Q_{\beta\beta} = m_N\left({}^{A}_{Z}X\right) - m_{N-2}\left({}^{A}_{Z+2}X'\right) - 2m_e \tag{3.3}$$

Here, m_e is the electron mass in the rest frame, $m_N\left({}^{A}_{Z}X\right)$ is the mass of the mother nucleus in the rest frame, and $m_{N-2}\left({}^{A}_{Z+2}X'\right)$ represents the mass of the daughter nucleus in the rest frame.

To understand why some nuclei are able to undergo $2\nu\beta\beta$ decay, we can look at the formula used to describe the mass of an atomic nucleus M in terms of atomic number A and proton number Z:

$$M = Zm_p + (A - Z)m_n - E(A, Z) \tag{3.4}$$

The constants m_p and m_n are the rest masses of the proton and neutron, and E is the binding energy given by the semi-empirical mass formula [23]:

$$E = a_v A - a_s A^{2/3} - a_C \frac{Z(Z-1)}{A^{1/3}} - a_A \frac{(N-Z)^2}{A} + \delta(N, Z) \tag{3.5}$$

The first term in Equation 3.5 is the volume term, which scales linearly with the number of nucleons and accounts for the strong force between nucleons and their nearest (and second nearest) neighbors. The second term, is referred to as the surface term, is a correction to the volume term to account for surface nucleons being less tightly bound. The third term accounts for the Coulomb repulsion between protons. The fourth term results from the Pauli exclusion principle, which imposes states of increasing energy for increasing number of nucleons. The final term, $\delta(N, Z)$, is the pairing energy parameterized by the constant $a_p \sim 12$ MeV and captures the effect of spin coupling. It's exact form is constrained by experimental data and is defined as:

$$\delta_p = \begin{cases} -a_p A^{-1/2} & \text{even-even nuclei} \\ 0 & \text{even-odd and odd-even nuclei} \\ +a_p A^{-1/2} & \text{odd-odd nuclei} \end{cases} \tag{3.6}$$

Nuclei with odd mass numbers have vanishing pairing energy, and thus one stable isobar. Those with even mass numbers have two possible pairing energies, separated by $2\delta_p$. Although an even-even nucleus has access to its lowest-energy isobar, sometimes single β^- decay is energetically forbidden or forbidden due to a large mismatch

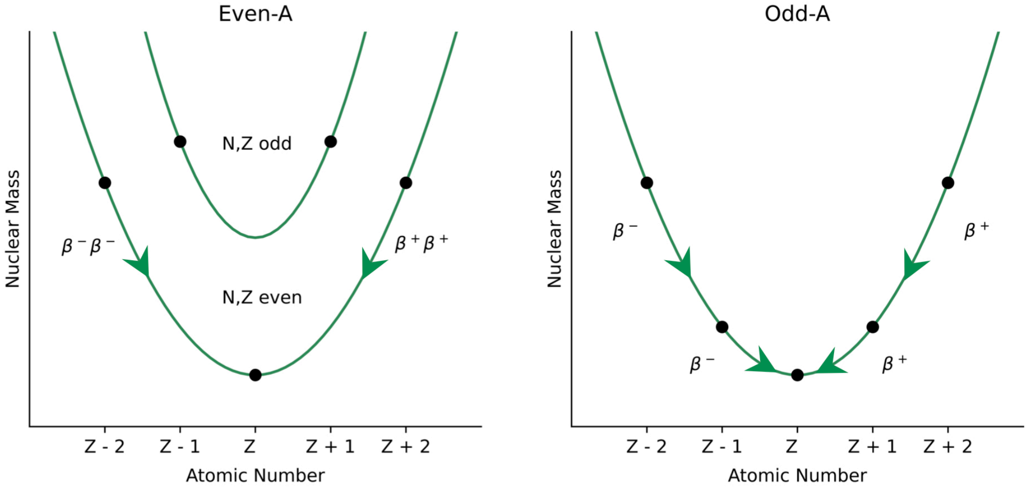

Figure 3·1: Ground state mass parabolas for nuclear isobars with even (left) and odd (right) mass number A.

in angular momentum between initial and final states. Nuclei with this restriction will decay directly to the next-available configuration through simultaneous emission of an additional β^-, which is illustrated by the mass parabolas in Figure 3·1.

The $2\nu\beta\beta$ decay rate (inverse half-life) can be calculated following Fermis golden rule. This can be written as:

$$\Gamma^{2\nu} = (T_{1/2}^{2\nu})^{-1} = G^{2\nu}(Q_{\beta\beta}, Z)\left|\mathcal{M}_{GT}^{2\nu} + \frac{g_V^2}{g_A^2}\mathcal{M}_F^{2\nu}\right|^2 \tag{3.7}$$

where $G^{2\nu}(Q_{\beta\beta}, Z)$ is the phase-space factor obtained by integrating over the phase space of the four final-state leptons, g_V is the vector coupling constant, g_A is the axial-vector coupling constant, $\mathcal{M}_F^{2\nu}$ is the Fermi matrix element, and $\mathcal{M}_{GT}^{2\nu}$ is the Gamow-Teller matrix element. The $\mathcal{M}_F^{2\nu}$ describes decays with no change in nuclear spin (a vector-type interaction) and the $\mathcal{M}_{GT}^{2\nu}$ describes decays with a change in spin of one unit (an axial-vector-type interaction). For the two-neutrino case, Fermi transitions are strongly suppressed, since the transition proceeds through a state with different isospin, $0^+ \rightarrow 1^+$.

While the phase-space factor $G^{2\nu}(Q_{\beta\beta}, Z)$ can be calculated rather accurately [24, 25], the matrix element rely on wave functions of the initial and final nuclear states, meaning that these calculations are often complex and require sophisticated nuclear models to determine. Given that the $2\nu\beta\beta$ is a second-order weak process, typical half-lives are on the order of 10^{19} - 10^{22} years [26]. It should also be noted that theoretical calculations of $\mathscr{M}^{2\nu}_{GT}$ are typically not in agreement with experimental measurements. Therefore, a kind of “renormalization" is often used to bring them into agreement. This is usually done in the form of $\mathscr{M}^{2\nu,\mathrm{eff}}_{GT} = (g^{\mathrm{eff}}_A/g_A)^2\mathscr{M}^{2\nu}_{GT}$, where g^{eff}_A is the effective vector-axial coupling which is *quenched* relative to the free nucleon value g_A. This quenching value differs widely among models and theoretical calculations and its origin is an active area of research in nuclear theory [27].

3.2 Neutrinoless Double Beta Decay ($0\nu\beta\beta$)

Soon after Maria Goeppert-Mayer’s landmark work on $2\nu\beta\beta$ decay, Wendell Furry proposed the process known as $0\nu\beta\beta$ decay [28]:

$$(Z,A) \rightarrow (Z+2,A) + 2e^- \tag{3.8}$$

Unlike the two-neutrino mode of double beta decay, the neutrinoless mode violates total lepton number conservation by two units and is therefore forbidden in the Standard Model. The $0\nu\beta\beta$ process is one of the best experimental probes for the Majorana nature of the neutrino since $0\nu\beta\beta$ decay can only occur if the neutrino is a Majorana particle.

In both the $2\nu\beta\beta$ and $0\nu\beta\beta$ decay modes, it should be noted that the final state leptons carry essentially all of the energy of the decay, with the nuclear recoil being negligible. Both $2\nu\beta\beta$ and $0\nu\beta\beta$ decay are second order weak processes, which means the decay rates are inherently slow as they are proportional to G_F^4 (where Fermi’s constant is $G_F \propto g^2/m_W^2$ with m_W being the mass of the W boson). Phase space considerations would generally give preference to the $0\nu\beta\beta$ mode if it weren’t for lepton number violation. The $2\nu\beta\beta$ process differs from the $0\nu\beta\beta$ in the sense that the kinetic energy of the final state electrons in the $2\nu\beta\beta$ process form a continuous distribution peaked somewhat below half the endpoint energy ($Q_{\beta\beta}/2$). The two electrons from $0\nu\beta\beta$ decay will have exactly the endpoint energy $Q_{\beta\beta}$ and will form a sharp peak at the end of the $2\nu\beta\beta$ energy distribution (with some smearing due to detector energy resolution).

The discussion surrounding $0\nu\beta\beta$ decay in this thesis will always directly refer to the process shown in Equation 3.8. Other neutrinoless decay modes, such as double positron emission ($\beta^+\beta^+$), single positron emission with single-electron capture (β^+EC), and double-electron capture (ECEC), are generally less favorable experimentally due to the smaller available phase space for these processes.

In general, in theories beyond the Standard Model, there may be several sources of total lepton number violation which can lead to $0\nu\beta\beta$. Nevertheless, irrespective of the mechanism, $0\nu\beta\beta$ necessarily implies Majorana neutrinos. This is called the Schechter-Valle (or 'black box') theorem [29]. The Schechter-Valle theorem says nothing about the physics mechanism driving the $0\nu\beta\beta$ rate that is large enough to be observable. The dominant mechanism could be directly connected to neutrino oscillations phenomenology, only indirectly connected, or not connected to it at all [30]. The case where the $0\nu\beta\beta$ mechanism is connected directly to neutrino oscillation phenomenology is referred to as *light Majorana neutrino exchange* and is the most popular mechanism that most sensitivity predictions are focused on. However, it is entirely possible that other physics mechanisms are responsible for $0\nu\beta\beta$ decay, should it exist in nature. Other possible physics mechanisms are outlined in Reference [30].

When considering the standard light Majorana neutrino exchange mechanism, $0\nu\beta\beta$ decay occurs when the parent nucleus emits a pair of virtual W bosons that exchange a Majorana neutrino to produce the outgoing electrons. This is illustrated in Figure 3·2 with $2\nu\beta\beta$ decay for comparison. The rate is non-zero only for massive, Majorana neutrinos. This can be described by the fact that the exchanged anti-neutrino is emitted

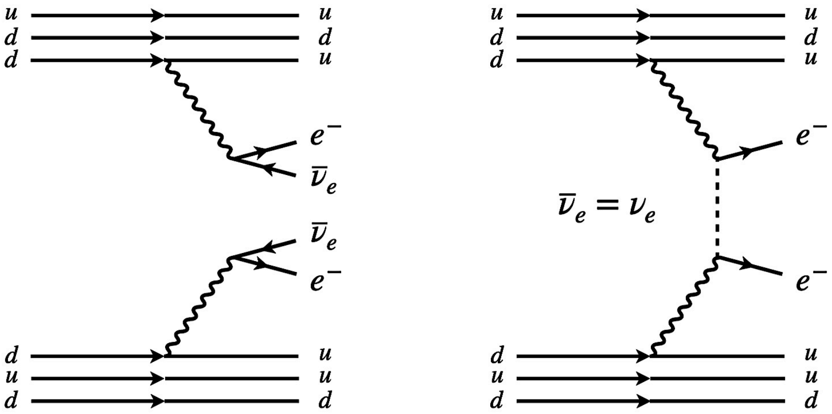

Figure 3·2: Feynmann diagrams for $2\nu\beta\beta$ decay (left) and $0\nu\beta\beta$ decay (right).

(in association with an electron) almost entirely with right-handed helicity, and its tiny, $\mathcal{O}(m/E)$, left-handed helicity component is absorbed in the other vertex by the Standard Model electroweak current.

In this case, the amplitude of the $0\nu\beta\beta$ decay is a sum over the contributions of the three light neutrino mass states ν_i, which are also proportional to the PMNS matrix elements U_{ei}^2. As such, we conclude that the amplitude for the $0\nu\beta\beta$ decay process must be proportional to the *effective neutrino Majorana mass* defined as:

$$m_{\beta\beta} = \left| \sum_{i=1}^{3} m_i \, e^{i\xi_i} \, U_{ei}^2 \right| \tag{3.9}$$

where the masses m_i correspond to the light neutrino mass states ν_i and $e^{i\xi_i}$ are the Majorana phases. In the case where light Majorana neutrino exchange is the dominant contribution, the decay rate (inverse half-life) for the process can be written in a simplified form as:

$$\Gamma^{0\nu} = \left(T_{1/2}^{0\nu}\right)^{-1} = G^{0\nu}\left(Q_{\beta\beta}, Z\right)\left|M^{0\nu}\right|^2 \left(\frac{m_{\beta\beta}}{m_e}\right)^2 \tag{3.10}$$

where $G^{0\nu}$ is the phase space factor obtained from integrating over all possible energies and angles of emitted particles and the m_e is the electron mass. The nuclear matrix element $M^{0\nu}$, often abbreviated (NME), contains the transition probabilities from the initial to final nuclear states. The NME for the light Majorana neutrino exchange is often expressed in terms of long-range and short-range contributions as $M^{0\nu} \rightarrow M^{0\nu}_{\text{light}} = g_A^4 \left(M^{0\nu}_{\text{long}} + M^{0\nu}_{\text{short}} \right)$. As stated earlier, g_A is axial-vector (long-range) coupling of the weak interaction to nucleons and is factor out. It represents the degrees of freedom used by many-body methods to calculate the NMEs. Furthermore, the short-range component of the NME is proportional to another two-nucleon coupling g^{NN}, which is not written explicitly here. Similar to the $2\nu\beta\beta$ decay case, the phase-space factors are quite accurately known for all relevant nuclei used in $0\nu\beta\beta$ decay experiments [24, 25]. However, despite recent progress, NMEs for $0\nu\beta\beta$ decay and some of their associated hadron couplings are still poorly known. This will be summarized in the following section.

To date, the best lower limit on $m_\beta\beta$, was achieved by the KamLAND-Zen experiment in 2024 with a half-life limit of $T^{0\nu}_{1/2} > 3.8 \times 10^{26}$ years at 90% confidence [31]. This corresponds to $m_{\beta\beta}$ in the range of 28-122 meV, where the width of the range is mostly due to the NMEs for ^{136}Xe, the $0\nu\beta\beta$ isotope used in KamLAND-Zen. This is shown in Figure 3·3 where $m_{\beta\beta}$ is plotted versus the lightest neutrino mass state. There are two possible regions where we would expect to see a signal, depending on whether the mass ordering is inverted (blue) or normal (magenta). KamLAND-Zen probed the inverted region for the first time but a large portion of the phase space remains unexplored.

3.3 $0\nu\beta\beta$ Nuclear Matrix Elements

The NMEs capture the effect of nuclear structure on $0\nu\beta\beta$ decay and are calculated using models of the initial and final nuclei, along with nuclear many-body approaches to evaluate the transition operator between them. When attempting to infer the value of $m_{\beta\beta}$ from the measured $T^{0\nu}_{1/2}$, the uncertainty on $m_{\beta\beta}$ is presently dominated the

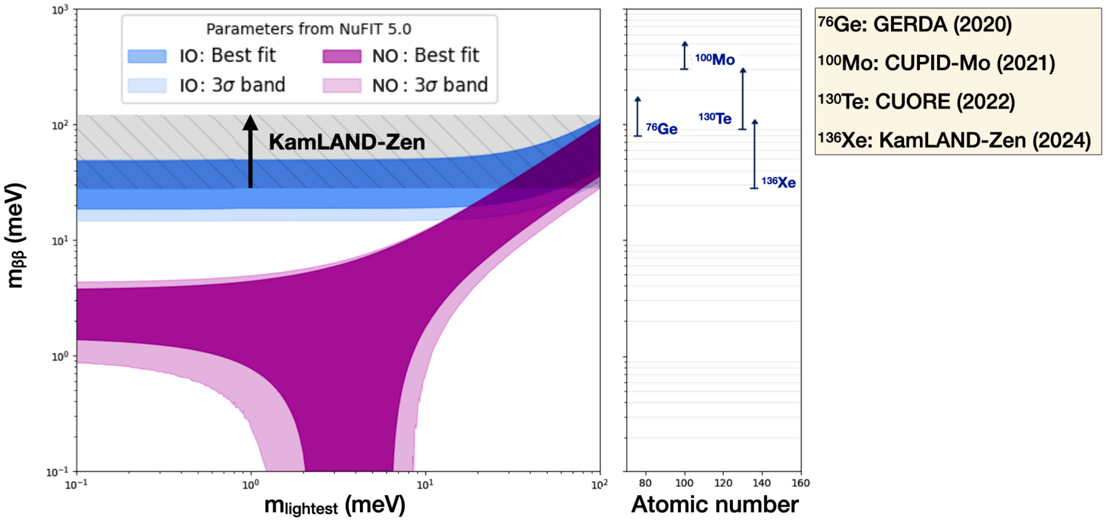

Figure 3·3: Possible Majorana masses for normal (magenta) and inverted (blue) mass orderings, calculated with mixing angles and mass differences from the PMNS matrix. The error bands come from uncertainties in the mixing parameters [32]. The KamLAND-Zen experimental limit on $m_{\beta\beta}$ is shown in gray. Present limits [33, 34, 35] for several other isotopes are shown in the panel on the right.

NMEs. This section will only briefly summarize the many-body models used to calculate the relevant NMEs. A detailed theoretical explanation of the NME calculations, uncertainties, and considerations for the light Majorana exchange mechanism, is well beyond the scope of this thesis. However, a more in-depth description of modern NME calculations can be found in Reference [36].

NME calculations typically proceed in two steps. The first and standard step creates a many-body Hamiltonian containing nucleon-nucleon interactions and the decay mechanism. The second step introduces a field that captures nuclear structure and higher-order interactions. To date, the most-often-used calculation methods are as follows:

- The **Nuclear Shell Model (NSM)** has historically been the main workhorse for the computation of nuclear structure. The basic assumption is that the nucleus as composed by nucleons that move independently of each other inside a spherically symmetric mean field with an additional strong spin-orbit term. This mean field potential represents the average interaction with the other nucleons and is often described by harmonic oscillator or Woods-Saxon potentials. The nucleons are considered to be non-relativistic point-like particles and as such without any internal structure. This produces single-particle states clustered in orbits, called shells, that are close in energy, which is especially true for observed nuclei with magic numbers of protons and/or neutrons (they more tightly bound than their neighbors). These energy levels are sufficiently separated from each other so that the nucleus can be regarded as an inert core, made up by filled shells and external nucleons, referred to as valence nucleons.

- The **Quasi Random Phase Approximation (QRPA)** is a many-body perturbation theory approach that spans a wide space of nuclear orbitals, using experimentally determined nucleon-nucleon potentials similar to the Hartree-Fock or Woods-Saxon potentials. It mainly builds on the mean-field picture of the nucleus and allows for a tractable treatment of complex nuclear correlations, particularly in medium and heavy nuclei, where fully microscopic methods like the nuclear shell model become computationally infeasible. The calculation relies heavily on proton-neutron interactions, parameterized by the constant g_{pp} which is the proton-neutron pairing interaction strength. The value of this parameter is determined from a fit of the model to the experimentally measured $2\nu\beta\beta$ rate, and this fitted value is then used to determine the $0\nu\beta\beta$ rate.
- **Energy Density Functional (EDF)** methods calculate nucleon interactions and properties using an energy density functional obtained from a standard nucleus. EDF methods generalize the idea of mean-field models and they also allow for a self-consistent description of a nucleus using functionals of densities, just like

density functional theory (DFT) in condensed matter or atomic physics. The functional comes from minimizing the Hamiltonian over local density functions (nucleon, spin, and current numbers, among others), with constraints from measurement. This provides a way to include complex correlations such as nuclear deformation, pairing, configuration mixing, and collective dynamics, which are essential for a reliable NME prediction in heavy, realistic nuclei. These models describe medium and heavy nuclei well, but produce some of the largest NME values, since they don't explicitly include proton-neutron pair correlations.

- The **Interacting Boson Model (IBM)** simplifies the description of medium and heavy nuclei by treating pairs of nucleons (proton-proton, neutron-neutron, or proton-neutron) as bosons. IBM provides a phenomenological but effective framework for calculating NMEs that are otherwise computationally challenging to obtain in fully microscopic models. Some of these approaches treat protons and neutrons identically, while others treat them separately. Specifically, the IBM has been extended to compute NMEs for double beta decay in its IBM-2 version, which distinguishes between proton and neutron bosons. Since these models require pairs of nucleons, they only apply to even-even nuclei (a fine application to double-β decay.)

- **Ab initio** (meaning from the beginning) methods use truly first-principles calculation of nuclei and involve solving the underlying theory, quantum chromodynamics (QCD), with quarks and gluon degrees of freedom. In contrast to other many-body approaches which have some phenomenological components, ab initio methods treat all nucleons in the nucleus and use unadjusted nuclear Hamiltonians that are mostly derived from chiral effective field theory. Moreover, the calculations are systematically improvable and the convergence of the results can be checked explicitly. Ab initio calculations reproduce the properties of nu-

> clei well in light and medium nuclei and some recent attempts have been made applying these calculations to heavier $0\nu\beta\beta$ nuclei. However, more theoretical work in this area is needed to improve the accuracy and applicability of these methods for heavy nuclei [37].

To give some context as to how the available NME calculations compare with each other, a summary of recent calculations for various $0\nu\beta\beta$ isotopes is shown in Figure 3·4. The spread in the magnitude of the NMEs varies by at least factors of 2-3.

3.4 Experimental Search for $0\nu\beta\beta$ Decay

As previously mentioned, the two electrons from $0\nu\beta\beta$ decay will have exactly the endpoint energy $Q_{\beta\beta}$ and will form a sharp peak at the end of the $2\nu\beta\beta$ energy distribution. Theoretically, the $0\nu\beta\beta$ peak should resemble a delta function positioned at $Q_{\beta\beta}$. However, the exact width of the measured signal, and its separation from the falling tail of

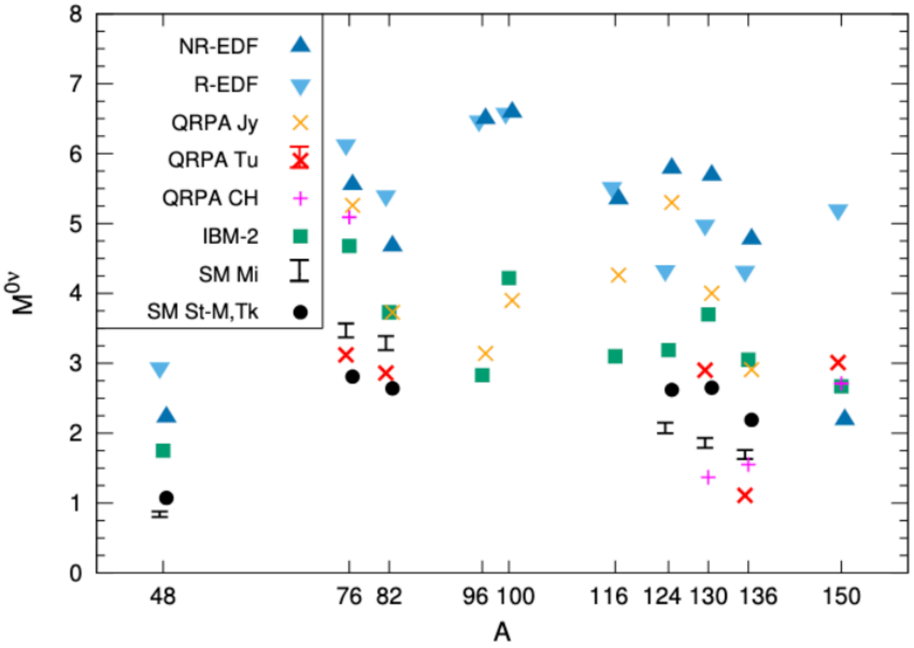

Figure 3·4: Nuclear matrix elements for $0\nu\beta\beta$ decay candidates as a function of mass number A, adapted from Reference [36].

the $2\nu\beta\beta$, depends on the detector energy resolution. This is illustrated in Figure 3·5 where the inset is zoomed into the region of interest (ROI) for the $0\nu]\beta\beta$ signal.

There are many naturally occurring even-even nuclei capable of undergoing double beta decay, but not all are equally well-suited for a $0\nu\beta\beta$ search. The selection of the double beta decay isotope is a critical aspect of detector design and significantly impacts the experiment's sensitivity to the $0\nu\beta\beta$ process. Choosing the optimal isotope involves considering several key factors. Firstly, a high Q-value is preferred in $0\nu\beta\beta$ searches because it enhances the decay rate through a larger phase-space factor, places the signal region above most natural radioactive backgrounds, and generally improves signal detection efficiency. Selecting a double beta decay isotope with high natural abundance is beneficial because it increases the number of target nuclei for a given mass, thereby enhancing sensitivity. While it's possible to enrich isotopes, this will typically increase the costs dramatically, a factor that can't be easily ignored. A relatively long $2\nu\beta\beta$ half-life improves the experiments sensitivity by reducing the $2\nu\beta\beta$ background. Also, calculations of the NMEs vary by isotope, influence the $0\nu\beta\beta$

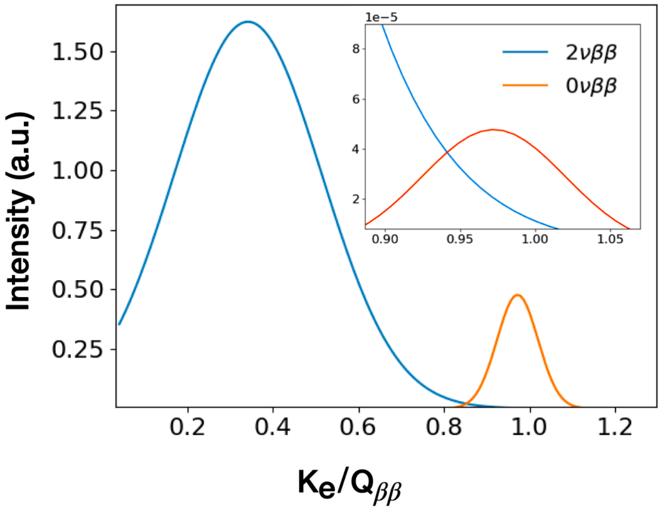

Figure 3·5: Illustration of the observable signal of $0\nu\beta\beta$ decay. The inset is zoomed into the ROI for $0\nu\beta\beta$ and gives a more accurate depiction of the $0\nu\beta\beta$ signal for a detector energy resolution of at least several percent or higher.

Isotope	Abundance (%)	$Q_{\beta\beta}$ (MeV)	$T^{2\nu}_{1/2}$ (10^{20} years)
^{48}Ca	0.187	4.268	$0.64^{+0.07}_{-0.06}$(stat)$^{+0.12}_{-0.09}$(syst)
^{76}Ge	7.75	2.039	20.22 ± 0.18(stat) ± 0.38(syst)
^{82}Se	8.82	2.998	0.869 ± 0.005(stat)$^{+0.009}_{-0.006}$(syst)
^{96}Zr	2.80	3.356	0.235 ± 0.014(stat) ± 0.016(syst)
^{100}Mo	9.744	3.034	0.0707 ± 0.0002(stat) ± 0.0011(syst)
^{116}Cd	7.512	2.814	$0.263^{+0.011}_{-0.012}$
^{130}Te	34.08	2.528	$8.76^{+0.09}_{-0.07}$(stat)$^{+0.14}_{-0.17}$(syst)
^{136}Xe	8.857	2.458	$21.6^{+6.2}_{-4.0}$(stat)$^{+4.0}_{-2.9}$(syst)
^{150}Nd	5.638	3.371	0.0934 ± 0.0022(stat)$^{+0.0062}_{-0.0060}$(syst)

Table 3.1: Summary of the most common $2\nu\beta\beta$-decay isotopes used to search for $0\nu\beta\beta$-decay. Abundances, $Q_{\beta\beta}$ values, and half-lives are taken from References [38] and [39]

half-life and observation prospects, with some isotopes offering both larger values and smaller theoretical uncertainties. Finally, technical limitations involving the instrumentation of the isotope in an experiment must also be considered. The most popular isotopes used for $0\nu\beta\beta$ experiments are summarized in Table 3.1.

When designing $0\nu\beta\beta$ experiments, we must also consider multiple sources in addition to the tail of the continuous $2\nu\beta\beta$ distribution. Natural radioactivity from materials both within and surrounding the detector, cosmic rays and their resulting radioactive spallation products, and other subtle particle physics processes could all mimic a $0\nu\beta\beta$ decay signal. As a result, a central challenge in designing a $0\nu\beta\beta$ experiment is maximizing the *signal-to-background ratio.*

Experiments searching for $0\nu\beta\beta$ decay are typically conducted deep underground to shield them from cosmic ray backgrounds, and they also prioritize the use of materials with exceptionally low levels of radioactive contamination. Smaller-scale (order of several hundred kilograms of isotope or less) experiments generally achieve lower background rates, superior energy resolution, and better spatial resolution, which aids in distinguishing signal from background. However, these detectors often find it challenging and very expensive to scale to multiple tons of $0\nu\beta\beta$ isotope. In contrast, while larger detectors can accommodate a tone or more of isotope, maintaining high energy resolution becomes more challenging. Additional factors such as safety and technical feasibility are crucial considerations in experimental design.

Experiments searching for $0\nu\beta\beta$ decay aim to measure the signal rate, which in turn provides a measure of the half-life. In the absence of any signal in the data, a limit on the half-life is obtained instead. To describe an experiment's sensitivity to the $0\nu\beta\beta$ half-life we start with the very familiar form of the radioactive decay law:

$$N(t) = N_0 e^{-\ln(2)\, t / T^{0\nu}_{1/2}} \tag{3.11}$$

where $N(t)$ is the number of nuclei remaining at a given time t, and N_0 is the initial number of $0\nu\beta\beta$ nuclei. Equation 3.11 can be written in terms of the number of decays $N_{\beta\beta}$ detected with efficiency ϵ:

$$\frac{N_{\beta\beta}}{\epsilon} = N_0 \left(1 - e^{-\ln(2)\, t / T^{0\nu}_{1/2}}\right) \tag{3.12}$$

Solving for the exponential gives:

$$\ln\left(\frac{N_0 - N_{\beta\beta}/\epsilon}{N_0}\right) = -\ln(2) \frac{t}{T^{0\nu}_{1/2}} \tag{3.13}$$

Using the approximation $\ln(1-x) \approx -x$ for small x, we arrive at the $0\nu\beta\beta$ half-life in terms of the number of the number of double-beta decay nuclei N, the number of $0\nu\beta\beta$ events detected $N_{\beta\beta}$, the livetime t, and the detection efficiency ϵ:

$$T^{0\nu}_{1/2} = \frac{\ln(2)\, t \epsilon N}{N_{\beta\beta}} \tag{3.14}$$

The number of double-beta decay nuclei N can be written more explicitly in terms of the following parameters:

$$N = \left(\frac{xaN_AM}{A}\right) \tag{3.15}$$

where x is the fraction of the total mass M made of the target element containing $0\nu\beta\beta$ isotope, a is the percent abundance of the isotope within the target element, N_A is Avogadro's number, and A is the molar mass. For example, suppose an experiment uses a 780 tons of liquid scintillator with 3.9 tons of natural tellurium, which has a 34% natural abundance of ^{130}Te. Then we would have the following: $M = 780\,\mathrm{kg}$, $x = 3.9\,\mathrm{tons}/780.0\,\mathrm{tons} = 0.005$, and $a = 0.34$.

Successful experiments maximize their sensitivity to $T^{0\nu}_{1/2}$. In the case where no clear signal exists in the data, we consider the maximum number of possible decays n_b hiding under the background:

$$n_b = \sigma_n\sqrt{tb\Delta_E M} \tag{3.16}$$

where b is the background rate, t is the exposure time, M is the total mass of the target material, Δ_E is the width of the ROI, and σ_n is the n-σ confidence level at which we believe the decays are hiding under the background.

Combining Equations 9.6, 3.15, and 3.16, we can write down the experimental sensitivity to the $0\nu\beta\beta$ half-life:

$$S^{0\nu}_{1/2} = \ln(2)\epsilon\frac{xaN_A}{A}\sqrt{\frac{Mt}{b\Delta_E}} \tag{3.17}$$

We can see that Equation (3.17) prescribes an ideal experiment: one with a minimal background rate b and fine energy resolution (small Δ_E), while maximizing the efficiency ϵ, live time t, and target material loading $(M, x$, and $a)$.

Chapter 4

The SNO+ Experiment

The Sudbury Neutrino Observatory (SNO) was constructed back in the late 1990s to measure solar neutrinos. In doing so, the experiment resolved the solar neutrino problem and confirmed for the first time ever that neutrinos must have mass. For this ground-breaking discovery, SNO, and the Super-Kamiokande experiment in Japan, won the 2015 Nobel Prize in Physics [4]. The experiment collected data until late 2006, and soon after it was shut down to begin major detector upgrades to enhance its sensitivity to new physics.

In May 2017, a series of initial upgrades were completed and the experiment began taking data while filled with ultra-pure water. Another major milestone was achieved after additional upgrades allowed the detector to be filled with pure liquid scintillator, followed by a period of low-background data acquisition. At the time of this thesis, the project is approaching it's final upgrade that will allow for a high-sensitivity search for $0\nu\beta\beta$ with ^{130}Te. Dubbed 'SNO+', this next-generation liquid scintillator detector supports a broad low-background physics program that aims to set new limits on $0\nu\beta\beta$ decay. This chapter details the major upgrades and detector components, along with each phase of detector operation.

4.1 Detector

SNO+ is located in Sudbury, Ontario, 2 km underground in SNOLAB. At this depth (6010 meters water equivalent), SNO+ has a cosmic μ^- rate of roughly four per hour, the lowest among its fellow detectors abroad[40]. SNOLAB, which houses several other neutrino and dark-matter experiments, also shares space with Vale's active Creighton nickel mine. To mitigate contamination from its surroundings and meet the experiment's strict radio-purity needs, the entire lab was built as a self-contained class-2000 clean room.

The detector target medium fills a 5.5-cm-thick spherical acrylic vessel (AV), which measures 6.0 m in diameter. A spherical array of 9,362 photomultiplier tubes (PMTs) with light concentrators are positioned ~ 2.35 m from the surface of the AV, pointing inward. Each PMT is mounted to a concentrically situated geodesic PMT support structure (PSUP) made out of stainless steel. The PSUP also contains 91, outward-facing PMTs without light concentrators, which help veto sources of light outside the detector. Much of the outside light comes from muons or ambient radiation from the surrounding rock. SNO+ has retained most of the original 8-inch Hamamatsu R1408 PMTs from the SNO era. As part of the upgrade, over 600 of these were repaired and reinstalled, while some were permanently removed to make room for the hold-down ropes. Finally, the entire cavity containing the detector is filled with roughly 7,000 metric tons of ultra pure water, which provides addtional shielding from trace radiation from the surrounding environment and PMTs. A diagram of the SNO+ detector is shown in Figure 4·1.

The original acrylic vessel (AV) from the SNO detector was built from 122 ultraviolet-transmitting acrylic panels and remains in place today. As mentioned above, the panels are nominally 5.5 cm thick, with the exception of ten panels around the equator being 11 cm thick to accommodate grooves for the supporting ropes. All the panels are glued together to form the 6-meter-diameter sphere[41]. This choice of acrylic has two advantages: it's easily cleansed of radiation, and its index of refraction matches that of the PMT faces.

Hold-up ropes once cradled the AV when it was filled with heavy water during

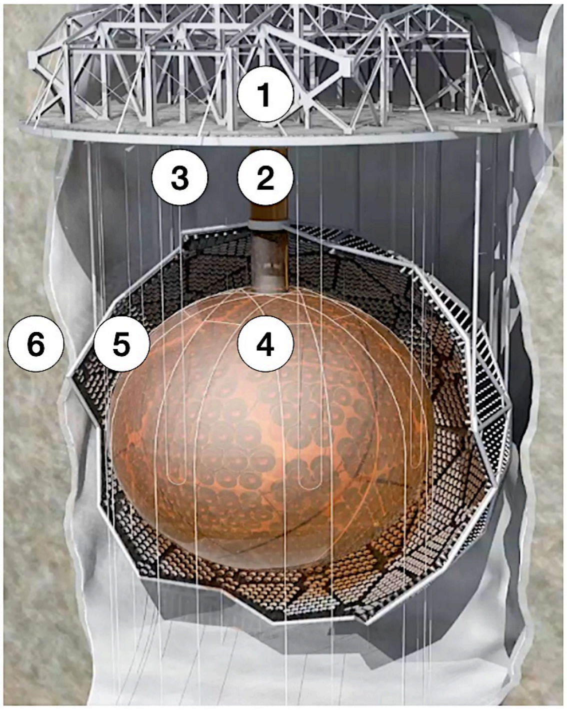

Figure 4·1: Diagram of the SNO+ detector's main structural components showing 1) the deck, 2) neck, 3) support ropes, 4) AV, 5) PSUP, and 6) rock cavern.

the SNO experiment, and they have since been upgraded to higher-purity Tensylon ropes[40]. These are 19 mm in diameter and hang from the cavity ceiling, where they suspend the AV by grooves in its equator. The hold-down ropes are a special addition to SNO+ and anchor the now-buoyant AV to the cavity floor (density of scintillator is $\sim 863\ \mathrm{kg/m^3}$ which is less than the density of water). The hold-down ropes are 38 mm in diameter and wrap around the top of the AV, at the base of its neck. The AV-rope system is illustrated in Figure 4·2.

This neck is a 7.0-meter-high tube with an inner diameter of 1.5 m. It is made of ultraviolet-absorbing acrylic and connects to the AV at a joint called the neck boss. The choice of ultraviolet absorbing acrylic was to mitigate the funneling unwanted light from the neck and down into the AV. The whole structure reaches to the level of the deck, where lab personnel can deploy internal calibration sources. This interface also

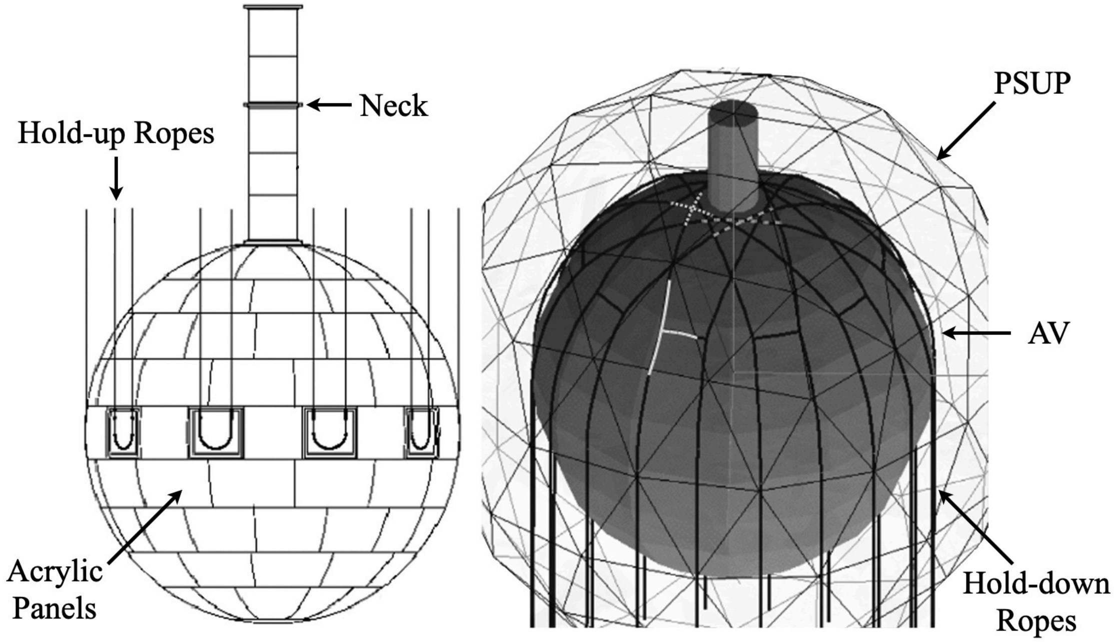

Figure 4·2: Left: The original SNO design and placement of the AV and hold-up ropes, adapted from Reference [41]. Right: Diagram of the PSUP, the AV (solid dark gray), and the hold-down ropes designed for SNO+ (thick black lines), adapted from Reference [40].

contains a system of pipes, which allow the loading and recirculation of detector media. A constantly flushed blanket of high-purity nitrogen seals the neck and cavern to protect each from radon.

The transitions between different detector media filled into the AV caused two artifacts. During the first transition, the AV was drained of heavy water, sitting exposed to air for several years before filling with ultra-pure water in 2017. Even with thorough cleaning before ultra-pure water filling, radon daughters remain embedded in the acrylic and constitute a major source of backgrounds, discussed in detail in Chapter 7. During the second transition in 2019, the AV started filling with pure liquid scintillator. This caused the AV's buoyancy to vary until the fill completed in 2021, causing an average vertical offset of 185.0 mm with respect to the PSUP. Both of these effects are monitored and included in simulation and event reconstruction described in Chapter 6.

4.2 PMTs and Data Acquisition Electronics

When particles interact inside the SNO+ detector medium, they create small bursts of Cherenkov and/or scintillation light. These optical photons are eventually detected by the PMTs, which have single-photon detection sensitivitiy. The PMTs convert the detected photons to an analog signal through the photoelectric effect. They generally consist of an anode at the base, and an evacuated glass dome which houses a grounded photocathode and a stack of dynodes. A positive high voltage (typically in the range of 1.6 – 2.3 kV) is applied to the anode, creating a strong electric field. Incident photons knock electrons from the photocathode. These electrons accelerate toward the first dynode, where they knock out secondary electrons in a cascade down the rest of the dynode stack. The number of electrons is amplified at each dynode stage until it reaches a factor of $\sim 10^7$. The resulting avalanche of electrons produces an analog signal that allows for two measured observables to be used in event reconstruction: charge and time.

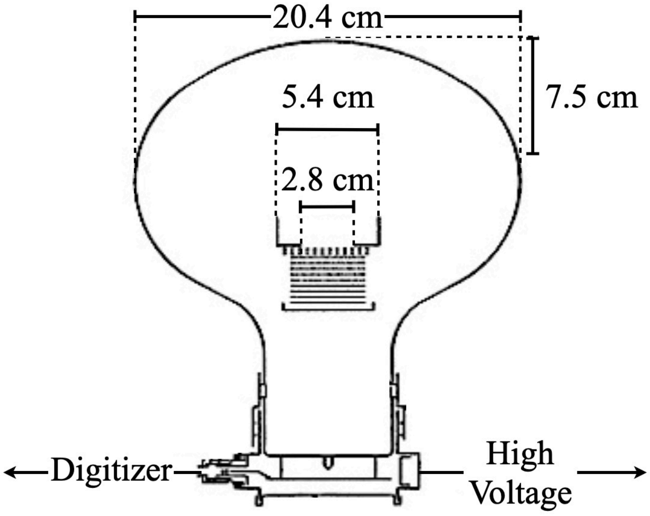

Figure 4·3: Diagram of the SNO Hamamatsu R 1408 PMTs adapted from Reference [41].

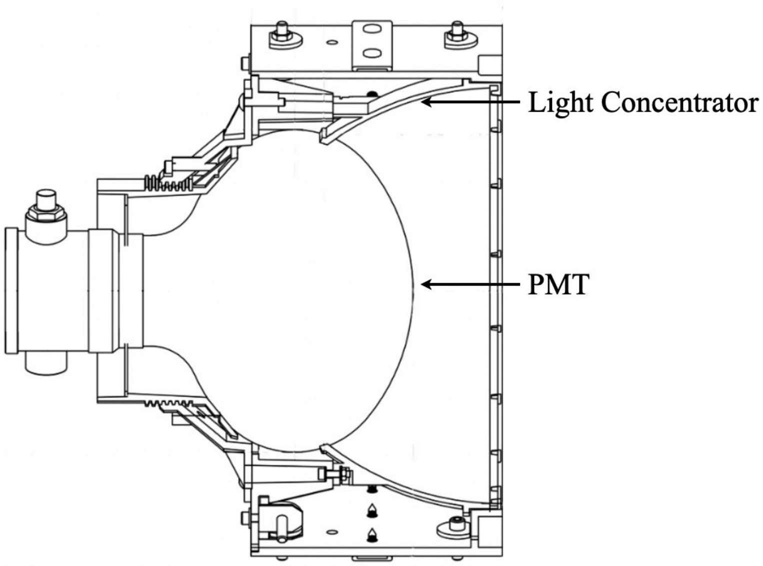

Figure 4·4: Diagram of a PMT shown with its casing and light-concentrator. Figure from [42].

Each inward-looking PMT sits in plastic casing, surrounded by a Winston cone (Figure 4·4). The Winston cones are concave arrangements of 18 curved, dielectric-coated aluminum sheets, which direct off-axis light toward the PMT faces. This dielectric coating contains a first layer of specular aluminum and a protective layer of titanium dioxide and praseodymium(III) oxide. The materials and thicknesses of these layers were chosen to maximize reflectance in water for the visible spectrum at several angles of incidence. The Winston cones enhance the overall photon collection efficiency resulting in a total photocoverage of 54%.

Calibration of each PMT requires knowledge of three intrinsic qualities: dark noise, wavelength efficiency, and transit time. Dark noise refers to the ambient amount of current generated at the anode without an incident photon. The SNO+ custom data acquisition electronics contain discriminators, which remove this current by only recording PMT signals with integrated charge above a certain threshold. The discriminator threshold is set to ~ 0.25 of the average photoelectron charge and is ~ 75% efficient in collecting only single-photoelectron signals.

Not every photon that hits the photocathode will produce a signal at the anode. PMTs have quantum and collection efficiencies that vary with wavelength. Quantum efficiency describes the initial photoelectron production and is the ratio of primary photoelectrons to incident photons. Collection efficiency is the fraction of these primary photoelectrons that reach the first dynode.

The path between photon incidence and production of maximum current at the anode determines the "transit time" of the PMT. Photoelectrons traversing a perfect path from photocathode to anode create the most intense prompt peak in the transit-time spectrum as illustrated in Figure 4·5. There are three other possible paths, though. Primary photoelectrons can originate from the *first dynode* when photons transmit through the photocathode. This path produces a subdominant peak $\sim 15\,\text{ns}$ earlier than the main peak. The other two possible paths are taken by photoelectrons that backscatter off the first dynode. In a single back scatter, the photoelectron scatters elastically, contributing to the electron cascade only when returned to the dynode by the electric field $\sim 10-40\,\text{ns}$ later. In a double backscatter, the photoelectron scatters in-elastically, contributing to the electron cascade twice. Its first collision contributes to the main peak, while its second collision contributes to the latest and broadest peak.

Measuring a photon-producing event inside SNO+ involves capturing light from as far away as 12.0 meters and converting it to a location and time. This requires several layers of electronics- one to supply voltage to the PMTs and others to process, discriminate, and record analog signals with high precision. SNO+ inherited most of its electronics from SNO, with the exception of a few upgrades made to accommodate an increased event rate in scintillator.

A waterproof coaxial cable connects each PMT to a PMT interface card (PMTIC), where it receives its high voltage, and to a front end card (FEC), where it outputs an analog signal for processing. The FEC has a discriminator, which digitizes the pulse and determines whether the amplitude crosses the discriminator threshold (again, this

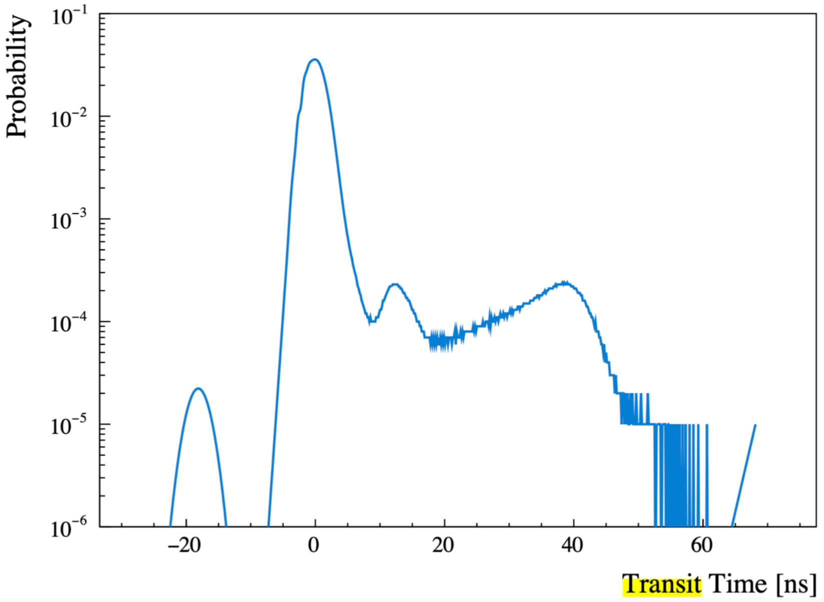

Figure 4·5: Transit time spectrum for a simulated R1408 SNO+ PMT adopted from Reference [43].

is set to ~ 0.25 the average PMT charge from a single photoelectron). If it does, the PMT is considered "hit", and the FEC goes on to calculate the hit time and total charge on the PMT.

All 9,362 PMT-PMTIC-FEC connections are divided among 19 crates, which sit on the detector deck. The discriminator pulses from the hit PMTs are summed for each crate, and these sums are passed to seven Master Trigger Card/Analog (MTC/A) boards. Each MTC/A board has its own pre-determined threshold and fires if the analog sum of the crate sums exceeds that threshold. Each of these thresholds correspond to unique trigger settings and are a helpful input to a physics analysis. For example, an analog sum that pass the NHIT100 trigger has an amplitude proportional to the number of PMTs hit within 100 nanoseconds.

The triggered pulses from the MTC/A boards reach the master trigger card/digital (MTC/D) which triggers the detector. This global trigger prompts the FECs to record all PMT data between 180 nanoseconds prior to and 220 nanoseconds after the global trigger. This 400-nanosecond period is called the "event window", and all PMT information in this window describes a single "event". The time to analog converter (TAC) gives the time between a PMT hit and its PMTIC's receipt of the global trigger. This is the relative time the PMT fired within the event window and is an input to event reconstruction. The total charge on each PMT also contributes to event reconstruction, and can be determined through short or long integration of the PMT pulse at high gain (QHS and QHL, respectively), or long integration of the PMT pulse at low gain (QLX).

Individual MTC/A boards can be masked in or out of the MTC/D requirements for a global trigger. The MTC/D can even be forced issue a global trigger during calibration campaigns. Each triggered event has a global trigger ID, and all are organized into hour-long runs. Lab personnel monitor detector conditions to assess data quality. Run type, material recirculation, and deck activity are some of the conditions noted, and only data taken under ideal running conditions are used for physics analyses.

4.3 Calibration

Reconstructing an event's time and energy from PMT information requires a thorough, and accurate, model of the detector response. Calibration campaigns provide model constraints and quantify features and artifacts. One such artifact is called the *time walk effect*, which refers to timing offsets between PMTs. These offsets arise from pulses traversing varying cable lengths and from low-amplitude pulses, which take the longest time to cross the discriminator threshold.

PMT-specific qualities affect signal efficiency. As discussed in Section 4.2, the probability of detecting a photon of a given wavelength depends on the PMT's quantum and collection efficiencies, which are known quantities provided by the manufacturer. Additionally, the signal efficiency depends on the discriminator threshold and resulting channel efficiency, as well as the angular response of the PMT and its light concentrator.

Other optical and scintillator quantities also affect light collection and event reconstruction. Scattering and absorption lengths of the scintillator, the AV, and the external water all determine how light propagates through the detector. Calculation of the time of flight of a photon from its origin to a PMT requires knowledge of the reflectivity of the Winston cones, and of the differences in refractive indices between scintillator, acrylic, and water. Energy reconstruction depends on the number of photons emitted by the scintillator per 1 MeV of deposited energy (the "light yield"). Finally, position reconstruction and particle identification depend on the timing response of the scintillator.

Most importantly, many of these quantities can be measured from the bench top, but change over time. For example, the reflectivity of the Winston cones and the angular response of the PMTs were both measured prior to their deployment in SNO. Over the detector's life, however, these components deteriorated in the ultra-pure water, changing the PMT response. In-situ calibrations provide the most realistic and current PMT response parameters and allow regular monitoring of scintillator stability.

4.3.1 Radioactive Sources

Deployed radioactive sources provide in situ energy calibration. By toggling the radiation type and energy (table 4.1), a deployed source calibration can determine the energy scale and resolution, as well as systematic uncertainty in reconstruction. During a deployment, a system of ropes lowers a sealed radioactive source down the neck and into the AV, where emitted particles deposit their energy in the detector media. This deposition produces light, which strikes the PMTs in its path, producing a signal for readout.

^{16}N High-Energy Gamma Source

When SNO+ was filled with ultra-pure water, energy calibration proceeded through a sealed^{16}N source. The β decay of ^{16}N produces a 6.1-MeV γ. During a source deployment, The β remained in the source container, while the γ escaped to the detector material. The β is first detected by a thin piece of plastic scintillator coupled to a calibration PMT. This prompt signal is used to tag the delayed 6.1-MeV γ. The ^{16}N source has been discontinued, as it (and β and α sources in general) cannot be sealed thoroughly enough to meet low-contamination requirements while also providing a source of radiation in the energy range relevant to the scintillator phase's physics goals[40].

Americium-Beryllium Neutron Source

The Americium-Beryllium (AmBe) source is a neutron and γ-emitter inherited from SNO. A Delrin-encapsulated stainless steel container holds a mixture of ^{241}Am and ^{9}Be powders. The α decay of ^{241}Am ($T_{1/2} = 432\,\mathrm{y}$) accompanies the emission of a 59.5 keV γ (branching ratio of 84.6%). ^{9}Be captures $\mathcal{O}(10^{-2})$% of these α particles to produce a neutron and the excited state ^{12}C*, which relaxes to its ground state and emits a 4.44 MeV γ. Within the next $\sim 200\,\mu$s , the neutron then thermally captures on H to produce a 2.2 MeV γ. This can be depicted in the following way:

$$\begin{aligned} {}^{9}\mathrm{Be} + \alpha \rightarrow n + {}^{12}\mathrm{C}^{*} \\ {}^{12}\mathrm{C}^{*} \rightarrow {}^{12}\mathrm{C} + \gamma \end{aligned} \tag{4.1}$$

The 4.44 MeV and 2.2 MeV γ's can be coincidence-tagged, and provide a high-energy calibration point, in addition to a measurement of neutron detection efficiency. This is important for physics studies involving antineutrinos that interact through inverse beta decay.

^{46}Sc Low-Energy Gamma Source

The ROI for $0\nu\beta\beta$-decay is centered around the Q-value of ^{130}Te at 2.528 MeV bounded by the on the left and right by $(-0.5\sigma, 1.5\sigma)$ of the hypothetical $0\nu\beta\beta$ peak. This highly asymmetric range was chosen to reduce as much of the $2\nu\beta\beta$-decay contamination as possible, but this choice depends heavily on accurate modeling of the $2\nu\beta\beta$ decay distribution. An increase of just 1% in the high-energy tail of the $2\nu\beta\beta$ spectrum could result in a 5% systematic error in $T^{0\nu}_{1/2}$[40].

The ^{16}N tagged source produces gammas well above the energies needed for a calibration near the $0\nu\beta\beta$ ROI. Another tagged source can be created by bombarding ^{45}Sc with neutrons to an activity of ~ 200 Bq. The resulting ^{46}Sc has a half-life of 83.39 d, and emits one β (111.8 keV) and two γ's (889.3 keV and 1120 keV). The β's have too low energy to reach the scintillator, and do so only 0.002% of the time. This presents minimal background to the γ's, which deposit all of their energy in the scintillator (98.5 ± 0.5)% of the time. The construction and principle of operation of this source are the same as for ^{16}N. Mainly, an enclosed PMT coupled to plastic scintillator detects the β's, providing a prompt tag for the γ's that escape to the scintillator.

Other Radioactive Sources

Intrinsic radioactivity can also provide a means of energy calibration. Radon daughters already present in the detector such as the ^{210}Po α, and the coincident decays of ^{212}Bi-Po and ^{214}Bi-Po, are all well-known and efficiently tagged, making them useful in monitoring energy scale and resolution over time. Other untagged gamma emitters are also being considered for future calibrations and are listed in Table 4.1.

4.3.2 Optical Calibration Sources

Optical calibrations inject the detector with photons of known wavelength, allowing the measurement of the target material's optical properties. Many of these properties,

Source	Particle	Energy (MeV)	Calibration
AmBe	n, γ	4.4 (2.2, 7.6, 9.0)	neutron capture response
^{16}N	γ	6.1	energy reconstruction
^{46}Sc	γ	0.9, 1.1	energy reconstruction
^{48}Sc	γ	1.0, 1.2, 1.3	energy reconstruction
^{57}Co	γ	0.122	energy reconstruction
^{137}Cs	γ	0.7	energy reconstruction

Table 4.1: List of SNO+ radioactive sources which have either been deployed or are being considered for deployment.

such as attenuation length, scattering, and absorption, change over time and need to be monitored in-situ. Asynchronous control of the optical calibration sources also allows the measurement of PMT timing offsets and overall response.

Laserball

Calibrations of SNO relied heavily on the laserball optical source, and it remained in use during the pure-water phase of SNO+ [42]. In this device, a bundle of optical fibers shuttles light from a nitrogen-pumped dye laser to a quartz sphere as illustrated in Figure 4·6. Here, the light diffuses isotropically through glass microbeads suspended in silicone gel. Just upstream from the laser's dye cells, a beam splitter diverts some of the light to a photodiode. A constant fraction discriminator processes the resulting signal into a pulse for the data acquisition system. These triggers serve as a clock against which absolute time delays and PMT time-walk can be measured. Calibrations with the laser ball source typically last 2.5 hours, during which the source delivers pulses at frequencies of 10 – 40 Hz in several locations inside the AV. Positioning the laserball at the center of the AV allows measurement of timing offsets between the PMTs, while the intensity of the pulses guarantees that they operate in the single-photoelectron emission mode.

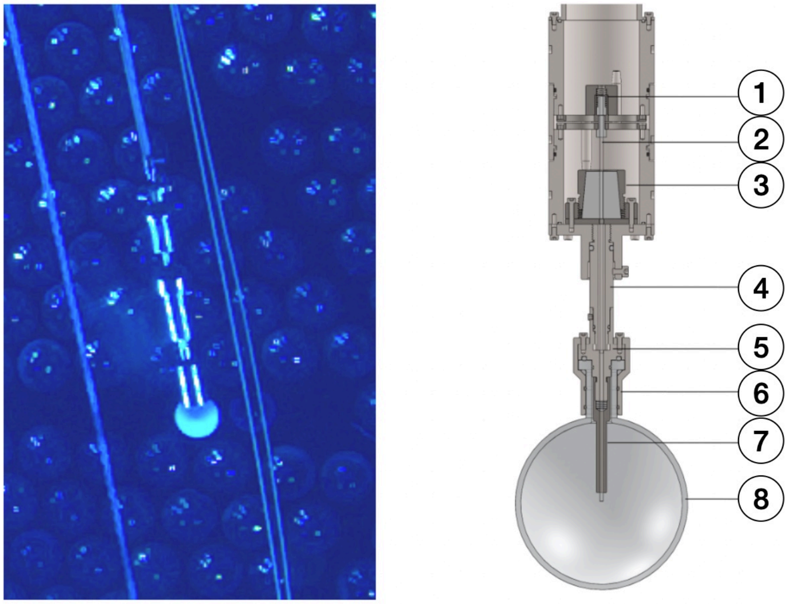

Figure 4·6: Left: the laserball source during deployment in the SNO+ water phase. Right: diagram of the upgraded laserball source. 1) cable fiber lock, 2) quartz rod light guide, 3) PTFE clamps, 4) neck tube, 5) neck bung, 6) neck clamp, 7) quartz support extension, 8) diffuser flask, filled with silicone gel and hollow glass microbeads. Figure adapted from Reference [40].

The laser ball has since been upgraded for the scintillator phase of SNO+. It has a slimmer neck, which reduces self-shadowing, and a movable quartz rod, which allows an adjustable light injection point. The latter feature improves the polar isotropy, and can be optimized, along with microbead concentration, to produce a quasi-isotropic light source with low temporal dispersion.

ELLIE

Deployed sources like the laser ball present an increased risk of contamination in SNO+. To mitigate this, SNO+ has been retrofitted with the Embedded LED/Laser Light Injection Entity (ELLIE). This system consists of optical fiber ends permanently affixed to the PSUP, and it operates under three modules, described here. A diagram of the system is shown in Figure 4·7.

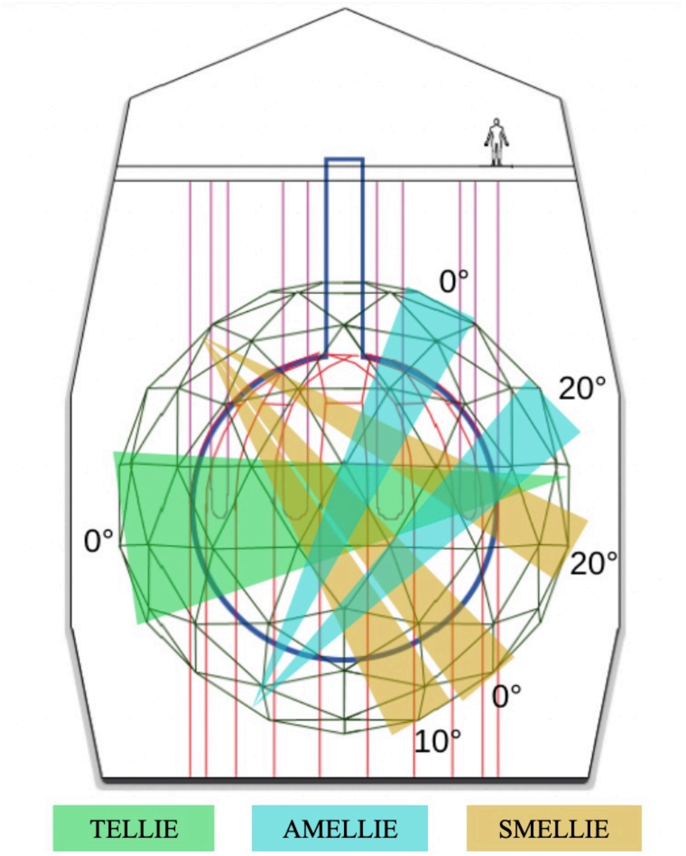

Figure 4·7: Diagram of the ELLIE system, adapted from Reference [40], showing a few beams from each ELLIE configuration. TELLIE points wide-angle light toward the center of the detector, while AMELLIE and SMELLIE point collimated beams at several angles relative to the center of the AV.

TELLIE

The timing module (TELLIE) injects wide beams of one wavelength from several points throughout the PSUP. It enables measurement of the gain, charge response, and timing offset of the PMTs. Accurate calibration of the detector timing affects the accuracy of photon time-of-flight, an important input for an event's position reconstruction. TELLIE's first job is to synchronize the PMTs' response to a global generation of a single photoelectron. Several factors mitigate this synchronization, like delays in the electronics between the readout and trigger system. TELLIE must determine and suppress these delays to within a nanosecond or better, since other irreducible timing uncertainties remain. The transit time spread of the PMTs ($1\sigma = 1.7$ ns) and the decay time of the scintillator, described in Section 5.2, contribute $\mathcal{O}(1-10)$ ns to the overall timing uncertainty.

TELLIE's second job is to calibrate the charge response and gain of the PMTs, both inputs to energy reconstruction. These need monitoring, because the discriminator threshold is set at ~ 25% of the average peak voltage for single photoelectron pulses. So, any change in the gain of a PMT changes the photoelectron detection efficiency of its channel.

To accomplish these tasks, TELLIE has 92 injection points distributed throughout the PSUP, each pointing toward the center of the AV. Duplex PMMA step-index optical fibers at each point inject light from an LED source. This light is 500 nm in wavelength, chosen from the range in which Te-loaded scintillator (described in Chapter 5) is most transparent, and where the PMTs are most quantum-efficient. Each fiber has a numerical aperture of 0.5 and provides a beam with a half-opening angle of 22° (measured in water for 450 nm), allowing thorough coverage. Each 45-meter-long fiber connects to a patch panel on the detector deck, where an additional two-meter long fiber connects it to to a driver box containing the LEDs. These LEDs deliver several optical pulses of width ~4 ns (FWHM), with diminishing statistical uncertainty. Under these conditions, TELLIE achieved a calibration of $\mathcal{O}(0.1)$ ns for each PMT.

SMELLIE

The scattering module (SMELLIE) injects collimated beams of several wavelengths from 15 locations on the PSUP. This calibration source measures the scattering properties of the detector media. Optical scattering of visible photons proceeds mainly through Rayleigh scattering, which has a clear and well-defined dependence on scattering angle and wavelength. SMELLIE exploits this using partially collimated beams of five wavelengths. Partial collimation precisely defines the initial angle of the photon, while still producing enough hit PMTs. Since scattered photons travel farther and longer than re-emitted or refracted photons, SMELLIE picks out genuine photon scatters.

Short-pulsed lasers inject light from five nodes. At each node, each of three Corning InfiniCor SXi telecommunications fibers sits at $-20°$, $0°$, $10°$ relative to the center of the AV. This configuration yields three sets of path lengths through which SMELLIE can separate the scattering properties of the external water from those of the detector media. The pulses themselves come from four diode lasers, and one super-continuum laser. The diode lasers produce fast ($50-200$ ps) pulses with wavelengths of 375, 405, 440, and 495 nm (FWHM in the range of $1.5-4$ nm). The super-continuum laser produces longer pulses (~ 2 ns) with a 10-nm bandwidth in the visible range.

AMELLIE

A photon can deviate from its initial path via absorption or scattering. These tendencies make up a wavelength-dependent material property called attenuation. Since a photon's path determines position and energy reconstruction, SNO+ needs a calibration to closely monitor optical attenuation over time. AMELLIE accomplishes this through injecting collimated beams of several wavelengths at different angles. The basic principle uses the timing and positions of hit PMTs along the path of the beam. PMTs aligned with the beam register prompt, unattenuated light, while those just outside the path register light delayed by scattering, reflection from boundaries, or absorption and re-emission by the detector media. The angle between the prompt and delayed PMTs, which AMELLIE can monitor over time, implies the attenuation length.

AMELLIE employs LED sources coupled to step-index multimode fibers, which have a numerical aperture of 0.22 and produce highly collimated beams ($9.47°$ half-opening angle at 450 nm in water). Eight of these 45-meter long fibers sit in four locations on the PSUP, two toward the north and south poles, and two close to the equator.

At each location sits one fiber pointing at 0° and a second at 20° relative to the center of the AV, providing a variety of path lengths. AMELLIE cannot separate contributions to attenuation in the detector medium, but can measure changes over time and serve as a crosscheck for SMELLIE data.

4.3.3 Calibration Hardware

Calibration sources deployed inside the detector require support and connections to the deck clean room. A central rope carries the weight of the source, while a 30-m umbilical cable runs optical fibers, thin coaxial gas tubing, and four hook-up cables between the source and the deck clean room. The umbilical cable can transmit power, signals, light pulses, and gas feed.

The umbilical can move about the central vertical axis, using the umbilical retrieval mechanism (URM). This system contains storage for the umbilical, an umbilical driver, and a rope driver as illustrated in Figure 4·8. In the storage system, the umbilical winds around two sets of six co-axial pulleys, which can lower or retract it. A pneumatic piston pulls on the pulleys, maintaining a constant tension. The umbilical retrieval mechanism also allows the source to move off axis, using side-ropes. Sources can also be lowered between the AV and PSUP, through any of six calibration guide tubes as shown in Figure 4·9. Since this region of the detector has lower radio-purity requirements, these calibrations can occur during all phases with minimal disturbance to the target material. These external calibrations were especially relied upon while the detector was actively filling with liquid scintillator.

The water phase of data taking (section 3.4.1) used the original source deployment system from SNO, but required updates for deployment in SNO+. Specifically, the drivers on the inside of the umbilical retrieval system had to be replaced with scintillator-compatible materials (low radon emanation and high grip in the slippery scintillator).

4.4 SNO+ Target Phases and Physics Goals

The terminal goal of SNO+ is a search for $0\nu\beta\beta$-decay, achievable with a target of ^{130}Te-loaded liquid scintillator. The transition to this medium from a drained and decommissioned SNO happened in phases: ultra-pure water, pure scintillator, and later, ^{130}Te-loading. This section briefly outlines these phases below, leaving a detailed discussion of processing and deployment for Chapter 5.

4.4.1 Water Phase

After decommission in 2007, SNO was drained of heavy water, cleaned, and re-filled with ultra-pure water. SNO+ resumed data-taking as a water Cherenkov detector from May 2017 until July 2019.

Cherenkov radiation is a polarization effect that occurs when a relativistic particle passes through a dielectric material. As the particle travels, molecules in its path polarize and then return to their ground states, emitting photons. If these photons have

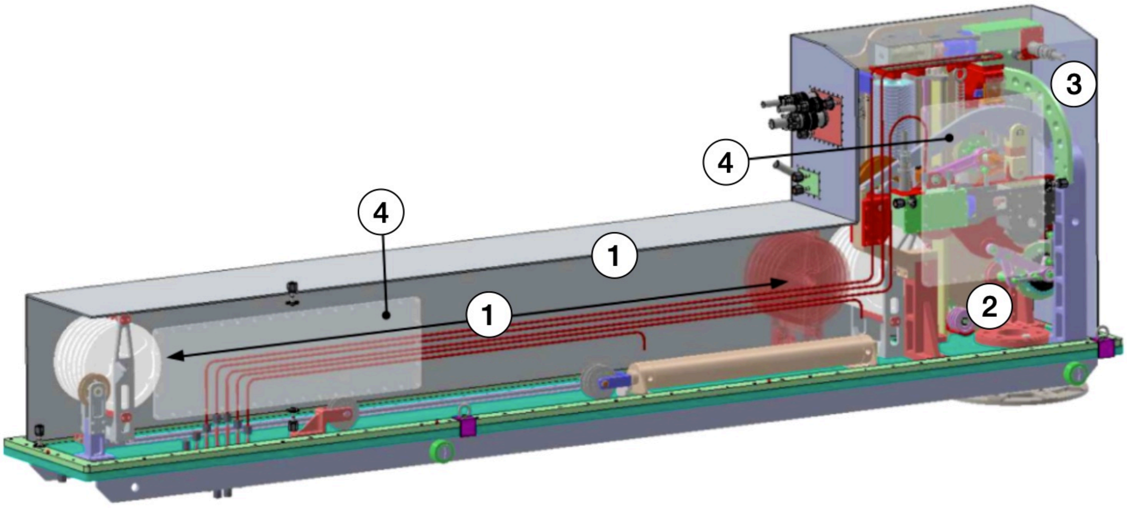

Figure 4·8: Diagram of the umbilical retrieval mechanism, adapted from Reference [40], showing 1) umbilical storage system, 2) drive system, 3) rope mechanism, and 4) acrylic viewing ports.

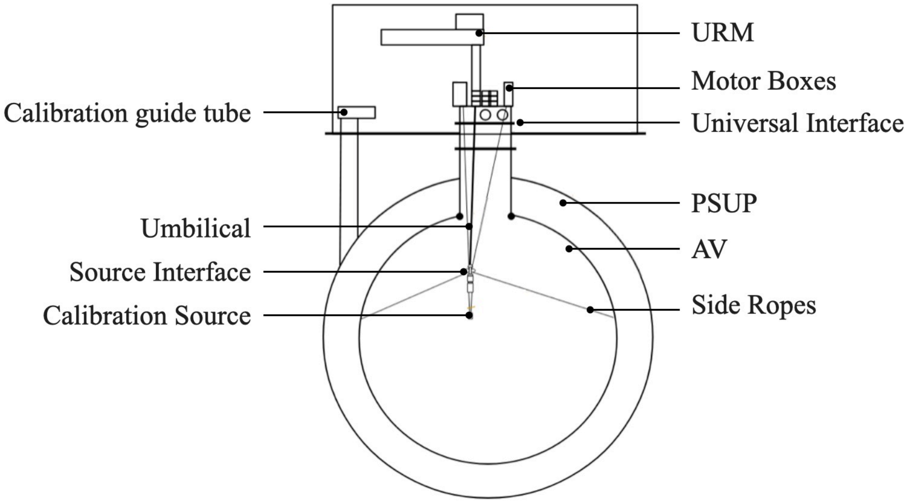

Figure 4·9: Figure adapted from Reference [40], showing the main components of the calibration source deployment system.

a slower velocity than the particle traveling through the medium, they emit in a cone along the particle's trajectory (similar to a sonic boom). This Cherenkov radiation has an opening angle θ , which is related to the particle's velocity v, the speed of light in vacuum c, and the material's wavelength-dependent index of refraction $n(\lambda)$:

$$\cos\theta = \frac{1}{n(\lambda)}\frac{c}{v} \tag{4.2}$$

A particle traveling fast enough produces Cherenkov radiation immediately, and the emission persists until the particle's speed v drops below c/n .

Two relations make Cherenkov radiation useful for particle identification and track reconstruction. The minimum kinetic energy (the "Cherenkov threshold") needed for a particle of mass m to produce Cherenkov radiation in a material with index of refraction n is:

$$E = m\left(\frac{1}{\sqrt{1-\frac{1}{n^2}}} - 1\right) \tag{4.3}$$

In water, the Cherenkov thresholds of the electron and muon differ by two orders of magnitude $\left(E_{e^-} = 0.26\,\text{MeV}, E_{\mu^-} = 55\,\text{MeV}\right)$. Moreover, calculation of the opening angle, along with measurement of the number of photons emitted N gives information about how long the track x is:

$$\frac{d^2N}{dxd\lambda} = \frac{2\pi\alpha Z^2}{\lambda^2}\sin^2\theta \tag{4.4}$$

Finally, since electrons tend to scatter more than muons, the shape of a trajectory reconstructed from Cherenkov radiation can also help distinguish between electrons and muons. This is a common technique used for particle ID in water Cherenkov detectors.

The water phase of SNO+ facilitated the search for invisible nucleon decay. Such a decay is "invisible" in the sense that, during this process, a nucleon decays into products that don't deposit energy in the detector. Since these products don't deposit energy, their presence has to be inferred. Ultra-pure water is well-suited to this task, since theoretically, neutron and proton decays in oxygen produce excited states that de-excite by emitting γ's in the following way:

$$\begin{aligned} {}^{16}\text{O} &\rightarrow \text{(invisible)} + {}^{15}\text{O}^* \\ {}^{15}\text{O}^* &\rightarrow {}^{15}\text{O} + \gamma \end{aligned} \tag{4.5}$$

$$\begin{aligned} {}^{16}\text{O} &\rightarrow \text{(invisible)} + {}^{15}\text{N}^* \\ {}^{15}\text{N}^* &\rightarrow {}^{15}\text{N} + \gamma \end{aligned} \tag{4.6}$$

Almost half the time, these de-excitation γ's deposit ~ 6 MeV of energy, well within the detector's sensitivity. The final analysis of 274.7 days of data produced the world's lowest limits on the proton and neutron lifetimes of 9.6×10^{29} and 9.0×10^{29} years, respectively [44]. It also set limits of 1.1×10^{29} and 6.0×10^{28} years on the pp and pn modes, each an improvement on limits set by Borexino and radiochemical experiments.

This period of data-taking also established the levels of backgrounds outside the target material, optical calibration of the external water, PMT response, and the state of the data acquisition system prior to scintillator deployment. During this time, the detector also showed a 50% neutron detection efficiency, the highest among water Cherenkov detectors. Additional physics results include the measurement of reactor antineutrinos, and of ^{8}B solar neutrinos down to 3.5 MeV [45, 46, 47, 48]. These recent measurements by SNO+ agree with earlier measurements made by other experiments and demonstrated the SNO+ detector's low intrinsic backgrounds and continued viability into its next phase of operation.

4.4.2 Scintillator Phase

Aromatic hydrocarbons scintillate through the de-excitation of electrons in a fully-conjugated π-system. Some of these materials are liquid at room temperature and can form the base of many neutrino detectors' targets, which are economically and safely scalable to hundreds of tons. In many ways, liquid scintillators are very clean, compared to water. Without ionic impurities, which are insoluble in non-polar materials like hydrocarbons, some liquid scintillators can achieve ^{238}U and ^{232}Th contamination as low as 10^{-18}g/g. Along with its high light yield, these qualities make liquid scintillator an ideal target for low-statistics analyses like a search for $0\nu\beta\beta$.

As SNO+ filled with liquid scintillator, described in more detail in the next Chapter, several analyses monitored intrinsic background levels, measured light yield over time, and completed a preliminary energy and timing calibrations. An image depicting the SNO+ detector partially filled with scintillator is shown in Figure 4·10. This period of data-taking also allowed for physics analyses. One of them demonstrated separation of Cherenkov and scintillation light at the event-level, taking advantage of the water and scintillator-filled halves of the target. This separation was able distinguish background events from ^{8}B solar neutrinos down to ~ 5 MeV [49]. The partial-fill period also saw the detector's first detection of reactor antineutrinos, thanks to efficient classification of the main background $^{13}\text{C}(\alpha, n)^{16}\text{O}$ [50].

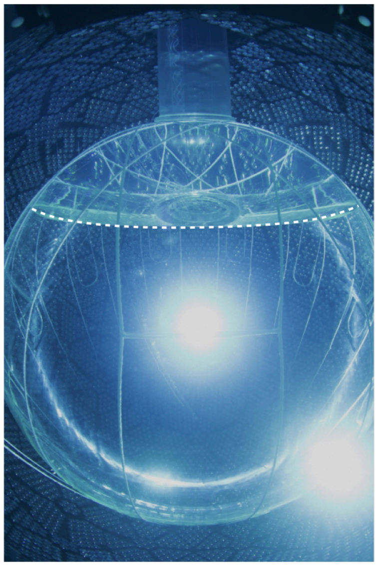

Figure 4·10: Image of the AV during the partial fill period, with the scintillator-water interface outlined by the white dashed line.

4.4.3 Tellurium-Loaded Scintillator Phase

The primary goal of the SNO+ experiment is to search for $0\nu\beta\beta$ in ^{130}Te. To enable this, the SNO+ collaboration has developed an innovative method for incorporating tellurium into the liquid scintillator, described in detail in Chapter 5. The tellurium phase of the experiment will begin once the scintillator has been loaded to a target concentration of 0.5% natural tellurium by mass, which corresponds to 3.9 tons of natural tellurium. About $\sim$ 34% of natural tellurium is ^{130}Te, the isotope that can undergo double-beta decay. To support SNO+ tellurium operations, two dedicated chemical processing facilities were constructed. The tellurium, initially acquired and transported underground in the form of telluric acid (TeA), will first be purified in a specialized TeA purification plant. Following purification, the TeA will undergo a condensation reaction with 1,2-butanediol (BD) in the BD synthesis plant, during which water produced in the reaction will be vaporized and removed. The resulting telluriumbutanediol (TeBD) complex will then be dissolved in linear alkylbenzene (LAB). The high

Q-value of ^{130}Te places the $0\nu\beta\beta$ region of interest outside of most of the low-energy backgrounds, but there remains contamination from ^{8}B solar neutrinos, daughters of the ^{238}U and ^{232}Th chains, and products from cosmic spallation of the tellurium.

Chapter 5

SNO+ Scintillator Upgrade

The primary goal of the SNO+ experiment is to perform a high-sensitivity search for $0\nu\beta\beta$ decay. To accomplish this task, the detector upgraded numerous subsystems that would support the delivery, purification, filling, and recirculation, of liquid scintillator into the AV. Following some initial data-taking with liquid scintillator in the AV, another dedicated system will allow for tellurium loading into the scintillator to enable the $0\nu\beta\beta$ search. The details of the scintillator filling campaign and the path towards tellurium loading is summarized here.

5.1 Scintillation Mechanism

A molecule's valence electrons dictate its chemical properties. In aromatic compounds, the electrons' wave functions overlap and delocalize around a ring of quasi-double bonds, which is depicted in Figure 5·1. After excitation by ionizing radiation, the electrons relax and isotropically emit photons over a time period of order tens of nanoseconds. This behavior is called *scintillation* and has been used in particle detection for decades. In large liquid detectors, the light emission should have rise times on the order of a nanosecond, or less, and decays times on the order of ten nanoseconds, or less. These timescales allow proper reconstruction of the location of an event and allows discrimination between signal and time-correlated backgrounds.

Aromatic hydrocarbons in particular have several other properties that make them ideal for low-background experiments like SNO+. First, aromatics are organic mean-

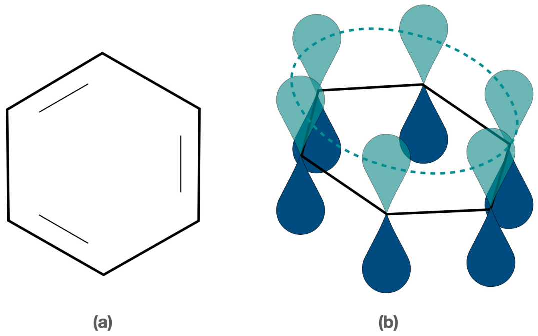

Figure 5·1: Benzene structures drawn with a) alternating single-double bonds, and b) with p-orbitals. Each vertex represents a CH bond. The p-orbitals in b) are bi-colored to represent the top and bottom sets of lobes, which overlap laterally (highlighted by the dotted line).

ing they consist of non-polar carbon-hydrogen bonds, making them relatively clean (eg., compared to water) in terms of ionic or metallic contaminants like potassium, uranium, and thorium. Second, they are generally non-reactive and thus safely scalable and deployable for many years. Finally, some aromatic hydrocarbons, such as linear alkylbenzene (LAB), are bases for common detergents, and thus cheap and readily available.

Scintillation can occur in organic and inorganic materials of all physical phases. In crystalline solids, like Tl-doped NaI, this proceeds in three steps: 1) absorption of energy and creation of electrons and holes (conversion), 2) migration of these electrons and holes toward dopant impurities or "luminescence centers" (transfer), and 3) the relaxation of these electron-hole pairs through radiative recombination (luminescence).

The process described above arises from the valence-conduction band structure of crystals. Organic scintillating materials, on the other hand, luminesce through excitation and relaxation of π-electrons. Like conduction-band electrons, π-electrons are also delocalized around the molecule. Ionizing radiation and electron scattering all excite π-electrons to higher-energy singlet states. The material fluoresces when these electrons relax directly to the ground state on a timescale of nanoseconds.

If excited to high enough energy, π-electrons can relax to a meta-stable triplet state via inter-system crossing. Triplet states brought into close enough contact annihilate. For example, in LAB, this close contact occurs through Van der Waals interactions between alkyl groups. That is, one molecule in a triplet state relaxes by transferring its energy to the second molecule, also in a triplet state. The re-excited molecule then relaxes directly to the ground singlet state, fluorescing $\mathcal{O}(10-100)$ ns later. Delayed fluorescence from triplet-triplet annihilation has an energy dependence. That is, the heavier the ionizing particle, the more delayed fluorescence it produces. Reconstruction algorithms can use this relationship to distinguish between types of particles. For example, the slow component in the timing profile of LAB is higher for α's (m_α =3.727 GeV) than for β's (m_β =0.511 MeV).

A scintillator must produce enough light for the detector to resolve the energies of incoming particles. It is important to note that only a fraction of scintillator de-excitations ("quantum efficiency") proceed by radiative emission. Moreover, the energy of the light emitted by the scintillator only makes up a small fraction of the energy originally deposited by the particle ("scintillator efficiency"). The quantity "light yield" captures both of these efficiencies and is defined as the number of photons produced per quantum of energy absorbed by the scintillator. Light yield obeys something known as Birks' law [51] which is written as follows:

$$\frac{dL}{dx} = S\frac{dE/dx}{1+k_B(dE/dx)} \tag{5.1}$$

where, L and S are the light yield and scintillator efficiency, respectively. The Birks' constant k_B is empirically determined through measurements and material-dependant. It describes the collective effect of different particle interactions that lead to non-radiative de-excitation ("ionization quenching"). For a given scintillator, the Birks' constant increases with the mass of the ionizing radiation. The energy loss per unit length in a material, dE/dx is given by the Beta-Bloch formula [39]:

$$\frac{dE}{dx} = \frac{z^2 e^4}{8\pi\epsilon_0^2 m_e v^2}\left(Z\frac{N_A}{A}\right)\log\frac{2m_e\beta^2}{E_{min}(1-\beta^2)} \tag{5.2}$$

where z is the charge of the particle, e and m_e are the charge and mass of the electron, v the speed of the particle, A and Z the atomic number and mass number of the material, N_A Avogadro's number, E_{min} the scintillator ionization potential, and $\beta = v/c$. Although light yield is generally defined in units of photons per energy, experiments like SNO+ often quote an analogous value which is the number of PMT hits per unit energy.

Non-radiative processes like vibrational relaxation produce a Stokes shift between absorption and emission spectra. That is, scintillators emit light at longer wavelengths than they absorb. For organic scintillators, these spectra aren't perfectly separated. This means that LAB reabsorbs a large portion of the light that it emits. The π-electrons excited by reabsorption then relax non-radiatively, decreasing the light yield. An additional fluor can recover this reabsorbed light by first absorbing either the energy from the non-radiative de-excitation, or some of the light from the radiative de-excitation. The fluor's π-electrons then de-excite radiatively, emitting light outside the absorption range of the solvent scintillator. An ideal fluor absorbs within the primary scintillator's emission spectrum, and it emits within the quantum efficiency range of the PMTs.

5.2 SNO+ Liquid Scintillator

The liquid scintillators chosen for neutrino experiments consist of solute-type fluors dissolved in a primary solvent scintillator. The Borexino and KamLAND experiments pioneered the use of liquid scintillator, each using 1,2,4-Trimethylbenzene as a solvent [52, 53]. Although successfully deployed for long-term data-taking, this material is ex-

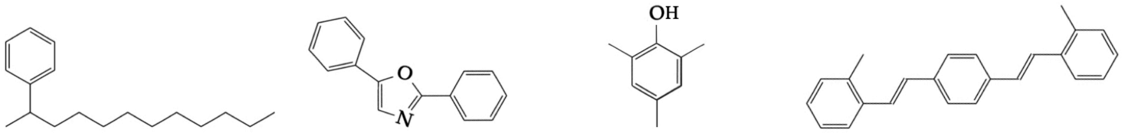

linear alkylbenzene (LAB)

2,5-diphenyloxazole (PPO)

butylated hydroxytoluene (BHT)

1,4-bis(2-methylstyryl)benzene (bis-MSB)

Figure 5·2: Primary components of the SNO+ liquid scintillator.

tremely volatile and toxic, and it requires careful and specialized handling to protect and the environment and lab personnel [54]. For this reason, the SNO+ collaboration explored alternative organic scintillators.

The clear winner, for several reasons, was LAB, a $C_{10} - C_{12}$ alkane containing a benzene group which is illustrated in Figure 5·2. This material is compatible with acrylic and has a high light yield (~600 PMT hits/MeV), a long attenuation length (29.9±2.0m), and a density (0.8 g/cm^3) close to that of water, making it the brightest, most transparent, and least buoyant of the six solvents originally considered [55]. LAB also has a high flash point (143 °C) and is certified for environmental release, making it the safest scintillator to transport, purify, and deploy underground. Amusingly, LAB's only documented danger to humans is choking [56].

The data discussed in this thesis were taken with LAB that contained 2.0 g/L of 2,5-diphenyloxazole (PPO), a fluor long used in liquid scintillator detectors. PPO has a wide Stokes shift, and its absorption spectrum overlaps with the emission spectrum of LAB [55]. Although listed in Figure 5·2, the fluor 1,4-bis(2-methylstyryl)benzene (bis-MSB) and stabilizer butylated hydroxytoluene (BHT) were not yet added to the scintillator during the data-taking period analyzed in this thesis. As of July 17th, 2023, the BHT has been nitrogen-purged, and water-extracted, and deployed in the AV. The loading of bis-MSB is nearly complete, and it will mitigate light absorption by tellurium during the $0\nu\beta\beta$ phase of data-taking [55].

5.3 Linear Alkylbenzene Fill

At a temperature of 12 °C, the AV contains ~ 780 tons of LAB. This volume had to be transported, stored, and loaded, in 20-ton batches as shown in Figure 5·3. First, tanker trucks transported the LAB (PETRALAB 500-Q LAB, purchased from CEPSA) to SNOLAB's tanker transfer facility, where a hose drew the scintillator into a 70-ton vessel (capacities quoted using density of water). From there, six 3-ton railcar tankers traveled back and forth, carrying the scintillator underground. Once there, the railcars were cleaned by hand, and then rolled into the lab to offload the LAB into one of two 60-ton storage tanks.

Concurrently, the LAB was purified and filled into the AV. During this process, the AV was constantly full of fluid; ultra-pure water drained from the bottom, as LAB filled from the top. Made possible by water's immiscibility with LAB, this technique ensured minimal contact between the AV and air. From the top of the neck to the bottom of the AV, the LAB fill lasted from 25 October, 2018 until 9 April, 2021.

5.3.1 Underground Purification Plant

All components of the SNO+ scintillator were synthesized above ground in commercial facilities like CEPSA, which supplied the experiment's LAB. They also traveled under ambient conditions above ground, before being rolled through a mine. To mitigate contamination from this long, multi-step journey, each conveyor underwent precleaning by citric acid passivation, and where possible, operated under a constant pressure of 99.999% pure N_2 gas. During the fill, LAB sat stored in polypropylene-lined holding tanks, each cleaned with a 1% Alconox and ultra-pure water solution.

These measures prevented degradation of the scintillator en route to the AV, but didn't remedy any contamination already sustained during synthesis, which took place under ambient conditions in an environment with much more lenient cleanliness re-

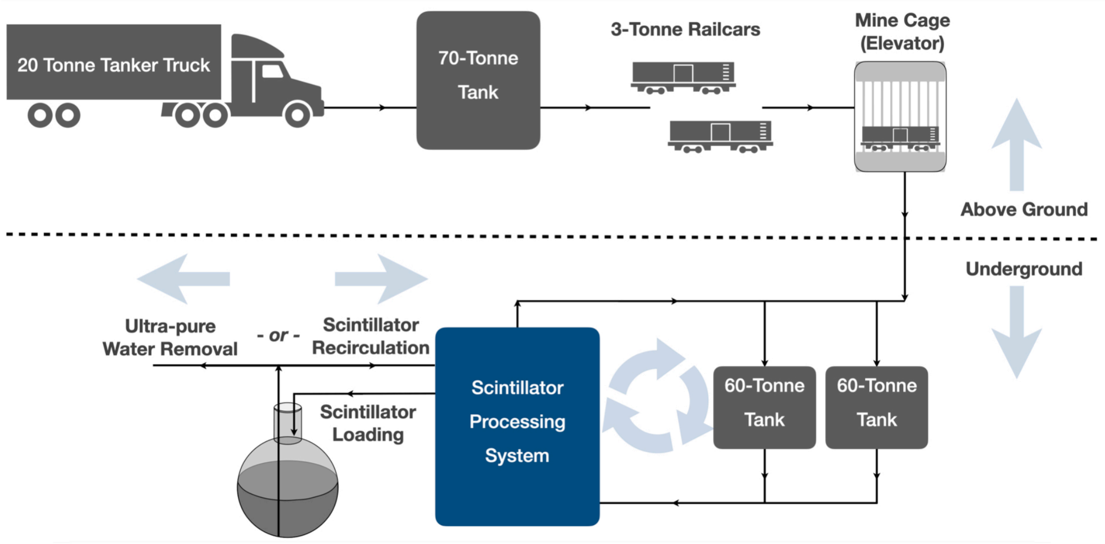

Figure 5·3: Block diagram of LAB's journey from the surface of SNOLAB to underground and into the AV.

quirements. To address this vulnerability, the collaboration designed a customized purification system to further clean and distill the scintillator underground as shown in Figure 5·4.

The underground purification plant consists of two distillation systems and a recirculation mechanism. The primary distillation cleaned the pure LAB in two stages: fractional distillation and gas stripping. Fractional distillation exploits differences in boiling points to separate liquids and dense solutes. During a fractional distillation, LAB enters the middle of a heated distillation column, allowing dense contaminants to sink, and the LAB to condense at the top. Repeating this process gradually strips dense contaminants, which collect at the bottom of the column for removal. Fractional distillation successfully removed metallic radioactive elements (Bi, K, Pb, Po, Ra, U, and Th), in addition to some non-volatile organic impurities that could have harmed the scintillator optical transparency.

The second stage of the primary distillation removed contaminants through gas

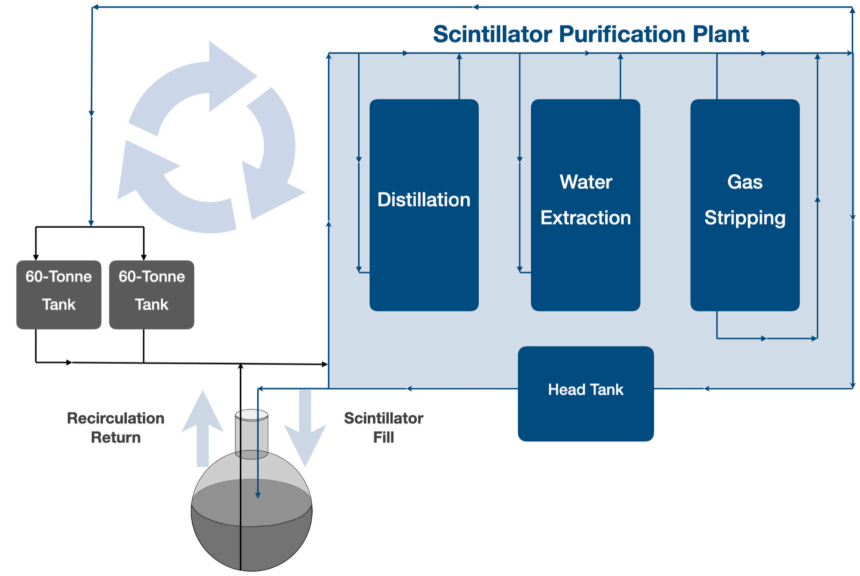

Figure 5·4: Block diagram of the SNO+ scintillator purification plant.

stripping, which exploits differences in vapor pressures. During gas stripping, the LAB enters the top of a column and flows through a porous, high-surface-area material at a rate of 4,000 kg/hr. At the same time, nitrogen enters from the bottom. The volatile impurities in the LAB then dissolve in the nitrogen, leaving through the top of the column by vacuum pump. This technique removed contaminants that have higher vapor pressures than LAB: O_2, Rn, Ar, and Kr, and water.

Originally, the primary distillation contained three stages, and the LAB was meant to go through a process called solvent-solvent extraction immediately after distillation. This technique is similar to gas stripping, in that it forces contaminants from one medium to dissolve in another medium with higher solubility. When brought into contact, ionic contaminants in the LAB would have dissolved in water and been carried away. This step never deployed, as fractional distillation removed a sufficient amount of U, Th, Ra, K, and Pb on its own. The machinery and column remain in place, ready to use, if desired during recirculation.

Unfortunately, fractional distillation didn't just effectively cleanse the LAB of metallic and organic impurities. It also thoroughly removed PPO, reducing its concentration by a factor of 100. This issue warranted the construction of a separate distillation system for the PPO. This loop, shown in Figure 5·5, proceeded in three steps: master solution preparation, water-water extraction, and flash distillation.

The PPO had to reach the AV already dissolved, since it was prone to clumping; once loaded, clumps could not have been heated, agitated, or easily removed for recirculation. To this end, a highly concentrated "master solution" conveyed the PPO through the secondary distillation system, and then to the loading tank for mixing and deployment. The master solution contained 400 L of distilled LAB and PPO at a concentration of 80 g/L. The master solution underwent several rounds of water-water extraction, losing 2.64% of its PPO. After this, the master solution underwent flash distillation, followed by in-line mixing with fractional-distilled LAB and several rounds more of water extraction and gas stripping.

Loading the powdered PPO into the secondary distillation system required some manual intervention. Dry and poorly mixed PPO is prone to clumping, and the line connecting the master solution to the rest of the purification system required regular unclogging. This extra layer of maintenance reduced the flow rate of master solution relative to the LAB fill. Because of this, the LAB needed ‘topping off’ with the master solution to reach the desired concentration of PPO. Moreover, the LAB had retained water that leached from the AV. To address these issues, the bulk scintillator was recirculated through the purification system, undergoing only gas stripping and addition of the PPO master solution. By the end of this recirculation, the AV had a humidity less than 5 ppm (parts-per-million) and a PPO concentration of 2.2 g/L. Although this process finished in April 2022, the scintillator can still be recirculated as needed.

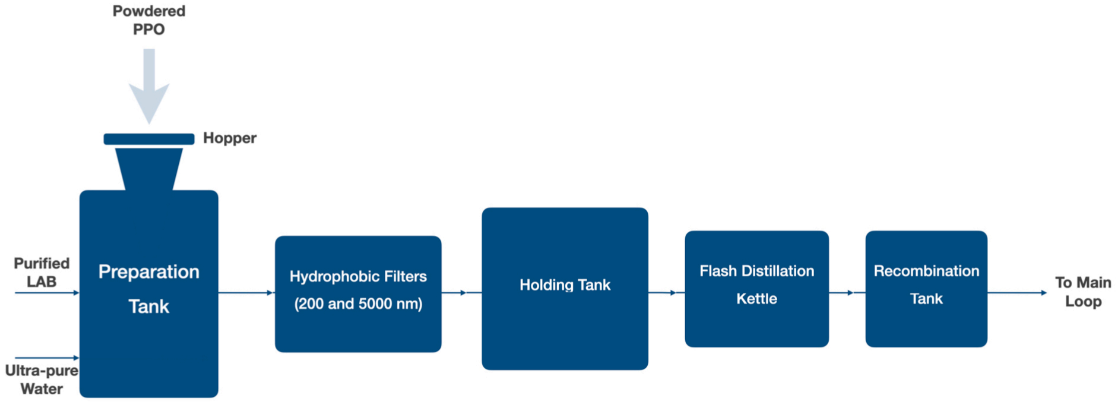

Figure 5·5: Block diagram of the secondary distillation loop used to purify the powdered PPO.

5.3.2 Ex-Situ Measurements

Scintillator deployment started with the LAB fill, concurrent with the ultra-pure water drain. This was followed by the addition of PPO. During each of these phases, samples of the scintillator cocktail were removed from the AV for external measurements of light yield, density, and absorption.

The emission spectra were measured with a Photon Technology International (PTI) QuantaMaster fluorescence spectrometer, with a 10-mCi ^{125}I excitation source. Light yield, calculated by integrating the spectrum from 350 – 425 nm (the quantum efficiency of the PMTs), was measured relative to that of a standard sample with 2 g/L PPO.

During the PPO top-off, samples were drawn from the AV at different times, and thus, different concentrations. The light yield measurements of these samples showed good agreement with 14 externally prepared standard samples, as well as with externally published values as shown in Figure 5·6. A fit of a double exponential to these data produced a relative light yield of $11,808 \pm 630$ photons/MeV. These measurements indicate that the scintillator plant achieved a 780-ton bulk scintillator with mixing and light yield comparable to that of reproducible, tightly-controlled, bench-top samples.

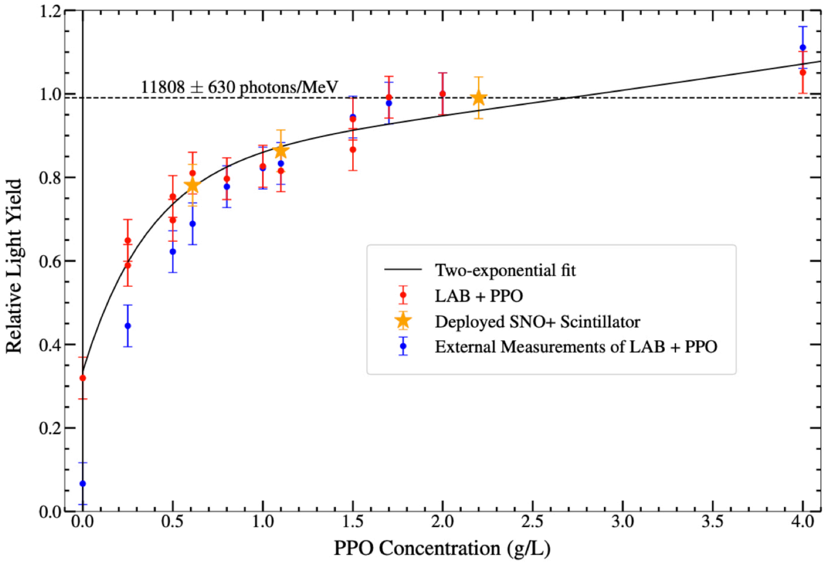

Figure 5·6: Relative light yield to a standard sample of LAB with 2.0 g/L PPO) versus concentration for externally prepared samples (red), scintillator drawn from the AV (gold), and values published in [57]. Figure adapted from Reference [58].

In addition, an Anton Paar DMA 25 Portable Density and Concentration Meter assessed the quality of the scintillator throughout the scintillator fill. Calibrated to 20 °C-ultra-pure water, the device measured the density of the scintillator at several temperatures. At a constant pressure of 0.12MPa, samples were either cooled in a fridge, and measured as they warmed, or heated and measured as they cooled. The data are shown to be in agreement with measurements from Daya Bay [59]. A linear fit of density as a function of temperature showed good agreement with externally published measurements, again indicating the thorough mixing and reproducibility as shown in Figure 5·7.

Finally, two separate campaigns measured the scintillator's absorption. One campaign was carried out on a Thermo Scientific Orion AquaMate 8000 UV-Vis spectrophotometer, and on a PerkinElmer Lambda 800 spectrometer. These measurements compared the absorption of the SNO+ scintillator to the LAB originally delivered by CEPSA

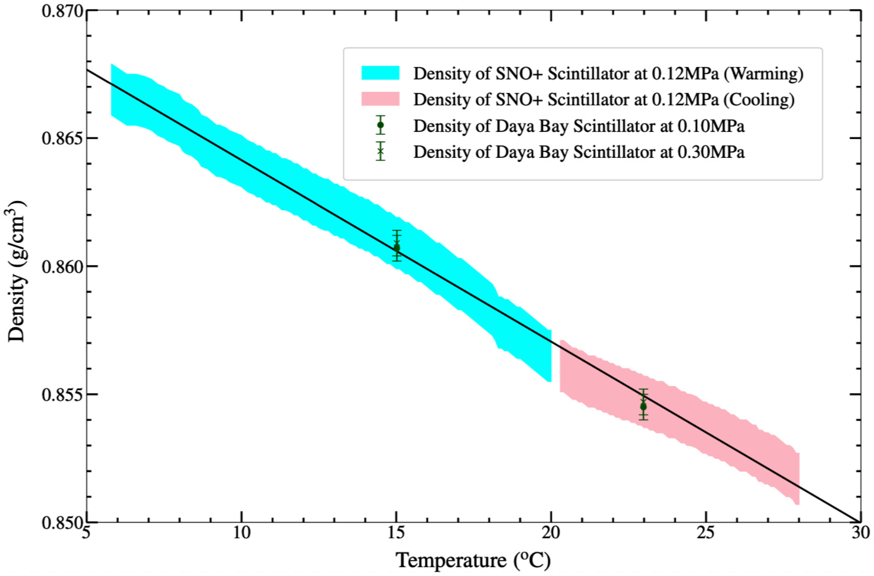

Figure 5·7: A plot of the scintillator's temperature-dependent density adapted from Reference [55]. The blue and pink bands indicate the errors on the respective density measurements listed in the legend.

and showed the expected decrease in self-absorption after the addition of PPO.

The second campaign assessed an incidentally discovered increase in absorption for certain wavelengths. Additional tests identified the absorbing contaminant as *p*-benzoquinone, which has an absorption peak too subdominant to that of PPO to be detected. This contaminant results from oxygen contamination and subsequent heating. Insufficient nitrogen purging, along with an unmitigated leak in one of the purification columns, caused oxygen to build up in the scintillator. The oxygenated scintillator then passed through the heated distillation columns in the course of recirculation. Increased nitrogen purging remedied these issues, and attenuation lengths at wavelengths over 420 nm were found to be roughly 20 m which are well in excess of the AV diameter of 12 m.

5.4 Tellurium-loaded Scintillator

As discussed previously in Chapter 3, the ideal double-beta decay isotope has a low half-life, a high endpoint, and a high natural abundance. These requirements ensure the experiment sees as many decays, in as low a background region, and for the cheapest cost possible. As shown in Table 3.1, real double-beta decay isotopes have many trade-offs.

Beyond these requirements, the final isotope-loaded scintillator must have a high light yield of at least several hundred PMT hits/MeV, and introduce minimal optical defects, such as increased scattering or absorption lines within the PMT quantum efficiency. The scintillator must be radiopure with less than $\mathcal{O}(10^{-15})$ grams of ^{238}U and ^{232}Th per gram of scintillator (commonly referred to as the unit 'g/g'). It must also be economical, and safely deployed and stored underground. Finally, the scintillator should retain all of these qualities for at least five years or longer, without any noticeable degradation to the AV.

5.4.1 Chemical Properties and Synthesis

The double-beta decay isotope of choice in SNO+ is ^{130}Te. It has a $2\nu\beta\beta$-decay half-life of $T_{1/2}^{2\nu} = 0.771^{+0.02}_{-0.021} \times 10^{21}$ years, a natural abundance of 34.08%, and an endpoint of 2.527 MeV. Tellurium metal is only soluble in some thiols, strong acids, and strong bases. This section covers the novel synthesis of an LAB-soluble tellurium compound.

Although tellurium is available in a pure metallic form, it is more commonly produced (and thus more economical) as a compound. The synthesis described here was adapted for orthotelluric acid (TA). Although this TA is also not soluble in LAB, it is soluble in water, and could thus be suspended in a surfactant. Linear alkylbenzene sulfonate (LAS) happens to be a readily available surfactant, commonly used in detergents. Although LAS stably suspends aqueous TA in LAB (by Calimulse proprietary sulfonation process), this method was abandoned as LAS cannot be easily distilled.

Luckily, glycols can also stably suspend TA in LAB, and are also commonly available for their use in cosmetics. The glycol 1,2-butanediol (BD) was chosen for its low molecular weight, sterile production, and ease of distilling. These qualities allow high tellurium loading, minimal contamination by living organisms (which introduce the well-known ^{14}C background), and high radiopurity.

The reaction of TA with BD proceeds in aqueous solution, and is simple to perform. First, powdered TA is dissolved in water and mixed with BD, forming a TeBD precursor. A diagram for this process is depicted in Figure 5·8. This compound consists of TA with BD bonds ("ligands") on all of the square-planar oxygen atoms. These organic ligands increase the solubility of TA in LAB. This reaction also produces water as a byproduct, a problem, since water promotes the backward reaction and is immiscible with LAB. To circumvent this issue, the aqueous solution is sparged with nitrogen, and the reaction is allowed to proceed at 70−80 °C until the water evaporates. As a bonus, this procedure promotes the formation of Te-Te complexes terminated with more BD ligands than the precursor which is depicted in Figure 5·9. The result is a mixture of Te complexes, all soluble in LAB.

The presence of even 300 ppm of water causes TeBD to hydrolize back to TeA and BD. This should be avoided at all costs since the TeA then immediately crashes out of solution. Since the scintillator was loaded as the water was drained, there is a risk

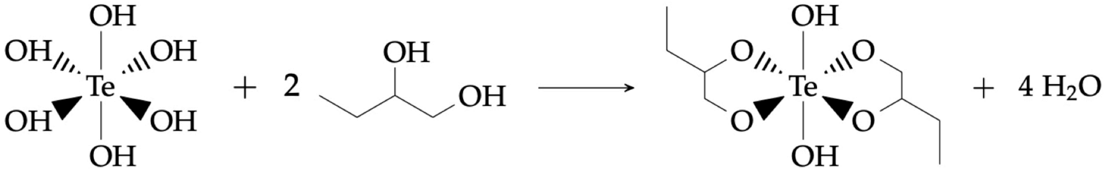

Figure 5·8: Diagram of the reaction of TA with BD which occurs in aqueous solution, forming TeBD, a tellurium complex soluable in LAB.

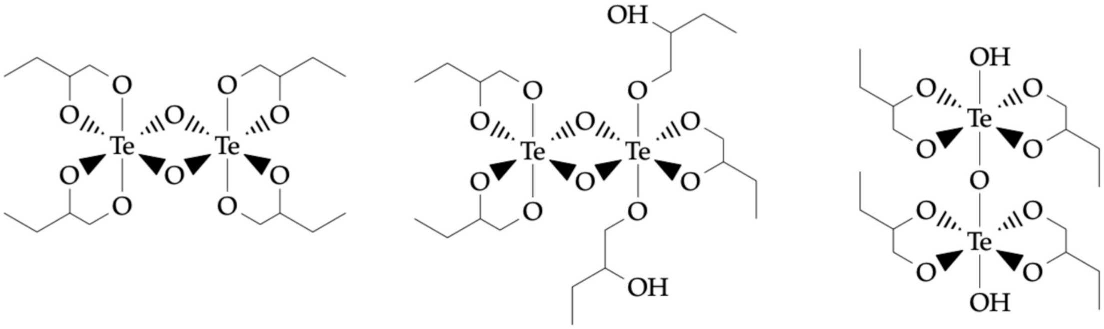

Figure 5·9: Diagram showing alternative TeBD complexes formed due to the presence of more BD ligands than the precursor TeBD.

that trace water in the walls of the AV can leach into the scintillator and damage the tellurium complexes. At the suggestion of Brookhaven National Laboratory (BNL), a long-chain amine will be added to the tellurium cocktail as a stabilizer. Although originally motivated by BNL's successful use of octylamine in a similar system, the chosen amine is N,N-Dimethyldodecylamine (DDA). This amine is much less toxic than octylamine and has a much higher flash point, making it safe for use underground. Moreover, its long chain makes it miscible with LAB. Of the seven amines investigated, only three stabilized the TeBD for longer than two months, and DDA had the highest flash point. A list of the amines that were investigated for use a stabilizers are shown in Table 5.1.

5.4.2 Bulk Tellurium Loading and Deployment

At the bench-top scale, the Te-loaded scintillating cocktail has been optimized for optical clarity, stability, and light yield. The bench-top synthesis of TeBD is easily scalable, and proceeds in the same way for bulk quantities. In a *hot* synthesis, aqueous TeA (30 – 50%w/w) will be added to BD in a 1:3 molar ratio. Heating and sparging with nitrogen drives the reaction by removing water. After achieving less than 1,000 ppm of water, the reaction yields mostly complexes with two tellurium centers, collectively called 'TeBD-I'. Following this procedure, DDA is subsequently added.

Amine	Stable for > 2 Months	Max. humidity before crashing (ppm)
Octylamine (OA)	no	-
1-Ethylhexylamine (EHA)	no	-
N,N-Dimethyldecylamine (DMDA)	yes	> 78000
N,N-Dimethyldodecylamine (DDA)	yes	> 78000
Tetradecylamine (TDA)	No	-
Hexadecylamine (HDA)	Yes	900
Dimethyltetradecylamine (DMTDA)	No	-

Table 5.1: List of amines investigated for the purpose of increasing stability of tellurium-loaded scintillator. Stability (second column) was assessed visually, while maximum humidity was calculated by titrating the cocktail with distilled water.

Loading can also proceed by a *cold* synthesis, whereby a single mixture of DDA, TeA, and BD is created in a 0.5:1:2 molar ratio. This reaction doesn't proceed in the aqueous phase and yields single-Te monomers dubbed 'TeBD-II'. These contain two bidentate BD ligands, and, if thoroughly sparged with nitrogen, are soluble in LAB.

Deployment will proceed in two steps: TA purification and TeBD synthesis, outlined in Figure 5·10. As of this writing, purification and synthesis plants are constructed and underground, awaiting final commissioning. The TA arrived at SNOLAB from ABSCO Ltd. in 2015. It remains underground today, allowing the decay of cosmogenic contaminants it incurred while above ground. The remainder of these will be removed in the TA purification plant, where the inherent U and Th contamination will need to be reduced by a factor of 10^2 to reach an acceptable level.

Similar to LAB purification, the TA purification will proceed in batches. A concentration of 0.5% (by mass) of tellurium requires 52 batches, each 200 kg. During purification, the amorphous solid TA will be dissolved in hot ultra-pure water and filtered to remove particulates. The dissolved TA will then be purified with two rounds of acid recrystallization, followed by one round of thermal recrystallization. The principle of operation is the same for each method; a solute is made to precipitate out of solution (either by the addition of a strong acid or by cooling), and the impurity-containing supernatant is removed by vacuum filtration. Other work has shown that one round of nitric acid recrystallization can reduce radio-impurities and cosmogenic contaminants by a factor of $\mathcal{O}(10^2 - 10^3)$, at the expense of only losing $\mathcal{O}(1)\%$ tellurium [60]. The final thermal recrystallization will cleanse the TA of any residual nitric acid. Separate from the TA purification, the BD purification will proceed in the main scintillator distillation column.

Following purification, the TeA and BD will be sent to the TeBD synthesis plant, where the reaction (shown in Figure 5·8) will proceed by both the hot and cold methods. Aqueous TA will flow to the plant and and react with BD in a flash chamber, which will remove the water (hot method). Previously distilled on the surface, DDA will be added directly to the plant (cold method), yielding TeBD diluted with LAB to a 1:1 ratio by mass. The TeBD will then go to the scintillator processing plant, where it will undergo further dilution and in-line mixing before deployment in the AV.

Every precaution has been taken to ensure material compatibility with nitric acid. To prevent leaching of metals into the scintillator, all vulnerable components in the purification and synthesis plants are made of polypropylene, teflon, or PFA. Other components were pre-leached with nitric acid and rinsed with ultrapure water. Preliminary tests demonstrate the tellurium processing plant's capability of achieving less than 1×10^{-13} g/g U, 5×10^{-14} g/g Th, and 7.5×10^{-13} Bq/kg ^{60}Co, all acceptable levels for the $0\nu\beta\beta$ decay search in SNO+.

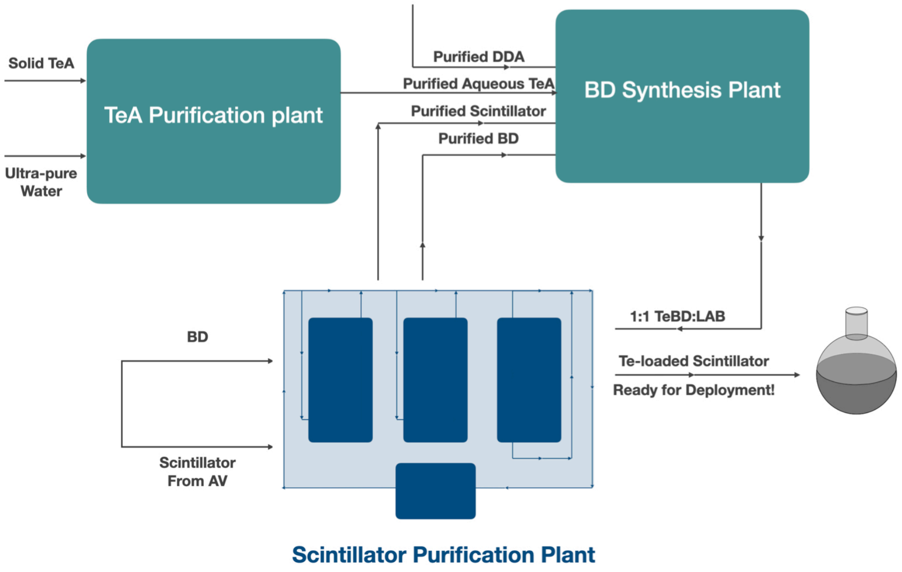

Figure 5·10: Block diagram showing the journey of TA, BD, and DDA to the AV. Aqueous TA will react with BD from the main scintillator purification plant. DDA, purified above ground, will be added directly to the sparged mixture (cold method). The resulting solution will be mixed with LAB in the main loop and deployed in the AV.

5.5 Quality Assurance and Analysis

A complex system like the underground scintillator purification plant is naturally vulnerable to small, easy-to-miss imperfections during construction and commissioning. One example are microscopic air leaks that could introduce oxygen and radon, which can be a danger to the scintillator quality or introduce backgrounds to the $0\nu\beta\beta$ search.

Oxygen reacts with LAB to form 1,4-benzoquinone, a compound that has absorption peaks within the quantum efficiency of the PMTs, as well as at wavelengths exceeding 500 nm. This absorption profile causes a kind of ‘yellowing’ of the scintillator, degrading its light yield and attenuation length. UV-vis measurements of oxygen-exposed LAB show minimal to no yellowing of samples stored at room temperature for many years. These periodic tests confirm that the reaction is slow at room temperature. Literature shows, however, that it is significantly accelerated through heating

[61], which happens during distillation. A microscopic air leak in the purification plant could cause the LAB to oxidize faster than it can be distilled.

To combat this, samples were taken from the scintillator during commissioning of the purification plant. Within a turnaround time of 10 minutes, scintillator undergoing purification could be drawn, and its absorption measured in a Thermo Scientific Orion AquaMate 8000 UV-Vis spectrophotometer. Periodically, these spectra were compared to 1) a perfect, bench-top distilled sample of LAB and 2) a paper-filtered sample of the fresh CEPSA delivery. The spectra for the drawn samples were checked against the perfect standard for anomolous peaks -especially at 368 nm and 389 nm, where 1,4-benzoquinone absorbs. The samples were also checked against the filtered CEPSA sample to gauge the effectiveness of the distillation. In one illustrative instance, two samples drawn 20 minutes apart showed growing peaks at 368 nm and 389 nm, indicating an oxygen leak. This lead to the discovery of a loose valve.

Two additional measurements were performed on samples periodically taken prior to filling: turbidity (on a Hach TL2310 LED Turbidimeter) and density (on an Anton Paar GmbH DMA 35 Portable Density and Concentration Meter). These measurements helped examine the scintillator for any dissolved contaminants not visible through UV-vis measurements.

Prior to each of these three measurements, a glass jar with a PTFE-lined cap first went through ultrasonic cleaning and was rinsed with methanol. A sample was then drawn through one of several methanol-cleansed ports, and the jar cap wrapped in several layers of teflon tape. All samples were placed in pre-cleaned bags for storage. All samples were then taken for QA analysis.

Samples were taken for QA analysis at each major stage of processing: arrival from CEPSA, storage, distillation, and deployment. After being primed with LAB, three consecutive QA analyses were performed for samples in closed-loop circulation. When the scintillator quality and quality improvement throughout the circulation were confirmed, a final QA analysis was performed before deployment in the AV. In this process, errors found at any point during closed-loop recirculation could be corrected (minutes to months, depending on the complexity of the failure) prior to filling.

QA analyses were also performed on the PPO in both the mixed and master solutions, with a few adaptations to accommodate the high concentration of PPO. First, the UV-vis spectra were standard-subtracted (using a perfect LAB sample) before being examined for anomalous absorption peaks. The concentration was too high to be measured in a similar fashion to LAB. Instead, it was determined by first preparing 20 samples with PPO concentrations from 0.5 – 6.0 g/L. A conversion from absorption to concentration was calculated by performing a third-degree polynomial fit to the 350-nm absorption versus concentration curve. Although this method reproduced the correct concentrations of additional prepared samples, it didn't extrapolate to the higher concentration of the master solution. Instead, to determine the concentration of the master solution, this fit method was carried out on 50:1 ratio dilutions.

Similarly to the LAB, the PPO was subject to these QA tests at each stage of addition: just prior to introduction, just after mixing, after each water extraction, just prior to distillation, and just prior to in-line mixing. Any solution batches that failed QA testing could either be thrown out or fully reprocessed in a new batch.

The Scintillator Counter Of Uranium and Thorium (SCOUT) was designed, shown in Figure 5·11, to check that the addition of PPO was improving the light yield of the scintillator. This small system contained an acrylic cylinder, into which master solution could be drawn. This was sandwiched between an array of four 3-inch PMTs and a 3.6 MBq ^{60}Co source, all encased by copper and lead shielding. The total charge on the PMTs after ten minutes was compared to that of an undistilled LAB sample. This measurement was taken for each batch of PPO, and at each juncture showed improved optical qualities relative to the undistilled LAB standard.

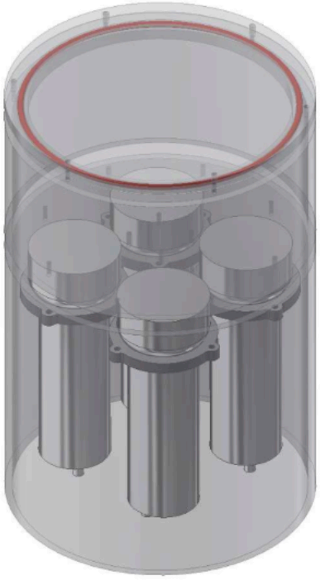

Figure 5·11: Rendering of SCOUT, showing the PMTs and scintillator volume encased in copper and lead shielding. Figure adapted from Reference [58].

5.5.1 Thousand Standard Test

As a followup to demonstrating the stabilizing effects of DDA, a study was conducted to find the optimal synthesis conditions and molar ratios needed to develop a stable cocktail. The Thousand Standard Test (TST) operated for 73 months and systematically tracked the quality of tellurium-loaded scintillator samples created from 150 TeBD synthesis variations. Each method differed in material suppliers, humidity exposure, LAB preparation, synthesis duration, and BD:Te ratio.

Each sample was strictly handled and prepared by diluting with LAB in a 0.5% ratio, under an oxygen-free, nitrogen-sparged environment. The dilutions were carried out with 30 mL of LAB that had been sparged with nitrogen for 15 minutes and filtered through a 0.2 μm filter. At least four samples per synthesis method were prepared, and all underwent visual monitoring for TeBD stability. Many more were subjected to extreme conditions. For example, a total of 88 samples were exposed to 6 hours of humidity and 240 samples were exposed to an oxygen-rich environment. All of the 1,136 samples were examined for signs of precipitation and degradation. The set of samples used in the stress test showed no precipitation with the presence of DDA, indicating the robustness of the stabilized cocktail. Additional measurements of light yield showed optimal BD:Te and DDA:Te molar ratios of $(\sim 2.5-3.0)$ and $(\sim 0.25-0.75)$, respectively. The final composition of the 0.5% tellurium-loaded scintillator is listed in Table 5.2.

5.5.2 Light Yield Monitoring of Tellurium-loaded Scintillator

Bench-top and small bulk studies have shown a 40% reduction in light yield after the addition of TeBD. This fluorescence quenching could have several causes. For example, in the presence of heavy atoms like tellurium, π-electrons preferentially excite to triplet states, which don't annihilate as frequently in the Te-doped cocktail, and otherwise decay non-radiatively. Moreover, the tellurium complexes are electronegative, due to the presence of several Te-O bonds. Electronegative species can partially localize π-electrons, inhibiting fluorescence.

Bench-top measurements of light yield were performed at Oxford University and The University of Pennsylvania, each producing similar results. The setup at Oxford used a ^{90}Sr source, which aimed β particles at a 1-cm (outer diameter) cylindrical cu-

Component	Concentration
LAB	-
PPO	2.2 g/L
bis-MSB	3.5 mg/L
BHT	5.4 mg/L
TeBD	0.5% by mass
DDA	0.25:1 molar ratio to Te

Table 5.2: Composition of the 0.5% tellurium-loaded SNO+ scintillator.

vette filled with 3 g of tellurium-loaded scintillator. A Hamamatsu H11432 PMT facing the cuvette was exposed for 5 minutes at a time, and its charge was accumulated and output by a multi-channel analyzer [43]. The light yield, reported in arbitrary analog-to-digital converter (ADC) units, was defined as the last 40 counts in each charge spectrum. This definition captures the maximum energy deposited in the scintillator. This metric is robust against changes in the shape of the charge spectra due to the β particles' variability in energy and path length, or to sample heterogeneity.

A study was also designed to measure the stability of tellurium-doped scintillator in bulk. The Long Term Test Tank, dubbed LT3, consisted of a cylindrical vessel, filled with 100 L of nitrogen-sparged, Te-doped liquid scintillator. The experimental setup is shown in Figure 5·12. The LT3 tank was fitted with a stand on top to hold three down-ward facing Hamamatsu PMTs, each connected to a CAEN multi-channel digitizer. This system triggered on cosmic muons, detected by two Cosmic Watch muon

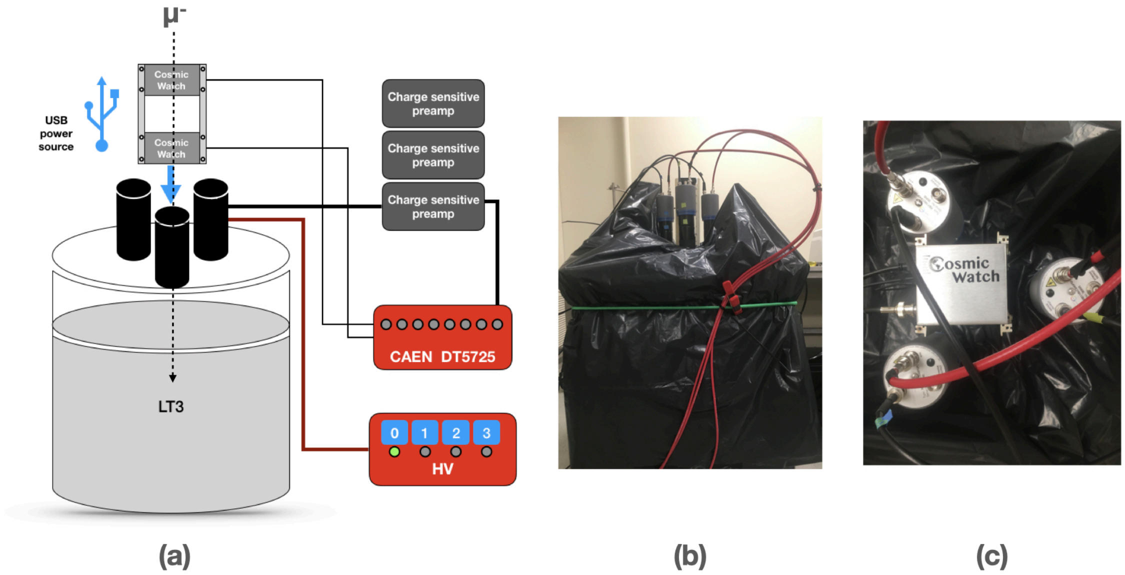

Figure 5·12: The setup of the long term test tank (LT3). a) Schematic of LT3, showing the Cosmic Watch triggers and three PMTs, all coupled to the scintillator-filled acrylic vessel. b) and c) side and top views of LT3, wrapped with a black plastic film, which serves as a dark box.

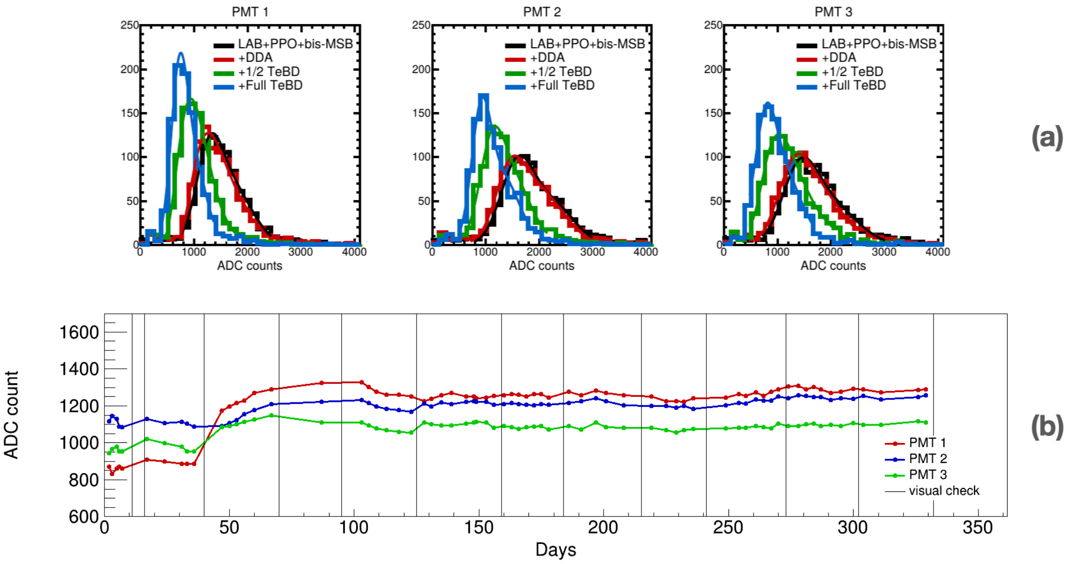

Figure 5·13: a) The light yield measured in LT3, using cosmic muons. Each distribution shows the data taken for each PMT. Light yield was measured for: plain (LAB, PPO, and bis-MSB) scintillator (black), after DDA addition (red), after half-Te loading (green), and after full-Te loading (blue). b) Plot of ADC count over time for each PMT under stable conditions, showing stable light yield for nearly a year (until decommissioning).

counters [62]. These counters were stacked on the center of the top face of LT3. Coincidences between the two watches and the three PMTs selected a sample of centrally located, downward-going muons. Charge on each PMT was collected in 2-3 day exposures, and light yield was defined as the peak of a gaussian fit to the ADC spectrum. This measurement was repeated at each stage of loading the DDA and TeBD, and the results are shown at the top of Figure 5·13. After full loading of 0.5% tellurium, the light yield decreased by 43.8 ± 3.0%.

Given the negative effect of TeBD on the light yield, observing an improvement could be an early indicator of TeBD crash or phase separation. As discussed, both effects happen in the presence of water. To demonstrate light yield as a useful monitoring tool, the Oxford group prepared 1% Te samples with DDA:Te ratios of 0, 0.25 and 0.5. The samples were then continuously exposed to humidity for 100 days. The light yield of the samples containing DDA modestly increased only ~ 10% over the course of the entire study. The light yield of the samples containing no DDA sharply increased ~ 40% within the first 20 days, as its tellurium either precipitated or phase separated out of solution. The light yield in LT3 was also monitored over a year. Under stable temperature, nitrogen sparging, and PMT placement, the light yield remained stable shown in at the bottom of Figure 5·13. We note that the observed shift in relative ADC count for all three PMTs on LT3 after about 40 days was due to a brief removal and reattachment of the PMTs on top of the acrylic tank, thereby changing the relative light collection efficiency of each PMT. This was the only time the PMTs were changed during operation of LT3. The ADC counts after this change show consistent outputs over measurement time. Periodic visual checks corroborated the findings above. Benchtop samples prepared with 0.5 DDA:Te molar fraction showed no precipitation after three years. Examinations of LT3 also did not show tellurium precipitation or adverse effects.

Chapter 6

Simulation and Reconstruction

Searches for extremely rare events like $0\nu\beta\beta$ decay require a precise accounting of backgrounds, and in turn, an accurate prediction of each event's energy and location inside the detector. From the predicted light yield of the tellurium-doped scintillator, and the Poisson fluctuations in the number of collected photoelectrons, the best achievable energy resolution at $Q_{\beta\beta}$ is 2.9%. The task of energy reconstruction is to get as close to this limit (the "Poisson limit") as possible. Detailed simulations of different events in the detector guide the development of reconstruction algorithms, which recover physics information from raw data.

6.1 Simulation

Physics analyses require detailed Monte Carlo (MC) simulations of all backgrounds in SNO+. These are handled by the Reactor Analysis Tool (RAT) software package [63], which contains bespoke event generators, as well as a detailed model of the entire SNO+ detector: its geometry, the PMTs, triggering electronics, detector component materials , and scintillator. The RAT framework was written in C++ with heavy reliance on the Geant4 [64] and GLG4sim [65] software packages. RAT also uses a database of detector and environment conditions during one-hour-long periods of data taking ('runs'), such as the number of PMTs turned on, downtime for on-deck maintenance work, mining operations ('blasts'), and occasional voltage breakdowns.

When RAT generates an event, such as solar neutrino scattering or a radioactive decay, Geant4 simulates all material interactions and GLG4sim handles the photon production and propagation through the detector. If a photon successfully hits a PMT photocathode, RAT simulates photoelectron creation, using the angle of photon incidence, wavelength, and individual PMT status. It then simulates the triggering system, and saves the event's information. This information consists of whether or not the event triggered the detector, and the PMT hit times and charge, which are converted by reconstruction algorithms to energy, position, and time. As with data, all of this information is stored in custom ROOT [66] files. In addition, these ROOT files also contain the true time, position, and energy of the generated event.

6.2 Reconstruction

The SNO+ detector reads out the PMTs that were hit and how much charge they produced. Reconstruction algorithms convert these observables to physics information like vertex position, time, energy. All of these values are critical to determining the type of event and the species involved. Event reconstruction is based on a Cartesian coordinate system centered at the origin of the PSUP. In this system, the xy-plane lies horizontally, and the positive z-axis extends upward through the neck of the detector. Given the detectors largely spherical geometry, the radial distance from the origin is frequently used as well. In some cases, the polar angle θ, defined relative to the z-axis, is also used, particularly when spherical symmetry is advantageous or when its important to account for the asymmetry introduced by the neck along the z-direction.

The SNO+ collaboration developed a suite of reconstruction algorithms within RAT, known as *ScintFitter*, to determine the energy, position, and time of events in scintillator. Because scintillator light is emitted isotropically, the signals are generally not sensitive to the original particles direction. All events are reconstructed assuming an electron, which generally does not affect the position but can influence the inferred energy. As such, the reconstructed energy is reported as the equivalent electron energy. The main reconstruction components of the *ScintFitter* are summarized below.

6.2.1 Vertex and Time Reconstruction

Several things happen during an event: a particle deposits its energy, this energy transfers to and excites the π-electrons of the scintillator, the π-electrons relax and emit light, this light propagates to some of the PMTs, and the SNO+ electronics register the hits. The correct reconstruction of an event's position and time requires the correct modeling of each of these processes, captured by a PMT-level quantity called the time residual:

$$t = t_{\mathrm{PMT}} - \tau - T_r \tag{6.1}$$

where t_{PMT} is the time the PMT was hit, t is the reconstructed time the event occurred, and τ is the time it took for the light to propagate from the reconstructed vertex to the PMT (time-of-flight). Written this way, one can see that the time residual T_r captures everything else that happens to the emitted light from its production and attenuation.

For a perfect reconstruction of t and accurate value of τ, the time residual exactly matches the timing profile of the scintillator, which is known from measurements taken inside and outside of the detector. So, reconstruction of event time and position proceeds through a log likelihood fit of the PMTs' time residuals to the scintillator timing profile, according to the following:

$$\log \mathscr{L} = \sum_{i=0}^{N} \log P(T_R^i) \tag{6.2}$$

where P is the probability of the ith PMT having a time residual T_R. The probability is taken from a time residual probability density function (PDF) stored in the RAT database, which combines the scintillator timing profile with detector effects. It is generated from Monte Carlo simulations using the true event position and time. This fit is referred to as the *PositionTimeLikelihood* [67], which floats the reconstructed position (contained in the time-of-flight τ) and reconstructed time t to produce the maximum likelihood $\mathscr{L}$.

We use simulations of electrons in SNO+ scintillator to quantify the performance of the vertex reconstruction. For each simulation, the differences between the electrons' true and reconstructed positions follow a Gaussian distribution with some mean and width, which we will call the 'bias' and 'resolution', respectively. Each of the points in the plots shown in Figure 6·1 represents a Gaussian fit to the distribution obtained from $\{x, y, z\}_{\text{reconstructed}} - \{x, y, z\}_{\text{truth}}$ of 10^4 electrons, with energy 2.0 MeV, in a radial shell of width 550.0 mm. A Gaussian fit to each distribution has a mean and width equal to the bias and resolution. The vertex resolution is observed to be within 9 – 12 cm, tending toward lower values as one moves away from the center of the detector.

The energy dependence of the vertex resolution is shown in Figure 6·2. Similarly, each point represents a distribution of $\{x, y, z\}_{\text{reconstructed}} - \{x, y, z\}_{\text{truth}}$ obtained for 10^4 electrons of a given energy, uniformly distributed within 5,500.0 mm of the AV center. A Gaussian fit to each distribution has a mean and width equal to the bias and resolu-

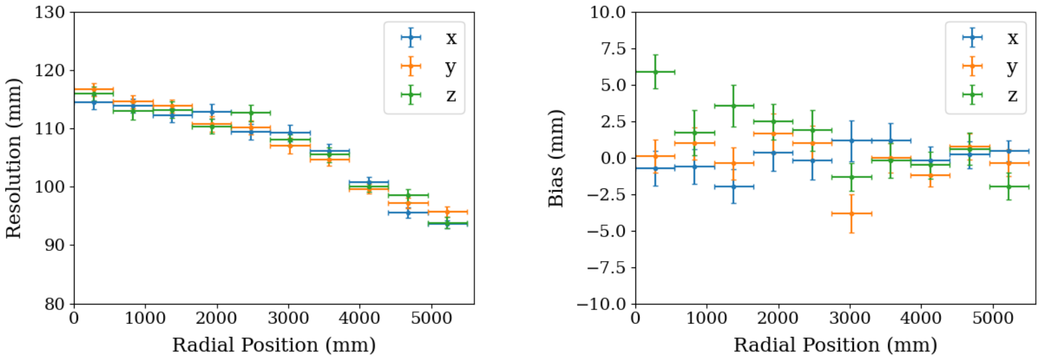

(a) **(b)**

Figure 6·1: The radial dependence of vertex a) resolution and b) bias. The horizontal error bars span the thickness of the shell, and the vertical error bars are the errors on each Gaussian fit.

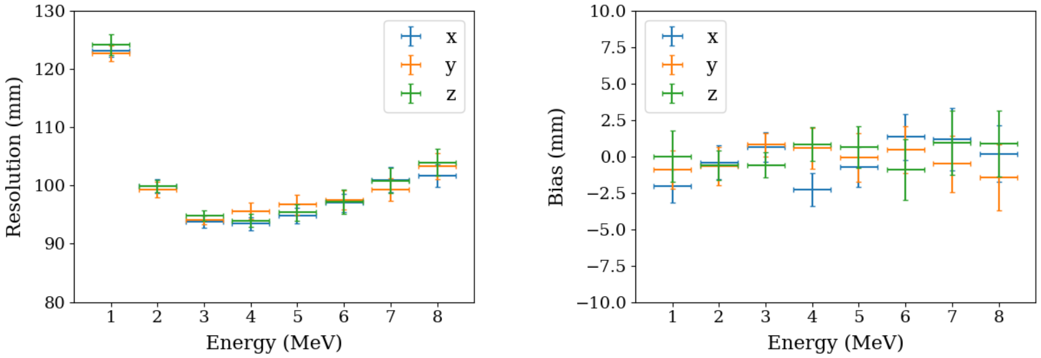

(a) **(b)**

Figure 6·2: The energy dependence of vertex a) resolution and b) bias. The vertical error bars are the errors on each Gaussian fit.

tion. The vertex resolution in each coordinate is finest at energy ~ 3.0 MeV and radial position ~ 5500.0 mm. This is expected, because the vertex reconstruction algorithm was developed for 2.5 MeV electrons (close to the $Q_{\beta\beta}$ for ^{130}Te). Across all energies and throughout the detector, the vertex bias is within resolution and is not considered to be problematic for the analysis in this thesis.

6.3 Energy Reconstruction

The software class responsible for energy reconstruction in SNO+ is referred to as the *EnergyRThetaFunctional* within *ScintFitter* and is outlined in Reference [68]. The number of scintillation photons produced in an event is proportional to the amount of energy deposited in the scintillator. At low energies, most hit PMTs are hit only once, and the number of scintillation photons is easily inferred. At high enough energies, the scintillator will produce enough photons to hit a PMT multiple times. Unfortunately, the SNO+ electronics cannot register more than one hit per PMT, nor can the PMTs' charge resolution distinguish between single and multiple photoelectrons. As a result, the number of scintillation photons becomes lost information at high energy, and the relationship between energy and integrated PMT charge becomes nonlinear.

A multi-hit correction method, developed in Reference [69] and applied within *EnergyRThetaFunctional*, recovers the linear conversion between number of photoelectrons and the number of scintillation photons at high energy. In lieu of registering multiple hits on each PMT, this approach groups the PMTs into segments with roughly the same hit probability, and calculates the average number of photoelectrons in each ensemble.

The multi-hit correction method firstly assumes that the probability of each PMT registering a hit follows a Poisson distribution centered at some expected number of photoelectrons, μ. Additionaly, we assume that the PMTs in a given ensemble have the same μ. Putting these assumptions together, the expected number of hits in an ensemble is:

$$\overline{N}_{\text{hits}} = \sum_{i=0}^{N_{\text{PMT}}} (1 - e^{-\mu^i}) = N_{\text{PMT}}(1 - e^{-\mu}) \tag{6.3}$$

where N_{PMT} is the number of PMTs in the ensemble, μ^i is the expected number of photoelectrons on a given PMT, and we have assumed that $\mu^i = \mu$ for each PMT. Solving for μ gives the expected number of photoelectrons μ for a particular ensemble, which can be written as:

$$\mu = -\log\left(1 - \frac{\overline{N}_{\text{hits}}}{N_{\text{PMT}}}\right) \tag{6.4}$$

We can then approximate the total number of photoelectrons in an event by adding μ for each ith ensemble. This results in a quantity H that scales linearly with energy:

$$H = -\sum_{i=1}^{S} N^i_{\text{PMT}} \log\left(1 - \frac{N^i_{\text{hits}}}{N^i_{\text{PMT}}}\right) \tag{6.5}$$

where $N^i_{\text{hits}}/N^i_{\text{PMT}}$ is the fraction of PMTs hit in the ith of S ensembles. The last piece of successful energy reconstruction is a correction to H for the positions of events that occur off-center. N^i_{hits} depends on the path lengths of the photons, which can undergo scattering and absorption. Moreover, hits close to the AV can totally internally reflect, or occur close to the north pole of the detector, where the photo-coverage is the smallest. Both of these effects can reduce the expected number of photoelectrons. These effects are modeled, and an empirical function of θ and r (parameters that are validated with calibration) relates the off-center correction to the equivalent correction at the center.

We assess the performance of the energy reconstruction by a method similar to the one used for vertex reconstruction. The distribution of $(E_{\text{truth}} - E_{\text{reconstructed}})/E_{\text{truth}}$ for 10^4 simulated electrons follows a Gaussian with mean and width equal to the bias and resolution, respectively. Figure 6·3 shows this calculation performed for simulations of 2.0 MeV electrons in 10 radial shells of thickness 550.0 mm, and Figure 6·4 for electrons of varying energy within 5,500.0 mm of the center of the AV. Across all spacial distributions and energies, the fractional bias is smaller than the fractional resolution.

Currently, the ratio of hit PMTs to electron energy is roughly 300/1.0 MeV, giving a Poisson limit on the resolution of $1/\sqrt{\text{N}_{\text{hits}}} = 1/\sqrt{3.0\text{ MeV} \times 300\text{ MeV}^{-1}} = 3.3\%$ at

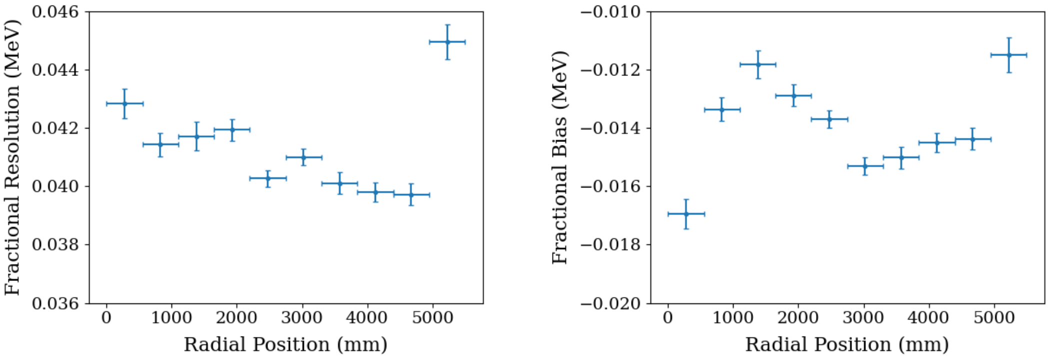

(a) **(b)**

Figure 6·3: The radial dependence of energy a) resolution and b) bias. The horizontal error bars span the thickness of the shell, and the vertical error bars are the errors on each Gaussian fit.

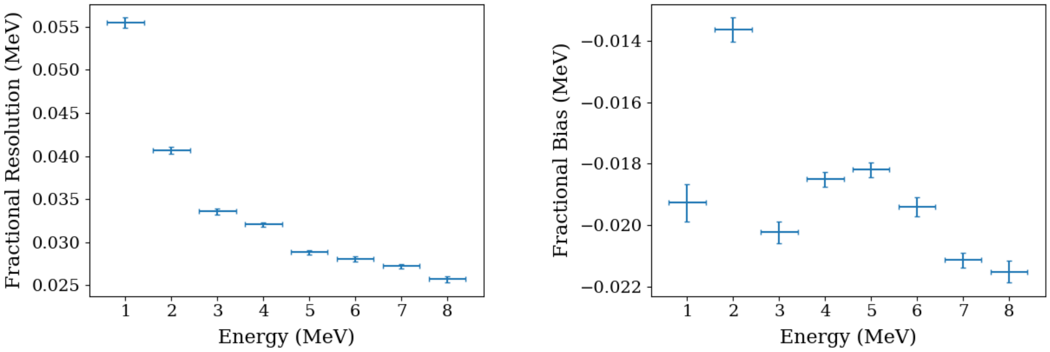

(a) **(b)**

Figure 6·4: The energy dependence of energy a) resolution and b) bias. The vertical error bars are the errors on each Gaussian fit.

3.0 MeV, and 4.1% at 2.0 MeV. Figure 6·4a shows resolutions 3.4% and 4.1% at these energies, indicating that reconstruction performs almost as well as is achievable over an energy range containing the ROI.

6.4 Classifiers

PMT-level information does not only provide the information needed for energy, time, and position reconstruction. It also enables particle classification by event type that leave unique hit patterns or occur with a signature timing profile, such as a set of coincident decays.

6.4.1 $\alpha-\beta$ Classifier

There are certain classes of correlated radioactive events that occur within very short time intervals of each other, at or near the same location inside the detector. Many such events come from two backgrounds known as ^{214}Bi – ^{214}Po and ^{212}Bi – ^{212}Po decays. The origin of these radioisotopes in SNO+ is discussed in more detail in Chapter 7. To briefly summarize, β-decays of ^{212}Bi and ^{214}Bi are immediately followed by the α-decays of ^{212}Po ($T_{1/2}$ = 0.299 μs) and ^{214}Po ($T_{1/2}$ = 164.3 μs), respectively.

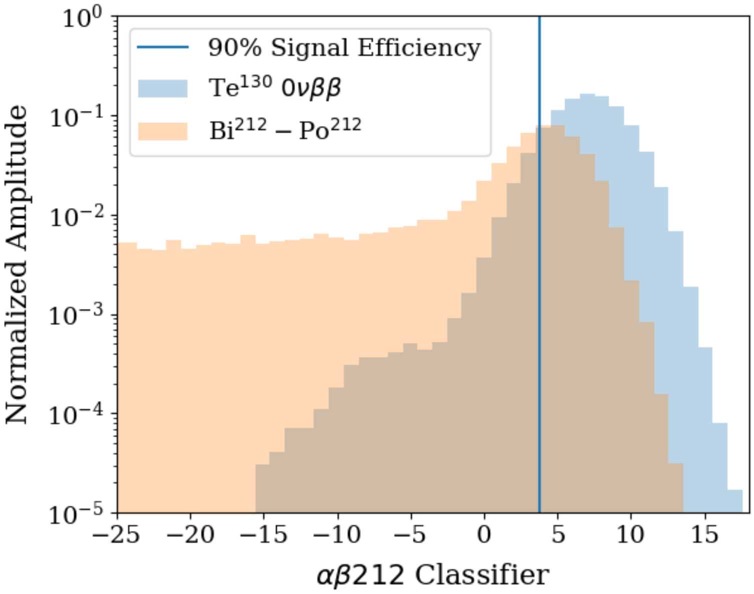

Figure 6·5: Distributions of the $^{212}(\alpha-\beta)$ classifier for the $0\nu\beta\beta$ signal and ^{212}Bi –^{212} Po background. The vertical line represents a cut at 90% signal efficiency, below which 99% of ^{212}Bi –^{212} Po events are rejected.

Since the half-life of ^{214}Po is much longer than the trigger window (400 ns), over 99% of ^{214}Bi – ^{214}Po decays can be rejected through coincidence tagging of the prompt β quickly followed by the delayed α. However, some tiny fraction of the ^{214}Bi –^{214}Po decays, and most of the ^{212}Bi –^{212}Po decays, happen within the trigger window as 'pileup' events.

There are two separate $\alpha-\beta$ classifiers, one for each of the in-window decays of ^{212}Bi –^{212}Po and ^{214}Bi – ^{214}Po. The PDFs are constructed from simulated time residuals of ^{214}Bi – ^{214}Po (or ^{212}Bi –^{212}Po) decays, and for $0\nu\beta\beta$ decays. Both classifiers perform a log likelihood fit of these pdfs to the time residuals of an event in data. The value of the classifier is likelihood ratio:

$$\log \mathscr{L} = \log P(0\nu\beta\beta) - \log P(^{214/212}\mathrm{Bi}-\mathrm{Po}) \tag{6.6}$$

The results of the ^{212}Bi –^{212}Po classifier performance against the simulated $0\nu\beta\beta$ signal is shown in Figure 6·5. A cut is chosen on the classifier output such that 90% of the $0\nu\beta\beta$ signal is accepted while rejecting about 99% the ^{212}Bi –^{212}Po background.

6.4.2 In-time ratio (ITR) and β_{14}

The ITR and β_{14} classifiers were originally developed to identify Cherenkov events in SNO. Cherenkov light is emitted early and produces anisotropic clusters of light (in comparison to other optical sources like flashers). The ITR classifier selects events with a high fraction of early PMT hits, and the β_{14} classifier selects events with PMTs hit within a small angular separation. The β_{14} and ITR classifiers have not been revised for the scintillator phase, but are examples of how particular PMT hit patterns and time residuals can be used to identify a process. Differences in timing and hit patterns are more subtle in scintillator than in water, but still resolvable. Instead of Cherenkov light, the latest event classifiers select low-energy radioactive decays by searching for deviations from point-like behavior in PMT hit patterns and time residuals.

6.4.3 Multi-site Classifiers in SNO+

A bimodal distribution of time residuals is just one example of a classifiable feature (in this case, for pileup of α–β decays). The width of a distribution of time residuals is also telling. This is because the PDFs used to fit position, given by Equation 6.1, are constructed from simulated 2.5 MeV electrons that behave as point-like events for which the fitter constructs only one vertex. Events with multiple energy depositions, such as the β-decay of ^{214}Bi and the ensuing gamma cascade (that undergo multiple Compton scatters), have broad time residuals, relative to single-site events, such as $0\nu\beta\beta$ decay. One of the recent multi-site classifiers developed for SNO+ exploits the differences in the shape of the time residuals to discriminate between single-site (signal-like) and multi-site (background-like) events. The details of this classifier are described in Reference [43].

Multi-site events can also display significant spacial anisotropy when they originate off-center. This topological classifier operates from the following principles: 1) γ's that travel in a straight line have the smallest energy depletion and are more likely to reconstruct in the fiducial volume (FV), and 2) the PMTs closest to the event origin have the earliest hits. This classifier has been demonstrated in simulation, and is most successful on external background events that reconstruct inside the FV. Details on the timing and topological multi-site classifiers can be found in [43].

The three classifiers described in this section each use either PMT timing or PMT hit patterns to identify an event. They also require the use of reconstructed quantities (eg. position). A new machine-learning framework called KamNet uses raw PMT timing and hit patterns simultaneously, and is the focus of Chapter 8.

Chapter 7

Neutrinoless Double Beta Decay Backgrounds

As detailed earlier, $0\nu\beta\beta$ decay experiments require extremely low-background detectors. These ultra-low backgrounds are realized with a clever selection of materials which are extremely radiopure, implementation of both passive shielding and active veto systems, a location in a very clean underground lab, and application dedicated data selection criteria. The most important backgrounds for the $0\nu\beta\beta$ search in SNO+ are detailed here.

7.1 Instrumental Backgrounds and Data Cleaning

The original SNO experiment encountered a class of bright backgrounds injected by its electronics — for unclear reasons, PMTs produced light that mimicked physics events. Thought to arise from electronic breakdown and static discharge, these events, called "flasher events" are an example of an instrumental background. Since SNO+ inherits all of its electronics from SNO, these backgrounds are still relevant today and have many of the same distinguishing features. The removal of instrumental background events is called "data cleaning", and the criteria used are the "data cleaning cuts". A full list and description of all data cleaning cuts developed through the ultra-pure water phase can be found in References [70] and [71]. Below is a brief description of the subset of data cleaning cuts used in this analysis and, where applicable, the upgrades they underwent to accommodate scintillator-phase data.

- **Zero-Zero Cut**: Collections of PMT hits that the DAQ doesn't correctly identify with an event are called 'orphans'. These are either wrapped into their own event or bleed into another event. One producer of orphans is the rollover of global trigger IDs; the zero-zero cut removes these.

- **Junk Cut**: This cut removes events in which a PMT errantly contributes two hits.

- **Muon Tag**: The cosmic muon rate in SNO+ is very low, only ~ 4 muons every hour. Nevertheless, they can induce spallation on the target material, producing nuclei with decays close to or in the ROI. The muon tag was originally developed for the ultra-pure-water phase, and underwent substantial changes for the scintillator phase. One of these is the addition of a set of cuts determined by the RMS positions and hit times of outward-looking (OWL) PMTs. The updated tag also requires an increased number of hit PMTs, which accommodates the scintillator's relative increase in light yield. Details on these updates, as well as the threefold coincidence tagging of ^{11}C they enable, can be found in Reference [72].

- **OWL Cut**: This cut removes events with ≥ 3 hit OWL PMTs.

- **TP Muon Follower-Short Cut:** This cut suppresses the decays of radioisotopes produced by cosmic muon spallation, by removing events that occur within 20 s of a tagged muon. Since the muon rate is low in SNO+, this cut results in only ~ 1m of data sacrifice.

- **Polling Cut**: PMTs are polled during each run to check for faulty channels. This polling produces electronics noise, which this cut removes.

- **Missing CAEN Data Cut**: The CAEN digitizer stores the sum of the ESUM trigger, so that its size and shape might be analyzed. Sometimes, the CAEN buffer can fill before the data are read out, resulting in events missing this information. This cut removes those events.

7.2 Internal

Internal backgrounds are radioactive decays that originate in the AV and scintillator (as opposed to "external" and "instrumental" backgrounds, which originate outside of the AV). Most of these decays are of isotopes in secular equilibrium along the uranium and thorium chains. These species found their way into the detector through ambient means; they are naturally present in the environment and detector materials. Accidental air leaks in particular disrupt secular equilibrium by introducing ^{222}Rn and ^{220}Rn, which feed the lower halves of these chains.

The daughters of ^{222}Rn and ^{220}Rn directly pollute the $0\nu\beta\beta$ ROI, and are discussed in more detail below. The scintillator and air leaks carry additional contaminants that, because of their low energy and absence from the $0\nu\beta\beta$ ROI, are only briefly mentioned here. The scintillator naturally contains trace amounts of ^{40}K, which has a half-life of 1.248×10^9 years, and undergoes both β-decay (Q-value = 1.311 MeV) and electron capture (1.460 MeV γ emission). Cosmic activation of ^{40}Ar in the atmosphere produces ^{39}Ar, a β-emitter (Q-value = 0.585 MeV) with a half-life of $T_{1/2}$ =269 years. Nuclear processing, weapons testing, and accidents have polluted the air with ^{85}Kr, a β-emitter with a half-life of $T_{1/2}$ =10.739 years. Its main decay mode (branching ratio = 99.57%) has a Q-value of 0.687 MeV, while its second decay mode emits a 173 keV β and a 514 keV γ with a half-life of 1.015 μs. Delayed coincidence tagging of this mode can help reject ^{85}Kr, but keeping the scintillator covered with pure nitrogen sufficiently mitigates most airborne contaminants.

7.2.1 Uranium-238 and Thorium-232 Decay Chains

Uranium-238 $\left(^{238}\text{U}\right)$ and Thorium-232 $\left(^{232}\text{Th}\right)$ are long-lived isotopes naturally present in rock, soil, and water. Each undergoes a chain of α and β decays to a stable isotope of lead, as show in Figure 7·1.

The decay daughters of ^{232}Th ($T_{1/2} = 1.4 \times 10^{10}$ years) are approximately in secular equilibrium, since the chain's second-longest half-life (5.75 years), belonging to ^{228}Ra, is smaller than its parent's by nine orders of magnitude. Moreover, the chain's only gaseous isotope ^{220}Rn, has a half-life of only $T_{1/2}$ =55.6 seconds, meaning that the only possible disruption to this equilibrium (ingress from an accidental air leak) is short-lived.

The most troublesome thorium-chain isotope is ^{212}Bi ($T_{1/2}$ = 60.6 minutes), which can β-decay to ^{212}Po (branching ratio 64%) or α-decay to ^{208}Tl (branching ratio 36%). The β-decay to ^{212}Po (half-life 0.299 μs) has a Q-value of 2.25 MeV, and is immediately followed by an α-decay with Q-value 8.95 MeV. Following quenching, this decay produces a peak below 1 MeV. Although delayed coincidence tagging can effectively remove some of these $^{212}\text{Bi} - ^{212}\text{Po}$ decays and place a constraint on the ^{232}Th rate in the detector, around half of them occur within the same trigger window. The energy of this trigger is the sum of the energies of the β and quenched α – perilously close to the $0\nu\beta\beta$ ROI. Exploiting the timing distributions of the PMTs hit during one of these 'in-window' events can effectively remove them. The α-decay of ^{212}Bi to ^{208}Tl is less a concern for the $0\nu\beta\beta$ search, since the energies of the quenched α (less than 1 MeV) and ensuing β-decay of ^{208}Tl (Q-value = 5.0 MeV) are well outside the ROI.

Like the ^{232}Th chain, the ^{238}U chain contains a gaseous isotope, ^{222}Rn , that can enter the detector through an accidental air leak and disrupt the secular equilibrium. Unlike the ^{232}Th chain, the ^{238}U chain can't quickly regain this equilibrium. This is because the half-life of ^{222}Rn ($T_{1/2}$ =3.82 days) is much longer than the half-life of ^{220}Rn, and because the remainder of the chain is interrupted by ^{210}Pb. The ^{210}Pb daughter has a half-life of $T_{1/2}$ =22.3 years, and so, is out of secular equilibrium with its quickly-decaying predecessors. Although the β-decay of ^{210}Pb has too low a Q-value to end up in the $0\nu\beta\beta$ ROI, its daughters ^{210}Bi and ^{210}Po span the entire low-energy solar-neutrino ROI.

In the event of ^{222}Rn ingress, the portion of the chain between ^{222}Rn and ^{210}Pb can regain secular equilibrium in around three weeks. After this period, coincident β-α decays of ^{214}Bi and ^{214}Po can be rejected from the $0\nu\beta\beta$ ROI and used to estimate the ^{238}U rate. Since the half-life of ^{214}Po is $T_{1/2}$ =164 μs (two orders of magnitude longer than the trigger window), and since the branching ratio of this decay mode is 99.979%, these decays can be tagged and rejected with over 99% efficiency.

7.2.2 Leaching from the AV

After decommissioning in 2007, SNO+ sat empty until 2016, with its AV exposed to air, and critically, ^{222}Rn. The decay daughters impinged on the acrylic, decaying to what is today a persistent layer of ^{212}Pb. This layer continuously pollutes the detector volume with ^{210}Bi and ^{210}Po. Although most of these species stay closest to the inner wall of the AV, confined to the region outside the FV, some of them can diffuse beyond it, carried away by eddies under fluctuating temperature.

The Q-value of the ^{210}Bi β- decay and the energy of the quenched α from the ^{210}Po decay are too low to threaten the $0\nu\beta\beta$ ROI, but can result in signal sacrifice by producing pileup signals that lead to false α–β tagging. Moreover, the α's from ^{210}Po contribute two additional backgrounds by reacting with the the Oxygen and Carbon atoms in the scintillator:

$$\begin{aligned} \alpha + {}^{13}\text{C} &\rightarrow {}^{16}\text{O} + n \\ \alpha + {}^{18}\text{O} &\rightarrow {}^{21}\text{Ne} + n \end{aligned} \tag{7.1}$$

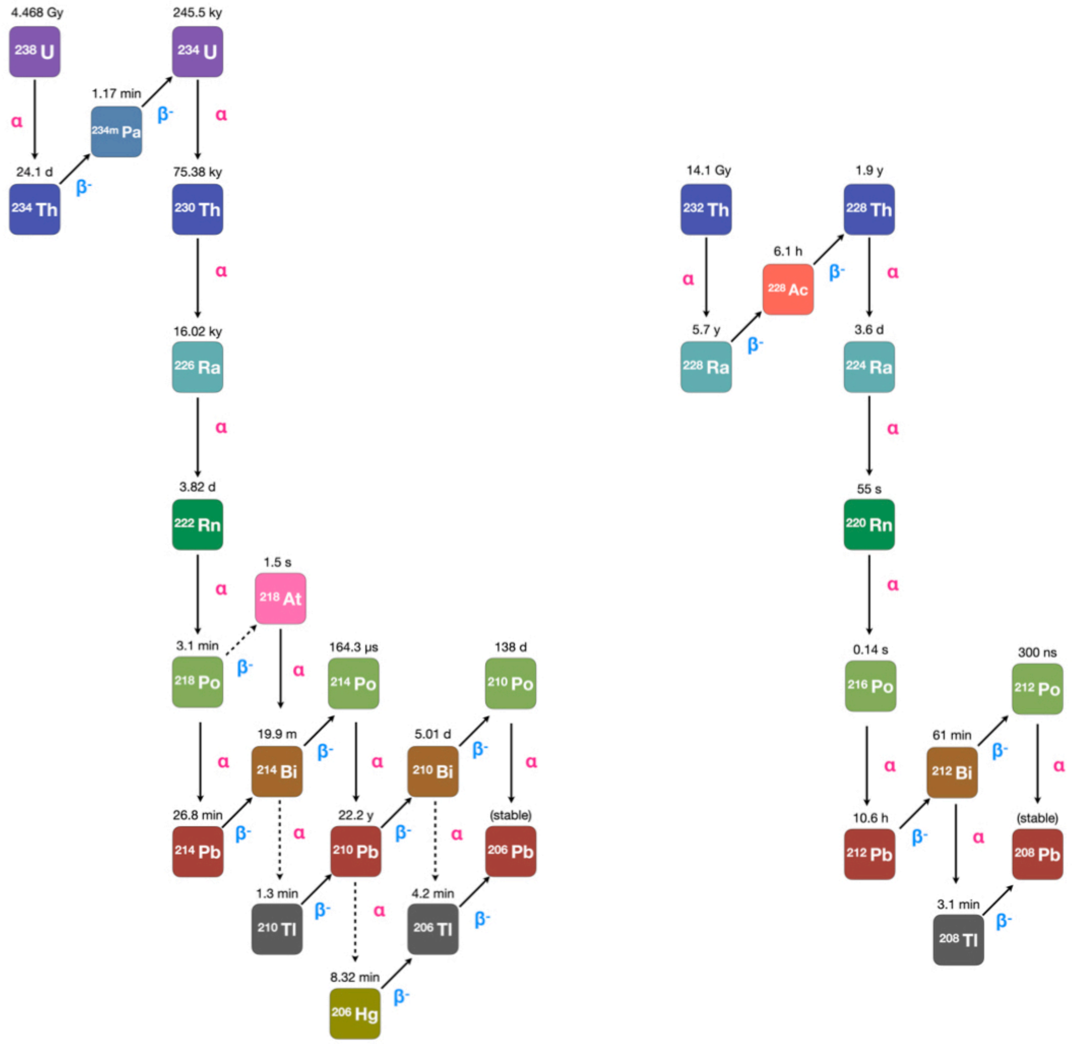

Figure 7·1: Decay chains of ^{238}U (left) and ^{232}Th (right). Half-lives are printed on top of the isotopes, while the primary and secondary decay modes are depicted with solid arrows and dashed arrows, respectively.

Isotopes shown on the right side of Equation 7.1 are produced in excited states, which promptly de-excite by emitting γ's. Under a time constant of $\sim 200\,\mu$s, the neutrons then capture on H nuclei, which de-excite by emitting a characteristic 2.22 MeV γ. Like the decays of ^{212}Bi–^{212}Po and ^{214}Bi–^{214}Po, these (α,n) reactions can also be rejected with prompt-delayed coincidence tagging.

7.3 Solar Neutrinos

Neutrinos that come from solar fusion cycles elastically scatter off atomic electrons in the target material, and these electrons recoil in directions highly correlated with the position of the sun. The SNO detector was able to resolve these trajectories using the Cherenkov light emitted along each track. In SNO+, this directional information is lost under isotropic scintillation light. Although efforts to separate and salvage Cherenkov from scintillation light are ongoing [73], this analysis treats solar neutrino elastic scattering as an irreducible background, with the expected rate given by:

$$R = \Phi n \int S_\nu(E) \left[P_{ee}(E)\sigma_e(E) + (1 - P_{ee})\,\sigma_{\mu,\tau}(E) \right] dE \tag{7.2}$$

where P_{ee} is the survival probability of ν_e, and the $\sigma_{e,\mu,\tau}$ are the electron-neutrino elastic scattering cross-sections for each neutrino flavor. The Φ and S_ν are the total and normalized energy spectrum of solar neutrino flux for a given solar fusion reaction. Between the proton-proton and carbon-nitrogen-oxygen solar fusion cycles, there are eight neutrino-producing reactions, and only the positron emission from ^{8}B produces a non-trivial flux of neutrinos with energies that span the ROI (Figure 7·2 A)). The number of electron targets n in the SNO+ scintillator, given the components' specifications from the manufacturers, is 3.357×10^{29} per tonne. With an observed total flux of ^{8}B solar neutrinos $\Phi = 5.16^{+2.5\%}_{-1.7\%} \times 10^6 \mathrm{cm}^{-2}\mathrm{s}^{-1}$ [6], numerical integration of Equation 7.2 over the entire ^{8}B solar neutrino energy range gives an expected rate of 1,630.0 events/kiloton-year. The energy spectrum of recoiled electrons was obtained through simulation.

7.4 External

External backgrounds originate in detector components outside of the scintillator: the AV, PMTs, rope systems, and ultra-pure water all contain trace amounts of radioisotopes. While the ultra-pure water attenuates the charged products of these decays, high-energy γ's, such as those from ^{40}K, ^{214}Bi, and ^{208}Tl, can propagate into the detector volume. Luckily, the number of events from such decays decreases exponentially toward the center of the detector, and rejecting data close to the AV can sufficiently remove them, leaving a clean FV for analysis. The rates of these backgrounds were estimated using data from the ultra-pure water phase; the complete analysis is detailed in Reference [**?**].

7.5 Cosmogenic Backgrounds

Cosmic rays interact with atoms in the scintillator cocktail, producing radioactive isotopes. Each component incurred most of these backgrounds during its above-ground manufacture and transport, and even underground, it remains vulnerable to spallation by the roughly 3-4 high-energy muons that make their way to the detector every hour. Some of these spallation isotopes quickly decay almost immediately, while others can remain in the scintillator much longer times of hours, days, or even years.

Spallation of the organic components' hydrogen, carbon, oxygen, and nitrogen atoms produces ^{14}C, ^{10}C, ^{7}Be and ^{11}Be. Although the Q-values of these isotopes are comfortably outside than the $0\nu\beta\beta$ ROI, their high rates (hundreds of Hz for ^{14}C) can lead to pileup. Scintillator purification efficiently ($> 99\%$) removes ^{7}Be and ^{14}C. The latter is also present naturally, and the total expected ^{14}C/^{12}C ratio is comparable to that of Borexino [75]. Rejecting data from the few minutes after a cosmic muon re-

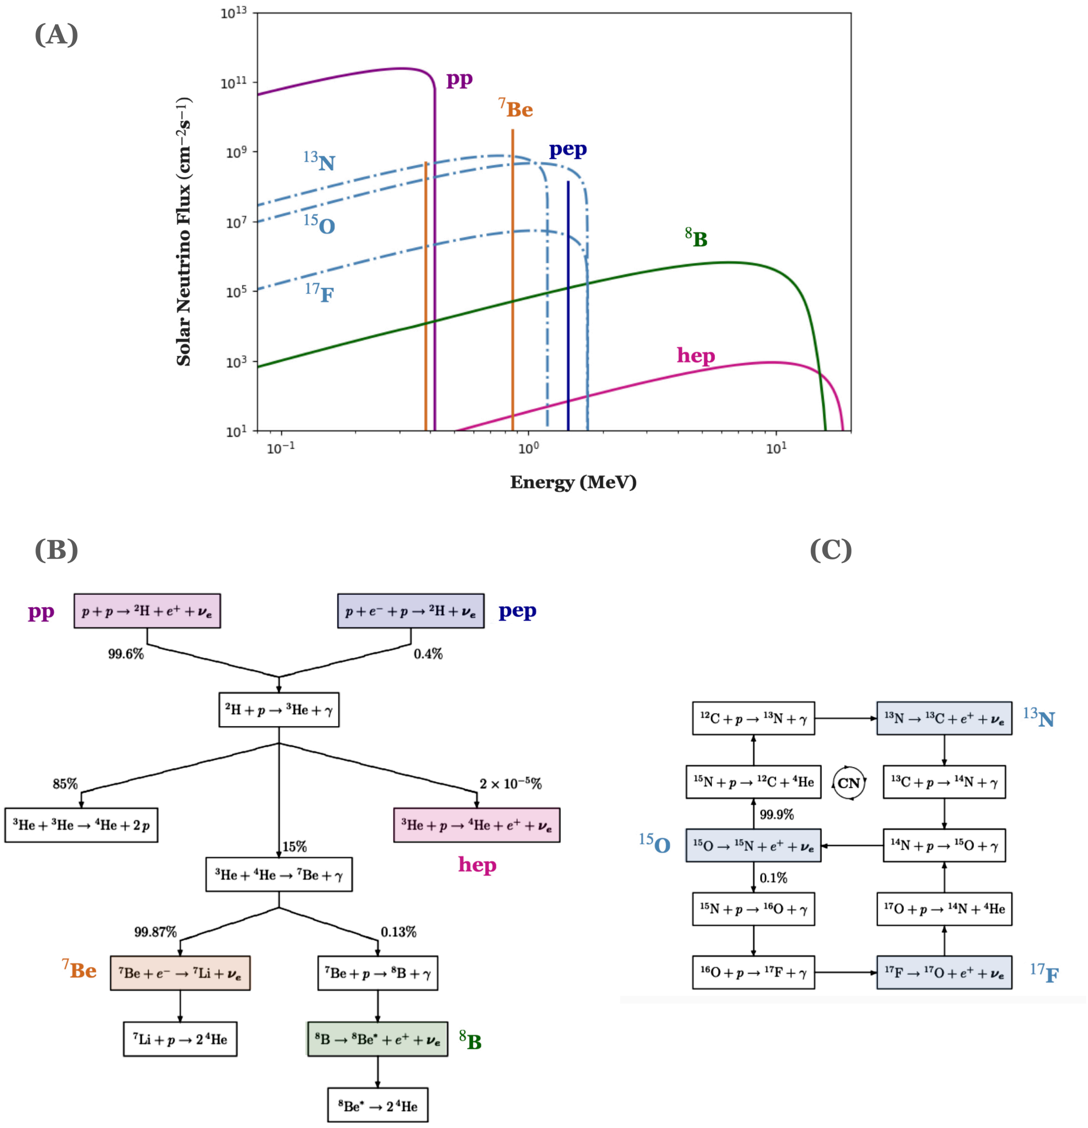

Figure 7·2: (A) The energy spectra of solar neutrinos from different solar fusion reactions (data retrieved from Reference [74]). The solid lines correspond to reactions from the proton-proton chain (B), and the dot-dashed lines correspond to those from the carbon-nitrogen-oxygen cycle (C). Flow charts adapted from Reference [6].

moves the short-lived ^{10}C ($T_{1/2}$ = 19.3 seconds) and ^{11}Be ($T_{1/2}$ = 13.8 seconds). Rejecting ^{11}C ($T_{1/2}$ = 20 minutes) takes a targeted approach like coincidence tagging, since a sufficient veto window of a few half-lives would sacrifice hours of data.

The cosmic spallation of ^{130}Te, threaten the $0\nu\beta\beta$ analysis the most. Many of the tellurium spallation products have Q-values within the ROI, as well as half-lives far exceeding months. Moreover, their concentrations scale with the addition of ^{130}Te. The telluric acid has been stored underground for several years (first batch arrival in 2015, third and last batch arrival in June 2018) to allow these contaminants to decay. Those remaining can be removed by a factor of at least 10^5 during purification. In addition, the charged current interaction of solar neutrinos with ^{130}Te creates stable and excited states of ^{130}I, which also decay close to the ROI as described in Reference [76]. The half-lives and Q-values of these Te-phase contaminants are listed in Table 7.1.

7.6 Two-Neutrino Double Beta Decay

As previously described in Chapter 3, the $0\nu\beta\beta$ signature is a distinct peak in the energy spectrum, at the Q-value of the relevant isotope. However, the energy resolution of the detector can smear the falling edge of the $2\nu\beta\beta$ decay peak thereby causing it to partially cover up the $0\nu\beta\beta$ signal. As a result, when reporting sensitivities to $0\nu\beta\beta$ decay, an asymmetric ROI is typically chosen to mitigate contamination from these two-neutrino events. Given the best limit on $T^{2\nu}_{1/2}$ for ^{130}Te (reported in Table 3.1), and a 0.5% loading of natural tellurium, the expected rate of the two-neutrino double-beta decay is 6.6×10^6 events/kiloton-year.

7.7 SNO+ Background Budget

A summary of the the background budget was calculated using a MC simulation with the best available information of the measured backgrounds during the initial water

Isotope	Q-Value (MeV)	$T_{1/2}$ (days)
^{22}Na	2.84	950.6
^{26}Al	4.00	2.62×10^{8}
^{42}K (direct and daughter of ^{42}Ar)	3.53	$0.51\left(1.20 \times 10^{4}\right)$
^{44}Sc (direct and daughter of ^{44}Tl)	3.65	$0.17\left(2.16 \times 10^{4}\right)$
^{46}Sc	2.37	83.79
^{56}Co	4.57	77.2
^{58}Co	2.31	70.9
^{60}Co (direct and daughter of ^{60}Fe)	2.82	$1.93 \times 10^{3}\left(5.48 \times 10^{8}\right)$
^{68}Ga (direct and daughter of ^{68}Ge)	2.92	$4.7 \times 10^{-2}(271)$
^{82}Rb (daugher of ^{82}Sr)	4.40	$8.75 \times 10^{-4}(271)$
^{84}Rb	2.69	32.8
^{88}Y (direct and daughter of ^{88}Zr)	3.62	106.63(83.4)
^{90}Y (direct and daughter of ^{90}Sr)	2.28	$2.67\left(1.05 \times 10^{4}\right)$
^{102}Rh (direct and daughter of ^{102m}Rh)	2.32	207.3
^{102m}Rh	2.46	1.37×10^{3}
^{106}Rh (daughter of ^{106}Ru)	3.54	$3.46 \times 10^{-4}(371.8)$
^{110m}Ag	3.01	249.83
^{110}Ag (daughter of ^{110m}Ag)	2.89	2.85×10^{-4}
^{124}Sb	2.90	60.2
^{126m}Sb (direct and daughter of ^{126}Sn)	3.69	$0.01\left(8.40 \times 10^{7}\right)$
^{126}Sb (direct and daughter of ^{126m}Sb)	3.67	12.35(0.01)
^{130}I	2.94	0.515
^{130}I	2.99	6.1×10^{-3}

Table 7.1: The Q-values, and half-lives of isotopes produced from ^{130}Te nuclei, either by charged-current interactions with solar neutrinos, or by spallation from cosmic protons and neutrons from Reference [77]. Multiple entries for Q-value and $T_{1/2}$ are for the isotope produced from the direct decay of a spallation isotope, and for the parent nuclide that feeds the reaction.

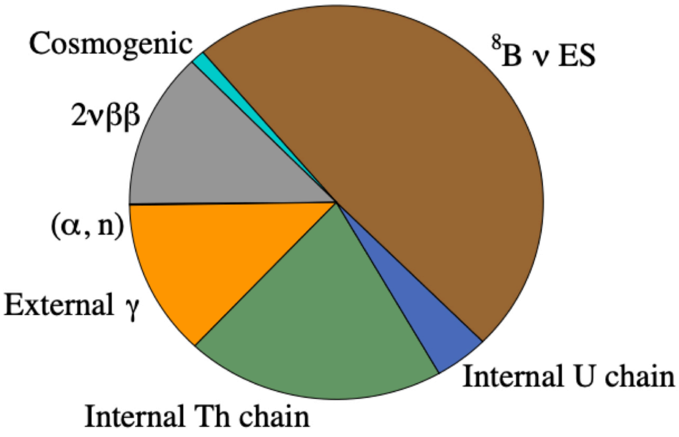

Figure 7·3: A pie chart indicating the relative conributions to the background budget inside the $0\nu\beta\beta$ ROI. Figure is adapted from Reference [79].

phase, backgrounds inferred from radioassays of materials, optical measurements of the scintillator components, and independent simulations of cosmogenic yield. Details on the background estimates and simulations are summarized in References [78] and [79]. The resulting number expected number of background counts in the ROI for a one-year exposure with the initial 0.5% Te-loaded scintillator is 9.47 events and the relative contributions from each of the backgrounds are shown in the form of a pie chart in Figure 7·3.

It's important to note that the ROI was chosen to be asymmetric such that is spans the region $-0.5\sigma < Q_{\beta\beta} < +1.5\sigma$, where σ is the standard deviation of the Gaussian-shaped $0\nu\beta\beta$ signal. This region correspond to the energy interval 2.42 – 2.56 MeV as estimated from simulation. As previously mentioned, this asymmetric ROI is chosen to minimize contribution from the rapidly falling $2\nu\beta\beta$ spectrum that spills into the ROI causing the $2\nu\beta\beta$ background to partially cover the signal. The simulated energy spectrum with all the estimated background contributions is shown in the form of a stacked histogram in Figure 7·4. The hypothetical $0\nu\beta\beta$ signal is shown in red for the

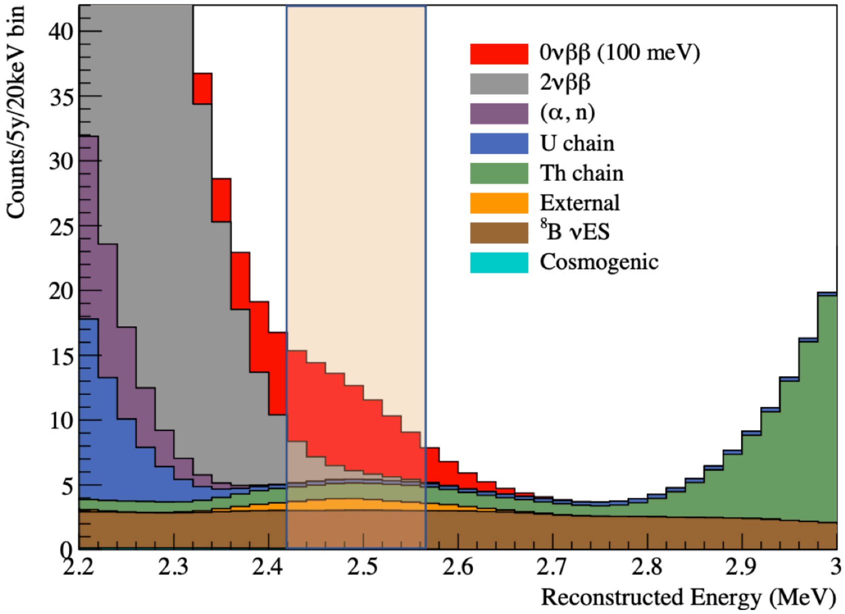

Figure 7·4: Simulated energy distribution of expected backgrounds in the $0\nu\beta\beta$ ROI after 5 years of tellurium phase data-taking with 0.5% loading of natural tellurium and a FV of 3.3 m. Figure adapted from Reference [79].

case of $m_{\beta\beta} = 100\,\text{meV}$ using the IBM-2 NME calculation ($M^{0\nu} = 4.03$) and a phase space factor $G^{0\nu} = 3.688 \times 10^{-14}\,\text{y}^{-1}$.

Using the background estimates reported in References [78] and [79], and the help of Equation 3.17, we can perform a simple counting analysis to predict the half-life sensitivity for SNO+ to be $S^{0\nu}_{1/2} > 2 \times 10^{26}$ years at 90% C.L. after three years of data-taking. This offers a kind of benchmark that SNO+ will aim for as the experiment gets closer to tellurium deployment. This sensitivity prediction will change as more information becomes available from new estimates of backgrounds in the ROI and evolving detector response due to the addition of scintillator components and tellurium.

Chapter 8

Background Rejection with KamNet

One of the main challenges for the $0\nu\beta\beta$ decay on SNO+ is background suppression. The KamLAND-Zen $0\nu\beta\beta$ decay experiment is very similar in many ways to SNO+, and they designed a novel deep learning algorithm to help them achieve extremely effective background suppression. The algorithm is known as KamNet [80] and while it was initially tailored to the experimental parameters of KamLAND-Zen, its general approach and architecture are transferable to other $0\nu\beta\beta$ decay experiments. The main goal of this thesis is to study the performance and application of KamNet on the SNO+ low-background scintillator data.

This Chapter briefly summarizes basic machine learning vocabulary and concepts, and then goes on to describe the components of KamNet network and its ability to separate single and multi-site events. After applying KamNet to SNO+ data and MC simulation, we use KamLand-Zen as a benchmark to quantify its performance while keeping in mind each experiments' differences in hardware and backgrounds. We find that the network produces similar background and signal separation on both experiments, a result consistent with their similar vertex resolutions. Finally, we compare KamNet's performance, using spaciotemporal data, to that of two existing classifiers that use only PMT time residuals. We find improved background and signal separation over each.

8.1 KamNet: A Spacio-temporal Deep Neural Network

KamNet represents a novel addition to a nearly 60-year-old class of machine learning models: Neural Networks. Although they have undergone several decades of performance improvement and resource optimization, all neural networks retain the same basic concepts of their common ancestor the perceptron, an algorithm used in pattern recognition. We will use this simplest case to illustrate the operation of a neural network and establish vocabulary used in the rest of this Chapter. A more thorough discussion of neural networks can be found in Reference [81].

A perceptron is a function that returns 1 if the weighted sum of its binary inputs $\vec{x}$ passes some threshold value b (the *bias*), and 0 otherwise. In essence, the weighted sum of the inputs produces some value, which 'turns on' a component in the next layer, if that value exceeds the corresponding bias. This can be expressed mathematically as follows:

$$f(\vec{x}) = H\left(b + \sum_i w_i x_i\right) \tag{8.1}$$

Qualitatively, the weights w_i represent how important an input is in determining the output. The machine "learns" to find the correct system of weights such that the output satisfies known conditions. A visual representation of a perceptron is shown in Figure 8·1.

One can initialize the network's system of weights and biases, and then pass it some input with a known output. The *loss function* gives the distance between the correct output and the machine's output. We train the network to recognize the correct answer by having it classify several known inputs from a *training dataset*, and minimize the loss function at each iteration over the entire dataset. An iteration over the entire dataset is called an *epoch* and is an example of a *hyperparameter*, which, unlike the system of weights, does not change during training. If a machine learning model is

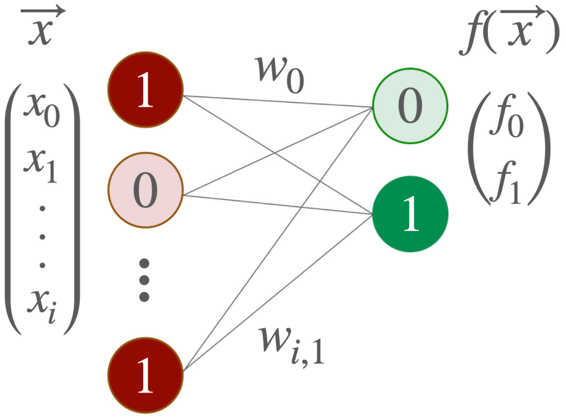

Figure 8·1: A visual representation of a perceptron, in which each input in $\vec{x}$ (red) has an associated weight w_i, and each output (green) has an associated bias b_i.

trained for too few epochs, it may *underfit* the data, thereby failing to capture underlying patterns and resulting in poor performance on both training and validation sets. If it is trained for too many epochs, it may *overfit*, meaning that it's capturing statistical fluctuations specific to the training set and will perform poorly on unseen data. Generally speaking, underfitting can also occur if, for example, the underlying model is too simple. Likewise, overfitting can sometimes be mitigated even with many epochs using regularization techniques, dropout, or early stopping based on validation performance.

Since perceptrons accept and return binary values, their ability to learn is somewhat limited as the length of $\vec{w}$ increases. By allowing f to be nonlinear (eg. like the sigmoid or rectified linear unit (ReLU) functions), and also allowing the inputs and outputs to have values between 0 and 1, a subtle change in each weight can produce a subtle change in the output. This altered perceptron is called a neuron, and these modifications result in a smooth parameter space for the loss function, which can then be minimized through gradient descent. The machine accomplishes this through an algorithm called backpropagation, which is beyond the scope of this thesis but thoroughly described in Reference [82].

Finally, layering these neurons produces produces a deep neural network. An arbitrary number of layers with an arbitrary number of neurons per layer allows the machine to efficiently produce a general and accurate model. However, a simple fully-connected neural network, like the example in Figure 8·1, has a few drawbacks when applied to time-ordered sets of three-dimensional images (i.e., events captured in a spherical detector). What follows is a brief description of these issues and the modifications KamNet uses to overcome them, namely, the combination of a spherical convolutional neural network, long short-term memory, and an attention mechanism.

8.1.1 Spherical Convolutional Neural Network

Traditional convolutional neural networks (CNN) perform two-dimensional convolutions to pick out certain patterns in closely-spaced regions of an image. Doing so reduces the dimensionality of the problem, making CNNs especially efficient in image recognition. In this structure, the input neurons (or, in the case of an image, the pixels) are grouped into local receptive fields. Each $l \times m$ receptive field (also called 'kernel' or 'filter') connects to a hidden neuron through $l \times m$ weights and biases. The value of each hidden neuron N_{ij} is the sum of the products of the overlapping filter and image pixels:

$$N_{ij} = A\left(b + \sum_{0}^{l}\sum_{0}^{m} w_{lm} I_{i+l,j+m}\right) \tag{8.2}$$

Equation 8.2 is a convolution of the filter and image, written for some activation function A and one filter shared over one layer of hidden neurons. The filter's weights w_{lm} and bias b are therefore shared among all of the hidden neurons. This system of shared weights and biases allows the machine to pick out a particular feature in different locations I_{xy} of the image. Rather than learning a system of unique weights in a fully-connected network, a CNN learns a system of filters. The layer of neurons following a convolution is called a feature map. In principle, a CNN contains several of these. Following the feature maps are pooling layers, which group the features into locations relative to other features, further reducing the number of parameters.

While flat image processing is handled easily by a CNN, it doesn't extend well to images with spherical symmetry. World map projections are a good example of the irregularities and inaccuracies encountered when trying to stretch a sphere into a flat surface. Moreover, a conventional CNN doesn't reflect the rotational invariance of an image. In the case of a spherical detector, one can imagine any unique arrangement of hit PMTs, such as a ring, compressed into a $\theta \times \phi$ grid, where θ and ϕ are polar coordinates of the PMTs. This ring can have completely different $\theta - \phi$ maps, depending on where it's located. KamNet incorporates a spherical convolutional neural network (sperical CNN) to eliminate spacial distortion and preserve rotational invariance of PMT hit maps. Instead of a two-dimensional filter sliding in cartesian space, a cpherical CNN uses a a three-dimensional filter, in which each cell represents a rotation in Euler angle space. The convolution proceeds through the $SO(3)$ Fourier Transform (SOFT) algorithm described in Reference [83].

8.1.2 Long Short-Term Memory

Generally, CNNs find and pool features in images with no knowledge or preservation of their order. However, patterns of hit PMTs aren't the only piece of information conveyed by an event in a scintillator detector. As discussed earlier, different classes of events proceed on different timescales. Different radioisotopes, such as ^{214}Po ($T_{1/2} = 1.64 \times 10^{-6}$ s) and ^{10}C ($T_{1/2} = 1.93 \times 10^{1}$ s), can have half-lives that differ by several orders of magnitude. Moreover, as with these examples, decays can proceed through different steps: ^{10}C decays directly to ^{10}B, while ^{214}Po decays in coincidence with ^{214}Bi. Most PMTs are fast enough to capture decays that happen within microseconds of one another , making coincidence tagging a reliable method of removing these backgrounds from data. There are more subtle timing differences within the trigger window that can be exploited, and are not visible through reconstructed quantities and coincidence tagging.

A recurrent neural network retains and uses the order of information in its learning and decision making by taking as its input both new information and recent output. Recurrent networks can be unwieldy and slow, however, suffering from the 'vanishing gradient problem' [84]. Long short-term memory (LSTM) is a type of recurrent neural network that skirts this problem by teaching the machine when to 'forget' information irrelevant to its task. More details can be found in Reference [85].

At this stage, we have described two key pieces of information contained in detector events and two long-established methods KamNet uses to incorporate them in its classification. However, there is one subtlety: these data represent movies, instead of isolated images or ordered lists of vectors. The machine's task is to consider all of this information simultaneously and classify a time-ordered set of images. In other words, the LSTM must take in, and pass, features obtained from convolutional layers. The convolutional LSTM (ConvLSTM) network was first proposed in 2015 [86] to predict rainfall intensity using a set of radar maps taken every 6-10 minutes. This type of network applies to other spaciotemporal inputs, such as the evolving PMT hit maps in an event.

8.1.3 Attention Mechanism

The final piece and novel contribution of KamNet to the deep learning community is the addition of an attention mechanism to the ConvLSTM network. Attention mechanisms have been around since 2016 and are widely used in language translation [87] . Their use in KamNet addresses a problem unique to scintillator detectors: the scintillator timing.

The most populated PMT hit maps fall on and around the scintillation peak. Although these maps contain most of the event's information -the most detailed images- the sparsely hit maps and their temporal locations can also hold clues about how the event unfolds. The attention mechanism allows the machine to focus on and weight the individual time slices that contain the most identifying information.

In practice, this ability allows the machine to capture different event topologies. In keeping with our example, ^{10}C doesn't only decay to stable ^{10}B by positron emission; the positron goes on to either find an electron and produce positronium, and eventually annihilates to produce a pair of γs, which Compton scatter. This series of events delay the energy deposition, and precede an additional emission of a 718-keV γ from an excited state of ^{10}B. This series of steps are lost in reconstruction, but appear to capture the attention of the machine, which places an 'attention score' on these time slices. The decay of ^{214}Bi presents an entirely different series of steps in which a gamma cascade follows the initial β^- decay by ~ 25 ns.

This example, analyzed with KamNet in Reference [80], illustrates the ability of the AttentionConvLSTM to distinguish between event topologies. Herein lies its classifying power, and relevance, to $0\nu\beta\beta$ decay: the separation of multi-site (background-like) and single-site (signal-like) events which is illustrated in Figure 8·2. The first category describes many of the radioactive backgrounds in SNO+, as β decays are often accompanied by γ's that veer off the β's path into multiple Compton scatters. The second category involves only the scattering of β's, which deposit their energy close to the event origin (within the vertex resolution) and are thus point-like.

8.2 Training

As mentioned previously, a neural network consists of layers of neurons linked to one another by a system of weights and biases. The output of one layer affects the output of the next layer, and so on. Ideally, the output of the final layer is some configuration

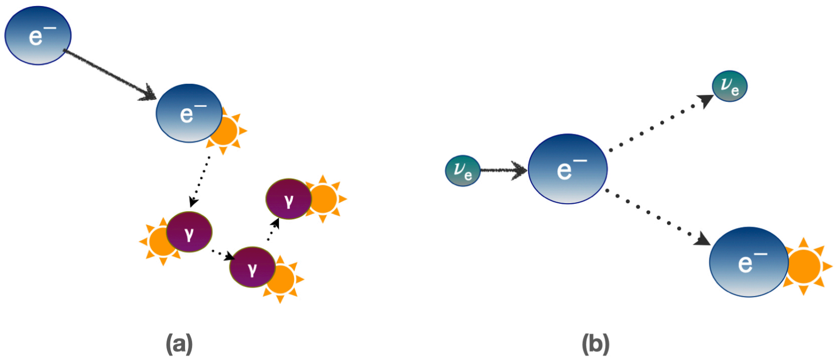

Figure 8·2: A cartoon of a) multi-site and b) single-site events.

of neurons representing the correct answer. When we 'train' a network, we ask the machine to find the system of weights and biases that consistently produce the correct answer for a variety of inputs. The machine 'learns' by minimizing the difference between its answer and a known answer. The machine does this iteratively over many known answers, and this set of known answers is called the training dataset.

8.2.1 Input Data Structure

Recall that PMTs can be hit at different times during an event. This feature is exploited by vertex and time reconstruction through a quantity called the time residual T_r, restated here:

$$t = t_{\mathrm{PMT}} - \tau - T_r \tag{8.3}$$

where t is the reconstructed time of the event, t_{PMT} is the time of the PMT hit, and τ is the scintillation photon's time-of-flight. The time residual accounts for the production and attenuation of the scintillation light and exactly matches the scintillator timing profile under perfect reconstruction. Some differences in event topology and event type can show up in the time residual, most notably the pronounced slow component of α's, relative to β's.

To build the spatiotemporal data structure, the time residuals are calculated for all PMTs hit during an event, and binned from -5.0 to 45.0 ns. The bins are 1.0 ns wide to reflect the timing of the PMTs. In each time bin is a 38×38 sparse matrix of PMT hits in $\theta \times \phi$. The final data structure is a $50 \times 38 \times 38$ hit map in t, θ, and ϕ. The input to KamNet is essentially a movie, rather than a picture, in that it is a time-ordered series of images. This is illustrated in Figure 8·3.

Plot a) in Figure 8·3 is a 2D histrogram providing a visual representation of the sparse matrices, with each pixel representing one of the $\theta \times \phi$ bins, and the cool-to-warm color gradient representing the fewest-to-most PMTs in each bin. Plot b) is a visualization of these bins' locations around the detector. In this example, the matrix pictured falls into the ninth time bin, highlighted by the solid color in Plot c). Plot c) is the hit time residual where the bins chosen for the time residual distribution reflect the timing of the SNO+ readout electronics, roughly 1.0 ns. The y-axis indicates the

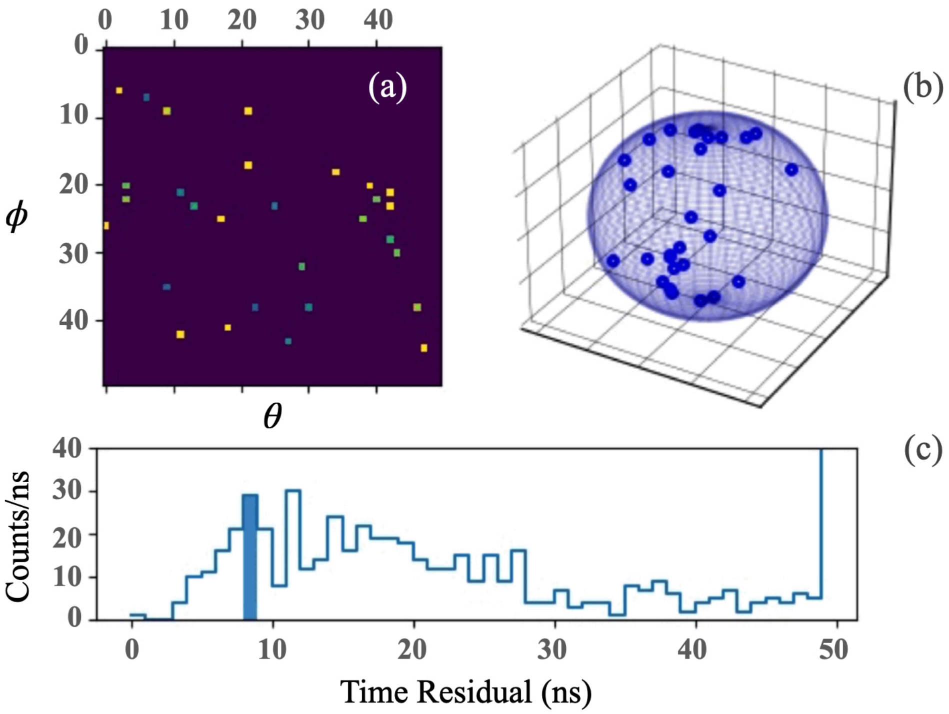

Figure 8·3: Example KamNet input showing the spatiotemporal structure of the data.

number of hit PMTs in that bin, or the sum of the bin contents in Plot a). The last bin of this hit time residual histogram is an overflow bin, chosen due to the lack of additional information gained at very late hit times.

8.2.2 Training Dataset

The training dataset used for this model consists of 20,000 simulated events each of ^{8}B solar neutrino elastic scatters and ^{214}Bi β^{-} decays. These events were simulated in a 3.3 m-FV of LAB and 2.2 g/L PPO, reflecting the conditions of the real dataset used in the final analysis described in Chapter 9. Care was taken to prevent the machine from distinguishing ^{214}Bi and ^{8}B solar neutrinos based on the number of hit PMTs alone. Early studies with KamNet revealed that when there is a significant difference in Nhit between signal and background events, KamNet tends to rely primarily on Nhit for classificationignoring other potentially useful features. To address this, we apply a technique called 'Nhit matching' before training. In this process, both signal and background events are sampled from a common Nhit distribution constructed from the region where their original Nhit distributions overlap. This encourages the network to learn more discriminating features beyond just Nhit.

The choice of ^{214}Bi β^{-} decays and ^{8}B solar neutrino scatters for the training dataset is motivated by the interpretability studies carried out during its most recent development for the KamLAND-Zen experiment, described in Reference [80]. To summarize the collaboration's finding: KamNet was evaluated on simulations of ^{136}Xe $0\nu\beta\beta$ decay, solar neutrino scattering, β^{-} decays of ^{214}Bi, β^{+} decays of ^{10}C, and several long-lived spallation isotopes. While retaining 90% of the $0\nu\beta\beta$ signal, KamNet was most able to reject the radioactive decays, and least able to reject solar neutrino scatters. Since single-β decays are accompanied by gamma cascades, each radioactive decay event consists of two or more separated vertices. Solar neutrino scatters and $0\nu\beta\beta$ decays, however, do not contain multiple vertices. This indicates that a significant amount of KamNet's classifying capability is in distinguishing multi-vertex and single-vertex events.

8.2.3 Validation and Figure of Merit

During the training, an additional 5,000 events were repeatedly classified using the iteration of KamNet present at the end of each training 'epoch', a cycle in which the network has seen all of the training data. Quizzing the network this way can reveal, and prevent, the use of a model that is either underfit (the validation and training accuracies are both low) or overfit (the validation accuracy is low and the training accuracy is high). An optimally trained network shows both minimal and comparable loss between the training and validation datasets. An optimally trained model obtained on the 14th epoch was chosen for the remainder of this analysis. The classifier value output by KamNet is the KamNet score, with higher score being signal-like. The normalized distributions of KamNet scores obtained by running the trained model over simulations of multi-site (^{214}Bi, ^{10}C) and single-site (^{130}Te$0\nu\beta\beta$, ^{8}B solar ν) events is shown in Figure 8·4. The shift in the multi-site distributions relative to the single-site distributions make the KamNet score a viable cut for data selection in a $0\nu\beta\beta$ decay analysis. The near-perfect overlap of the single-site distributions alongside the slight offset between the multi-site distributions reflects the network's use of γ activity, and/or the slightly elongated hit-time distribution from ortho-positronium coming from the β^+ decay of ^{10}C, in its decision making. The KamNet score cut value chosen for the remainder of this work is the value above which 90% of the validation signal is retained.

8.3 Benchmarking and Model Validation

To quantify KamNet's performance on SNO+, we compare the SNO+ classification ability its original evaluation on KamLAND-Zen data and Monte Carlo as a benchmark. The trained KamNet model was run over SNO+ simulations of ^{130}Te $0\nu\beta\beta$ decay, ^{8}B so-

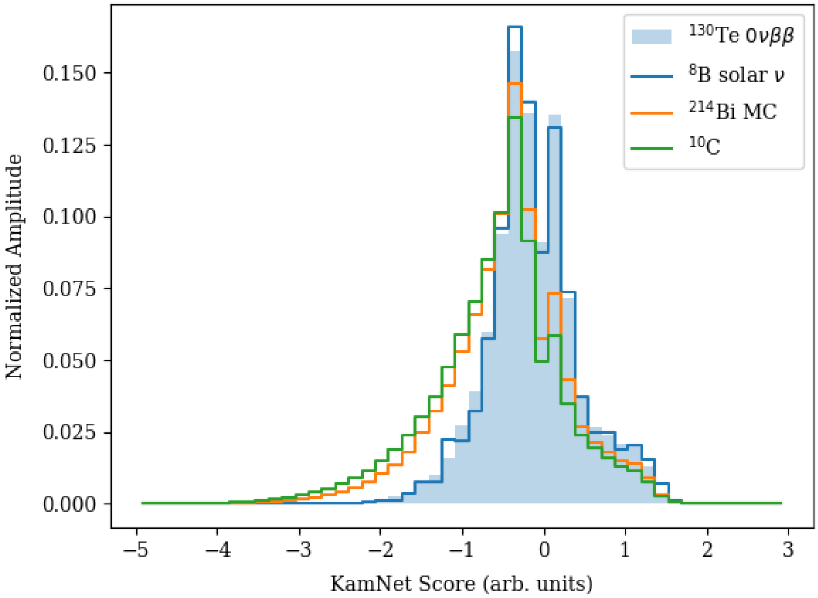

Figure 8·4: Normalized distributions of KamNet scores obtained by running the trained model over simulations of multi-site (^{214}Bi, ^{10}C) and single-site (^{130}Te $0\nu\beta\beta$ decay, ^{8}B solar ν) events.

lar ν, and the β decays of ^{214}Bi and ^{10}C. As mentioned earlier, these simulations represent two event classes: single-site and closely-spaced multi-site. At 90% signal acceptance, KamNet showed rejection efficiencies similar to those obtained for KamLAND-Zen [80]. A summary of these results is shown in Table 8.1 where the AUC is the area under the receiver operating characteristic curve. The higher this value, the better the separation between signal and background. The Rejection efficiencies ϵ are quoted for 90% acceptance of $0\nu\beta\beta$ events. For the SNO+ values, all event types were simulated in a 3.3 m FV, and the reported rejection efficiencies were calculated from events with energy between 2.0 and 3.0 MeV.

This behavior is expected, because both KamLAND and SNO+ have similar vertex resolution and thus produce input images of roughly the same clarity [88]. From a careful γ calibration campaign, KamLAND determined their vertex resolution to be $13.7\,\mathrm{cm}/\sqrt{\mathrm{MeV}}$. In absence of a similar calibration, this work used the MC study to

Isotope	Event Type	KamLAND AUC	SNO+ AUC	KamLAND ϵ	SNO+ ϵ
$0\nu\beta\beta$	signal	0.5	0.5	10%	10%
^{8}B solar ν	e^{-}	0.49	0.49	9.5%	9.8%
^{10}C	$\beta^{+}+\gamma$	0.72	0.68	40.0%	34.5%
^{214}Bi	$\beta^{-}+\gamma$	0.65	0.63	27.0%	26.2%

Table 8.1: Performance summary of KamNet in terms of the AUC and rejection efficiency ϵ. Event Type indicates the multi-site ($\beta^{+/-}+\gamma$) or single-site (e^{-}) nature of the event.

estimate a SNO+ vertex resolution of $12.6 \pm 0.7\,\text{cm}/\sqrt{\text{MeV}}$.

As discussed in Chapter 6, the vertex resolution in SNO+ has some position dependence throughout the AV with slightly better performance at larger radial distance from the center. This behavior was also present in KamNet's performance as demonstrated in Figure 8·5, which shows the z versus ρ^2 distributions for simulated ^{8}B and ^{10}C events. Each set of events were subjected to a cut of 90% ^{8}B acceptance. For both sets of events, KamNet performed slightly better at classifying events as multi-site compared to single-site at locations farther from the center of the AV. This is consistent with the slightly improved vertex reconstruction resolution at locations farther from the center of the AV.

To show that KamNet's performance on simulation agrees with its performance on data, we ran the network over β^{-} decays of ^{214}Bi. Because these decays are followed by α decays of ^{214}Po, they can be efficiently isolated from data through coincidence tagging. In this study, the inclusive coincidence tagging criteria in Table 9.1 were applied to both data and MC simulations of ^{214}Bi–^{214}Po. An additional cut in energy of $(2.0, 3.0)$ MeV was applied to select only the ^{214}Bi β^{-} decays . The trained network was then run over each of these event classes and on new simulation of ^{8}B solar ν, producing the KamNet score distributions in Figure 8·6. A cut defined at 90% acceptance of the solar ν rejects 29.0% and 28.2% ^{214}Bi β^{-} decays in simulation and data, respectively, such

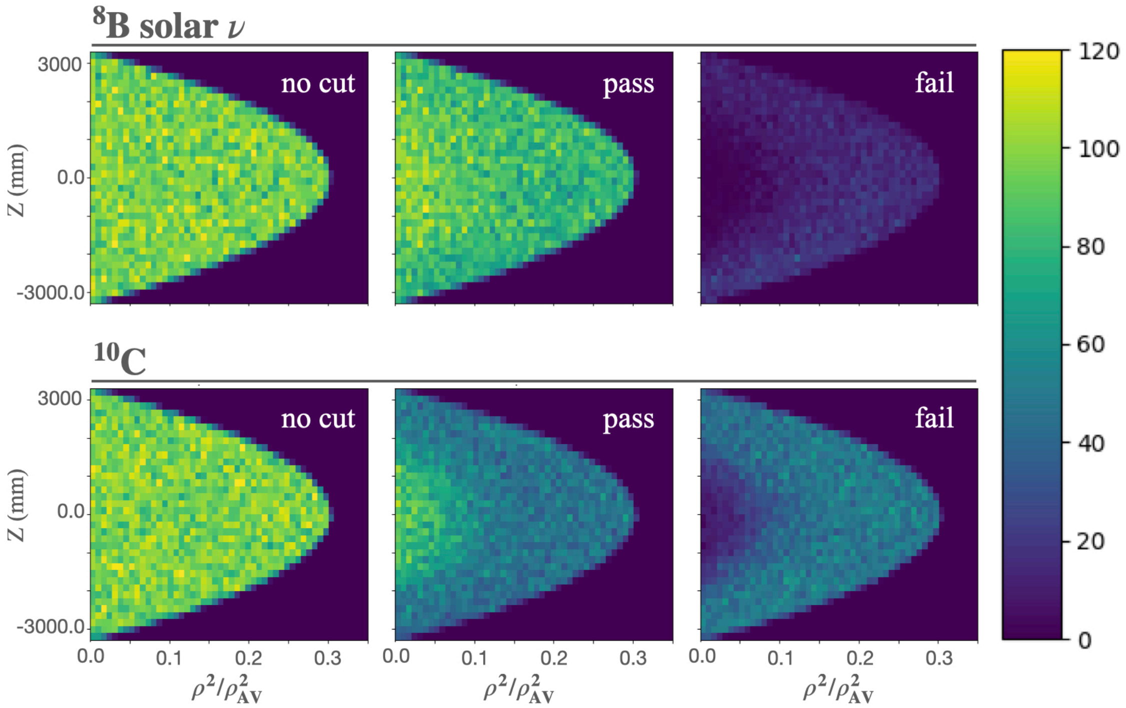

Figure 8·5: KamNet performance shown for z versus ρ^2 distributions for simulated ^{8}B and ^{10}C events. The far left plots show the total number of generated events inside the volume. The central, and right, plots show the number of events passing the KamNet score cut and failing to pass the KamNet score cut, respectively. The color gradient indicates the number of events per unit volume.

that they agree within 2.8%.

8.4 Comparison to Existing Event Classifiers

The in-time ratio (ITR) classifier described in Chapter 6 is the fraction of PMTs hit early in an event, with time residuals in the range $[-2.5, 5.0]$ ns. Originally optimized to identify early Cherenkov light in SNO, this window also provides a loose measure of goodness-of-fit for timing and vertex reconstruction in the scintillator phase by selecting time residual distributions peaked at zero. Some events have distinct timing profiles with features obviously captured by ITR. For instance, α decays produce a noticeable broadening of the slow component of the scintillator timing profile and have low values of ITR. However, the low-energy γ cascades that accompany the many non-

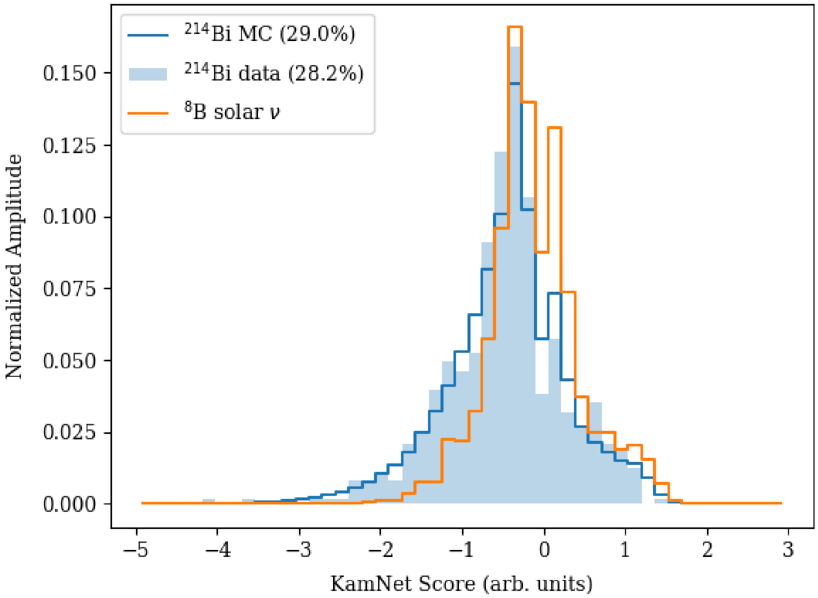

Figure 8·6: The overlap of KamNet score distributions, showing agreement between rejection efficiences of MC simulated ^{214}Bi and coincidence-tagged ^{214}Bi in real data at a 90% acceptance of ^{8}B MC.

coincident β^- decay backgrounds in SNO+ don't always produce such obvious features, and their identification requires more information.

The study in Figure 8·7 was carried out on simulations of ^{214}Bi β^- decay and ^{130}Te $0\nu\beta\beta$ decay. In order to avoid considering energy-driven decisions, equal numbers of events with between 535 and 545 PMT hits were considered. At a cut of 90% ^{130}Te $0\nu\beta\beta$ acceptance, KamNet produced a $10.5 \pm 2.1\%$ increase in rejection of ^{214}Bi compared to ITR. This substantial increase is expected given that ITR is not optimized for event classification in scintillator.

Another dedicated classifier used to identify multi-site events was recently developed for use on SNO+ [89, 43]. This multi-site classifier is based on a likelihood approach and uses more information than a just a fraction of early PMT hits. This classifier takes advantage of the subtle differences in the shapes of the timing residuals for multi and single-site events. The classifier is linear with a discriminant of the form:

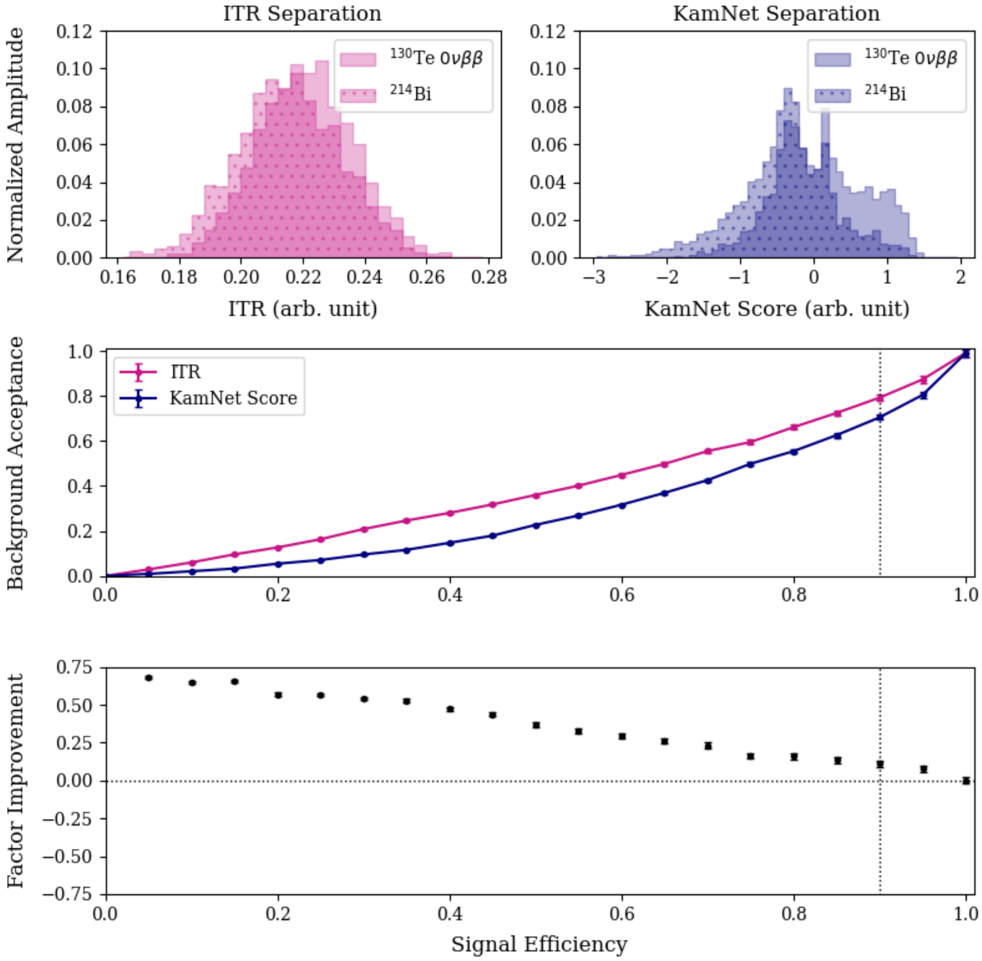

Figure 8·7: Distributions in ITR (pink) and KamNet score (blue) for ^{130}Te $0\nu\beta\beta$ decay and ^{214}Bi are shown at the top. Background versus signal acceptances for each classifier, with 90% acceptance marked with a vertical line is shown in the middle plot. This value is also marked on the bottom plot, which shows the factor change in background acceptance, (ITR – KamNetScore/ITR). The horizontal line indicates no change, while values above indicate improvement in signal and background separation. The error bars in the middle and bottom plots represent the statistical error and are difficult to see due to the large number of MC events used in the study.

$$D = \vec{w} \cdot \vec{T} \tag{8.4}$$

where D is the classifier value, $\vec{w}$ is some vector of weights determined from MC, and $\vec{T}$ is the vector of the bin contents in a classified event's time residual distribution. The most powerful choice of weight is defined as:

$$w^i = \log\left(\frac{\mu^i_{0\nu}}{\mu^i_{\text{multi-site}}}\right) \tag{8.5}$$

The μ^i represent the probability of a PMT hit in time residual bin i for a signal (0ν) and background (multi-site), respectively. In this form, the discriminant becomes the log-likelihood ratio:

$$D = \Delta \log \mathscr{L} = \sum_{i=0}^{Nhits} \log\left(\frac{\mu^i_{0\nu}}{\mu^i_{\text{multi-site}}}\right) \tag{8.6}$$

According to the Neyman-Pearson lemma, the log-likelihood ratio provides the optimal discriminant when the PDFs of both signal and background are completely known. However, for multi-site events, the average PDF may not accurately represent individual events. For instance, longer-traveling events have more late hits and fewer near the peak, creating inter-bin correlations that the log-likelihood ignores. An alternative approach is to use the Fisher discriminant which include a covariance matrix to account for correlations:

$$\vec{w} = (\Sigma_{\text{multi-site}} - \Sigma_{0\nu}) \cdot \left(\vec{\mu}_{0\nu} - \vec{\mu}_{\text{multi-site}}\right) \tag{8.7}$$

where $\vec{\mu}$ are the same as before (vectors of mean PMT hit probability from the PDFs) and the Σ are the covariance matrices. Studies described in Reference [43] found that the Fisher discriminant was not as powerful, and it was hypothesized that the Fisher discriminant approach was limited by statistics ($\mathcal{O}(10^2$ events)) available for the creation of the background covariance matrix. For the purpose of comparison to KamNet we continue with the multi-site classifier define based on the weights in Equation 8.5.

To compare with KamNet's performance, a new model was created using a training dataset of ^{60}Co and ^{130}Te $0\nu\beta\beta$ decays within a 4-m FV to match the simulation used in the multi-site classifier's original development. The results of the network's performance on additional ^{60}Co and ^{130}Te $0\nu\beta\beta$ MC is summarized in Figure 8·8. The values for the multi-site classifier were obtained from Reference [89]. At a cut of 90% signal acceptance, KamNet provides a slight but noticeable improvement of $4.2 \pm 2.1\%$. Future studies, with higher statistics MC and real datasets, could provide a more robust comparison between the two methods.

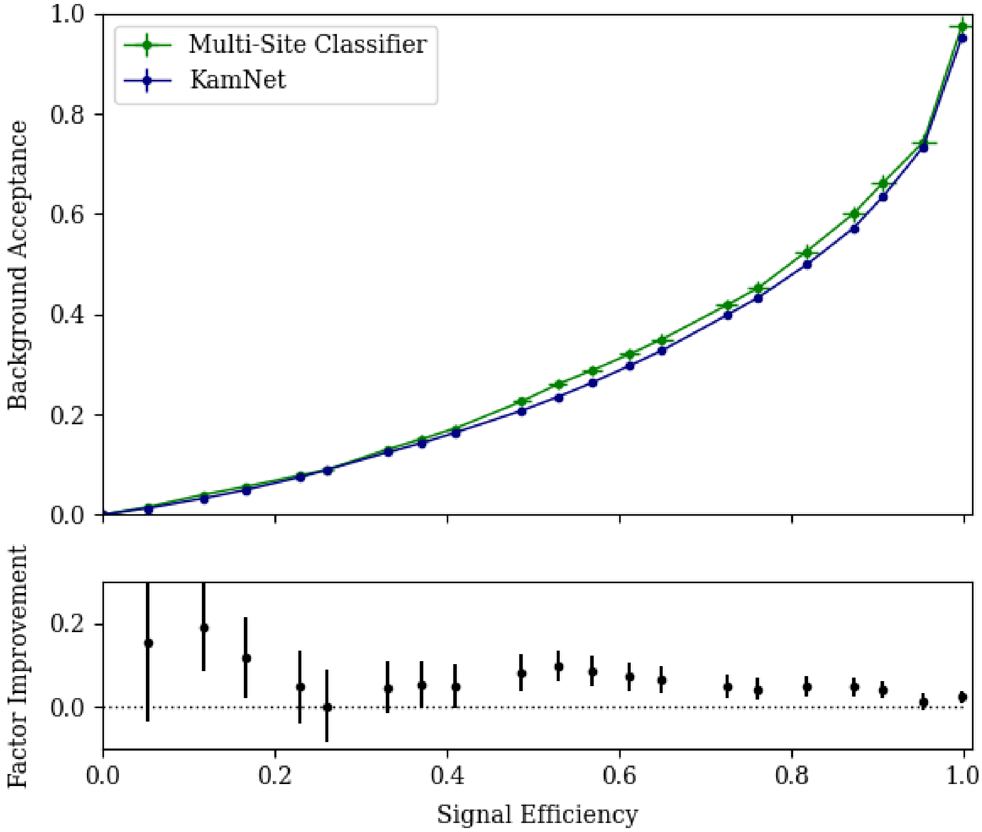

Figure 8·8: Plot of background versus signal acceptance for the multi-site (green) and KamNet score (blue) classifiers for background acceptance at 90% signal acceptance. The error bars in the bottom plot are statistical and come from a limited number of MC events.

Chapter 9

$0\nu\beta\beta$ Sensitivity Analysis with KamNet

After training and validating KamNet to classify $0\nu\beta\beta$ decay signal and backgrounds, it will now be applied to the real low-background scintillator data in SNO+. This Chapter summarizes the low-background analysis of a select period of scintillator data-taking, including the data selection, traditional cuts applied to the data, and the likelihood fitting procedure used to estimate the scintillator background rates. The performance of KamNet is evaluated by comparing the fitted background rates in the scintillator data before, and after, KamNet is applied. Finally, the impact of applying KamNet to the $0\nu\beta\beta$ phase of SNO+ is studied with sensitivity calculations, assuming the present background rates are maintained after tellurium loading.

9.1 Data Selection, Analysis Cuts, and Livetime

The SNO+ liquid scintillator fill was completed on April 29, 2022, and first physics data was collected immediately after scintillator recirculation (starting with run number 300000). The end of the data-taking period in this thesis is March 10, 2023 (corresponding to run number 309277). Each data-taking run is roughly an hour in length. Out of the collection of runs selected for the analysis presented in this thesis, a subset were selected for ideal running conditions: appropriate physics triggers applied for low-background data acquisition, AV recirculation turned off, no personnel on deck, quiet mine conditions (no ‘rock blasts’), and all electronics crates turned on (details in [90]). Hereafter, these runs are referred to as *golden runs*, and they were subjected to the analysis cuts described in the following section. Data cleaning cuts were made to remove neutron followers, special classes of instrumental backgrounds, and artifacts from electronics. An additional background called a re-trigger occurs when a single event, especially those that produce PMT pileup, triggers the detector twice. Events that occur with $\Delta t = 460\,\mathrm{ns}$ of one another are considered re-triggers and are cut from the data.

The first physics data-taking in the scintillator phase began with some residual ^{222}Rn contamination from the distillation system. The data used in this analysis was taken 869 runs ($\sim$ 36 days) later, after the radioactivity cooled to secular equilibrium ($T_{1/2} = 3.82$ days). In general, radon contamination from both the ^{238}U and ^{232}Th chains can be monitored from measurements of ^{214}Bi and ^{212}Bi decays, respectively. This is because each of these isotopes decays in coincidence with an α-decay of $^{214/212}$Po, a feature that shows up either in the trigger window in the form of altered time residuals (see $\alpha-\beta$ classifiers in Chapter 6), or as separate triggers that can be selected with the coincidence tagging criteria.

The only coincidence tagging deployed for this analysis is for ^{214}Bi–^{214}Po, since most ^{212}Bi–^{212}Po decays occur within the trigger window and are removed with the $\alpha-\beta$ classifier. Looping through events in a single run, unique pairs of events that individually pass ^{214}Bi and ^{214}Po cuts within a spacial separation Δr and time interval Δt are tagged as ^{214}Bi–^{214}Po coincidences and can be removed from data. To efficiently capture nearly all of the possible ^{214}Bi–^{214}Po decay pairs, the time interval Δt was chosen to be 10 half-lives of ^{214}Po. The delayed energy E_d was chosen to exclusively select the ^{214}Po α energy region, and the prompt energy E_p was chosen to select most of the β^- distribution of ^{214}Bi down to a lower bound of 1.25 MeV, which is sufficiently high-enough to avoid accidentally selecting events inside the ^{214}Po α energy region. The tagging criteria used in this analysis are summarized in Table 9.1. The requirements of data cleaning and valid fits were also applied to this data and are described in Chap-

^{214}Bi	^{214}Po	Coincidence
$r < 5.5\,\mathrm{m}$	$r < 5.5\,\mathrm{m}$	$\Delta t < 1.64 \times 10^6\,\mathrm{ns}$
$1.25\,\mathrm{MeV} < E_p < 5.00\,\mathrm{MeV}$	$0.70\,\mathrm{MeV} < E_d < 1.10\,\mathrm{MeV}$	$\Delta r < 1.0\,\mathrm{m}$
Valid Fit	Valid Fit	
Passes Data Cleaning	Passes Data Cleaning	

Table 9.1: List of criteria used to tag ^{214}Bi–^{214}Po coincidences.

ters 6 and 7, respectively.

These cuts were optimized for data-MC agreement and efficiency in the ROI. This is demonstrated by the plots in the left column of Figure 9·1, which shows the overlapping energy and Δr distributions obtained from all simulated ^{214}Bi–^{214}Po decays, and from the subset of simulated decays tagged by the coincidence algorithm. Additionally, a fit of the half-life equation to the Δt distribution also produces the correct half-life of ^{214}Po = 164.3 μs, indicating accurate and efficient coincidence tagging. The plots in the right column of Figure 9·1 show the results of scaling normalized (to unity) spacial and energy distributions of tagged MC events to match the tagged data. Similarly, a fit to the Δt distribution reproduces the correct half-life, and good agreement in Δr and Δt. The selection efficiency achieved in the ROI is 99.0%. While the removal efficiency is quite high for ^{214}Bi–^{214}Po decays, it was realized at the time of this thesis that these selection criteria could likely be further improved in future analyses to select more than 99.9% of the events. However, due to limited time and computing resources, the current selection criteria is used in the remainder of this analysis.

The spatial distribution of the tagged ^{214}Bi–^{214}Po in the golden data is shown in Figure 9·2. This distribution is quite complex and can vary over time due to decay of short-lived radioactivity, internal convection, and temperature changes inside the inner detector. A hot spot can be seen at the bottom of the AV while the cleanest region

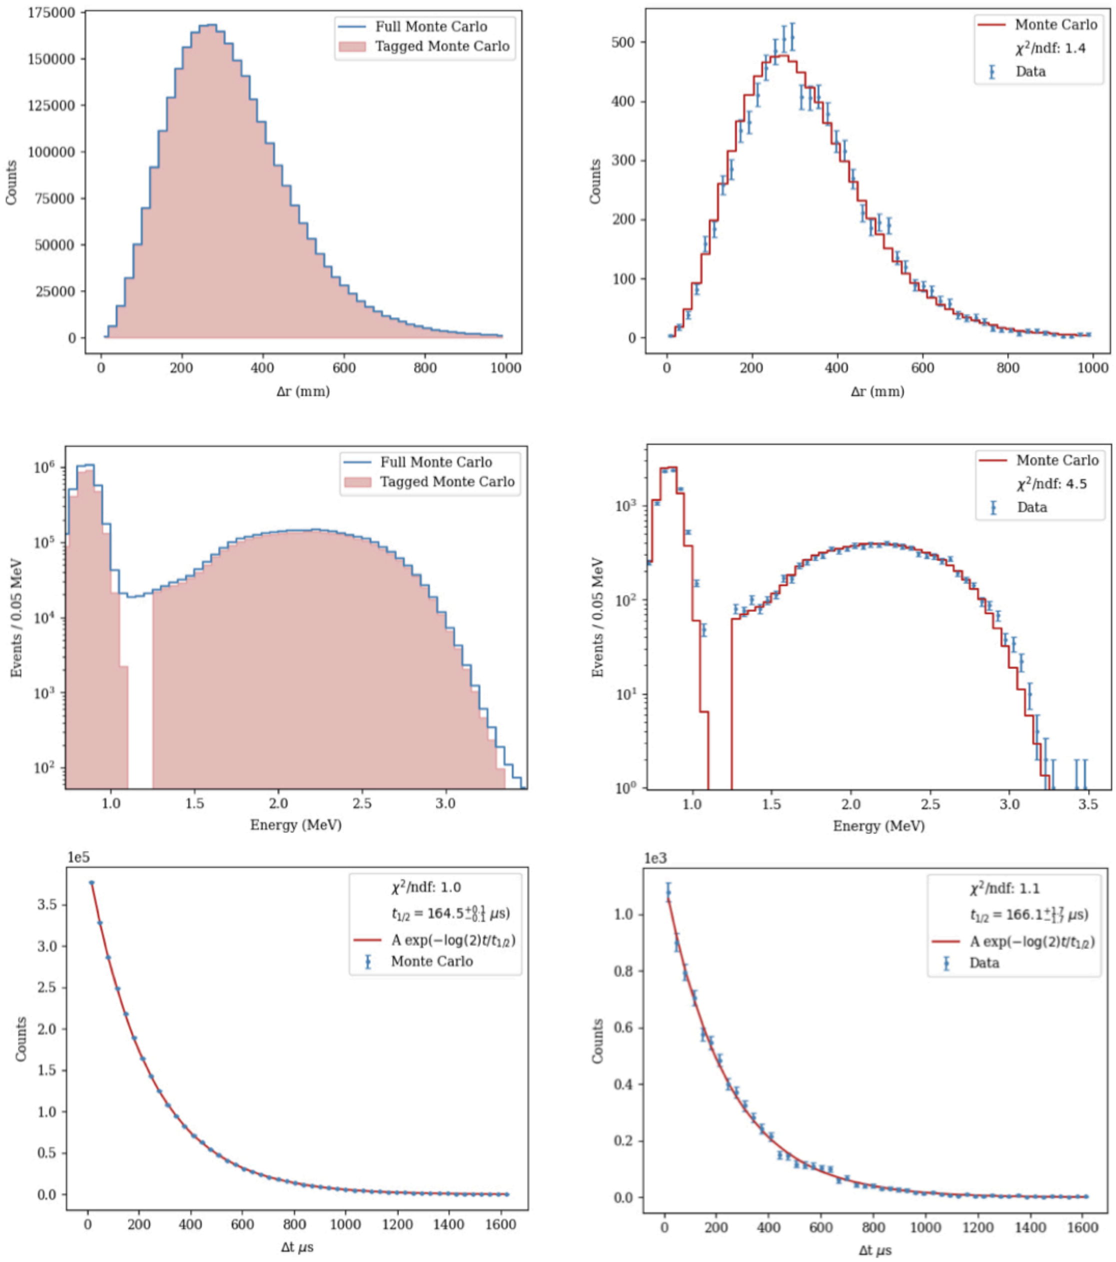

Figure 9·1: Demonstration of the MC tuning of ^{214}Bi–^{214}Po decay selection criteria is shown in the left column of plots, with the ^{214}Po half-life reproduced by the fit of Δt between the ^{214}Bi prompt ^{214}Po delayed events at the very bottom. The application of the ^{214}Bi–^{214}Po selection criteria on real data is shown in the right column of plots.

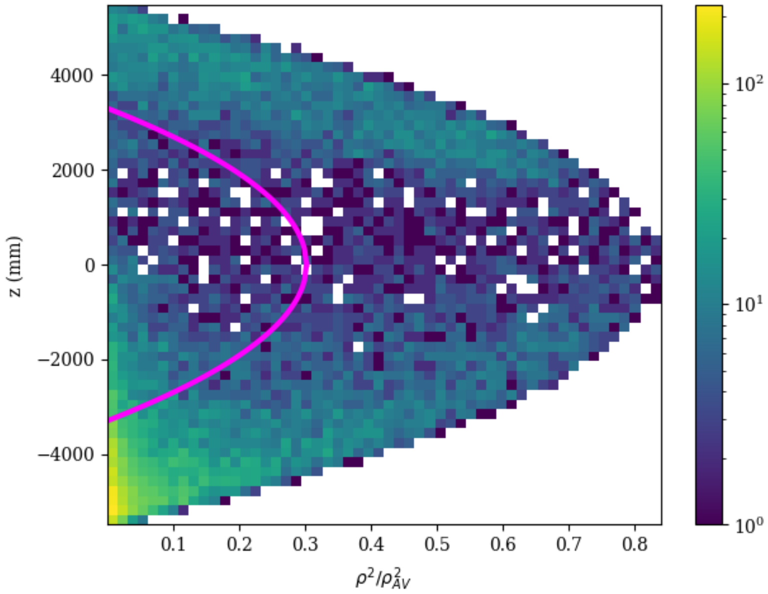

Figure 9·2: Plot of z-position versus $\rho^2/\rho^2_{\mathrm{AV}}$ showing spacial distribution of tagged ^{214}Bi–^{214}Po coincidences. The 3.3 m FV shown in magenta suppresses external backgrounds from the outer regions of the detector.

lies between $-1.0\,\mathrm{m} < z < +2.0\,\mathrm{m}$. The FV of 3.3 m used for this analysis was chosen to be centered on the cleaner part of the region (shown by the magenta line in Figure 9·2) and far from external backgrounds from the outer detector materials. The chosen FV could be optimized in future analyses by allowing the FV selection criteria to vary for different data-taking periods.

To illustrate the manifestation of the varying ^{214}Bi–^{214}Po background during data-taking, we can take the data and plot the number of tagged ^{214}Bi events as a function of z position versus run number. This is shown in Figure 9·3, where the tagged ^{212}Bi events from ^{212}Bi–^{212}Po decays were also included with the help of the $\alpha-\beta(212)$ classifier. This classifier is summarized in great detail in Reference [91]. It tags and removes $\sim$ 98% of the ^{212}Bi – ^{212}Po coincidences occurring within the trigger window, using a cut on the value $\alpha-\beta(212) > -5.0$. The highlighted portions of Figure 9·3 show the

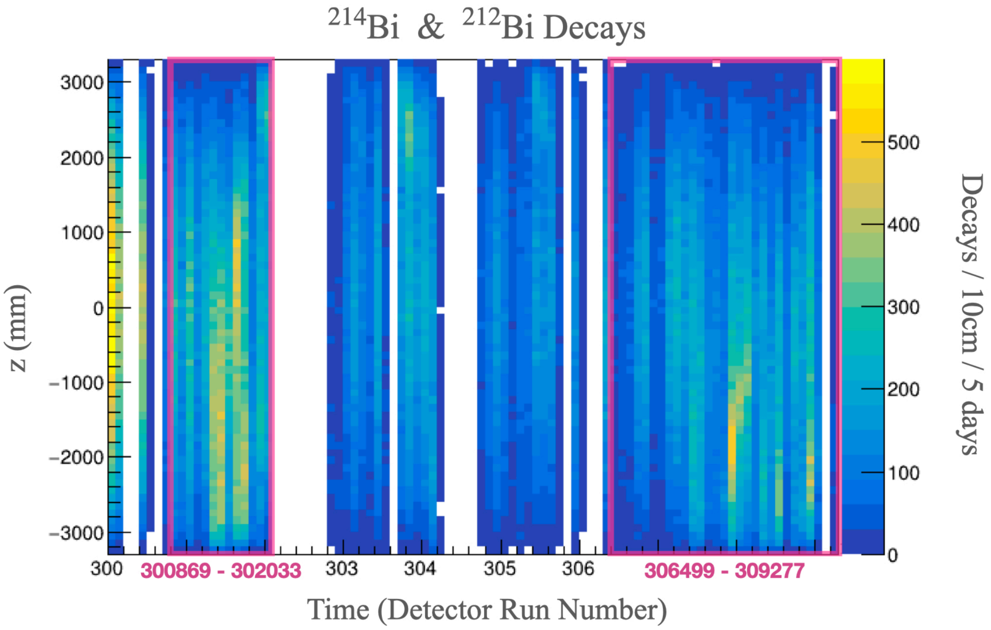

Figure 9·3: The number of ^{214}Bi and ^{212}Bi decays throughout the detector over time, by run number.

run ranges of golden data used in this thesis. Runs missing from this dataset were either contaminated after completion of the fill (runs 300000 - 300869) or simply missed during the data downloading process. The left-most bins show a concentration of decays at the center of the detector, indicating high radon contamination. After confirming this with an independent radon monitor, data from a period of ten ^{222}Rn half-lives were removed from the analysis. The period after showed lower, but still elevated numbers of ^{214}Bi and ^{212}Bi decays, which were removed from the dataset with coincidence tagging criteria and $\alpha - \beta$ classifier.

Within the remaining data, there are further cuts to remove ^{10}C decays and mysterious "hot spots", or events of high light output (> 5000 PMTs triggered). The first background is a spallation isotope and easily removed by the 60-second muon follower cut in data cleaning. Hot spots are bursts of several PMT hits, often without an obvious cause. All events that occur within 20 seconds of an event with more than 5000 hit PMTs are removed.

Prompt	Delayed	Coincidence
$r < 5.5\,\mathrm{m}$	$r < 5.5\,\mathrm{m}$	$\Delta t < 2.0\,\mathrm{ms}$
$0.90\,\mathrm{MeV} < E_p < 8.00\,\mathrm{MeV}$	$1.85\,\mathrm{MeV} < E_d < 2.50\,\mathrm{MeV}$	$\Delta r < 2.5\,\mathrm{m}$
Valid Fit	Valid Fit	
Passes Data Cleaning	Passes Data Cleaning	

Table 9.2: Table of coincidence cuts used to remove events with possible neutron followers. The energy and coincidence cuts were optimized in Reference NO+-doc-8288-v1:SNO+.

Neutron followers from $\alpha - n$ reactions and inverse beta decay (IBD) candidates are a special class of backgrounds, as they contribute to an as yet unexplained peak in the energy spectrum of around 2.2 MeV, discussed in the following section. The following IBD coincidence cuts are a subset of those listed in Reference NO+-doc-8288-v1:SNO+, and are thus too broad to obtain a pure sample of IBD events. They are, however, able to reject any events with neutron followers that might contribute to the mysterious 2.2 MeV peak with greater than 99% efficiency. Table 9.2 lists these coincidence criteria.

The final livetime obtained from the run selection, data cleaning, coincidence tagging, and hot spot removal described here is 88.58 days. To study the performance of KamNet, we compared the background rates of this final dataset with, and without, an additional cut on KamNet score > -0.61. The value of the KamNet cut was obtained during training and is defined at 90% acceptance of the validation signal as described in Chapter 8.

9.2 Binned Log-Likelihood Analysis

This analysis takes a likelihood approach to estimating the contribution of each expected background to the dataset. The Poisson likelihood of observing the set of n independent data points $\vec{x} = \{x_1, x_2, ..., x_i, ...x_n\}$, given some combination of m expected parameters $\vec{\theta} = \{\theta_1, \theta_2, ..., \theta_j, ..., \theta_m\}$ is given by:

$$\mathscr{L}(\vec{x}|\lambda, \vec{\theta}) = e^{-\lambda} \frac{\lambda^n}{n!} \prod_i^n p(x_i|\vec{\theta}) \tag{9.1}$$

where λ is the expected total number of events and p the probability of observing the i^{th} data point x_i. The parameters $\vec{\theta}$ are scaling factors for the normalized probability density functions (PDFs) obtained from MC simulation which has been tuned to accurately reproduce physics data.

Binning these data in a histogram significantly reduces the dimensionality of the problem, and we can rewrite the likelihood equation in terms of bin contents b_{ij} for the i^{th} bin of the j^{th} PDF. The expected number of events λ in each bin, is expressed in the following way:

$$\lambda_i = \sum_j \mu_j b_{ij} \tag{9.2}$$

where the μ_j are the scaling factors for each PDF. The λ_i show up in the total likelihood, now written as the product of the Poisson likelihoods of observing k_i events in the i^{th} bin:

$$\mathscr{L} = \prod_i \frac{\lambda_i^{k_i} e^{-\lambda_i}}{k_i!} \tag{9.3}$$

The task of this analysis is to maximize the likelihood $\mathscr{L}$, or find the combination of μ_i for which the data are most likely observed.

For ease of computation, we make two adjustments. The first is to consider the logarithm of the likelihood, which removes the factorial factor of k_i in the denominator and simplifies Equation 9.3 in the following way:

$$\log \mathscr{L} = \sum_i (\lambda_i - k_i \log \lambda_i) \tag{9.4}$$

The second adjustment is to consider the negative log likelihood, so that we may take advantage of existing minimization algorithms. The minimizer used for these fits is the limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm described in Reference [92].

Finally, fit parameters can be constrained with values that are determined either from independent analyses or external measurements. These values and their associated uncertainties are added to the likelihood function in the form of Gaussian constraints, producing the following log-likelihood function:

$$\log \mathscr{L} = \sum_i (\lambda_i - k_i \log \lambda_i) + \sum_k \frac{(\lambda_k - \overline{\lambda_k})^2}{2\sigma_k^2} \tag{9.5}$$

where k indexes the set of parameters that are constrained with measured value $\overline{\lambda}_k$ and associated uncertainty σ_k. All ^{238}U and ^{232}Th-chain backgrounds, with the exception of $^{234\text{m}}$Pa and ^{208}Tl, were constrained based on separate measurements of ^{212}Bi and ^{214}Bi concentration by Reference [93]. The constraints on $\alpha - n$ prompt signal came from an independent measurement reported in Reference [**?**]. Fits were performed both with the ^{8}B solar ν rate constrained to its expected value, and without constraint. For the case where ^{8}B solar ν is constrained we use the value reported in Reference [78], which is based on new solar models in Reference [39]. The comparison between fits with ^{8}B solar ν floated and constrained is discussed later. A summary of the parameters that were floated and/or constrained are shown in Table 9.3.

The PDFs used in this analysis were prepared from MC simulated classes of events that had exactly the same cuts applied as the real data (with the exception of data cleaning, which is not run on MC). Importantly, two sets of PDFs were created: one set that was additionally cut for KamNet score, and one that was not. Moving forward, we'll refer to the PDFs in the latter set as the 'Primary' analysis and the former set as the 'KamNet' analysis. All PDFs are energy spectra normalized to unity and binned from 1.5 – 5.0 MeV in bins of width 0.05 MeV.

Parameter	Decay Chain	Fit Inclusion	Primary Constraint (# events)	KamNet Constraint (# events)
^{228}Ac	^{232}Th	Constrain	66.22	50.82
$\alpha-$n delayed	-	Float	-	-
$\alpha-$n prompt	-	Constrain	0.12	0.06
^{8}B solar ν	-	Float/Constrain	15.54	13.62
^{212}Bi $-$ ^{212}Po	^{232}Th	Constrain	1.85	1.29
^{214}Bi $-$ ^{214}Po	^{238}U	Constrain	3.10	1.00
^{11}C	-	Constrain	3.89	2.73
^{234m}Pa	^{238}U	Float	-	-
^{208}Tl	^{232}Th	Float	-	-
^{210}Tl	^{238}U	Constrain	0.13	0.09

Table 9.3: Table of all backgrounds included in Primary and KamNet fits.

9.3 Fit Validation and Background Estimation

The fitting framework was tested and validated using an Asimov dataset, as shown in Figure 9·4. An Asimov dataset is a deterministic dataset constructed by summing all average background spectra, with each component normalized to its expected event rate [94]. This approach avoids statistical fluctuations inherent in pseudo-experiments and provides a representative dataset that reflects the median behavior of the experiment under the null hypothesis. As such, it enables sensitivity estimation with significantly reduced computational cost compared to generating and analyzing a full ensemble of MC simulations.

In this validation of the fitting framework, the Asimov dataset is constructed as the sum of all PDFs used in the fit, each of them scaled by their expected rates as listed in Table 9.5.

Species	Injected (# events)	Fitted (# events)
^{228}Ac	272.85	$269.06^{+22.4}_{-14.6}$
α-n delayed	592.93	$598.54^{+52.2}_{-69.2}$
α-n prompt	0.49	$0.49^{+0.8}_{-0.8}$
^{8}B solar ν	147.43	$131.9^{+56.3}_{-21.4}$
^{212}Bi–^{212}Po	7.62	$7.65^{+3.0}_{-3.0}$
^{214}Bi–^{214}Po	13.02	$12.62^{+3.8}_{-4.1}$
^{11}C	16.03	$15.8^{+4.6}_{-4.3}$
^{234m}Pa	8618.25	$8626.70^{+285.2}_{-292.3}$
^{130}Te 0ν	0.0	$1.02^{+6.6}_{-10.0}$
^{130}Te 2ν	51435.32	$51444.23^{+278.1}_{-276.0}$
^{208}Tl	448.51	$460.45^{+20.5}_{-43.0}$
^{210}Tl	0.55	$0.56^{+0.8}_{-0.8}$

Table 9.4: The resulting fit values for an Asimov dataset with the above species injections.

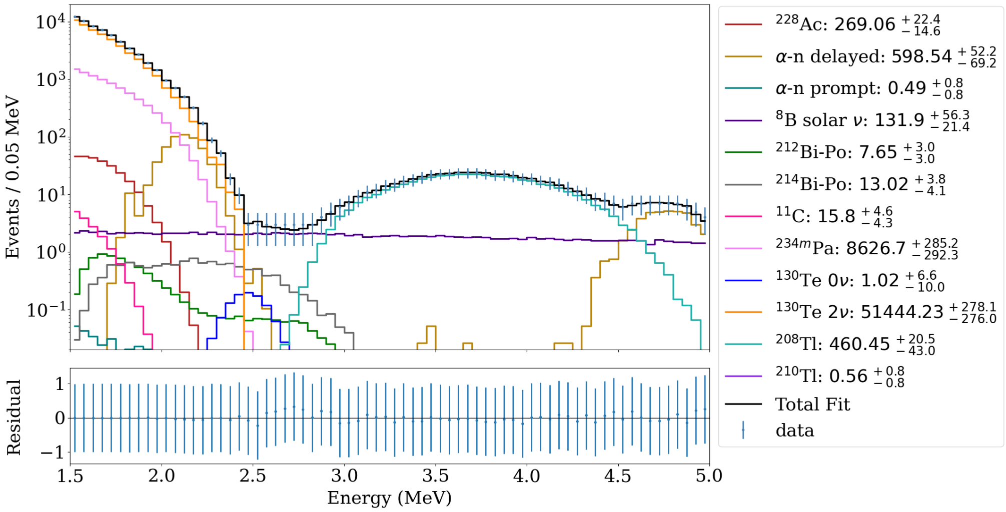

Figure 9·4: A fit to an Asimov dataset, showing agreement between injected and fitted values.

A Poisson error is assumed for each histogram bin, reflecting the statistical uncertainty that would be expected in real data. The negative log-likelihood function is then minimized using the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS-B) algorithm, a quasi-Newton method well-suited for problems with a large number of parameters and bounded constraints. The results of the fit show that the recovered (fitted) parameter values are consistent with the injected (true) values within 1σ standard deviation. This agreement confirms the internal consistency of the fitting framework and demonstrates its robustness for extracting physics parameters from the expected dataset composition.

With the fitting framework validated we now move onto the application to the real detector dataset. The log-likelihood fit was performed on both the Primary and KamNet datasets. The first application of the fits with the parameters listed in Table 9.3 is shown in Figures 9·5 and 9·6, where the ^{8}B solar ν background was constrained to its expected value. The fit residuals, per i^{th} bin defined as $(\text{data}_i - \text{total fit}_i)/\sqrt{\text{data}_i}$, are all within 1σ of zero, indicating a good fit to the dataset. Both sets of fit residuals also exhibit a slight bump around 2.0 MeV. This behavior persists through the significant, and expected, drop in ^{214}Bi–^{214}Po in the KamNet dataset. This behavior is also consistent with other analyzers' datasets, which have undergone even more aggressive ^{214}Bi–^{214}Po tagging, using algorithms of 99.96% and 99.997% efficiencies ([58], [95]).

Very careful inspection shows that the fits in Figures 9·6 and 9·5 systematically underestimate the number of events falling between 2.25 and 2.85 MeV. This range completely contains the ROI, which presents a problem pertaining to sensitivity studies assuming this particular background model. Toy MC datasets generated from these fitted rates will underestimate the backgrounds in the ROI. Therefore, the sensitivity to $0\nu\beta\beta$ will be overestimated if these toy MC datasets are used in the sensitivity calculation. To address the underestimate of the background in the ROI, the ^{8}B solar ν spectrum was allowed to float as it was one of the only backgrounds that could accom-

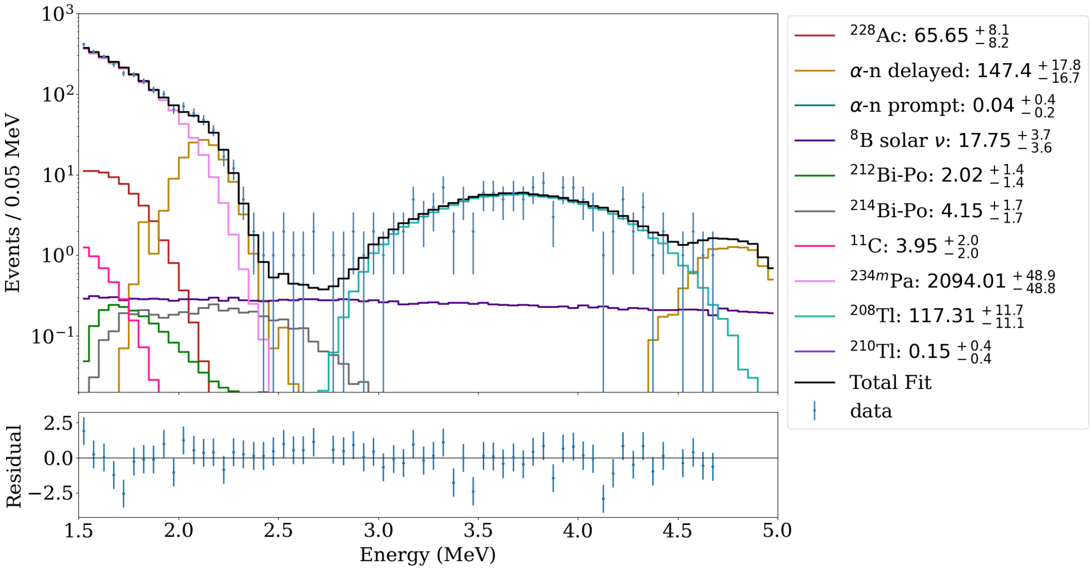

Figure 9·5: A fit of the backgrounds listed in Table 9.3 to the Primary dataset with the ^{8}B solar ν rate constrained.

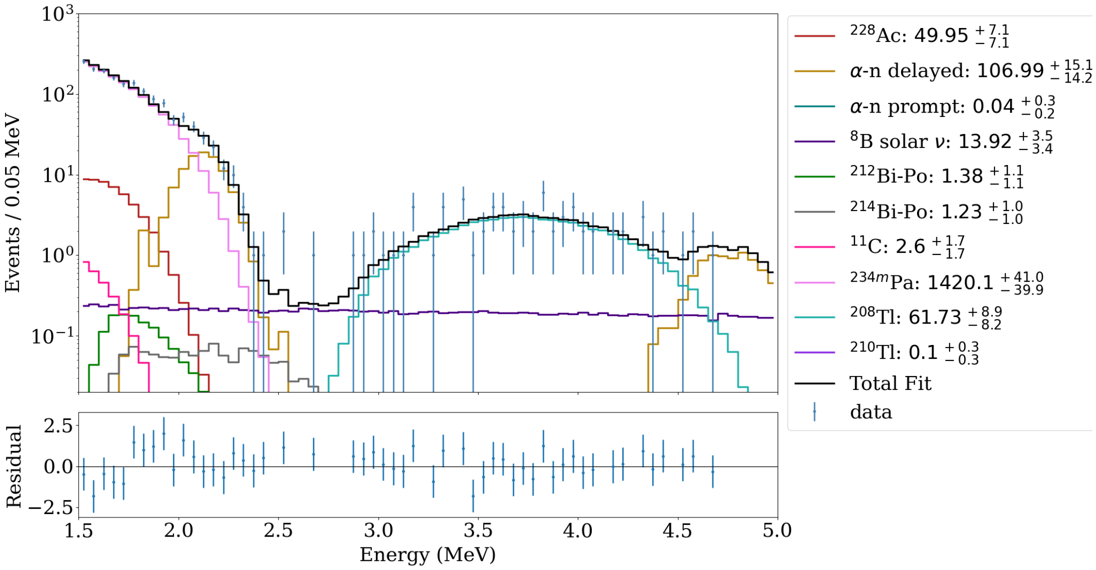

Figure 9·6: A fit of backgrounds listed in Table 9.3 to the KamNet dataset with the ^{8}B solar ν rate constrained.

modate the excess in the ROI without diminishing the overall fit quality. The resulting fits in Figures 9·7 and 9·8 show a slightly better fit to the data in the ROI region, with the improvement most pronounced in the Primary fit. The fitted rates for the ^{8}B solar ν spectra are within 1-2σ of the predicted value, using the flux reported in [78].

The percent change of the fitted ^{8}B solar ν rate between constrained and unconstrained fits is 102% for the Primary fit, and 17% for the KamNet fit. These differences suggest that the excess of backgrounds between 2.25 and 2.85 MeV could be multi-site in nature and removed by KamNet. Future background studies could investigate the possibility of alternative background sources not included in the current fits to explain this slight excess, or more stringent cuts could be used on the data to further remove the excess events. Higher statistics would also provide more insight into the subtle features in the background. Since the purpose of these backgrounds studies is to study the performance of KamNet on the data, the background model from the Primary and KamNet fits with the ^{8}B solar ν rate allowed to float is assumed in the sensitivity calculations that follow.

9.4 Toy MC

To evaluate the analysis sensitivity, at least 2,000 toy MC datasets (often referred to as 'fake' datasets) were generated and analyzed. The full MC sample was first divided into two statistically independent subsets: one was used to construct PDFs, and the other was reserved for generating the toy datasets.

Each toy dataset consists of the sum of background spectra, with each background component independently fluctuated to simulate statistical variations. Specifically, for a given background source Θ_i with an expected rate r_i determined from fits to data, a fluctuated rate r'_i is drawn from a Poisson distribution Pois($\lambda = r_i$). A 'Poisson-fluctuated' background sample is then generated by randomly selecting r'_i events from

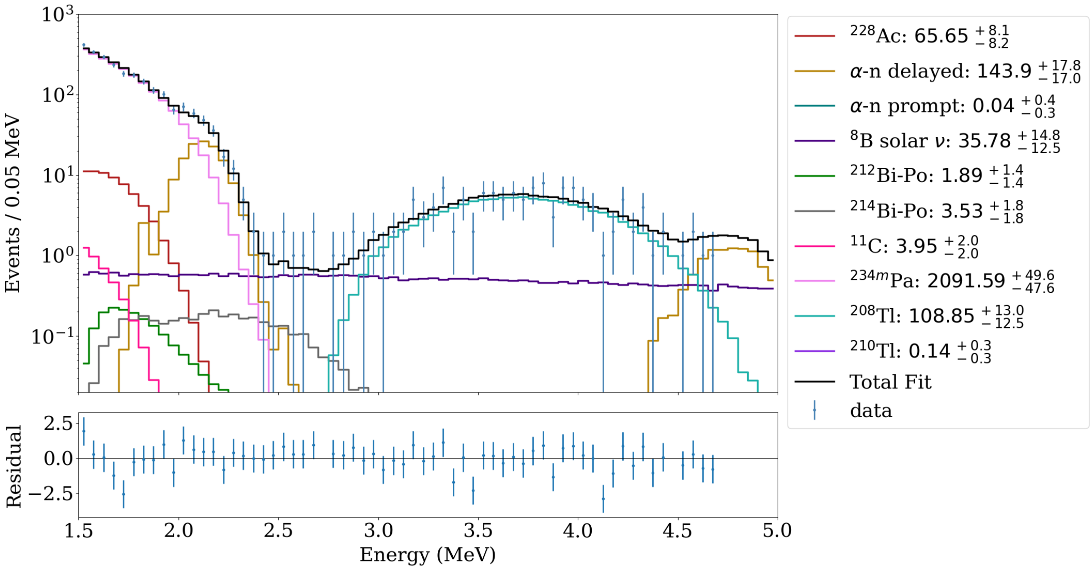

Figure 9·7: A fit of the backgrounds listed in Table 9.3 to the Primary dataset with the ^{8}B solar ν rate allowed to float.

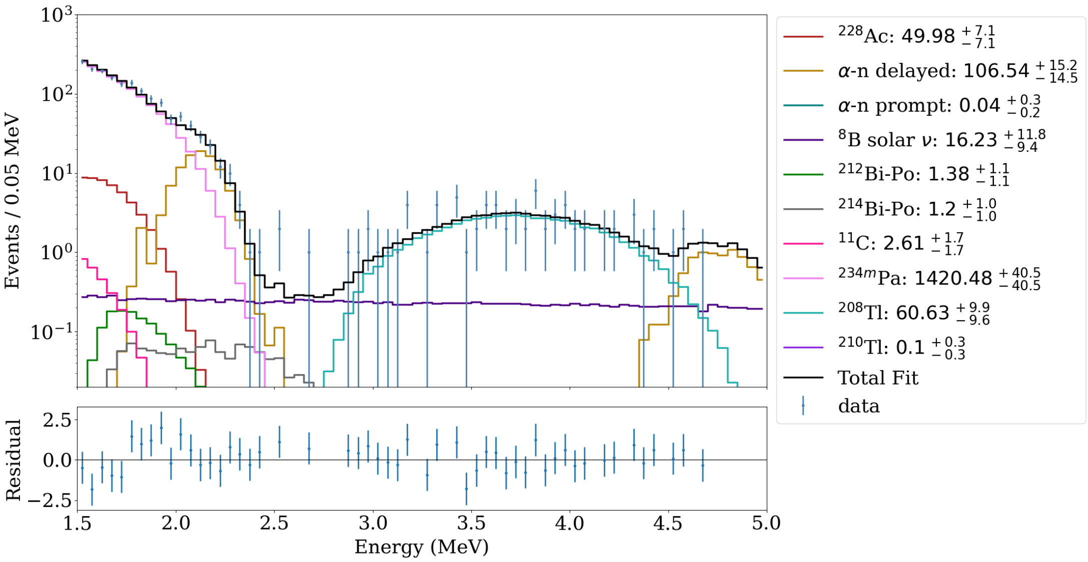

Figure 9·8: A fit of backgrounds listed in Table 9.3 to the KamNet dataset with the ^{8}B solar ν rate allowed to float.

the corresponding MC simulation of Θ_i. This process is repeated for each background component, and the combined set of events forms the complete toy dataset. Each toy dataset is then binned into a histogram with a fixed bin width of 0.05 MeV, and subsequently used in the sensitivity studies.

Figure 9·9 shows an overlay of two energy spectra: one of the original primary dataset, and the other of a toy MC dataset generated using the primary fit values. The close agreement between the two spectra, evidenced by their substantial overlap across the entire energy range, demonstrates that the toy dataset accurately reproduces the original data. This agreement confirms that the underlying assumptions in the modeling and fluctuation process are consistent with the observed data.

The SNO+ projected sensitivities given each set of fitted backgrounds in the Primary and KamNet datasets were calculated for 88.58 days of livetime at 0.5% tellurium loading by weight, assuming no signal. A ^{130}Te $2\nu\beta\beta$ spectrum was added to each toy dataset by Poisson fluctuating the corresponding PDF first scaled by its expected rate in the SNO+ FV. This expected rate was calculated using the CUORE experiment's most recent measurement of the half-life [96]. An example energy spectrum of a final toy MC dataset with the assumed 0.5% tellurium loading is shown in Figure 9·10.

As an additional validation of the toy MC generation, two crosschecks were conducted on the Primary and KamNet datasets. The first crosscheck confirmed the correct toy generation for each background. To do this, 2,000 fake spectra were produced for a given background and summed together. The accumulated histogram from summing the individual fake spectra was then normalized to unity and overlaid with its corresponding PDF to check for agreement in the shape of the overall distribution. Moreover, a distribution of the injected rates was created, and its average rate over the generation of many fake datasets was noted.

Example results from these crosschecks for one example background, ^{11}C, are shown in Figures 9·11 and 9·12. For every background used in the toy MC dataset, the overlaid

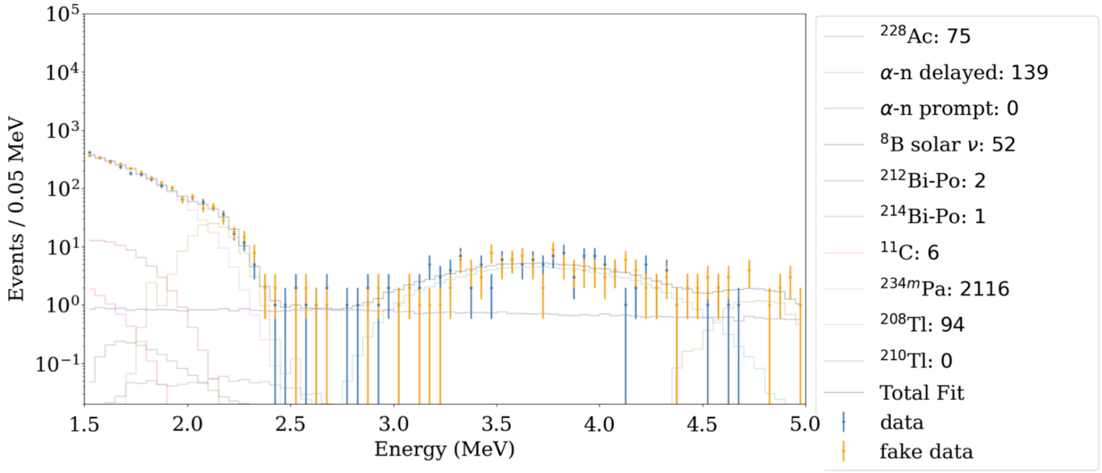

Figure 9·9: Energy spectra of a real dataset (blue) and a toy MC dataset (orange). The numbers of events of each background injected into the toy MC dataset is listed in the legend on the right.

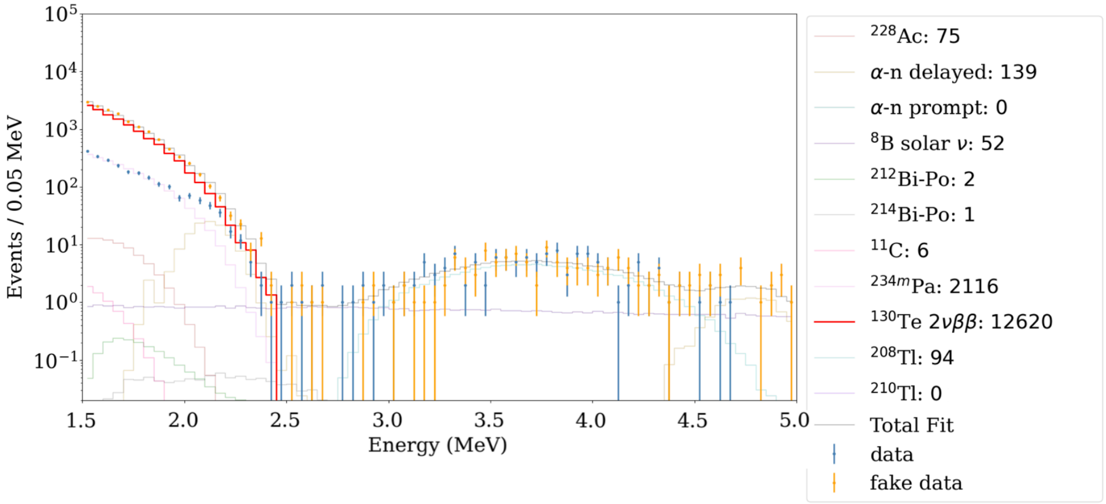

Figure 9·10: Energy spectra of real dataset (blue) and a toy MC dataset (orange), with the $2\nu\beta\beta$ injected (red). The numbers of events of each background injected into the toy MC dataset is listed in the legend on the right.

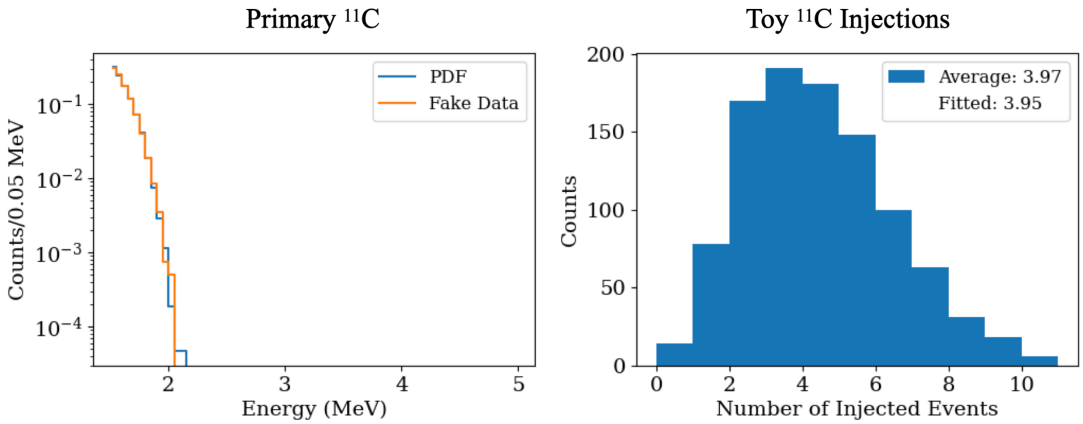

Figure 9·11: Example validation of a Primary toy MC background. The shape comparison between the accumulated fake data and the PDF used to generate the fake data is on the left. The distribution of injected rates are shown on the right.

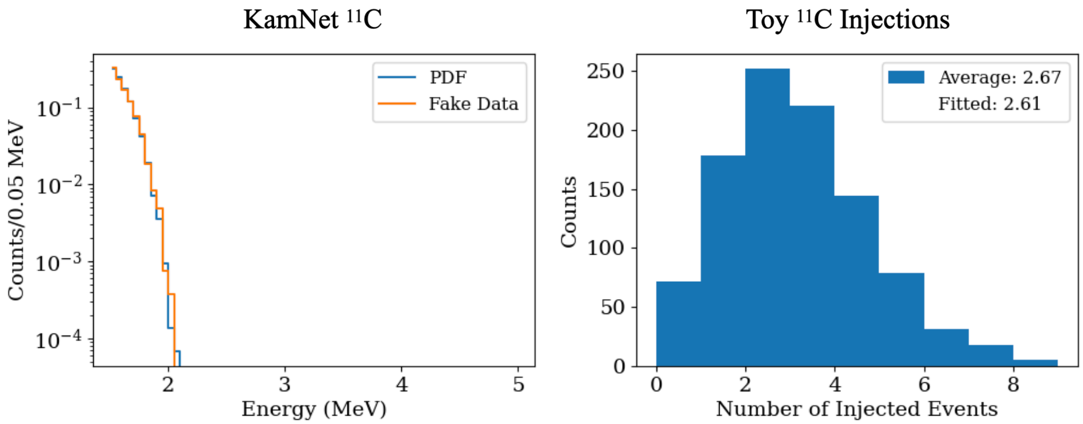

Figure 9·12: Example validation of a KamNet toy MC background. The shape comparison between the accumulated fake data and the PDF used to generate the fake data is on the left. The distribution of injected rates are shown on the right.

spectra showed the good agreement in the shapes. Furthermore, each distribution of injected rates of the fake datasets followed a Poisson distribution centered about the original fitted rate from the real dataset. In all cases, the average injected rates from fake datasets and corresponding fitted rates from real datasets agreed to within several percent.

For the second crosscheck of the toy MC dataset, over 2,000 toy MC datasets were accumulated and the distributions of the number of events generated in the ROI were recorded. These datasets were generated with the $2\nu\beta\beta$ spectrum included and without the $0\nu\beta\beta$ spectrum . As expected, the KamNet toy datasets delivered ~ 27% fewer events on average in the ROI compared to the Primary toy datasets. The average numbers of events in the Primary and KamNet toy datasets were also found to be within 1σ of the numbers observed in the corresponding real dataset, given the relatively large statistical uncertainties in the real datasets.

After validating and crosschecking the toy MC datasets, they were used to calculate the SNO+ sensitivity to the ^{130}Te $0\nu\beta\beta$ decay half-life, assuming the background model of the 88.58 days of livetime for the real Primary and KamNet datasets. Each sensitivity calculation was carried out on an ensemble of toy MC datasets, assuming

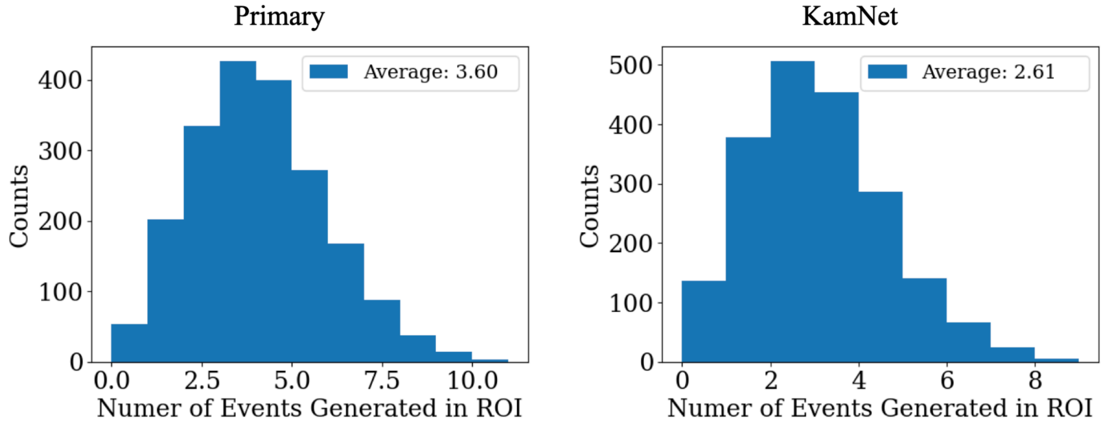

Figure 9·13: Distributions showing number of events in the ROI over 2,000 Primary (left) and KamNet (right) toy datasets.

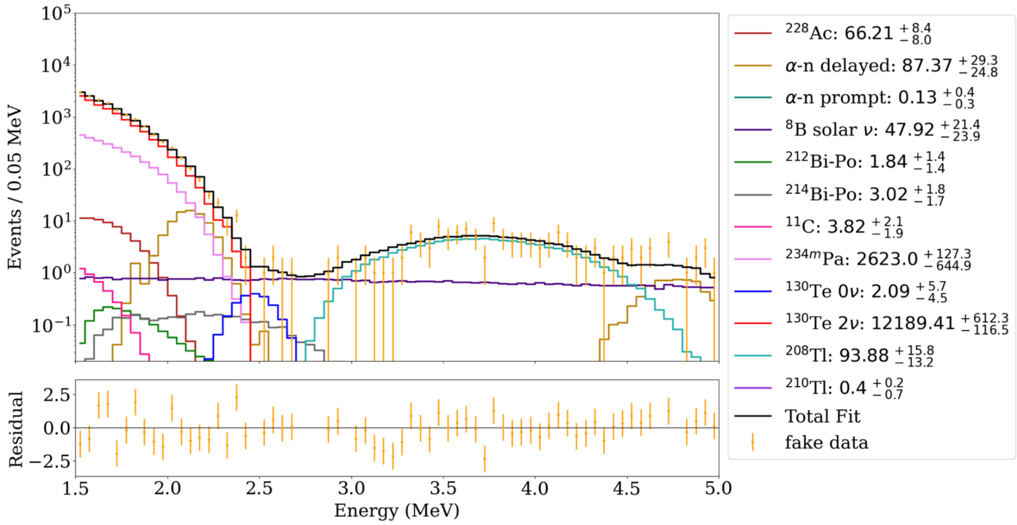

Figure 9·14: The fit of $0\nu\beta\beta$ decay signal to the Primary toy MC dataset shown in Figure 9·10. These fitted rates are compared to the injected rates in the toy MC dataset in Table 9.5.

Species	Injected (# events)	Fitted (# events)
^{228}Ac	75	$66.21^{+8.4}_{-8.0}$
α-n delayed	139	$87.37^{+29.3}_{-24.8}$
α-n prompt	0	$0.13^{+0.4}_{-0.3}$
^{8}B solar ν	52	$47.92^{+21.4}_{-23.9}$
^{212}Bi–^{212}Po	2	$1.84^{+1.4}_{-1.4}$
^{214}Bi–^{214}Po	1	$3.02^{+1.8}_{-1.7}$
^{11}C	6	$3.82^{+2.1}_{-1.9}$
$^{234\text{m}}$Pa	2116	$2623.45^{+127.3}_{-644.9}$
^{130}Te 0ν	0	$2.09^{+5.7}_{-4.5}$
^{130}Te 2ν	12620	$12189.41^{+612.3}_{-116.5}$
^{208}Tl	94	$93.88^{+15.8}_{-13.2}$
^{210}Tl	0	$0.4^{+0.2}_{-0.7}$

Table 9.5: The resulting fit values for an example toy MC dataset compared with the rates injected into the toy MC dataset.

no $0\nu\beta\beta$ decay signal. The 2,000 fits of the ^{130}Te $0\nu\beta\beta$ decay signal were performed on Primary and KamNet toy MC datasets created without injected ^{130}Te $0\nu\beta\beta$ decay events. Of these fits, 376 failed to converge for the KamNet toy MC dataset, and 241 failed to converge for the Primary toy MC dataset. The failed fits were discarded from the final sensitivity calculation. An example of one fit to a Primary toy MC dataset is shown in Figure 9·14, and a summary of the fitted and injected rates is listed in Table 9.5 for comparison.

Using the fit of the $0\nu\beta\beta$ signal to a toy MC dataset we follow a similar approach to what is described in Chapter 3. From Equation 9.6, we can see that the number of detected $0\nu\beta\beta$ decay events $N_{\beta\beta}$ can be written as:

$$N_{\beta\beta} = \frac{\ln(2)\, t\epsilon N}{T^{0\nu}_{1/2}} \tag{9.6}$$

where t is the livetime in years, ϵ is the signal detection efficiency, and N is the number of ^{130}Te nuclei in the FV. Therefore, we can reverse the logic and use the 90% C.L. on the fitted $0\nu\beta\beta$ decay rate $C_{90\%}$ in the toy MC dataset and extract a lower limit on the $0\nu\beta\beta$ half-life:

$$T^{0\nu}_{1/2} = \frac{\ln 2}{C_{90\%}} t\epsilon N \tag{9.7}$$

where 90% of the fitted $0\nu\beta\beta$ rates in the toy MC ensemble fall below $C_{90\%}$. With the current level of backgrounds, at 0.5% tellurium loading and 88.58 days of livetime, the use of KamNet as a background rejection tool increases the SNO+ projected sensitivity by 98.5% (nearly double). This is illustrated in a side-by-side comparison of the sensitivity calculations for the Primary toy MC dataset and the KamNet toy MC dataset shown in Figure 9·15.

As described earlier in Chapter 7, the background budget for SNO+ to reach a sensitivity of $S^{0\nu}_{1/2} > 2 \times 10^{26}$ years is 9.47 events/year in the ROI (assuming a simple counting analysis). This would scale down to about 2.3 events in the 88.58 days of livetime used

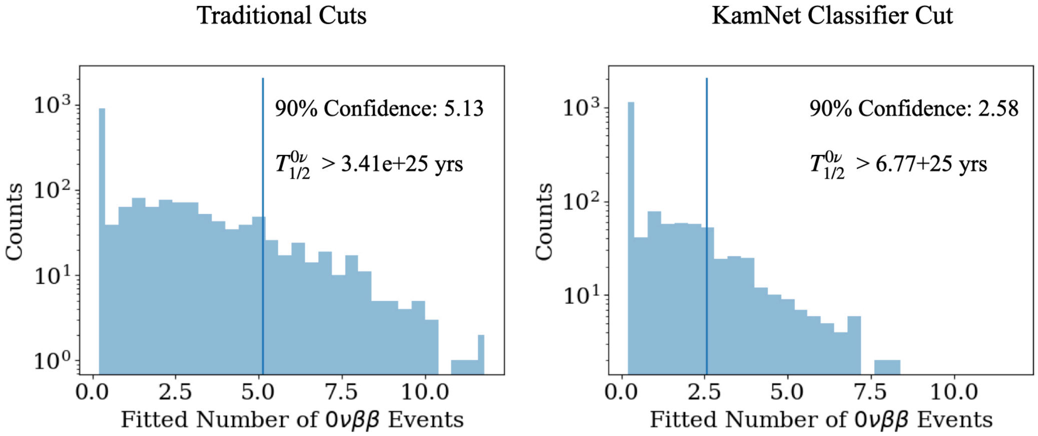

Figure 9·15: Histograms of number of fitted $0\nu\beta\beta$ counts in the Primary (left) and KamNet (right) toy MC analyses, showing KamNet's sensitivity improvement of 98.5%.

for this analysis. In this dataset, there were five events in the ROI. KamNet removed two events, bringing the number of events in the ROI within 1σ of the target number. It should be noted that this is a very early analysis with limited data, so these results should be taken with a grain of salt. Future improvements on the traditional cuts applied to the data to reject backgrounds, and the fitting methods, could boost the limits reported here to be within (or surpass) the estimated half-life sensitivity reported for the simple SNO+ counting analysis.

To put the KamNet result into context, recall from Chapter 3 that the sensitivity of an experiment to $0\nu\beta\beta$ decay scales according to the following experimental parameters:

$$S^{0\nu}_{1/2} \propto \sqrt{\frac{mt}{B\Delta E}} \tag{9.8}$$

where m is the mass of the isotope, t is detector livetime, B is the background rate, and ΔE is the detector energy resolution. To achieve its best sensitivity, SNO+ will run for several years at progressively higher tellurium loading. Figure 9·16 shows the half-life sensitivity versus livetime for 0.5% and 1.5% tellurium loading, and 0.5% tellurium

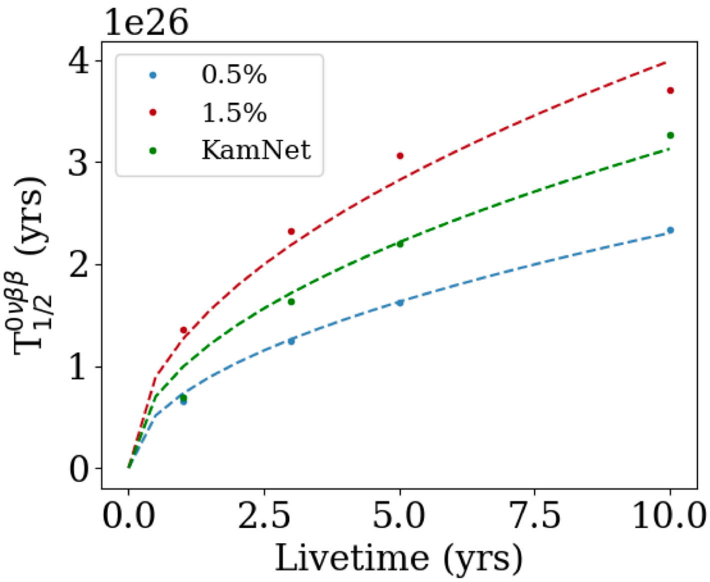

Figure 9·16: Half-life sensitivities versus livetime for various isotope loading. Each point represents a fit, assuming no signal, to an Asimov dataset generated at the given livetime and isotope loading.

loading with KamNet applied. Each data point represents a fit to an Asimov dataset of the given livetime and loading, assuming no $0\nu\beta\beta$ signal, while the curves represent a fit of the function $f(t) = a\sqrt{t}$ to the data points where a is a floating constant. Under the assumed background model in this thesis, with three years of livetime, an analysis of 0.5% tellurium-loaded data using KamNet is projected to produce a half-life limit of $T_{1/2}^{0\nu} > 1.63\times10^{26}$ years. A traditional analysis with the same amount of Te would require two additional years of livetime. Alternatively, 1.5% Te loading could achieve this limit in under two years, *but at the additional cost of* \$2.6*M of natural tellurium!* This, in turn, shows the power of AI/ML tools like KamNet for background classification in low-background particle detectors.

Chapter 10

Summary and Conclusions

This work examined the deep learning neural network KamNet as a promising background rejection tool for the $0\nu\beta\beta$ decay search on SNO+. Originally developed for the KamLAND-zen experiment, the network was hypothesized to be able to distinguish multi-site backgrounds from single-site backgrounds with similar performance on other similar spherical liquid scintillator detectors. This hypothesis was found to be true for SNO+. As a test case, KamNet was trained on ^{214}Bi background and ^{8}B solar ν signal datasets, and it went on to reject ^{214}Bi and ^{10}C with efficiencies similar to those achieved in KamLAND-Zen, while accepting 90% of the signal. These similarities are reflected by the KamLAND-Zen and SNO+ detectors' comparable vertex resolutions, since PMT hit-timing information is one of several important factors going into KamNet's classification power. An additional validation study demonstrated that the KamNet score distribution of a pure, tagged sample of ^{214}Bi data agreed with that of ^{214}Bi MC and produced rejection efficiencies similar to within 2.8%.

KamNet is not the only event classifier available on SNO+. Two existing classifiers were also discussed in this thesis, and both of them use only time residuals in their decision making. The ITR cut is a vestige from SNO and the water phase of SNO+, and it was intended to identify Cherenkov light by the characteristic sharp rising-edge it produces in water. This classifier was retired briefly under the assumption that it wouldn't be useful in a detector in which Cherenkov light is overwhelmed by scintillation light. It was recently found, however, that different kinds of events produce changes in time residuals clear enough for ITR to identify. A comparison of ITR to KamNet on equal numbers of simulated $0\nu\beta\beta$ and ^{214}Bi in a narrow range of total PMT hits revealed that KamNet rejects 10.5% more background than ITR at 90% signal efficiency. This increase in rejection efficiency demonstrates the power that KamNet has through its additional use of spacial information. In a second study, KamNet was trained on ^{60}Co and ^{130}Te $0\nu\beta\beta$ decay in a 4.0-m FV, to perform a direct comparison to the most recently developed log-likelihood multi-site classifier. At 90% signal efficiency, KamNet also produced a slight overall improvement of $4.2 \pm 2.1\%$.

After validating KamNet's performance with a tagged background in real data and demonstrating its superior performance relative to other available multi-site classifiers, the remainder of this analysis focused on the effect of KamNet's background rejection on the $0\nu\beta\beta$ decay sensitivity of SNO+. The experimental sensitivity was assessed using three approaches, each applied independently to a Primary dataset and a KamNet dataset: toy MC, Asimov datasets, and a counting analysis. An ensemble of toy MC consisted of 2,000 fake energy spectra, assuming 0.5% tellurium loading, no $0\nu\beta\beta$ signal, and 88.58 days of livetime. These fake data were generated from simulated events, according to the procedure described in Chapter 9. Negative log-likelihood fits containing the $0\nu\beta\beta$ spectrum were performed, and the 90th percentile of fit $0\nu\beta\beta$ events was used to calculate the half-life sensitivity $T^{0\nu}_{1/2}$ assuming no background. Half-life limits of $T^{0\nu}_{1/2} > 3.41 \times 10^{25}$ years and $T^{0\nu}_{1/2} > 6.77 \times 10^{25}$ years were calculated for the Primary and KamNet datasets, respectively, showing an improvement of 98.5% when KamNet is applied.

Earlier predictions of SNO+ sensitivity relied on a simple counting analysis in the ROI, using a background budget informed by ongoing measurements leading to tellurium loading. The counting analysis produced a half-life sensitivity of $T^{0\nu}_{1/2} > 2 \times 10^{26}$ years in three years of livetime and 0.5% tellurium loading. The target ROI rate for the counting was calculated to be 9.47 events/year, corresponding to 2.3 events in the 88.58 days of data analyzed in this thesis. Although the KamNet-cut dataset has 3 events in the ROI, it is important to note that the analysis in this thesis came from a low-statistics dataset, and that the target and measured rates are still within 1σ of one another. It is also valuable to note that the Primary dataset contains 5 events in the ROI - more than 1σ out of the background budget. One potential cause for the higher background could be contamination from external γ sources in the data, and it is recommended that future analyzers continue to improve cuts to remove such backgrounds, even if KamNet is eventually applied to the data.

To interpret the practical implications of the improvement observed by applying KamNet to the data, additional fits to Asimov datasets were performed, assuming 0.5% and 1.5% loading at 1, 3, 5, and 10 years of livetime. This study showed, at a benchmark of three years, that KamNet could achieve a sensitivity of $T^{0\nu}_{1/2} > 1.63 \times 10^{26}$ years. Without KamNet, the same tellurium loading would take roughly two more years to reach the same limit. Loading to 1.5% tellurium could achieve the same limit in under two years, but at an additional cost of ~ \$2.6M, thereby showing the power of AI/ML tools like KamNet.

Bibliography

[1] M. Grunewald et al. Precision electroweak measurements on the Z resonance. *Physics Reports*, 427(5):257--454, 2006.

[2] E. Fermi. Tentativo di una Teoria Dei Raggi β. *Nuovo Cimento*, 1, 1934.

[3] NobelPrize.org. The Nobel Prize in Physics 1979. `https://www.nobelprize.org/prizes/physics/1979/summary/`. Accessed: 2025-01-24.

[4] NobelPrize.org. The Nobel Prize in Physics 2015. `https://www.nobelprize.org/prizes/physics/2015/summary/`. Accessed: 2025-01-24.

[5] B. Pontecorvo. Mesonium and anti-mesonium. *Soviet Journal of Experimental and Theoretical Physics*, 6:429, 1957.

[6] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 1 2007.

[7] Alessandro Paoloni. Status of the JUNO experiment. *Proceeding of Science*, HQL2023:017, 2024.

[8] Ali Ajmi. Status of the Hyper-Kamiokande Experiment. *Proceeding of Science*, HQL2023:098, 2024.

[9] Inés Gil-Botella. The Deep Underground Neutrino Experiment (DUNE) Program. *Proceeding of Science*, ICHEP2024:165, 2025.

[10] Ivan Esteban, M. C. Gonzalez-Garcia, Michele Maltoni, Ivan Martinez-Soler, João Paulo Pinheiro, and Thomas Schwetz. NuFit-6.0: Updated Global Analysis of Three-Flavor Neutrino Oscillations. *Journal of High Energy Physics*, 2024(12), 2024.

[11] M. Aker et al. Direct Neutrino-Mass Measurement With Sub-Electron Volt Sensitivity. *Nature Physics*, 18(2):160--166, 2022.

[12] A. Ashtari Esfahani, S. Böser, N. Buzinsky, M. C. Carmona-Benitez, C. Claessens, L. de Viveiros, P. J. Doe, M. Fertl, J. A. Formaggio, J. K. Gaison, L. Gladstone, M. Grando, M. Guigue, J. Hartse, K. M. Heeger, X. Huyan, J. Johnston, A. M. Jones, K. Kazkaz, B. H. LaRoque, M. Li, A. Lindman, E. Machado, A. Marsteller, C. Matthé, R. Mohiuddin, B. Monreal, R. Mueller, J. A. Nikkel, E. Novitski, N. S. Oblath, J. I.

Peña, W. Pettus, R. Reimann, R. G. H. Robertson, D. Rosa De Jesús, G. Rybka, L. Saldaña, M. Schram, P. L. Slocum, J. Stachurska, Y.-H. Sun, P. T. Surukuchi, J. R. Tedeschi, A. B. Telles, F. Thomas, M. Thomas, L. A. Thorne, T. Thümmler, L. Tvrznikova, W. Van De Pontseele, B. A. VanDevender, J. Weintroub, T. E. Weiss, T. Wendler, A. Young, E. Zayas, and A. Ziegler. Tritium Beta Spectrum Measurement and Neutrino Mass Limit from Cyclotron Radiation Emission Spectroscopy. *Physical Review Letters*, 131:102502, Sep 2023.

[13] Eleonora Di Valentino, Stefano Gariazzo, and Olga Mena. Most constraining cosmological neutrino mass bounds. *Physical Review D*, 104:083504, Oct 2021.

[14] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Physical Review Letters*, 13:321--323, Aug 1964.

[15] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Physical Review Letters*, 13:508--509, Oct 1964.

[16] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global Conservation Laws and Massless Particles. *Physical Review Letters*, 13:585--587, Nov 1964.

[17] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of Neutrinos. *Physical Review*, 109:1015--1017, Feb 1958.

[18] Pedro A. N. Machado. Neutrino Properties and Interactions. `https://arxiv.org/abs/2206.13449`, 2022.

[19] E. Majorana. Symmetric Theory of Electrons and Positrons. *Il Nuovo Cimento*, 14:201233, 1937.

[20] M. Fukugita and T. Yanagida. Barygenesis Without Grand Unification. *Physics Letters B*, 174(1):45--47, 1986.

[21] Mikola, Veera. Results from the T2K+NOvA Joint Analysis. *European Journal of Physics Web of Conferences.*, 312:02002, 2024.

[22] M. Goeppert-Mayer. Double Beta-Disintegration. *Physical Review*, 48:512--516, Sep 1935.

[23] C. F.v. Weizsäcker. Zur Theorie der Kernmassen. *Zeitschrift für Physik*, 96(7-8):431--458, 7 1935.

[24] J. Kotila and F. Iachello. Phase-space factors for double-β decay. *Physical Review C*, 85:034316, Mar 2012.

[25] Mihai Horoi and Andrei Neacsu. Shell Model Study of Using an Effective Field Theory for Disentangling Several Contributions to Neutrinoless Double-β Decay. *Physical Review C*, 98:035502, Sep 2018.

[26] B. Pritychenko and V.I. Tretyak. Comprehensive Review of 2β decay half-lives. *Atomic Data and Nuclear Data Tables*, 161:101694, 2025.

[27] Jouni Suhonen and Joel Kostensalo. Double Decay and the Axial Strength. *Frontiers in Physics*, 7, 2019.

[28] W. H. Furry. On Transition Probabilities in Double Beta-Disintegration. *Physical Review*, 56:1184--1193, Dec 1939.

[29] J. Schechter and J. W. F. Valle. Neutrinoless double-β decay in SU(2)×U(1) theories. *Physical Review D*, 25:2951--2954, Jun 1982.

[30] WERNER RODEJOHANN. NEUTRINO-LESS DOUBLE BETA DECAY AND PARTICLE PHYSICS. *International Journal of Modern Physics E*, 20(09):1833--1930, 2011.

[31] S. Abe, T. Araki, K. Chiba, T. Eda, M. Eizuka, Y. Funahashi, A. Furuto, A. Gando, Y. Gando, S. Goto, T. Hachiya, K. Hata, K. Ichimura, S. Ieki, H. Ikeda, K. Inoue, K. Ishidoshiro, Y. Kamei, N. Kawada, Y. Kishimoto, M. Koga, A. Marthe, Y. Matsumoto, T. Mitsui, H. Miyake, D. Morita, R. Nakajima, K. Nakamura, R. Nakamura, R. Nakamura, J. Nakane, T. Ono, H. Ozaki, K. Saito, T. Sakai, I. Shimizu, J. Shirai, K. Shiraishi, A. Suzuki, K. Tachibana, K. Tamae, H. Watanabe, K. Watanabe, S. Yoshida, S. Umehara, K. Fushimi, K. Kotera, Y. Urano, B. E. Berger, B. K. Fujikawa, J. G. Learned, J. Maricic, Z. Fu, S. Ghosh, J. Smolsky, L. A. Winslow, Y. Efremenko, H. J. Karwowski, D. M. Markoff, W. Tornow, S. Dell'Oro, T. O'Donnell, J. A. Detwiler, S. Enomoto, M. P. Decowski, K. M. Weerman, C. Grant, O. Penek, H. Song, A. Li, S. N. Axani, M. Garcia, and M. Sarfraz. Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset. `https://arxiv.org/abs/2406.11438v1`, 6 2024.

[32] Ivan Esteban, M. C. Gonzalez-Garcia, Michele Maltoni, Thomas Schwetz, and Albert Zhou. The fate of hints: updated global analysis of three-flavor neutrino oscillations. *Journal of High Energy Physics*, 09:178, 2020.

[33] M. Agostini, G. R. Araujo, A. M. Bakalyarov, M. Balata, I. Barabanov, L. Baudis, C. Bauer, E. Bellotti, S. Belogurov, A. Bettini, L. Bezrukov, V. Biancacci, D. Borowicz, E. Bossio, V. Bothe, V. Brudanin, R. Brugnera, A. Caldwell, C. Cattadori, A. Chernogorov, T. Comellato, V. D'Andrea, E. V. Demidova, N. Di Marco, E. Doroshkevich, F. Fischer, M. Fomina, A. Gangapshev, A. Garfagnini, C. Gooch, P. Grabmayr, V. Gurentsov, K. Gusev, J. Hakenmüller, S. Hemmer, R. Hiller, W. Hofmann, J. Huang, M. Hult, L. V. Inzhechik, J. Janicskó Csáthy, J. Jochum, M. Junker, V. Kazalov, Y. Kermaïdic, H. Khushbakht, T. Kihm, I. V. Kirpichnikov, A. Klimenko, R. Kneißl, K. T. Knöpfle, O. Kochetov, V. N. Kornoukhov, P. Krause, V. V. Kuzminov, M. Laubenstein, A. Lazzaro, M. Lindner, I. Lippi, A. Lubashevskiy, B. Lubsandorzhiev, G. Lutter, C. Macolino, B. Majorovits, W. Maneschg, L. Manzanil-

las, M. Miloradovic, R. Mingazheva, M. Misiaszek, P. Moseev, Y. Müller, I. Nemchenok, K. Panas, L. Pandola, K. Pelczar, L. Pertoldi, P. Piseri, A. Pullia, C. Ransom, L. Rauscher, S. Riboldi, N. Rumyantseva, C. Sada, F. Salamida, S. Schönert, J. Schreiner, M. Schütt, A.-K. Schütz, O. Schulz, M. Schwarz, B. Schwingenheuer, O. Selivanenko, E. Shevchik, M. Shirchenko, L. Shtembari, H. Simgen, A. Smolnikov, D. Stukov, A. A. Vasenko, A. Veresnikova, C. Vignoli, K. von Sturm, T. Wester, C. Wiesinger, M. Wojcik, E. Yanovich, B. Zatschler, I. Zhitnikov, S. V. Zhukov, D. Zinatulina, A. Zschocke, A. J. Zsigmond, K. Zuber, and G. Zuzel. Final Results of GERDA on the Search for Neutrinoless Double-β Decay. *Physical Review Letters*, 125:252502, Dec 2020.

[34] E. Armengaud, C. Augier, A. S. Barabash, F. Bellini, G. Benato, A. Benoît, M. Beretta, L. Bergé, J. Billard, Yu. A. Borovlev, Ch. Bourgeois, V. B. Brudanin, P. Camus, L. Cardani, N. Casali, A. Cazes, M. Chapellier, F. Charlieux, D. Chiesa, M. de Combarieu, I. Dafinei, F. A. Danevich, M. De Jesus, T. Dixon, L. Dumoulin, K. Eitel, F. Ferri, B. K. Fujikawa, J. Gascon, L. Gironi, A. Giuliani, V. D. Grigorieva, M. Gros, E. Guerard, D. L. Helis, H. Z. Huang, R. Huang, J. Johnston, A. Juillard, H. Khalife, M. Kleifges, V. V. Kobychev, Yu. G. Kolomensky, S. I. Konovalov, A. Leder, P. Loaiza, L. Ma, E. P. Makarov, P. de Marcillac, R. Mariam, L. Marini, S. Marnieros, D. Misiak, X.-F. Navick, C. Nones, E. B. Norman, V. Novati, E. Olivieri, J. L. Ouellet, L. Pagnanini, P. Pari, L. Pattavina, B. Paul, M. Pavan, H. Peng, G. Pessina, S. Pirro, D. V. Poda, O. G. Polischuk, S. Pozzi, E. Previtali, Th. Redon, A. Rojas, S. Rozov, C. Rusconi, V. Sanglard, J. A. Scarpaci, K. Schäffner, B. Schmidt, Y. Shen, V. N. Shlegel, B. Siebenborn, V. Singh, C. Tomei, V. I. Tretyak, V. I. Umatov, L. Vagneron, M. Velázquez, B. Welliver, L. Winslow, M. Xue, E. Yakushev, M. Zarytskyy, and A. S. Zolotarova. New Limit for Neutrinoless Double-Beta Decay of ^{100}Mo from the CUPID-Mo Experiment. *Physical Review Letters*, 126:181802, May 2021.

[35] D. Q. Adams, C. Alduino, K. Alfonso, F. T. Avignone, O. Azzolini, G. Bari, F. Bellini, G. Benato, M. Beretta, M. Biassoni, A. Branca, C. Brofferio, C. Bucci, J. Camilleri, A. Caminata, A. Campani, L. Canonica, X. G. Cao, S. Capelli, C. Capelli, L. Cappelli, L. Cardani, P. Carniti, N. Casali, E. Celi, D. Chiesa, M. Clemenza, S. Copello, O. Cremonesi, R. J. Creswick, A. D'Addabbo, I. Dafinei, F. Del Corso, S. Dell'Oro, S. Di Domizio, S. Di Lorenzo, V. Dompè, D. Q. Fang, G. Fantini, M. Faverzani, E. Ferri, F. Ferroni, E. Fiorini, M. A. Franceschi, S. J. Freedman, S. H. Fu, B. K. Fujikawa, S. Ghislandi, A. Giachero, A. Gianvecchio, L. Gironi, A. Giuliani, P. Gorla, C. Gotti, T. D. Gutierrez, K. Han, E. V. Hansen, K. M. Heeger, R. G. Huang, H. Z. Huang, J. Johnston, G. Keppel, Yu. G. Kolomensky, R. Kowalski, R. Liu, L. Ma, Y. G. Ma, L. Marini, R. H. Maruyama, D. Mayer, Y. Mei, S. Morganti, T. Napolitano, M. Nastasi, J. Nikkel, C. Nones, E. B. Norman, A. Nucciotti, I. Nutini, T. O'Donnell, M. Olmi, J. L. Ouellet, S. Pagan, C. E. Pagliarone, L. Pagnanini, M. Pallavicini, L. Pattavina, M. Pavan, G. Pessina, V. Pettinacci, C. Pira, S. Pirro, S. Pozzi, E. Previtali, A. Puiu, S. Quitadamo, A. Ressa, C. Rosenfeld, S. Sangiorgio, B. Schmidt,

N. D. Scielzo, V. Sharma, V. Singh, M. Sisti, D. Speller, P. T. Surukuchi, L. Taffarello, F. Terranova, C. Tomei, K. J. Vetter, M. Vignati, S. L. Wagaarachchi, B. S. Wang, B. Welliver, J. Wilson, K. Wilson, L. A. Winslow, S. Zimmermann, and S. Zucchelli. New Direct Limit on Neutrinoless Double Beta Decay Half-Life of ^{128}Te with CUORE. *Physical Review Letters*, 129:222501, Nov 2022.

[36] Jonathan Engel and Javier Menéndez. Status and Future of Nuclear Matrix Elements for Neutrinoless Double-Beta Decay: A Review. *Reports on Progress in Physics*, 80(4):046301, mar 2017.

[37] A. Belley, C. G. Payne, S. R. Stroberg, T. Miyagi, and J. D. Holt. Ab Initio Neutrinoless Double-Beta Decay Matrix Elements for ^{48}Ca, ^{76}Ge, and ^{82}Se. *Physical Review Letters*, 126:042502, Jan 2021.

[38] Matteo Agostini, Giovanni Benato, Jason A. Detwiler, Javier Menéndez, and Francesco Vissani. Toward the discovery of matter creation with neutrinoless $\beta\beta$ decay. *Reviews of Modern Physics*, 95:025002, May 2023.

[39] S. Navas et al. Review of particle physics. *Physical Review D*, 110(3):030001, 2024.

[40] The Sno+ Collaboration, V. Albanese, R. Alves, M. R. Anderson, S. Andringa, L. Anselmo, E. Arushanova, S. Asahi, M. Askins, D. J. Auty, A. R. Back, S. Back, F. Bar o, Z. Barnard, A. Barr, N. Barros, D. Bartlett, R. Bayes, C. Beaudoin, E. W. Beier, G. Berardi, A. Bialek, S. D. Biller, E. Blucher, R. Bonventre, M. Boulay, D. Braid, E. Caden, E. J. Callaghan, J. Caravaca, J. Carvalho, L. Cavalli, D. Chauhan, M. Chen, O. Chkvorets, K. J. Clark, B. Cleveland, C. Connors, D. Cookman, I. T. Coulter, M. A. Cox, D. Cressy, X. Dai, C. Darrach, B. Davis-Purcell, C. Deluce, M. M. Depatie, F. Descamps, F. Di Lodovico, J. Dittmer, A. Doxtator, N. Duhaime, F. Duncan, J. Dunger, A. D. Earle, D. Fabris, E. Falk, A. Farrugia, N. Fatemighomi, C. Felber, V. Fischer, E. Fletcher, R. Ford, K. Frankiewicz, N. Gagnon, A. Gaur, J. Gauthier, A. Gibson-Foster, K. Gilje, O. I. González-Reina, D. Gooding, P. Gorel, K. Graham, C. Grant, J. Grove, S. Grullon, E. Guillian, S. Hall, A. L. Hallin, D. Hallman, S. Hans, J. Hartnell, P. Harvey, M. Hedayatipour, W. J. Heintzelman, J. Heise, R. L. Helmer, B. Hodak, M. Hodak, M. Hood, D. Horne, B. Hreljac, J. Hu, S. M.A. Hussain, T. Iida, A. S. Inácio, C. M. Jackson, N. A. Jelley, C. J. Jillings, C. Jones, P. G. Jones, K. Kamdin, T. Kaptanoglu, J. Kaspar, K. Keeter, C. Kefelian, P. Khaghani, L. Kippenbrock, J. R. Klein, R. Knapik, J. Kofron, L. L. Kormos, S. Korte, B. Krar, C. Kraus, C. B. Krauss, T. Kroupová, K. Labe, F. Lafleur, I. Lam, C. Lan, B. J. Land, R. Lane, S. Langrock, P. Larochelle, S. Larose, A. Latorre, I. Lawson, L. Lebanowski, G. M. Lefeuvre, E. J. Leming, A. Li, O. Li, J. Lidgard, B. Liggins, P. Liimatainen, Y. H. Lin, X. Liu, Y. Liu, V. Lozza, M. Luo, S. Maguire, A. Maio, K. Majumdar, S. Manecki, J. Maneira, R. D. Martin, E. Marzec, A. Mastbaum, A. Mathewson, N. McCauley, A. B. McDonald, K. McFarlane, P. Mekarski, M. Meyer, C. Miller, C. Mills, M. Mlejnek, E. Mony, B. Morissette, I. Morton-Blake, M. J. Mottram, S. Nae, M. Nirkko, L. J. Nolan, V. M.

Novikov, H. M. O'Keeffe, E. O'Sullivan, G. D. Orebi Gann, M. J. Parnell, J. Paton, S. J.M. Peeters, T. Pershing, Z. Petriw, J. Petzoldt, L. Pickard, D. Pracsovics, G. Prior, J. C. Prouty, S. Quirk, S. Read, A. Reichold, S. Riccetto, R. Richardson, M. Rigan, I. Ritchie, A. Robertson, B. C. Robertson, J. Rose, R. Rosero, P. M. Rost, J. Rumleskie, M. A. Schumaker, M. H. Schwendener, D. Scislowski, J. Secrest, M. Seddighin, L. Segui, S. Seibert, I. Semenec, F. Shaker, T. Shantz, M. K. Sharma, T. M. Shokair, L. Sibley, J. R. Sinclair, K. Singh, P. Skensved, M. Smiley, T. Sonley, A. Sörensen, M. St-Amant, R. Stainforth, S. Stankiewicz, M. Strait, M. I. Stringer, A. Stripay, R. Svoboda, S. Tacchino, B. Tam, C. Tanguay, J. Tatar, L. Tian, N. Tolich, J. Tseng, H. W.C. Tseung, E. Turner, R. Van Berg, E. Vázquez-Jáuregui, J. G.C. Veinot, C. J. Virtue, B. Von Krosigk, J. M.G. Walker, M. Walker, J. Wallig, S. C. Walton, J. Wang, M. Ward, O. Wasalski, J. Waterfield, J. J. Weigand, R. F. White, J. R. Wilson, T. J. Winchester, P. Woosaree, A. Wright, J. P. Yanez, M. Yeh, T. Zhang, Y. Zhang, T. Zhao, K. Zuber, and A. Zummo. The SNO+ experiment. *Journal of Instrumentation*, 16(08):P08059, 8 2021.

[41] J. Boger, R. L. Hahn, J. K. Rowley, A. L. Carter, B. Hollebone, D. Kessler, I. Blevis, F. Dalnoki-Veress, A. Dekok, J. Farine, D. R. Grant, C. K. Hargrove, G. Laberge, I. Levine, K. McFarlane, H. Mes, A. T. Noble, V. M. Novikov, M. O'Neill, M. Shatkay, C. Shewchuk, D. Sinclair, E. T.H. Clifford, R. Deal, E. D. Earle, E. Gaudette, G. Milton, B. Sur, J. Bigu, J. H.M. Cowan, D. L. Cluff, E. D. Hallman, R. U. Haq, J. Hewett, J. G. Hykawy, G. Jonkmans, R. Michaud, A. Roberge, J. Roberts, E. Saettler, M. H. Schwendener, H. Seifert, D. Sweezey, R. Tafirout, C. J. Virtue, D. N. Beck, Y. D. Chan, X. Chen, M. R. Dragowsky, F. W. Dycus, J. Gonzalez, M. C.P. Isaac, Y. Kajiyama, G. W. Koehler, K. T. Lesko, M. C. Moebus, E. B. Norman, C. E. Okada, A. W.P. Poon, P. Purgalis, A. Schuelke, A. R. Smith, R. G. Stokstad, S. Turner, I. Zlimen, J. M. Anaya, T. J. Bowles, S. J. Brice, Ernst Ingo Esch, M. M. Fowler, Azriel Goldschmidt, A. Hime, A. F. McGirt, G. G. Miller, W. A. Teasdale, J. B. Wilhelmy, J. M. Wouters, J. D. Anglin, M. Bercovitch, W. F. Davidson, R. S. Storey, S. Biller, R. A. Black, R. J. Boardman, M. G. Bowler, J. Cameron, B. Cleveland, A. P. Ferraris, G. Doucas, H. Heron, C. Howard, N. A. Jelley, A. B. Knox, M. Lay, W. Locke, J. Lyon, S. Majerus, M. Moorhead, and M. Omori. Sudbury neutrino observatory. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 449(1), 2000.

[42] M. R. Anderson, S. Andringa, M. Askins, D. J. Auty, F. Barão, N. Barros, R. Bayes, E. W. Beier, A. Bialek, S. D. Biller, E. Blucher, M. Boulay, E. Caden, E. J. Callaghan, J. Caravaca, M. Chen, O. Chkvorets, B. Cleveland, D. Cookman, J. Corning, M. A. Cox, C. Deluce, M. M. Depatie, F. Di Lodovico, J. Dittmer, E. Falk, N. Fatemighomi, V. Fischer, R. Ford, K. Frankiewicz, A. Gaur, K. Gilje, O. I. González-Reina, D. Gooding, C. Grant, J. Grove, A. L. Hallin, D. Hallman, J. Hartnell, W. J. Heintzelman, R. L. Helmer, J. Hu, R. Hunt-Stokes, S. M.A. Hussain, A. S. Inácio, C. J. Jillings, T. Kaptanoglu, P. Khaghani, H. Khan, J. R. Klein, L. L. Kormos, B. Krar, C. Kraus,

C. B. Krauss, T. Kroupová, I. Lam, B. J. Land, A. Latorre, I. Lawson, L. Lebanowski, C. Lefebvre, A. Li, J. Lidgard, Y. H. Lin, Y. Liu, V. Lozza, M. Luo, A. Maio, S. Manecki, J. Maneira, R. D. Martin, N. McCauley, A. B. McDonald, M. Meyer, C. Mills, I. Morton-Blake, S. Nae, M. Nirkko, L. J. Nolan, H. M. O'Keeffe, G. D. Orebi Gann, J. Page, W. Parker, J. Paton, S. J.M. Peeters, T. Pershing, L. Pickard, G. Prior, P. Ravi, A. Reichold, S. Riccetto, R. Richardson, M. Rigan, J. Rose, J. Rumleskie, I. Semenec, F. Shaker, M. K. Sharma, P. Skensved, M. Smiley, R. Stainforth, R. Svoboda, B. Tam, J. Tseng, E. Turner, S. Valder, E. Vázquez-Jáuregui, J. G.C. Veinot, C. J. Virtue, J. Wang, M. Ward, J. J. Weigand, J. R. Wilson, A. Wright, J. P. Yanez, M. Yeh, S. Yu, T. Zhang, Y. Zhang, K. Zuber, and A. Zummo. Optical calibration of the SNO+ detector in the water phase with deployed sources. *Journal of Instrumentation*, 16(10):P10021, 10 2021.

[43] Tereza Kroupová. *Improving the Sensitivity to Neutrinoless Double Beta Decay in SNO+*. Doctoral dissertation, University of Oxford `https://snoplus.phy.queensu.ca/theses/tkroupova.pdf`, 2021.

[44] A. Allega, M. R. Anderson, S. Andringa, M. Askins, D. J. Auty, A. Bacon, N. Barros, F. Barão, R. Bayes, E. W. Beier, T. S. Bezerra, A. Bialek, S. D. Biller, E. Blucher, E. Caden, E. J. Callaghan, S. Cheng, M. Chen, O. Chkvorets, B. Cleveland, D. Cookman, J. Corning, M. A. Cox, R. Dehghani, C. Deluce, M. M. Depatie, J. Dittmer, K. H. Dixon, F. Di Lodovico, E. Falk, N. Fatemighomi, R. Ford, K. Frankiewicz, A. Gaur, O. I. González-Reina, D. Gooding, C. Grant, J. Grove, A. L. Hallin, D. Hallman, J. Hartnell, W. J. Heintzelman, R. L. Helmer, J. Hu, R. Hunt-Stokes, S. M. A. Hussain, A. S. Inácio, C. J. Jillings, T. Kaptanoglu, P. Khaghani, H. Khan, J. R. Klein, L. L. Kormos, B. Krar, C. Kraus, C. B. Krauss, T. Kroupová, I. Lam, B. J. Land, I. Lawson, L. Lebanowski, J. Lee, C. Lefebvre, J. Lidgard, Y. H. Lin, V. Lozza, M. Luo, A. Maio, S. Manecki, J. Maneira, R. D. Martin, N. McCauley, A. B. McDonald, M. Meyer, C. Mills, I. Morton-Blake, S. Naugle, L. J. Nolan, H. M. O'Keeffe, G. D. Orebi Gann, J. Page, W. Parker, J. Paton, S. J. M. Peeters, L. Pickard, P. Ravi, A. Reichold, S. Riccetto, R. Richardson, M. Rigan, J. Rose, J. Rumleskie, I. Semenec, P. Skensved, M. Smiley, R. Svoboda, B. Tam, J. Tseng, E. Turner, S. Valder, J. G. C. Veinot, C. J. Virtue, E. Vázquez-Jáuregui, J. Wang, M. Ward, J. J. Weigand, J. D. Wilson, J. R. Wilson, A. Wright, J. P. Yanez, S. Yang, M. Yeh, S. Yu, T. Zhang, Y. Zhang, K. Zuber, and A. Zummo. Improved Search for the Invisible Modes of Nucleon Decay in Water With the SNO + Detector. *Physical Review D*, 105:112012, Jun 2022.

[45] M. R. Anderson, S. Andringa, M. Askins, D. J. Auty, N. Barros, F. Barão, R. Bayes, E. W. Beier, A. Bialek, S. D. Biller, E. Blucher, R. Bonventre, M. Boulay, E. Caden, E. J. Callaghan, J. Caravaca, D. Chauhan, M. Chen, O. Chkvorets, B. Cleveland, M. A. Cox, M. M. Depatie, J. Dittmer, F. Di Lodovico, A. D. Earle, E. Falk, N. Fatemighomi, V. Fischer, E. Fletcher, R. Ford, K. Frankiewicz, K. Gilje, D. Gooding, C. Grant, J. Grove, A. L. Hallin, D. Hallman, S. Hans, J. Hartnell, P. Harvey,

W. J. Heintzelman, R. L. Helmer, D. Horne, B. Hreljac, J. Hu, A. S. M. Hussain, A. S. Inácio, C. J. Jillings, T. Kaptanoglu, P. Khaghani, J. R. Klein, R. Knapik, L. L. Kormos, B. Krar, C. Kraus, C. B. Krauss, T. Kroupova, I. Lam, B. J. Land, A. LaTorre, I. Lawson, L. Lebanowski, E. J. Leming, A. Li, J. Lidgard, B. Liggins, Y. H. Lin, Y. Liu, V. Lozza, M. Luo, S. Maguire, A. Maio, S. Manecki, J. Maneira, R. D. Martin, E. Marzec, A. Mastbaum, N. McCauley, A. B. McDonald, P. Mekarski, M. Meyer, C. Mills, I. Morton-Blake, S. Nae, M. Nirkko, L. J. Nolan, H. M. O'Keeffe, G. D. Orebi Gann, M. J. Parnell, J. Paton, S. J. M. Peeters, T. Pershing, L. Pickard, G. Prior, A. Reichold, S. Riccetto, R. Richardson, M. Rigan, J. Rose, R. Rosero, P. M. Rost, J. Rumleskie, I. Semenec, F. Shaker, M. K. Sharma, K. Singh, P. Skensved, M. Smiley, M. I. Stringer, R. Svoboda, B. Tam, L. Tian, J. Tseng, E. Turner, R. Van Berg, J. G. C. Veinot, C. J. Virtue, E. Vázquez-Jáuregui, S. C. Walton, J. Wang, M. Ward, J. J. Weigand, J. R. Wilson, P. Woosaree, A. Wright, J. P. Yanez, M. Yeh, T. Zhang, Y. Zhang, K. Zuber, and A. Zummo. Measurement of Neutron-Proton Capture in the SNO+ Water Phase. *Physical Review C*, 102:014002, Jul 2020.

[46] A. Allega, M. R. Anderson, S. Andringa, M. Askins, D. M. Asner, D. J. Auty, A. Bacon, F. Barão, N. Barros, R. Bayes, E. W. Beier, A. Bialek, S. D. Biller, E. Blucher, E. Caden, E. J. Callaghan, M. Chen, S. Cheng, B. Cleveland, D. Cookman, J. Corning, M. A. Cox, R. Dehghani, J. Deloye, M. M. Depatie, F. Di Lodovico, C. Dima, J. Dittmer, K. H. Dixon, M. S. Esmaeilian, E. Falk, N. Fatemighomi, R. Ford, A. Gaur, O. I. González-Reina, D. Gooding, C. Grant, J. Grove, S. Hall, A. L. Hallin, D. Hallman, W. J. Heintzelman, R. L. Helmer, C. Hewitt, B. Hreljac, J. Hu, P. Huang, R. Hunt-Stokes, S. M. A. Hussain, A. S. Inácio, C. J. Jillings, S. Kaluzienski, T. Kaptanoglu, J. Kladnik, J. R. Klein, L. L. Kormos, B. Krar, C. Kraus, C. B. Krauss, T. Kroupová, C. Lake, L. Lebanowski, C. Lefebvre, V. Lozza, M. Luo, A. Maio, S. Manecki, J. Maneira, R. D. Martin, N. McCauley, A. B. McDonald, G. Milton, D. Morris, M. Mubasher, S. Naugle, L. J. Nolan, H. M. O'Keeffe, G. D. Orebi Gann, J. Page, K. Paleshi, W. Parker, J. Paton, S. J. M. Peeters, L. Pickard, B. Quenallata, P. Ravi, A. Reichold, S. Riccetto, J. Rose, R. Rosero, I. Semenec, J. Simms, P. Skensved, M. Smiley, R. Svoboda, B. Tam, J. Tseng, E. Vázquez-Jáuregui, C. J. Virtue, M. Ward, J. R. Wilson, J. D. Wilson, A. Wright, S. Yang, M. Yeh, Z. Ye, S. Yu, Y. Zhang, K. Zuber, and A. Zummo. Measurement of the ^{8}B Solar Neutrino Flux Using The Full SNO+ Water Phase Datset. *Physical Review D*, 110:122003, Dec 2024.

[47] M. Anderson, S. Andringa, S. Asahi, M. Askins, D. J. Auty, N. Barros, D. Bartlett, F. Barão, R. Bayes, E. W. Beier, A. Bialek, S. D. Biller, E. Blucher, R. Bonventre, M. Boulay, E. Caden, E. J. Callaghan, J. Caravaca, D. Chauhan, M. Chen, O. Chkvorets, B. Cleveland, C. Connors, I. T. Coulter, M. M. Depatie, F. Di Lodovico, F. Duncan, J. Dunger, E. Falk, V. Fischer, E. Fletcher, R. Ford, N. Gagnon, K. Gilje, C. Grant, J. Grove, A. L. Hallin, D. Hallman, S. Hans, J. Hartnell, W. J. Heintzelman, R. L. Helmer, J. L. Hernández-Hernández, B. Hreljac, J. Hu, A. S. Inácio, C. J. Jillings, T. Kaptanoglu, P. Khaghani, J. R. Klein, R. Knapik, L. L. Kormos, B. Krar,

C. Kraus, C. B. Krauss, T. Kroupova, I. Lam, B. J. Land, R. Lane, A. LaTorre, I. Lawson, L. Lebanowski, E. J. Leming, A. Li, J. Lidgard, B. Liggins, Y. Liu, V. Lozza, M. Luo, S. Maguire, A. Maio, S. Manecki, J. Maneira, R. D. Martin, E. Marzec, A. Mastbaum, N. McCauley, A. B. McDonald, P. Mekarski, M. Meyer, M. Mlejnek, I. Morton-Blake, S. Nae, M. Nirkko, H. M. O'Keeffe, G. D. Orebi Gann, M. J. Parnell, J. Paton, S. J. M. Peeters, T. Pershing, L. Pickard, D. Pracsovics, G. Prior, A. Reichold, R. Richardson, M. Rigan, J. Rose, R. Rosero, J. Rumleskie, I. Semenec, K. Singh, P. Skensved, M. I. Stringer, R. Svoboda, B. Tam, L. Tian, J. Tseng, E. Turner, R. Van Berg, J. G. C. Veinot, C. J. Virtue, E. Vázquez-Jáuregui, J. Wang, J. J. Weigand, J. R. Wilson, P. Woosaree, A. Wright, J. P. Yanez, M. Yeh, K. Zuber, and A. Zummo. Measurement of the ^{8}B Solar Neutrino Flux in SNO+ With Very Low Backgrounds. *Physical Review D*, 99:012012, Jan 2019.

[48] A. Allega, M. R. Anderson, S. Andringa, J. Antunes, M. Askins, D. J. Auty, A. Bacon, N. Barros, F. Barão, R. Bayes, E. W. Beier, T. S. Bezerra, A. Bialek, S. D. Biller, E. Blucher, E. Caden, E. J. Callaghan, S. Cheng, M. Chen, B. Cleveland, D. Cookman, J. Corning, M. A. Cox, R. Dehghani, J. Deloye, C. Deluce, M. M. Depatie, J. Dittmer, K. H. Dixon, F. Di Lodovico, E. Falk, N. Fatemighomi, R. Ford, K. Frankiewicz, A. Gaur, O. I. González-Reina, D. Gooding, C. Grant, J. Grove, A. L. Hallin, D. Hallman, W. J. Heintzelman, R. L. Helmer, J. Hu, R. Hunt-Stokes, S. M. A. Hussain, A. S. Inácio, C. J. Jillings, S. Kaluzienski, T. Kaptanoglu, P. Khaghani, H. Khan, J. R. Klein, L. L. Kormos, B. Krar, C. Kraus, C. B. Krauss, T. Kroupová, I. Lam, B. J. Land, I. Lawson, L. Lebanowski, J. Lee, C. Lefebvre, J. Lidgard, Y. H. Lin, V. Lozza, M. Luo, A. Maio, S. Manecki, J. Maneira, R. D. Martin, N. McCauley, A. B. McDonald, C. Mills, I. Morton-Blake, S. Naugle, L. J. Nolan, H. M. O'Keeffe, G. D. Orebi Gann, J. Page, W. Parker, J. Paton, S. J. M. Peeters, L. Pickard, P. Ravi, A. Reichold, S. Riccetto, R. Richardson, M. Rigan, J. Rose, R. Rosero, J. Rumleskie, I. Semenec, P. Skensved, M. Smiley, R. Svoboda, B. Tam, J. Tseng, E. Turner, S. Valder, C. J. Virtue, E. Vázquez-Jáuregui, J. Wang, M. Ward, J. R. Wilson, J. D. Wilson, A. Wright, J. P. Yanez, S. Yang, M. Yeh, S. Yu, Y. Zhang, K. Zuber, and A. Zummo. Evidence of Antineutrinos from Distant Reactors Using Pure Water at SNO+. *Physical Review Letters*, 130:091801, Mar 2023.

[49] A. Allega, M. R. Anderson, S. Andringa, J. Antunes, M. Askins, D. J. Auty, A. Bacon, J. Baker, N. Barros, F. Barão, R. Bayes, E. W. Beier, T. S. Bezerra, A. Bialek, S. D. Biller, E. Blucher, E. Caden, E. J. Callaghan, M. Chen, S. Cheng, B. Cleveland, D. Cookman, J. Corning, M. A. Cox, R. Dehghani, J. Deloye, M. M. Depatie, F. Di Lodovico, J. Dittmer, K. H. Dixon, E. Falk, N. Fatemighomi, R. Ford, A. Gaur, O. I. González-Reina, D. Gooding, C. Grant, J. Grove, S. Hall, A. L. Hallin, W. J. Heintzelman, R. L. Helmer, C. Hewitt, B. Hreljac, V. Howard, J. Hu, R. Hunt-Stokes, S. M. A. Hussain, A. S. Inácio, C. J. Jillings, S. Kaluzienski, T. Kaptanoglu, P. Khaghani, H. Khan, J. R. Klein, L. L. Kormos, B. Krar, C. Kraus, C. B. Krauss, T. Kroupová, C. Lake, L. Lebanowski, J. Lee, C. Lefebvre, Y. H. Lin, V. Lozza, M. Luo, A. Maio, S. Ma-

necki, J. Maneira, R. D. Martin, N. McCauley, A. B. McDonald, C. Mills, G. Milton, I. Morton-Blake, M. Mubasher, A. Molina Colina, D. Morris, S. Naugle, L. J. Nolan, H. M. O'Keeffe, G. D. Orebi Gann, J. Page, K. Paleshi, W. Parker, J. Paton, S. J. M. Peeters, L. Pickard, P. Ravi, A. Reichold, S. Riccetto, M. Rigan, J. Rose, R. Rosero, J. Rumleskie, I. Semenec, P. Skensved, M. Smiley, J. Smith, R. Svoboda, B. Tam, J. Tseng, S. Valder, E. Vázquez-Jáuregui, C. J. Virtue, J. Wang, M. Ward, J. R. Wilson, J. D. Wilson, A. Wright, J. P. Yanez, S. Yang, M. Yeh, Z. Ye, S. Yu, Y. Zhang, K. Zuber, and A. Zummo. Event-by-event Direction Reconstruction of Solar Neutrinos in a High Light-Yield Liquid Scintillator. *Physical Review D*, 109:072002, Apr 2024.

[50] A. Allega et al. Initial measurement of reactor antineutrino oscillation at SNO+. *The European Physical Journal C*, 85(1):17, 2025. [Erratum: Eur.Phys.J.C 85, 296 (2025)].

[51] J B Birks. Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations. *Proceedings of the Physical Society. Section A*, 64(10):874, oct 1951.

[52] C. Arpesella, G. Bellini, J. Benziger, S. Bonetti, B. Caccianiga, F. Calaprice, F. Dalnoki-Veress, D. D'Angelo, H. de Kerret, A. Derbin, M. Deutsch, A. Etenko, K. Fomenko, R. Ford, D. Franco, B. Freudiger, C. Galbiati, S. Gazzana, M. Giammarchi, M. Goeger-Neff, A. Goretti, C. Grieb, S. Hardy, G. Heusser, Aldo Ianni, Andrea Ianni, M. Joyce, G. Korga, D. Kryn, M. Laubenstein, M. Leung, E. Litvinovich, P. Lombardi, L. Ludhova, I. Machulin, G. Manuzio, A. Martemianov, F. Masetti, K. McCarty, E. Meroni, L. Miramonti, M. Misiaszek, D. Montanari, M. E. Monzani, V. Muratova, L. Niedermeier, L. Oberauer, M. Obolensky, F. Ortica, M. Pallavicini, L. Papp, L. Perasso, A. Pocar, R. S. Raghavan, G. Ranucci, A. Razeto, A. Sabelnikov, C. Salvo, S. Schönert, H. Simgen, O. Smirnov, M. Skorokhvatov, A. Sonnenschein, A. Sotnikov, S. Sukhotin, Y. Suvorov, V. Tarasenkov, R. Tartaglia, G. Testera, D. Vignaud, S. Vitale, R. B. Vogelaar, F. von Feilitzsch, M. Wojcik, O. Zaimidoroga, S. Zavatarelli, and G. Zuzel. First real time detection of 7Be solar neutrinos by Borexino. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 658(4):101--108, 1 2008.

[53] K. Eguchi, S. Enomoto, K. Furuno, J. Goldman, H. Hanada, H. Ikeda, K. Ikeda, K. Inoue, K. Ishihara, W. Itoh, T. Iwamoto, T. Kawaguchi, T. Kawashima, H. Kinoshita, Y. Kishimoto, M. Koga, Y. Koseki, T. Maeda, T. Mitsui, M. Motoki, K. Nakajima, M. Nakajima, T. Nakajima, H. Ogawa, K. Owada, T. Sakabe, I. Shimizu, J. Shirai, F. Suekane, A. Suzuki, K. Tada, O. Tajima, T. Takayama, K. Tamae, H. Watanabe, J. Busenitz, Z. Djurcic, K. McKinny, D. M. Mei, A. Piepke, E. Yakushev, B. E. Berger, Y. D. Chan, M. P. Decowski, D. A. Dwyer, S. J. Freedman, Y. Fu, B. K. Fujikawa, K. M. Heeger, K. T. Lesko, K. B. Luk, H. Murayama, D. R. Nygren, C. E. Okada, A. W.P. Poon, H. M. Steiner, L. A. Winslow, G. A. Horton-Smith, R. D. McKeown, J. Ritter, B. Tipton, P. Vogel, C. E. Lane, T. Miletic, P. W. Gorham, G. Guillian, J. G.

Learned, J. Maricic, S. Matsuno, S. Pakvasa, S. Dazeley, S. Hatakeyama, M. Murakami, R. C. Svoboda, B. D. Dieterle, M. DiMauro, J. Detwiler, G. Gratta, K. Ishii, N. Tolich, Y. Uchida, M. Batygov, W. Bugg, H. Cohn, Y. Efremenko, Y. Kamyshkov, A. Kozlov, Y. Nakamura, L. de Braeckeleer, C. R. Gould, H. J. Karwowski, D. M. Markoff, J. A. Messimore, K. Nakamura, R. M. Rohm, W. Tornow, A. R. Young, and Y. F. Wang. First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance. *Physical Review Letters*, 90(2):6, 2003.

[54] Edwin Cartlidge. Gran Sasso lab to shut down controversial experiments. *Physics World*, 32(2):12, 2 2019.

[55] M. R. Anderson, S. Andringa, L. Anselmo, E. Arushanova, S. Asahi, M. Askins, D. J. Auty, A. R. Back, Z. Barnard, N. Barros, D. Bartlett, F. Barao, R. Bayes, E. W. Beier, A. Bialek, S. D. Biller, E. Blucher, R. Bonventre, M. Boulay, D. Braid, E. Caden, E. J. Callaghan, J. Caravaca, J. Carvalho, L. Cavalli, D. Chauhan, M. Chen, O. Chkvorets, K. J. Clark, B. Cleveland, D. Cookman, C. Connors, I. T. Coulter, M. A. Cox, D. Cressy, X. Dai, C. Darrach, B. Davis-Purcell, C. Deluce, M. M. Depatie, F. Descamps, J. Dittmer, F. Di Lodovico, N. Duhaime, F. Duncan, J. Dunger, A. D. Earle, D. Fabris, E. Falk, A. Farrugia, N. Fatemighomi, V. Fischer, E. Fletcher, R. Ford, K. Frankiewicz, N. Gagnon, A. Gaur, K. Gilje, O. I. González-Reina, D. Gooding, P. Gorel, K. Graham, C. Grant, J. Grove, S. Grullon, E. Guillian, S. Hall, A. L. Hallin, D. Hallman, S. Hans, J. Hartnell, P. Harvey, M. Hedayatipour, W. J. Heintzelman, J. Heise, R. L. Helmer, D. Horne, B. Hreljac, J. Hu, S. M.A. Hussain, T. Iida, A. S. Inácio, C. M. Jackson, N. A. Jelley, C. J. Jillings, C. Jones, P. G. Jones, K. Kamdin, T. Kaptanoglu, J. Kaspar, K. Keeter, C. Kefelian, P. Khaghani, L. Kippenbrock, J. R. Klein, R. Knapik, J. Kofron, L. L. Kormos, S. Korte, B. Krar, C. Kraus, C. B. Krauss, T. Kroupova, K. Labe, F. Lafleur, I. Lam, C. Lan, B. J. Land, R. Lane, S. Langrock, A. Latorre, I. Lawson, L. Lebanowski, G. M. Lefeuvre, E. J. Leming, A. Li, J. Lidgard, B. Liggins, Y. H. Lin, X. Liu, Y. Liu, V. Lozza, M. Luo, S. Maguire, A. Maio, K. Majumdar, S. Manecki, J. Maneira, R. D. Martin, E. Marzec, A. Mastbaum, J. Mauel, N. McCauley, A. B. McDonald, P. Mekarski, M. Meyer, C. Miller, C. Mills, M. Mlejnek, E. Mony, I. Morton-Blake, M. J. Mottram, S. Nae, M. Nirkko, L. J. Nolan, V. M. Novikov, H. M. O'Keeffe, E. O'Sullivan, G. D. Orebi Gann, M. J. Parnell, J. Paton, S. J.M. Peeters, T. Pershing, Z. Petriw, J. Petzoldt, L. Pickard, D. Pracsovics, G. Prior, J. C. Prouty, S. Quirk, A. Reichold, S. Riccetto, R. Richardson, M. Rigan, A. Robertson, J. Rose, R. Rosero, P. M. Rost, J. Rumleskie, M. A. Schumaker, M. H. Schwendener, D. Scislowski, J. Secrest, M. Seddighin, L. Segui, S. Seibert, I. Semenec, F. Shaker, T. Shantz, M. K. Sharma, T. M. Shokair, L. Sibley, J. R. Sinclair, K. Singh, P. Skensved, M. Smiley, T. Sonley, R. Stainforth, M. Strait, M. I. Stringer, R. Svoboda, A. Sörensen, B. Tam, J. Tatar, L. Tian, N. Tolich, J. Tseng, H. W.C. Tseung, E. Turner, R. Van Berg, J. G.C. Veinot, C. J. Virtue, B. Von Krosigk, E. Vázquez-Jáuregui, J. M.G. Walker, M. Walker, S. C. Walton, J. Wang, M. Ward, O. Wasalski, J. Waterfield, J. J. Weigand, R. F. White, J. R. Wilson, T. J. Winchester,

P. Woosaree, A. Wright, J. P. Yanez, M. Yeh, T. Zhang, Y. Zhang, T. Zhao, K. Zuber, and A. Zummo]deceased. Development, characterisation, and deployment of the SNO+ liquid scintillator. *Journal of Instrumentation*, 16(05):P05009, 5 2021.

[56] CEPSA Qimica. Material Safety Data Sheet: PETRELAB$^{®}$ 550-Q. Technical report, CEPSA, 2009.

[57] J. B. Cumming, S. Hans, and M. Yeh. Improving light yield measurements for low-yield scintillators. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 925:1--5, 5 2019.

[58] Benjamin Tam. *Enabling Neutrinoless Double Beta Decay in the SNO+ Experiment through the Deployment and Study of Liquid Scintillator*. PhD thesis, Queen's University `https://snoplus.phy.queensu.ca/theses/btam.pdf`, 2023.

[59] X. Zhou, Q. M. Zhang, Q. Liu, Z. Y. Zhang, Y. Y. Ding, L. Zhou, and J. Cao. Densities, isobaric thermal expansion coefficients and isothermal compressibilities of linear alkylbenzene. *Physica Scripta*, 90(5):055701, 4 2015.

[60] S. Hans, R. Rosero, L. Hu, O. Chkvorets, W. T. Chan, S. Guan, W. Beriguete, A. Wright, R. Ford, M. C. Chen, S. Biller, and M. Yeh. Purification of telluric acid for SNO+ neutrinoless double-beta decay search. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 795:132--139, 9 2015.

[61] Kamal Niaz, Fazlullah Khan, and Muhammad Ajmal Shah. Analysis of quinonoids. *Recent Advances in Natural Products Analysis*, pages 749--766, 1 2020.

[62] S. N. Axani, K. Frankiewicz, and J. M. Conrad. The CosmicWatch Desktop Muon Detector: a self-contained, pocket sized particle detector. *Journal of Instrumentation*, 13(03):P03019, 3 2018.

[63] S. Seiber et al. RAT. `https://rat.readthedocs.io/en/latest/index.html`. Accessed April 2025.

[64] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J.J. Gómez Cadenas, I. González, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F.W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata,

M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampén, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O'Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M.G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. Di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J.P. Wellisch, T. Wenaus, D.C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschiesche. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250--303, 2003.

[65] G. Horton-Smith. GLG4sim Webpage. `https://www.phys.ksu.edu/personal/gahs/GLG4sim/`. Accessed April 2025.

[66] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389:81--86, 1997. See also "ROOT" [software], Release v6.32/08, Nov 14, 2024.

[67] Ian T. Coulter. *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation.* Doctoral dissertation, University of Oxford `https://snoplus.phy.queensu.ca/theses/icoulter.pdf`, 2013.

[68] Matthew Mottram. Updated functional form energy reconstruction: EnergyRThetaFunctional. Internal Document: SNO+ DocDB 3488, 2015.

[69] Jack Dunger. *Topological and Time Based Event Classification for Neutrinoless Double Beta Decay in Liquid Scintillator.* Doctoral dissertation, University of Oxford `https://snoplus.phy.queensu.ca/theses/jdunger.pdf`, 2018.

[70] Aobo Li Teal Pershing Christopher Grant Diana Gooding, Morgan Askins. Data Cleaning Documentation. Internal Document: SNO+ DocDB 6593, 2020.

[71] John Walker. *Study of Invisible Mode Nucleon Decay in the SNO+ Detector.* PhD thesis, University of Liverpool `https://snoplus.phy.queensu.ca/theses/jwalker.pdf`, 2016.

[72] Lorna Nolan. *Cosmogenic Muon Induced Background Tagging in the SNO+ Detector.* PhD thesis, Queen Mary University of London `https://snoplus.phy.queensu.ca/theses/lnolan.pdf`, 2023.

[73] Josephine Paton. *Directional Reconstruction in Liquid Scintillator Neutrino Detectors.* Doctoral dissertation, University of Oxford `https://snoplus.phy.queensu.ca/theses/jpaton.pdf`, 2023.

[74] John Bahcall. Solar neutrino software and data. `https://www.sns.ias.edu/~jnb/SNdata/sndata.html`. Accessed: 2025-04-18.

[75] G. Alimonti, G. Angloher, C. Arpesella, M. Balata, G. Bellini, J. Benziger, S. Bonetti, L. Cadonati, F. P. Calaprice, G. Cecchet, M. Chen, N. Darnton, A. De Bari, M. Deutsch, F. Elisei, F. Von Feilitzsch, C. Galbiati, F. Gatti, M. G. Giammarchi, D. Giugni, T. Goldbrunner, A. Golubchikov, A. Goretti, T. Hagner, F. X. Hartmann, R. Von Hentig, G. Heusser, A. Ianni, J. Jochum, M. Johnson, M. Laubenstein, P. Lombardi, S. Magni, S. Malvezzi, I. Manno, G. Manuzio, F. Masetti, U. Mazzucato, E. Meroni, M. Neff, A. Nostro, L. Oberauer, A. Perotti, A. Preda, R. S. Raghavan, G. Ranucci, E. Resconi, M. Ruscitti, R. Scardaoni, S. Schönert, O. Smirnov, R. Tartaglia, G. Testera, P. Ullucci, R. B. Vogelaar, S. Vitale, and O. Zaimidoroga. Measurement of the 14C abundance in a low-background liquid scintillator. *Physics Letters B*, 422(1-4):349--358, 3 1998.

[76] H. Ejiri and S. R. Elliott. Charged current neutrino cross section for solar neutrinos, and background to $\beta\beta$ (0 ν) experiments. *Physical Review C - Nuclear Physics*, 89(5):055501, 5 2014.

[77] V. Lozza and J. Petzoldt. Cosmogenic activation of a natural tellurium target. *Astroparticle Physics*, 61:62--71, 2 2015.

[78] Mark Chen, Valentina Lozza, Helen O'Keefe. Expected radioactive backgrounds in SNO+. Internal Document: SNO+ DocDB 507-v43, 2021.

[79] Tanner Kaptanoglu. Double Beta Official Plots. Using Te+DDA Optics. Internal Document: SNO+ DocDB 5681-v3, 2019.

[80] A. Li, Z. Fu, C. Grant, H. Ozaki, I. Shimizu, H. Song, A. Takeuchi, and L. A. Winslow. KamNet: An integrated spatiotemporal deep neural network for rare event searches in KamLAND-Zen. *Physical Review C*, 107(1):014323, 1 2023.

[81] Michael A. Nielsen. Neural Networks and Deep Learning. `http://neuralnetworksanddeeplearning.com`, 2015.

[82] Charu C. Aggarwal. Neural Networks and Deep Learning: A Textbook. *Neural Networks and Deep Learning: A Textbook*, pages 1--529, 1 2023.

[83] Taco S. Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical CNNs. *Proceedings of the International Conference on Learning Representations, 2018*, 2 2018.

[84] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning Long-Term Dependencies with Gradient Descent is Difficult. *IEEE Transactions on Neural Networks*, 5(2):157--166, 1994.

[85] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735--1780, 11 1997.

[86] Xingjian SHI, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun WOO. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. *Advances in Neural Information Processing Systems*, 28, 2015.

[87] Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 9 2014.

[88] S Abe, S Asami, M Eizuka, S Futagi, A Gando, Y Gando, T Gima, A Goto, T Hachiya, K Hata, S Hayashida, K Hosokawa, K Ichimura, S Ieki, H Ikeda, K Inoue, K Ishidoshiro, Y Kamei, N Kawada, Y Kishimoto, M Koga, M Kurasawa, N Maemura, T Mitsui, H Miyake, T Nakahata, K Nakamura, R Nakamura, H Ozaki, T Sakai, H Sambonsugi, I Shimizu, J Shirai, K Shiraishi, A Suzuki, Y Suzuki, A Takeuchi, K Tamae, K Ueshima, H Watanabe, Y Yoshida, S Obara, A K Ichikawa, D Chernyak, ğ A Kozlov, K Z Nakamura, S Yoshida, Y Takemoto, S Umehara, K Fushimi, K Kotera, Y Urano, B E Berger, B K Fujikawa, J G Learned, J Maricic, S N Axani, J Smolsky, Z Fu, L A Winslow, Y Efremenko, H J Karwowski, D M Markoff, W Tornow, S Dell'oro, T O'donnell, J A Detwiler, S Enomoto, M P Decowski, C Grant, A Li, and H Song. Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen. *Physical Review Letters*, 130(7):51801, 2023.

[89] Tereza Kroupová. Multisite Classifer Update 2p2 Scintillator. Internal Document: SNO+ DocDB 8448, 2024.

[90] Clara Dima. Instructions for Analysers. GitHub Repository `https://github.com/claradima/rat/tree/reconstruction_pmt_selection`, 2025.

[91] Ziping Ye. Analysis of BiPo backgrounds for the search of Te130 0vbb. Internal Document: SNO+ DocDB 7981, 2023.

[92] Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1-3):503--528, 8 1989.

[93] Rafael Hunt-Stokes. Multisite Conference Plots. Internal Document: SNO+ DocDB 8278, 2024.

[94] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2):1554, 2011.

[95] Sam Naugle. Collaboration Meeting BB Update - March 2024. Internal Document: SNO+ DocDB 8185, 2024.

[96] D. Q. Adams, C. Alduino, K. Alfonso, F. T. Avignone, O. Azzolini, G. Bari, F. Bellini, G. Benato, M. Biassoni, A. Branca, C. Brofferio, C. Bucci, J. Camilleri, A. Caminata, A. Campani, L. Canonica, X. G. Cao, S. Capelli, L. Cappelli, L. Cardani, P. Carniti, N. Casali, D. Chiesa, M. Clemenza, S. Copello, C. Cosmelli, O. Cremonesi, R. J. Creswick, A. D'Addabbo, I. Dafinei, C. J. Davis, S. Dell'Oro, S. Di Domizio, V. Dompè, D. Q. Fang, G. Fantini, M. Faverzani, E. Ferri, F. Ferroni, E. Fiorini, M. A. Franceschi, S. J. Freedman, S. H. Fu, B. K. Fujikawa, A. Giachero, L. Gironi, A. Giuliani, P. Gorla, C. Gotti, T. D. Gutierrez, K. Han, K. M. Heeger, R. G. Huang, H. Z. Huang, J. Johnston, G. Keppel, Yu. G. Kolomensky, C. Ligi, L. Ma, Y. G. Ma, L. Marini, R. H. Maruyama, D. Mayer, Y. Mei, N. Moggi, S. Morganti, T. Napolitano, M. Nastasi, J. Nikkel, C. Nones, E. B. Norman, A. Nucciotti, I. Nutini, T. O'Donnell, J. L. Ouellet, S. Pagan, C. E. Pagliarone, L. Pagnanini, M. Pallavicini, L. Pattavina, M. Pavan, G. Pessina, V. Pettinacci, C. Pira, S. Pirro, S. Pozzi, E. Previtali, A. Puiu, C. Rosenfeld, C. Rusconi, M. Sakai, S. Sangiorgio, B. Schmidt, N. D. Scielzo, V. Sharma, V. Singh, M. Sisti, D. Speller, P. T. Surukuchi, L. Taffarello, F. Terranova, C. Tomei, K. J. Vetter, M. Vignati, S. L. Wagaarachchi, B. S. Wang, B. Welliver, J. Wilson, K. Wilson, L. A. Winslow, S. Zimmermann, and S. Zucchelli. Erratum: Measurement of the $2\nu\beta\beta$ Decay Half-Life of ^{130}Te with CUORE [Phys. Rev. Lett. 126, 171801 (2021)]. *Physical Review Letters*, 131:249902, Dec 2023.

Curriculum Vitae

Diana M. Gooding
dgooding@bu.edu

Education

2017–present **PhD Candidate, Boston University**
(Expected May 2024) Department of Physics

2017–2020 **Boston University**
M.A. in Physics

2014–2016 **Tufts University**
M.S. in Chemistry

2010–2014 **Transylvania University**
B.A. in Physics and Mathematics

PhD Thesis Research

2017–present **Boston University**

As a graduate student with Prof. Christopher Grant, I have worked on:

- Applying a deep learning framework to classify low-energy processes that occur in the SNO+ neutrino detector
- Implementing algorithms written in C++/ROOT to quantify radioactive contamination and infer changes to detector environment and data quality
- Performing one of the first log-likelihood fits of simulated backgrounds to clean scintillator data to determine the experiment's projected sensitivity to neutrinoless double-beta decay

Additional Recent Research

2015–2018 **Massachusetts Institute of Technology**

As a member of Prof. Lindley Winslow's research group, I applied my training as a chemist to set up a lab space and help develop a quantum-dot-based liquid scintillator for use in a novel neutrino detector.

In addition to acting as safety coordinator, I contributed by:

- Performing chemical synthesis of CdSe-core and lead-perovskite quantum dots, and dispersing them in liquid scintillator
- Testing their performance in a bench-top prototype detector
- Presenting results in an invited talk at the 2017 Division of Nuclear Science conference and coauthoring two peer-reviewed papers

Publications and Presentations

- *Light yield of Perovskite nanocrystal-doped liquid scintillator* E. Graham, **D. Gooding**, J. Gruszko, C. Grant, B. Naranjo and L. Winslow. *JINST* P11024 (2019)
- *Detecting Cherenkov light from 12 MeV electrons in linear alkylbenzene* J. Gruszko, B. Naranjo, B. Daniel, A. Elagin, **D. Gooding**, C. Grant, J. Ouellet and L. Winslow *JINST* P02005 (2019)
- Invited speaker at the 2017 APS Division of Nuclear Physics Meeting: *Quantum Dots in Liquid Scintillator.*
- *Influence of Environment on the Measurement of Rates of Charge Transport across* Ag^{TS}*/SAM//*Ga_2O_3*/ EGaIn Junctions.* Jabulani R. Barber, Hyo Jae Yoon, Carleen M. Bowers, Martin M. Thuo, Benjamin Breiten, **Diana M. Gooding**, and George M. Whitesides. Chem. Mater. (2014), 26, 13, 39383947
- *Electrets.* Kaufman, G.K. and **Gooding, D.M**. In Encyclopedia of Inorganic and Bioinorganic Chemistry, R.A. Scott (Ed.). (2019) DOI: 10.1002/9781119951438.eibc2240.pub2
- *Tribocharging and the Triboelectric Series.* **Gooding, D.M**. and Kaufman, G.K. In Encyclopedia of Inorganic and Bioinorganic Chemistry, R.A. Scott (Ed.). (2019) DOI: 10.1002/9781119951438.eibc2239.pub2