Calibration of the SNO+ Detector using Natural Radioactivity and First Demonstration of Multi-Site Discrimination for Solar Neutrinos

Rafael Hunt-Stokes
Saint Catherine's College

A thesis submitted for the degree of
Doctor of Philosophy
Trinity Term 2024

Abstract

A high-precision measurement of the ^{8}B solar neutrino flux in the 2.0 to 6.0 MeV transition region is a key objective in neutrino physics. This region, where the ν_e survival probabilities are highly sensitive to physics beyond the Standard Model, such as non-standard neutrino interactions or additional neutrino flavors, remains largely unexplored. However, challenges arise due to the high rates of internal ^{208}Tl radioactive backgrounds in liquid scintillator detectors. This thesis introduces multisite event discrimination as an effective method to enhance the precision of ^{8}B solar flux measurements in this background-dominated energy regime. A Bayesian optimization calibration method was developed using internal ^{214}Bi-^{214}Po coincidences to extract clean α and β event samples. These samples were used to build precise models of the scintillation response, leading to the creation of high-accuracy probability density functions (PDFs) for ^{8}B solar flux fitting. For the first time, multisite event discrimination was applied to liquid scintillator data, yielding a 30% improvement in measurement precision when combined with energy PDF information, as compared to fits performed in energy alone.

Sensitivity projections for the upcoming SNO+ bisMSB phase, featuring improved light yields and faster timing, show a 17% increase in precision over the current full-fill phase, with position resolution enhanced to 6 cm.

While this work marks a significant step towards achieving a high-precision measurement of the ^{8}B solar flux in the transition region, future analyses will integrate multisite, energy, and ^{212}Bi-^{208}Tl coincidence tagging into the likelihood functions, maximizing the information extracted to further refine measurement precision.

Acknowledgements

First of all, I'd like to thank my parents. From the bottom of my heart, I couldn't have wished for a more supportive family. You've always backed me in whatever direction I pursued, and been there to pick up the pieces on the occasions it's not worked out. I will always appreciate that. Dad, for giving me curiosity about the natural world and a specific knowledge of the indigenous butterflies of the United Kingdom. Mum, for being so frequently hilarious, all the sage advice over the years, and giving me a working knowledge of sorcery. Things have definitely taken an unexpected turn, but I'm excited to see what the future holds for everyone.

To the guys, the bros, you four that have been with me forever: Captain Caswell, Slim Tim, and MacroRat. It is impossible to put a price on the camaraderie we've shared over the years. However, Tom, if you consult splitwise, you will see you owe me approximately £14.62 for that curry. I will always fondly remember screaming CY*A B*Y* down the mic into the wee hours, and relentlessly rushing B, despite the protestations of The Captain. For all the hikes, NYE celebrations, and the sheer quantity of banter we've shared - I would've gone insane without you.

A particular shoutout to MacroRat: well, we dreamed about it, we thought about it and we achieved it - the founders of INSIP house. What an honour to share an abode, and so many absolutely massive dinners (naturally followed by a couple episodes of 'spys'). I miss you already. Your academic excellence has always sparked a healthy competitive spirit in me, and I think it's safe to say I wouldn't have tried half as hard if I hadn't had you to look up to.

John - thanks for being there for the past 8 (?!) years, through all the highs and lows. The image of you reading the FT in the club whilst dressed as a history teacher continues to live rent free in my mind. Where would DJ Shadowbiscuit be without Gregory Quasar?

I have been fortunate to work with so many good people in the SNO+ collaboration. I would like to express my gratitude to my supervisor, Steven Biller, for their guidance and support throughout my PhD journey, along with Armin Reichold and Jeff Tseng. My thanks to Tereza, who helped get me started when the world was in the grip of a pandemic and I was a lost first year. I'd also like to thank you for leaving me that bag of Huel: little did I know when I started that I would never eat solid food again. And those I've had the privilege of sharing the office with: Josie, Daniel, Cal, Gulliver, Jasmine, Po-Wei, Ben, Will and Ana Sofia. Good stuff.

Katharine - for brightening my Sudbury experience, and sharing the freezing mornings on the UK apartment balcony alongside a heady portion of Russian literature. The trips we took were epic.

To The Bishop and The Big Bad Beast: my sincere thanks for always being available to chat about anything, work or otherwise, specifically over a few pints. Do you remember when we went to see that massive apple and you guys (not me) ate that piglet's anus?

ASCII - for always having my back and believing in me, oftentimes when I didn't believe in myself. You taught me so much in the ways of science, and I owe you a hell of a lot. Your mentorship in the last couple of years has been utterly invaluable, as has your friendship.

Finally, to my friends at St. Catz. That pandemic was so strange, but overall I'm glad I was locked up with you lot: Ana, Peter, Grace and Roma.

Roma - for all the coffees and market lunches, the scifi novels recommended, and the Stewart Lee gigs attended.

Statement of Authorship

Chapter 1: This chapter provides the necessary theoretical foundation for relevant areas of neutrino physics, including neutrino oscillations, solar neutrinos, and neutrino mass. The information is compiled from research papers and textbooks. The second half of the chapter explains the basic components of a liquid scintillator detector, such as SNO+. This material is also drawn from a variety of textbooks, as referenced.

Chapter 2: This chapter provides an overview of the SNO+ detector, including the data and simulations used in the subsequent parts of this thesis. The information was obtained from published SNO+ detector papers.

Chapter 3: The third chapter discusses the use of internal Bi-Po coincidence decays for calibration and for measuring the ^{238}U and ^{232}Th concentrations. While the use of Bi-Po coincidences is an established technique, the application to full-fill SNO+ data and the subsequent ^{238}U/^{232}Th measurements are the original work of the author. The grid-search timing calibration using Bi-Po events is an existing method, but the specific calibration of the full-fill phase was led by the author. The tuning of light yield and α-quenching was carried out in collaboration with Serena Riccetto, who performed an independent analysis. The development of the Bayesian optimizer for timing tuning was solely the work of the author, building on established techniques from the machine learning field.

Chapter 4: The use of multisite event discrimination builds on the previous work of Tereza Kroupova, Steve Biller, and Jack Dunger, as referenced. The author developed the fitting framework independently and applied the classifier to data for the first time. All stages of the data analysis, Asimov studies, and PDF creation were led by the author, with valuable contributions from Ana

Sofia Inacio, who provided expertise in clean data selection cuts. The bisMSB phase predictions rely on the bisMSB-tuned emission time model, created using the techniques developed by the author in Chapter 3 and implemented by Po-Wei Huang. Other components of the bisMSB phase model were calibrated by Serena Riccetto, and the reconstruction algorithms were coordinated by Will Parker (position) and Tereza Kroupova (energy).

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Since their first proposal by Wolfgang Pauli in 1930 as a 'desperate remedy' to the apparent violation of conservation of energy and momentum in nuclear β-decay [1], the neutrino, or 'little neutral one', has been a fascinating and enigmatic particle. Dubbed the 'ghost particle', the neutrino was theorised to be a massless, electrically neutral spin-1/2 particle in the Standard Model, with an interaction cross-section so small it was feared undetectable. Pauli himself supposedly stated "I have done a terrible thing: proposed a particle that cannot be detected" [2]. It took twenty six years for the Cowen-Reines experiment to prove him both right and wrong, with the first confirmed observation of electron antineutrinos from the Savannah River nuclear power station [3].

Since then, neutrinos have provoked surprise and provided a window into exciting new physics, beyond the Standard Model (BSM). In particular, solar neutrinos, the focus of this thesis, have been an invaluable cosmological tool to probe the nature of neutrinos and take steps towards answering the most pressing questions remaining in fundamental physics. This introduction is split into three major parts: first, an overview of the current knowledge concerning solar neutrinos, including their discovery, the solar neutrino problem and the flavour oscillations that provided a solution. Second, a summary of the open questions in neutrino physics is given, including possible mechanisms for generating the neutrino mass, neutrinoless double beta decay ($0\nu\beta\beta$), and theoretical BSM neutrino interactions, which may reveal themselves through precise measurements of the ν_e survival probability. Lastly, large liquid scintillator detectors are presented as powerful tools to answer these most pressing questions in neutrino physics. A step by step description of the production, propagation and detection of scintillation light is detailed, providing a rationale for their use in rare event searches.

1.1 The Discovery of Solar Neutrinos, the Solar Neutrino Problem and its Solution

The sun, being powered by nuclear fusion, produces an extremely large flux of electron neutrinos. Stellar physicists, particularly Bahcall, were able to use constraints on the solar mass, age and isotopic abundances to make predictions of this flux [4]. Ray Davis, a radiochemist, set out to directly measure the ν_e flux via an ingenious detector design. Based at the Homestake Mine in South Dakota, Davis filled a 380 m^2 tank with perchloroethylene (a common dry-cleaning fluid) and placed it 1478 m underground [5]. The large tank ensured the detector would capture a measurable quantity of solar ν_e interactions, and the underground location provided the necessary shielding from cosmic muons.

These interactions were detected via the inverse β-decay reaction between chlorine atoms in the medium and solar ν_e:

$$\nu_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar}^{+} + e^{-} \tag{1.1}$$

By carefully counting the number of argon atoms produced, and recording the exposure time, Davis was able to precisely measure the solar neutrino flux at his detector, with the first results reported in 1968 [6]. However, there was a problem: the measured flux was consistently 1/3 of that predicted by Bahcall's theoretical calculations.

This prompted great excitement amongst theorists and experimentalists alike. As new experiments came online, similar deficits in the expected neutrino fluxes were reported, pointing to new physics 7–9 . The discrepancy became known as the 'solar neutrino problem' (SNP) [10].

Ultimately, the SNP was resolved by the SNO collaboration's 2002 publication, which demonstrated the first observation of ν_e to $\nu_{\mu,\tau}$ oscillations for solar neutrinos [11]. These results showed ν_e produced in the sun oscillate into $\nu_{\mu,\tau}$ as they traversed space, and thus, Ray Davis' experiment, being sensitive only to ν_e, recorded a deficit relative to the expectation.

The SNO experiment, a 6 m radius acrylic vessel filled with heavy water (D_2O), and instrumented with approximately 10 000 PMTs, was unique in that it was able to detect all flavours of neutrino [12]. This was accomplished via three observable reactions: charged current (CC) between ν_e and deuterons, neutral current (NC) between all flavours of neutrino and the deuteron, and elastic scattering (ES) from all flavours of neutrino with atomic electrons. These ES may occur via CC and NC for ν_e, but only NC for ν_μ and ν_τ.

The rates of each process are proportional to the flux of each respective neutrino flavour, $\phi(\nu_l)$:

$$\text{CC rate} \propto \phi(\nu_e) \tag{1.2}$$

$$\text{NC rate} \propto \phi(\nu_e) + \phi(\nu_\mu) + \phi(\nu_\tau) \tag{1.3}$$

$$\text{ES rate} \propto \phi(\nu_e) + 0.154\,[\phi(\nu_\mu) + \phi(\nu_\tau)] \tag{1.4}$$

By exploiting differences in the angular and energy distributions of the Cherenkov rings from the ES, NC and CC interactions, the rates of each individual process were obtained. The interaction cross sections of each neutrino flavour were used to convert these rates to flux measurements of the ν_e and $\nu_e + \nu_\mu + \nu_\tau$, allowing a comparison of the total solar neutrino flux from the sun to the ν_e flux.

The total solar neutrino flux from all flavours was in good agreement with the ν_e flux expected from the Standard Solar Model, and demonstrated that the solar neutrino flux at Earth must contain a large ν_μ and ν_τ component. Since ν_μ and ν_τ cannot be produced in any of the known solar fusion reactions in the Standard Solar Model, the SNO experiment provided clear evidence that ν_e neutrinos were changing flavour as they travelled. This transformation, or 'oscillation', had been postulated by Pontecorvo in 1957 in the context of neutrino-antineutrino oscillations [13], while the framework for flavor oscillations was formulated by Maki, Nakagawa, and Sakata in 1962 [14]. The SNO results provided conclusive evidence for the flavor oscillation solution to the solar neutrino problem. Neutrino oscillations remain the only accepted form of BSM physics to date: in order for neutrinos to oscillate and change state, they must have mass, in direct conflict with the massless neutrinos of the Standard Model.

1.1.1 Neutrino Oscillations

Neutrino oscillations provided a compelling solution to the solar neutrino problem. This section outlines the basic theoretical framework underpinning these oscillations, beginning with vacuum oscillations, before going on to consider the impact of matter on neutrino survival probabilities. These matter effects are needed to explain the solar neutrino measurements at different energy scales, as well as allowing the determination of the sign of the mass difference between two of the neutrino mass eigenstates.

Vacuum Oscillations

In the absence of matter, neutrinos propagate as eigenstates of the free particle Hamiltonian according to the time-dependent Schrodinger equation.

For neutrino oscillations to occur, two conditions must be met:

1. The weak eigenstates and the mass eigenstates must be distinct, and mixed according to a unitary transformation
2. The mass eigenstates have unique, and non-zero, mass eigenvalues of the free particle Hamiltonian

The second condition enfolds two key requirements for oscillations, namely neutrinos possessing a finite mass and these masses not being degenerate. The non-zero mass requirement is a consequence of special relativity: without mass, neutrinos would not experience duration, and therefore would have no time to change state. The fact the mass states cannot be degenerate arises from the need for phase differences between each mass eigenstate as the neutrino propagates.

For three neutrino flavours, the weak eigenstates are related to the mass eigenstates via the unitary transformation:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \hat{U} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.5}$$

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.6}$$

Where the unitary matrix, $\hat{U}$ is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [14].

When an electron neutrino is created in a CC weak interaction, it is produced in a definite flavour eigenstate, $|\nu_e\rangle$. The PMNS matrix is used to express this flavour eigenstate as a linear superposition of the mass states:

$$|\nu_e\rangle = U_{e1}|\nu_1\rangle + U_{e2}|\nu_2\rangle + U_{e3}|\nu_3\rangle \tag{1.7}$$

At a later space-time point, $(\vec{x}, t)$, the neutrino state has evolved according to:

$$|\Psi(\vec{x}, t)\rangle = U_{e1}e^{-i\phi_1}|\nu_1\rangle + U_{e2}e^{-i\phi_2}|\nu_2\rangle + U_{e3}e^{-i\phi_3}|\nu_3\rangle \tag{1.8}$$

Where $\phi_k = E_k t - \vec{p_k} \cdot \vec{x}$ is the phase of each mass eigenstate in the superposition, in natural units. The probability of transition from $|\nu_e\rangle$ to $|\nu_\mu\rangle$ is given by:

$$P(\nu_e \to \nu_\mu) = |\langle \nu_\mu | \Psi(\vec{x}, t)\rangle|^2 \tag{1.9}$$

To calculate this, each mass state in equation 1.8 is re-expressed in terms of the flavour states, taking advantage of the unitarity of the PMNS matrix to invert equation 1.6:

$$P(\nu_e \to \nu_\mu) = |U_{e1}U^*_{\mu 1}e^{-i\phi_1} + U_{e2}U^*_{\mu 2}e^{-i\phi_2} + U_{e3}U^*_{\mu 3}e^{-i\phi_3}|^2 \tag{1.10}$$

Where the orthogonality of flavour states has been used.

At this point, the requirement for each mass eigenstate to have a distinct mass can be seen. Given the unitarity of the PMNS:

$$\hat{U}\hat{U}^\dagger = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.11}$$

This yields 9 constraints on the products of the PMNS matrix elements. Taking the first row and second column of equation 1.11:

$$U_{e1}U^*_{\mu 1} + U_{e2}U^*_{\mu 2} + U_{e3}U^*_{\mu 3} = 0 \tag{1.12}$$

Given the above, if $m(\nu_1) = m(\nu_2) = m(\nu_3)$, then $\phi_1 = \phi_2 = \phi_3$ and the transition probability is zero, which is at odds with experiment. Thus, it is the phase differences, arising from non-degenerate masses, which allow vacuum oscillations to occur. Equation 1.10 may be expanded into a form that more clearly showcases the oscillation physics:

$$P(\nu_e \to \nu_\mu) = -4\sum_{i>j}(U_{ei}U_{\mu_i}U_{ej}U_{\mu j})sin^2\left(\frac{\Delta m^2_{ij}L}{4E}\right) - 2\sum_{i<j}(U_{ei}U^*_{\mu i}U^*_{\alpha j}U_{\beta j})sin\left(\frac{\Delta m^2_{ij}L}{2E}\right) \tag{1.13}$$

Where $\Delta m^2_{ij} = (m^2_i - m^2_j)$ is the mass splitting between the mass eigenstates, L is the propagation length and E is the neutrino energy.

This equation shows the probability of observing an oscillation is dependent on the length scale, L, the neutrino energy, E, the mass splitting between each of the neutrino mass eigenstates, Δm^2_{ij}, and the PMNS matrix elements. The PMNS elements determine the amount of mixing between each mass and flavour eigenstate, and are fundamental components of neutrino oscillations. Accurately determining the values of each element in the PMNS matrix is the primary goal of many historical and next-generation neutrino experiments. Canonically, the PMNS matrix is parameterised in terms of four variables: three mixing angles, θ_{ij}, and 1 complex, possibly CP-violating, phase, δ.

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23}c_{23} & \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.14}$$

Where $c_{ij} = cos(\theta_{ij})$ and $s_{ij} = sin(\theta_{ij})$. It is also possible to include two additional 'Majorana phases' in the PMNS definition, which do not alter the behaviour of neutrino oscillations. However, they have important consequences for neutrino mass generation mechanisms, with more details given in section 1.2.1.

Over time, neutrino oscillation experiments have measured the parameters of the PMNS matrix,

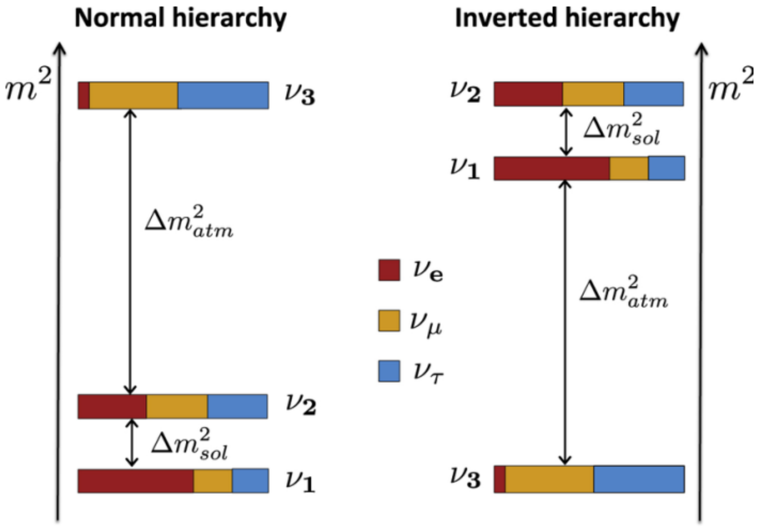

Figure 1.1: The two possible orderings of the neutrino masses, given only the sign of Δm^2_{21} is known to be positive. Taken from [16].

each with varying precision [15]. By placing detectors at different $\frac{L}{E}$ scales from neutrino sources, it is possible to focus measurements on particular PMNS parameters. Depending on the scale, different terms in the transition probability equation dominate or become negligible.

One important result from solar neutrino measurements was determining the sign of Δm^2_{21}. Reactor oscillation experiments alone are not sensitive to the sign, only the magnitude of the mass splittings, due to symmetry in the oscillation probability equation. However, solar neutrino measurements are sensitive to the sign, due to the effect of dense stellar matter on oscillations. From these experiments, it was determined $\Delta m^2_{21} > 0$, for reasons explained below. The sign of Δm^2_{32} remains unknown. Therefore, there are two distinct possibilities for the neutrino mass orderings, shown in figure 1.1:

- Normal Ordering: $m_1 < m_2 < m_3$
- Inverted Ordering: $m_3 < m_1 < m_2$

Oscillations in Matter

While neutrino oscillations in principle explain the solar neutrino problem, subsequent experiments obtained solar ν_e flux measurements that were highly dependent on the energy threshold, as shown

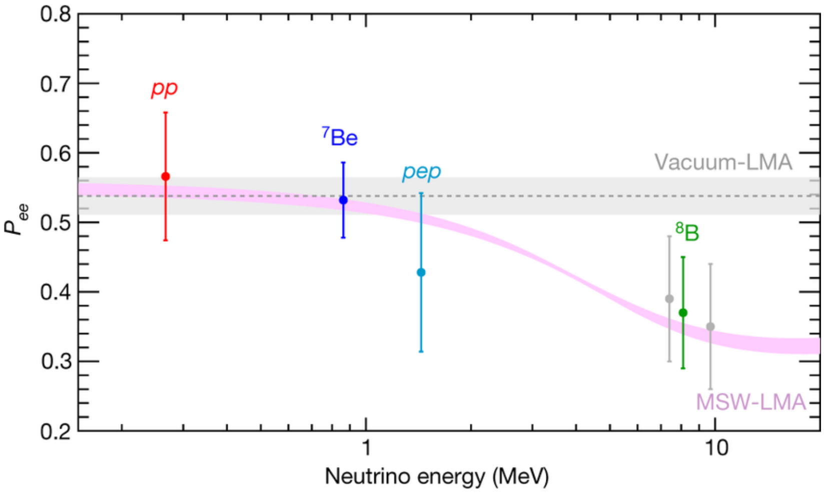

Figure 1.2: Solar ν_e survival probability as a function of neutrino energy, as measured by the Borexino collaboration [17]. Shaded regions represent the theoretical uncertainties for vacuum oscillations (grey) and MSW-LMA (pink).

in figure 1.2. The oscillation probability for ν_e appeared to have an additional energy dependence not accounted for by the vacuum oscillation theory alone. This is especially apparent for higher energy ^{8}B neutrinos ($E \geq 5$ MeV), where the survival probabilities significantly deviated from the vacuum oscillation expectation [17].

The energy dependence of the survival probabilities is explained by considering the environment in the stellar core, where the ν_e are born. Instead of vacuum, the ν_e propagate in an environment rich in electrons. These stellar electrons provide two interaction pathways for ν_e, that of the CC and NC. However, ν_μ and ν_τ are kinematically forbidden from interacting via CC. Thus, there exists an effective potential, V_{eff}, proportional to the electron density in the Sun, and experienced solely by the ν_e flavour neutrinos due to their additional CC interactions. The interaction potential has the form:

$$V_{eff} = V_e - V_\mu = 2\sqrt{2} G_F N_e E \tag{1.15}$$

Where G_F is the Fermi constant, N_e is the electron density of the medium and E is the neutrino energy.

The presence of this effective potential changes the form of the Hamiltonian. The vacuum mass states are no longer eigenstates of this new Hamiltonian, and the neutrinos propagate instead with new 'effective' masses. The additional potential experienced by ν_e flavour neutrinos alters the oscillation probabilities, and is known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect [18].

Due to the dependence of the interaction potential on neutrino energy, the degree to which the MSW effect changes ν_e to ν_μ oscillations is also energy dependent. This accounts for the different flux measurements as a function of energy scale measured by solar neutrino experiments. At lower energies, the interaction potential is negligible, and the neutrino oscillation probabilities recover the vacuum results. At higher energies, the MSW effect dominates, and the ν_e to ν_μ oscillation probability is enhanced.

The MSW effect allows the determination of the sign of Δm^2_{21}. In a two flavour approximation, the vacuum oscillation parameters may be related to the effective mixing angle, θ_m:

$$sin(2\theta_m) = \frac{sin(2\theta)}{\sqrt{\left(\frac{V_{eff}}{\Delta m^2} - cos(2\theta)\right)^2 + sin^2(2\theta)}} \tag{1.16}$$

Taking $\Delta m^2 \rightarrow -\Delta m^2$, the value of the effective mixing angle changes. In fact, with negative Δm^2 values, the MSW effect would suppress $\nu_e \rightarrow \nu_\mu$ oscillations as energy increases, which is entirely at odds with the measured results.

1.1.2 Solar Neutrinos

Given both their historical impact on the development of neutrino oscillation theory, and that this thesis is primarily concerned with solar neutrinos, it bears thinking briefly about their production mechanisms. Solar neutrinos are produced in the Sun via two main fusion processes: the proton-proton (pp) chain ($\sim 99\%$) and the carbon-nitrogen-oxygen (CNO) cycle ($\sim 1\%$). These processes are summarised in figure 1.3.

Due to the different binding energies and number of fusion products produced, each distinct neutrino-emitting process emits electron neutrinos according to different energy spectra. Focussing on the pp-chain, there are five reactions producing neutrinos, with the overall chain fusing hydrogen to produce helium:

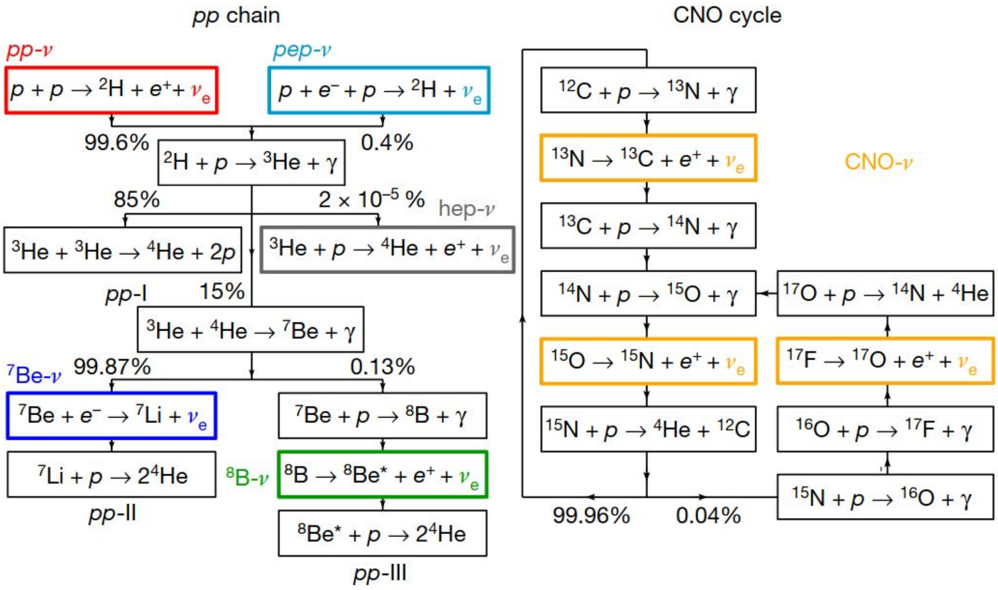

Figure 1.3: The fusion processes comprising the proton-proton (pp) chain and CNO cycles, which produce the solar neutrino flux measured on Earth. Taken from [17].

1. $p + p \to {}^2\text{H} + e^+ + \nu_e$: pp neutrinos
2. $2p + e^- \to {}^2\text{H} + \nu_e$: pep neutrinos
3. ${}^3\text{H} + p \to {}^4\text{He} + e^+ + \nu_e$: hep neutrinos
4. ${}^7\text{Be} + e^- \to {}^7\text{Li} + \nu_e$: ^{7}Be neutrinos
5. ${}^8\text{B} \to 2\ {}^4\text{He} + e^+ + \nu_e$: ^{8}B neutrinos

With the relative branching ratios given in figure 1.3. All but the hep neutrinos have been identified in nature [19–21], particularly by the Borexino collaboration, suggesting reaction (3) in the pp-chain has an extremely small branching ratio.

As shown in figure 1.4, decays (1), (3) and (5) are three body decays, and produce a continuous energy spectra. However, (2) and (4) are two body decays, producing mono-energetic neutrinos.

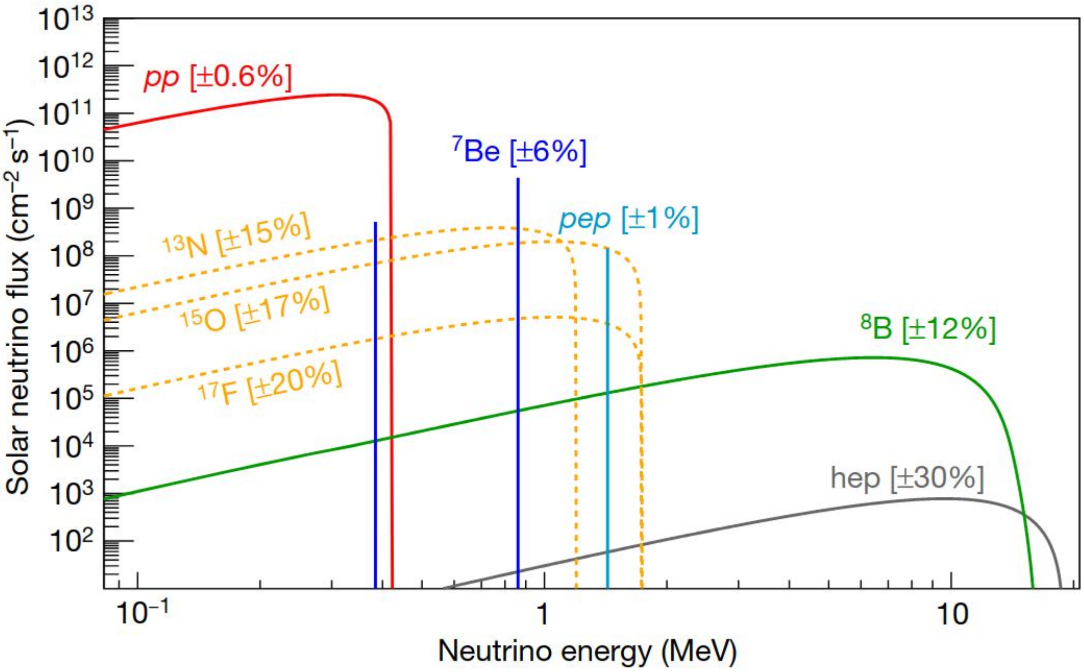

Figure 1.4: Solar neutrino energy distributions for the pp-chain and CNO cycle. Taken from [17].

The CNO cycle is the second process that produces solar neutrinos, however it is very subdominant to the pp-chain[1]. Despite this, Borexino was able to directly observe CNO neutrinos with high statistical significance, paving the way for further measurements of the CNO cycle [22]. Since the CNO neutrino flux is directly proportional to the isotopic abundance of the fusion materials, they offer a direct probe of the composition of the sun.

As shown in figure 1.3, two linked reaction pathways produce CNO neutrinos. These depend on the β-decays of nitrogen, oxygen and fluorine atoms, produced within each cycle. The neutrino emitting reactions are:

- Cycle I:

    1. $^{13}\mathrm{N} \rightarrow {}^{13}\mathrm{C} + e^{+} + \nu_e$
    2. $^{15}\mathrm{O} \rightarrow {}^{15}\mathrm{N} + e^{+} + \nu_e$

---

[1] Although the CNO cycle is expected to contribute the majority of energy in stars greater than 1.3 times the solar mass

- Cycle II:

  1. $^{17}\text{F} \rightarrow {}^{17}\text{O} + e^{+} + \nu_e$

1.2 Open Questions in Neutrino Physics

As described above, solar neutrino oscillations revealed the first chink in the SM's armour, and catalysed an era of progressively larger and more precise detectors. Despite this, there remains many open questions in both stellar and particle physics. These include: what is the hierarchy of the neutrino mass states? By what mechanism do neutrinos obtain their mass? Are neutrinos Majorana particles? Are there more than 3 flavours of neutrino? Do neutrinos undergo non-standard, BSM interactions? And in astrophysics, what is the metallicity of the sun? What follows is a brief overview of these questions, and the experimental methods that may yield answers.

1.2.1 The Origins of Neutrino Mass

While oscillations show neutrinos must be massive particles, they provide no information on the absolute neutrino mass, and this remains an open question in neutrino physics. The most stringent constraint comes from the KATRIN experiment, which has reported an absolute neutrino mass of $m_\nu < 0.45$ eV from studying the end-point of tritium nuclear β decays [23].

The apparent mass scale of neutrinos is significantly smaller than those of the other fermions: around a million times lighter than the electron. This difference in scale opens up further questions not just about the absolute value of the neutrino mass, but the very mechanism by which they obtain their mass, and whether it is different from that of the other fermions.

Massive particles in the SM obtain their masses via coupling to the Higgs field, with Dirac mass terms of the form:

$$\mathcal{L}_D = -m_D\left(\bar{\nu_R}\nu_L + \bar{\nu_L}\nu_R\right) \tag{1.17}$$

Where $\nu_{L,R}$ indicate left or right-handed neutrino fields and m_D is the Dirac mass term. Since neutrinos are always observed with left-handed helicity, the existence of right-handed neutrinos implies they must be sterile, interacting only through gravitational forces. This constraint allows neutrinos to obtain their mass via Higgs coupling, whilst remaining consistent with the V-A structure of the weak interaction.

Since neutrinos are neutral particles, the particle states might be equivalent to their antiparticles, being so-called ‘Majorana particles’ [24]. The Majorana condition is:

$$\psi^C = \hat{C}\hat{P}\psi = i\gamma^2\gamma^0\psi^* \tag{1.18}$$

This opens up alternative mechanisms for neutrinos to obtain their mass, without the need for invoking sterile right-handed particles. One such mechanism is the ‘see-saw’ mechanism [25]. A Majorana mass term can be added to the Dirac mass lagrangian in equation 1.17, without breaking the gauge invariance of the Standard Model. The most general mass term in the lagrangian combines both the Dirac and Majorana mass terms:

$$\mathcal{L}_{DM} = -\frac{1}{2}\begin{pmatrix} \bar{\nu_L} & \bar{\nu_R^C} \end{pmatrix}\begin{pmatrix} 0 & m_D \\ m_D & M \end{pmatrix}\begin{pmatrix} \nu_L^C \\ \nu_R \end{pmatrix} + h.c. \tag{1.19}$$

Where m_D is the Dirac mass, M is the Majorana mass and *h.c.* stands for the Hermitian conjugate.

The physical neutrino states correspond to the eigenvalues of the mass-matrix in equation 1.19:

$$m_{\pm} = \frac{M \pm \sqrt{M^2 + \frac{4m_D^2}{M^2}}}{2} \tag{1.20}$$

Assuming $M >> m_D$, there are two possible states, light (-) and heavy (+):

$$m_+ \sim M \tag{1.21}$$

$$m_- \sim \frac{m_D^2}{M} \tag{1.22}$$

By making the Majorana mass term sufficiently large, the see-saw mechanism hypothesises that the Dirac mass is on the order of the other fermions, but the observed light neutrino state is orders of magnitude smaller, due to the large Majorana mass.

While the see-saw mechanism provides a compelling explanation for the different scales of the neutrino absolute mass relative to the other fermions, it is only a theory, and contingent on neutrinos being Majorana particles. The Majorana nature of neutrinos, or otherwise, is another open area of investigation, with one of the most compelling tests for the Majorana nature of neutrinos being the search for neutrinoless double beta decay ($0\nu\beta\beta$).

1.2.2 Neutrinoless Double Beta Decay

Heavy, unstable nuclei may β-decay to lighter, more stable nuclei. In this case, the atomic number of the nucleus changes by +1.

$$(Z, A) \rightarrow (Z+1, A) + e^{-} + \bar{\nu_e} \tag{1.23}$$

Where Z is the number of protons and A the total nucleons in the nucleus.

It is energetically favourable for like nucleons to pair, leading to nuclei with even numbers of both protons and neutrons possessing lower binding energies than nuclei with odd nucleon counts. Because of this, isotopes exist where a single β^{-} decays would result in daughter nuclei acquiring *greater* mass than the parents, due to the increased binding energy going from an even to odd nucleus. In this case, it may be possible for the nucleus to instead decay via a second-order weak-interaction, the simultaneous double beta decay [26]:

$$(Z, A) \rightarrow (Z+2, A) + 2e^{-} + 2\bar{\nu_e} \tag{1.24}$$

Alternatively, it may be possible for double-beta allowed isotopes to decay similarly to equation 1.24 , but without the emission of neutrinos. This lepton-number violating process is neutrino*less* double beta decay, and to date has never been observed. It is possible only if neutrinos are Majorana particles [28], whereby neutrinos are there own anti-particles and may annihilate. Thus, the observation of $0\nu\beta\beta$ would be conclusive evidence for their Majorana nature, and support the see-saw mechanism's explanation for the absolute scale of the neutrino mass. Figure 1.5 shows the Feynman diagrams of $2\nu\beta\beta$ and $0\nu\beta\beta$.

The canonical mechanism behind $0\nu\beta\beta$ decay is the exchange of a light Majorana neutrino within the nucleus [29]. Two virtual W bosons are created, one of which produces an electron neutrino and electron. If neutrinos are Majorana particles, $\nu = \bar{\nu}$, and the neutrino may be absorbed by the

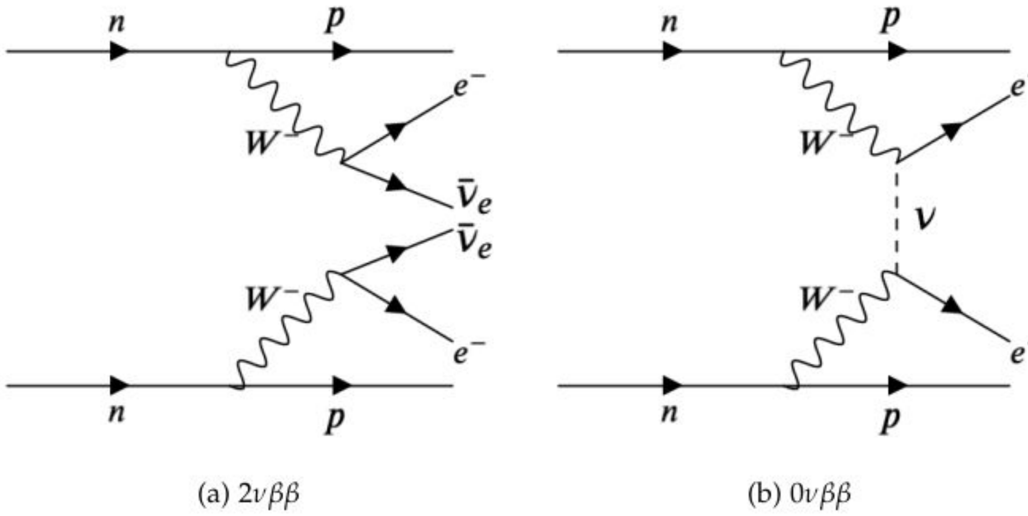

Figure 1.5: Feynman diagrams for a) SM double beta decay and b) BSM neutrinoless double beta decay, which is possible only if neutrinos are Majorana particles. Taken from [27].

second W, producing the second electron in $0\nu\beta\beta$ decay.

The primary goal of the SNO+ experiment is this direct observation of $0\nu\beta\beta$ from ^{130}Te nuclei loaded into liquid scintillator [30]. Since no neutrinos are emitted, all the decay energy is observable. Thus, the observable signature of $0\nu\beta\beta$ is a sharp peak at the end-point of the $2\nu\beta\beta$ energy spectrum, as shown in figure 1.6.

In addition to determining whether neutrinos are Majorana particles, the observed rate provides indirect information on the neutrino mass hierarchy, as shown in figure 1.7. The $0\nu\beta\beta$ rate is proportional to:

$$\Gamma \propto |M_{\beta\beta}|^2 \tag{1.25}$$

Where $|M_{\beta\beta}|$ is the effective Majorana mass, which is related to the PMNS matrix elements and the neutrino masses via:

$$M_{\beta\beta} = \sum_i U_{ei}^2 m_{\nu_i} \tag{1.26}$$

Thus, the effective Majorana mass is directly related to the neutrino mass eigenstates and PMNS matrix elements. Figure 1.7 shows the allowed values of $m_{\beta\beta}$ as a function of the lightest neutrino, $m_{lightest}$ (ν_1 or ν_3). By measuring the $0\nu\beta\beta$ rate and converting to the effective Majorana mass,

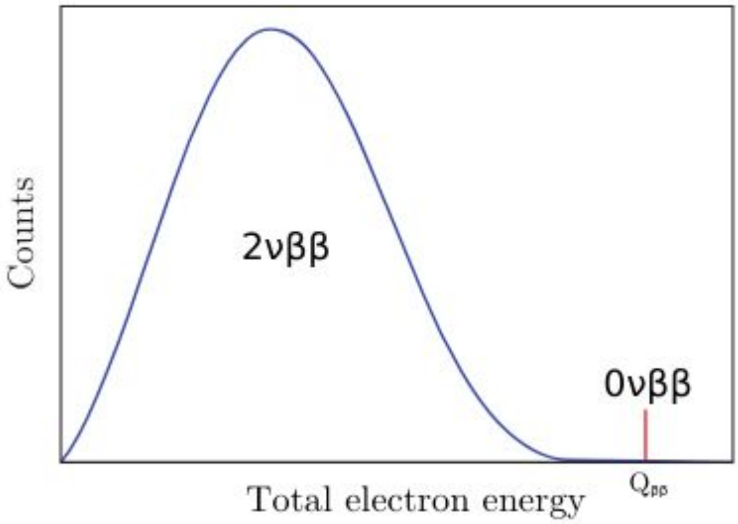

Figure 1.6: Schematic representation of the observable signature of $0\nu\beta\beta$ decay, which appears as a discrete peak at the Q-value of the $2\nu\beta\beta$ decay. In reality, the delta-peak will be smeared by detector resolution. Figure taken from [31].

horizontal limits can be placed on figure 1.7. In principle, with a sensitive enough detector, it is possible to use $0\nu\beta\beta$ to probe the neutrino mass hierarchy.

1.2.3 Non-Standard Neutrino Interactions

Non-standard neutrino interactions (NSIs) are theoretical models that allow neutrinos to interact in additional ways, beyond the SM [34]. In particular, the neutral current weak interaction lagrangian may be re-written to include NSIs between neutrinos and other SM fermions, as in equation 1.27:

$$\mathcal{L}_{NSI} = -\sum_{\alpha,\beta} 2\sqrt{2} G_F \epsilon_{\alpha\beta}^{ff'C} (\bar{\nu_\alpha} \gamma^\mu P_L \nu_\beta)(\bar{f} \gamma_\mu P_C f') \tag{1.27}$$

Where α and β denote the neutrino flavour, f and f' denote quarks of the same charge, but not necessarily the same flavour, C is the chirality and ϵ is a dimensionless coupling, which parametrises the strength of the NSI.

By providing additional, flavour dependent, interaction pathways, NSIs may alter the effective potential experienced by solar neutrinos as they propagate through the Sun. In the case of flavour

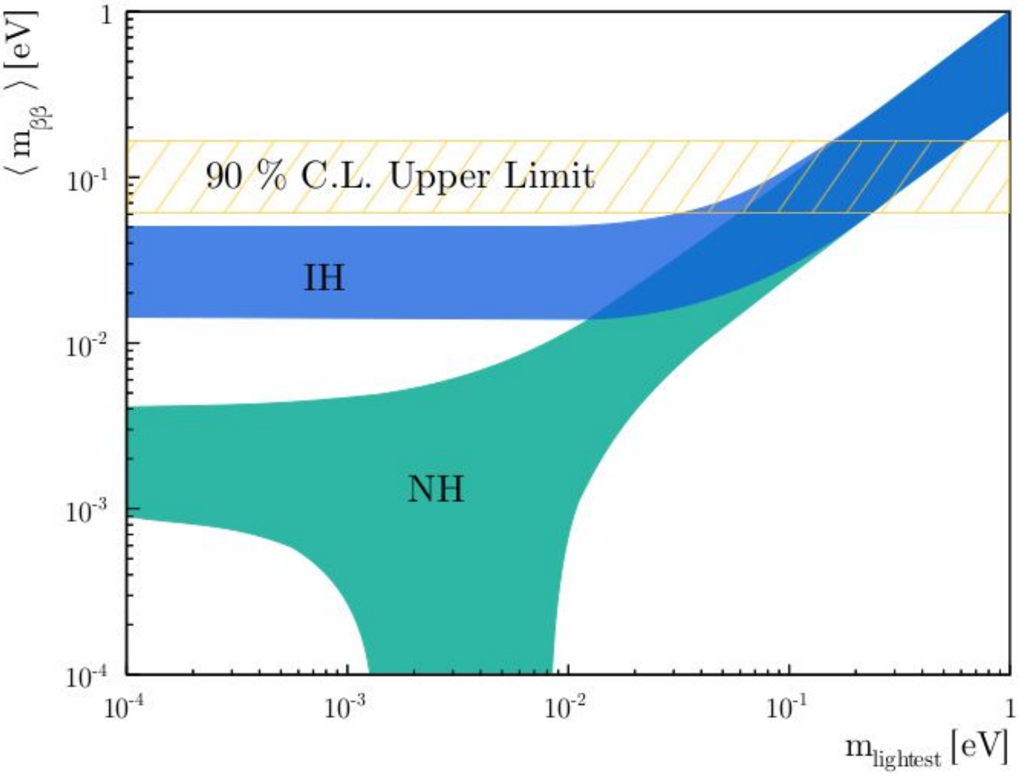

Figure 1.7: The effective $M_{\beta\beta}$ as a function of the lightest neutrino mass eigenstate for the normal (NH) and inverted hierarchy (IH). The upper limit is excluded by latest limits obtained from [32] and the figure is taken from [33].

conserving NSIs, the SM weak coupling constants, $g_{L,R}$, are shifted:

$$g_R \rightarrow g'_R = g_R + \epsilon^R_\alpha \tag{1.28}$$

$$g_L \rightarrow g'_L = g_L + \epsilon^L_\alpha \tag{1.29}$$

This shift in the coupling constants leads to a change in the effective potential experienced by ν_e neutrinos as they propagate through the Sun.

$$V_{eff} \rightarrow V^{NSI}_{eff} = (1 - \epsilon')V_{eff} \tag{1.30}$$

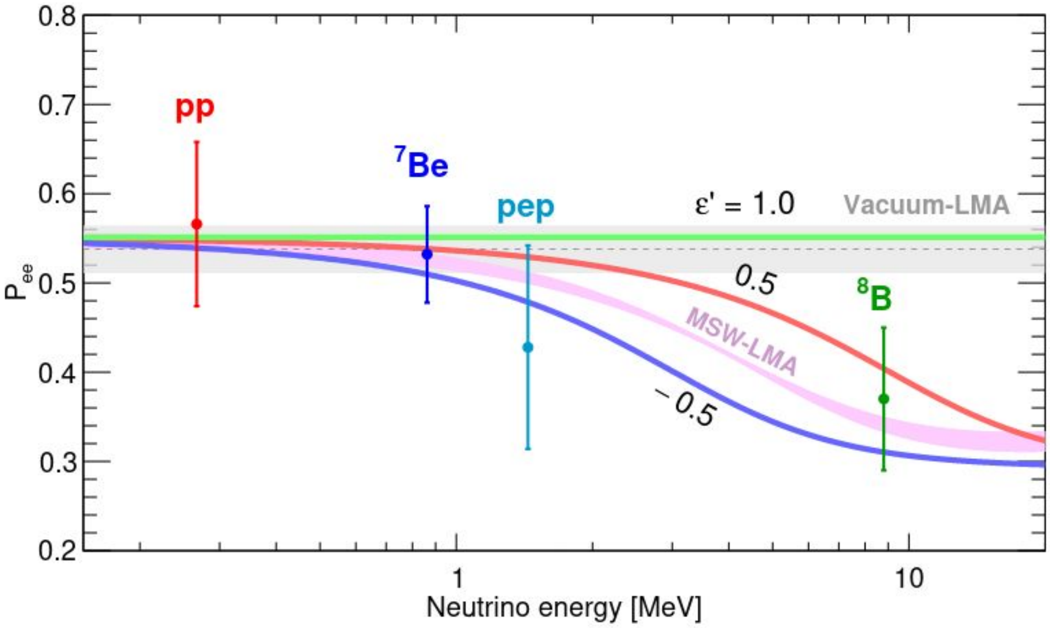

Figure 1.8: Solar ν_e survival probabilities as a function of energy, taken from [34]. Measurements are included from [17], and predicted survival curves for $\epsilon' = 1.0, 0.5, -0.5$ are included. The current measurements are not sensitive enough to determine if NSIs are present in solar oscillations. It is therefore highly desirable to perform a high precision measurement of the solar ν_e flux in the transition region, between 2 to 5 MeV.

Where ϵ' is defined in equation 1.31.

$$\epsilon' = \epsilon_\tau^V sin^2(\theta_{23}) - \epsilon_e^V \tag{1.31}$$

$$\epsilon_\alpha^V = \epsilon_\alpha^L + \epsilon_\alpha^R \tag{1.32}$$

Due to the change in effective potential, the neutrino oscillation parameters are further altered from those discussed in section 1.1.1. This gives a sensitivity to the strength of NSI couplings in the measured ν_e survival probabilities, as shown in figure 1.8 for three different values of ϵ'. As $\epsilon' \rightarrow 1$, the survival probability returns to the vacuum expectation and matter effects are removed. This is understood by referring to equation 1.30, which shows the effective potential becomes weaker with ϵ'. Conversely, as ϵ' grows smaller, the survival probability is increasingly suppressed.

Between 2 to 5 MeV, there is expected to be a significant decrease in the ν_e survival probability, due to the Mikheyev-Smirnov-Wolfenstein (MSW) effect. This transition region, between vacuum dominated ($\leq$ 2 MeV) and matter-dominated ($\geq$ 5 MeV) oscillations, is particularly sensitive to NSIs. By measuring the shape of the ^{8}B neutrino spectrum in the transition region and comparing it to theoretical predictions including new physics, we may find neutrinos hold yet more surprises.

1.3 Large Liquid Scintillator Detectors

Large liquid scintillator detectors are uniquely suited to answering many of the open questions described above. This section describes the key attributes of a large liquid scintillator experiment, such as the SNO+ detector, alongside the core physics behind the generation, propagation and detection of scintillation light.

Liquid scintillator experiments may be scaled to very large target masses, with the SNO+ experiment featuring a 780t target, and the next generation JUNO detector set to deliver 20kt [35]. Coupled with high light yields and optical clarity, liquid scintillator detectors offer high statistic measurements with significantly lower energy threshold and reconstruction resolution than water Cherenkov detectors. For example, the SNO+ detector is currently operating with a $\sim$ 0.3 MeV threshold and percent-scale energy resolution. Thus, these experiments are well suited for rare event searches, such as solar neutrino measurements and $0\nu\beta\beta$ decay searches.

A high precision measurement of the ν_e survival probability in the transition region is of major interest to the community, for the reasons outlined above. Large liquid scintillators, if internal backgrounds are sufficiently suppressed, provide the necessary fiducial volumes and energy thresholds to facilitate such a measurement.

For $0\nu\beta\beta$ searches, liquid scintillator detectors offer ease of loading and scaling the $0\nu\beta\beta$ isotope mass, such that detecting a statistically significant signal is feasible on the time scale of the experiment.

Liquid scintillator cocktails may be tuned to produce scintillation photons with different emission time profiles. Depending on the energy transfer mechanism between the species comprising the cocktail, fast or slow responses may be obtained. With fast response times, liquid scintillator experiments are able to perform accurate particle identification, which is important for background rejection, such as in the multisite analysis presented in this thesis. Conversely, slow scintillators allow good separation between directional Cherenkov and isotropic scintillation light, which allows lower threshold directional studies to be performed.

In order to provide large light yields, liquid scintillators are commonly composed of a cocktail of different molecules. The next section describes the action of a trinary (3 component) scintillator cocktail, consisting of a solvent, primary fluor and secondary wavelength shifter.

1.3.1 Liquid Scintillator Cocktails

Liquid scintillators are chemical cocktails that convert the energy deposited by charged particles traversing the medium into scintillation light [36]. The chemical cocktails usually consist of at least two species: a solvent and a dissolved fluor. The solvent molecules are excited by the incident charged particles, and transfer this excitation energy to the fluor, which subsequently relaxes to the ground state by emitting scintillation light.

In order for a liquid scintillator to be considered for rare event searches, it must satisfy the following constraints:

1. Be feasible to deploy in large quantities in underground laboratories
2. Produce a large amount ($\mathcal{O}(100)$ photons MeV^{-1}) of scintillation light
3. Produce scintillation light primarily in the wavelength range matching the peak quantum efficiency of an experiment's photodetectors
4. Be optically transparent

The most important feature of a liquid scintillator cocktail is the observed light yield, which is defined as the fraction of scintillation photons that survive propagation to a detector's photodetectors, per unit energy deposited in the medium. This is maximised by ensuring the scintillator is transparent to scintillation light in the dominant emission wavelengths (commonly the optical domain) and through the use of a secondary, wavelength shifting, fluor. These secondary fluors reduce losses due to self-absorption of scintillation light, whereby the medium re-absorbs emitted photons. By transferring energy to a lower concentration secondary fluor, with an emission spectra shifted relative to the solvent and primary, the probability of self-absorption losses is minimised. Secondary fluors are commonly used to wavelength shift the primary fluor's emissions to better match the peak quantum efficiencies of any photodetectors. It is also important to consider the losses of the fluors themselves, and fluor molecules should emit scintillation light with high quantum efficiency, with minimal non-radiative quenching.

In addition to these physical constraints, it is also important to consider practical issues relating to the deployment of large quantities of the scintillator, often in underground laboratories. Thus, the LS must:

1. Be safe to handle, particularly in terms of volatility and flammability
2. Be cheap and readily available

The following section examines the physics of liquid scintillators, using the example of a trinary cocktail such as that used in the SNO+ scintillator phase.

1.3.2 Excitations of the Solvent

The solvent molecules are the dominant population of molecular species in a liquid scintillator cocktail. As such, cases where the primary or secondary fluors are directly excited by incident charged particle interactions are neglected, and instead only the case where solvent molecules receive the initial input energy are considered.

Organic solvents scintillate due to the presence of delocalised π-electron systems [37]. In its ground state, carbon has the electronic configuration $1s^2 2s^2 2p^2$. During carbon-carbon bonding, it is energetically favourable for the $2s$ and $2p$ orbitals to 'hybridise', promoting one of the $2s$ electrons into the empty $2p$ orbital. The three p orbitals and the s orbital become 'mixed', creating one of four kinds of hybrid sp^n orbital.

Aromatic solvents undergo sp^2 hybridisation. In this case, two of the p orbitals and the s orbital are mixed, to create three identical sp^2 orbitals. These three orbitals lie in the $x - y$ plane at 120°. The remaining electron exists in the unmixed p_z orbital, at 90° to the sp^2 hybrid orbitals. It is sp^2 hybridisation that creates the hexagonal ring structure of benzene and other polycyclic aromatic

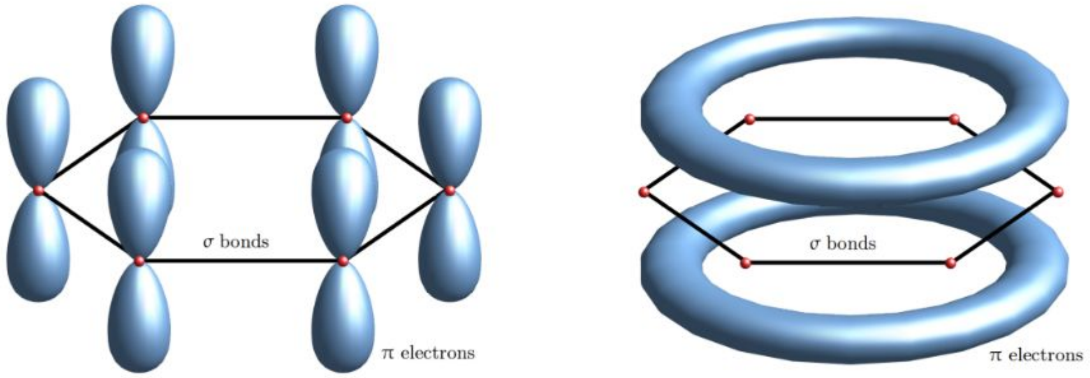

Figure 1.9: Figure shows the overlapping of p-orbitals (left) in an aromatic ring to produce a delocalised π-electron system (right), responsible for the phenomenon of scintillation. Figure taken from [33].

hydrocarbons, whereby the sp^2 orbitals between adjacent carbons and hydrogens overlap to form planar σ-bonds, and the remaining six electrons in the p_z orbitals overlap to form a fully delocalised π-electron system above and below the molecular plane, as shown schematically in figure 1.9. The electrons delocalised in this way are known as π-electrons, and it is excitations of these π-electrons which are responsible for scintillation in organic liquid scintillators.

There are four possible excitation mechanisms in the solvent:

1. Excitation of π-electrons into π-electron excited singlet states[2]
2. Excitation of other electrons, e.g. σ-electrons and those in carbon 1s orbitals
3. Ionisation of π-electrons
4. Ionisation of other electrons

Process 1 is the primary mechanism by which scintillation light is produced; ionisation of π-electrons (process 3), followed by recombination, may also produce scintillation light. However, recombination predominantly results in longer lived triplet states, producing a delayed component to the scintillation response compared with the fast $S_1 \rightarrow S_0$ singlet transitions as in process 1. The excitation

---

[2] It is unlikely π-electrons undergo the spin-forbidden direct transition to triplet states

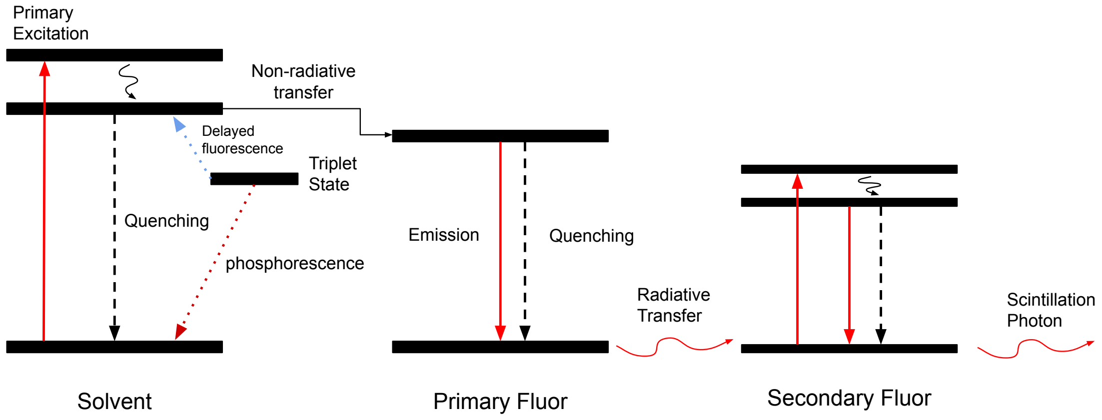

Figure 1.10: Dominant processes for a trinary scintillator cocktail. Radiative emission of solvent considered negligible compared to non-radiative transfer to primary fluor and is emitted.

energy of process 2 is dissipated non-radiatively and process 4 results in either temporary or permanent damage to the solvent molecules. Temporary damage results in 'ionisation quenching' along the particle track. This reduces the scintillation efficiency until recombination restores the solvent. Permanent ionisation damage creates 'quenching centres' within the scintillation cocktail.

Figure 1.10 shows the main energy transfer processes in a trinary scintillator cocktail, such as that used by the SNO+ detector. Each process is described in turn below.

1.3.3 Energy Transfer Mechanisms

There are various pathways energy may take through a trinary scintillation cocktail [38]. The primary excitation energy is deposited in the scintillator via elastic scattering between an incident charged particle and the delocalised π-electrons of the solvent. This generally excites the solvent to higher order electronic singlet states. These relax to the ground vibrational and first excited electronic state via internal conversion, without the emission of a photon. At this point, the solvent may relax via: a) radiative decay to its ground electronic state, with photon emission, b) internal quenching to its ground electronic state, without emission of a photon or c) non-radiatively transfer its excitation energy to the primary fluor. Since the solvent may absorb its own emission with high probability, processes a) and b) lead to efficiency losses. In a well optimised cocktail, non-radiative transfer to the primary fluor dominates.

From here, the fluor generally de-excites radiatively to the ground state, with high efficiency. The secondary fluor is chosen such that it strongly absorbs the primary's emission. In addition, the secondary's emission spectrum is such that it shifts the wavelength out of the primary's absorption band, thus suppressing the self-absorption probability of the cocktail.

As discussed above, it is also possible to produce scintillation light following ionisation and recombination of delocalised π electrons. This is especially relevant for incident particles with high charge densities, such as α-particles. Ionisation-recombination produce metastable triplet states approximately 75% of the time. These states may de-excite via: a) internal quenching to the ground singlet state, without photon emission, b) phosphorescence photon emission to the ground state or c) delayed fluorescence, with the emission of a photon. The latter occurs when the triplet state is converted back into the first excited singlet state, and from then is indistinguishable from the normal scintillation processes. However, the transition to and from a metastable triplet state means the scintillation light is emitted according to a different emission time profile, and gives rise to a slow component in the observed time distributions.

The ratio of the slow component is directly related to the charge density of the incident particles: thus, the observed emission time distributions of α particles are 'slower' than those of electrons. This has important consequences in large liquid scintillator detectors, allowing time-based α-β discrimination, which is important for background discrimination in many analyses.

1.3.4 Scintillation Emission Timing

Following excitation by ionising radiation, fluorescence emission by the terminal fluor is in accordance with a probability distribution over emission times [39]. The shape of the emission time PDF is dependent on the mass and charge of the incident particles and the respective rates of the energy transfer mechanisms detailed in section 1.3.3. The rates of these processes primarily depend on the concentrations of the solutes and the overlap of the absorption and emission spectra of each component in the energy transfer chain.

Despite the complexity of the scintillator physics, the empirical model in equation 1.33 may be used to well describe the observed emission time response of ternary scintillator cocktails:

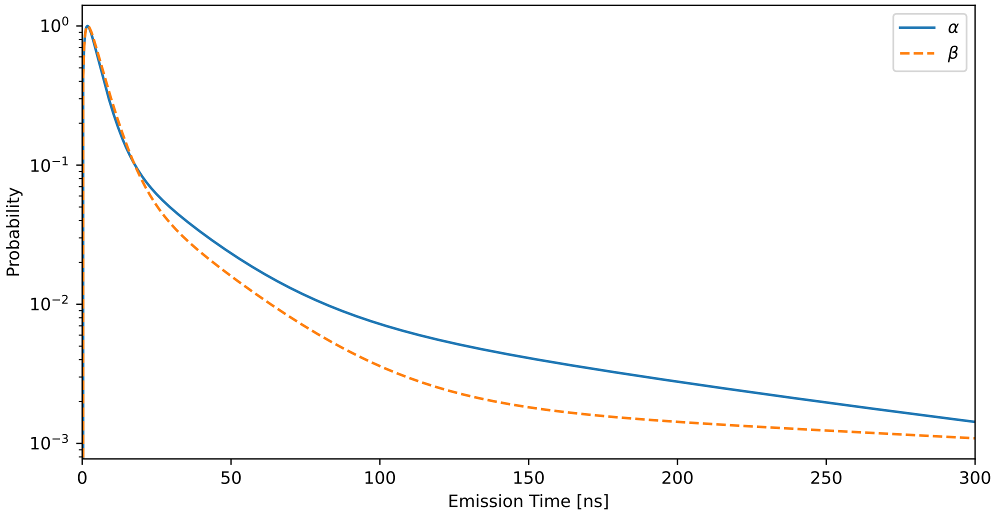

Figure 1.11: Comparison of peak-normalised scintillation emission timing distributions for α and β particles in the liquid scintillator. Distributions obtained by using equation 1.33 with tuned constants for SNO+ full-fill phase. Greater quantity of ionisation-recombination into metastable triplet states for αs leads to greater slow component.

$$f(t) = \sum_i A_i \frac{e^{-\frac{t}{ti}} - e^{-\frac{t}{t_r}}}{t_i - t_r} \tag{1.33}$$

Each term in the expansion is characterised by an amplitude, A_i, a decay constant, t_i, and a common rise time, t_r. In addition, the amplitudes satisfy,

$$\sum_i A_i = 1. \tag{1.34}$$

This empirical model is shown to provide a good description of scintillation emission timing, as demonstrated by the successful timing tuning calibration for the SNO+ scintillator phase described in chapter 3.

Each term in the expansion may be thought of as accounting for a different excitation energy transfer process, alongside a weighting factor to account for the respective rate. The rise time accounts for the finite lifetime of the excited state prior to excitation transfer. This model accurately accounts for both the fast and slow components of the emission time PDF. A comparison of the scintillation timing models for α and β particles, used in the SNO+ full-fill phase, is shown in figure 1.11.

1.3.5 Intrinsic Light Yield

The intrinsic light yield of a liquid scintillator is given by Birks' Law [40]:

$$\frac{dN}{dx} = L_0 \frac{\frac{dE}{dx}}{1 + K_B \frac{dE}{dx}} \tag{1.35}$$

Where N is the number of primary scintillation photons emitted per unit energy absorbed by the cocktail, x is the step size in mm, L_0 is the scintillation efficiency, $\frac{dE}{dx}$ is the energy loss per unit step along the particle trajectory, and K_B is the Birks constant. This constant in the denominator is dependent on the ionisation density of the incident particles, and accounts for the ionisation quenching effects (processes 2 and 4 in section 1.3.2) which reduce the amount of scintillation photons generated.

1.3.6 Photon Propagation

Emitted scintillation photons propagate through the medium and, ideally, trigger a photodetector. However, scattering and self-absorption processes conspire against this.

Rayleigh Scattering

The dominant scattering process in organic liquid scintillators, such as that used in SNO+, is Rayleigh scattering. Rayleigh scattering is an elastic $\frac{1}{\lambda^4}$ wavelength dependent process, where the incident scintillation light may interact electromagnetically with the electron cloud of over-dense groupings of solvent molecules, which occurs due to random thermal motion of the molecules [41]. This interaction causes the electrons to oscillate, which leads to the emission of scattered light of the same wavelength as the incident light. The angular dependence of the scattered light is anisotropic, preferring the forward and backwards directions acording to [42]:

$$I(\theta) \propto 1 + \frac{1-\delta}{1+\delta} cos^2(\theta) \tag{1.36}$$

Where $I(\theta)$ is the intensity of the scattered light, θ is the scattering angle relative to the incident light and δ is the depolarisation ratio of the medium.

For liquid scintillator detectors, the quantity of interest is the mean distance a scintillation photon travels before scattering. This quantity is the Rayleigh scattering length, L_{scat}, given empirically by [43]:

$$L_{scat} = \left[\frac{8\pi^3}{3\lambda^4}\left(\frac{(n^2-1)(2n^2+0.8n)}{n^2+0.8n+1}\right)^2 k_B T k_T \frac{6+3\delta}{6-7\delta}\right]^{-1} \tag{1.37}$$

Where n is the refractive index of the scintillator, λ is the wavelength, T is temperature, k_B is Boltzmann's constant, k_T is the isothermal compressibility and δ is the depolarisation ratio.

It is desirable to have very long scattering lengths, on the order of tens of metres, for the emitted scintillation light. Excess scattering delays the detection of light, making reconstruction of the ionising radiation's event vertex more challenging. In addition, deviations from straight line paths from the vertex to the photodetectors may increase the distance traversed by photons in the scintillator. This increases the probability of self-absorption, and thus the likelihood of internal quenching processes removing the light. Therefore, scattering reduces the observed light yield of the scintillator, and increases the resolution of position reconstruction.

Absorption

As scintillation photons propagate through the medium, they may be re-absorbed by the scintillator. The probability of absorption depends on the wavelength of the scintillation photons, as well as the absorption spectra of the scintillator components. In principle, any component of the scintillator may be responsible for absorption, however the most likely candidate is the solvent, based on molecular abundance arguments. Low concentration wavelength shifting secondary fluors, as discussed above, are therefore crucial to shift emitted scintillation light to longer wavelengths, out of the solvent's absorption spectrum.

Where absorption does occur, the excitation energy will again be subject to the energy transfer mechanisms detailed above. Since the energy may be internally quenched, excess levels of absorption lead to severely curtailed light yields. Thus, it is imperative scintillator cocktails are carefully chosen to maximise the separation between the absorption spectrum of the solvent, and the emission spectrum of the secondary fluor.

Practically, the combination of absorption and scattering effects may be observed through measurements of the extinction length, L_{ex}:

$$L_{ex} = \frac{L_{abs} L_{scat}}{L_{scat} + L_{abs}} \tag{1.38}$$

The extinction length is the mean distance an optical photon travels before it is absorbed or scattered. Since it is difficult to separate contributions from absorption and scattering, the extinction length is the main bulk property of the scintillator accessible to measurements.

1.3.7 Photon Detection

In liquid scintillator experiments, scintillation light is typically detected by arrays of photomultiplier tubes (PMTs) [44]. PMTs are a common choice, providing a cost effective way to provide large area coverage, while maintaining good time responses and high efficiencies. Typical PMTs comprise of photocathode, a set of 'dynodes' and a charge collecting anode [44], as shown in figure 1.12. In addition, the PMTs may be placed within a concentrator 'bucket', which increases the PMT coverage of a detector, by reflecting indirect light onto the photocathode.

An incident scintillation photon striking the photocathode may induce the release of a photoelectron via the photoelectric effect. This photoelectron is focussed by a strong electric field onto the first dynode, causing a number of secondary electrons to be ejected from the surface. A strong electric field focusses these secondary electrodes onto the second dynode, which further amplifies the number of electrons produced. This process continues through a series of focussing and amplification steps at each dynode, before the charge is collected on the anode.

The cascade of electrons produced by collisions with each dynode serves to amplify the charge collected at the anode, ensuring even single photoelectron emissions from single photon hits produce observable signals.

To detect low energy physics events (< 1 MeV), it is desirable for tubes to feature high quantum efficiencies, to maximise the observed light yield, and a fast time response, to ensure good position resolution. A fast time response is also important for applications seeking to resolve differences between Cherenkov and scintillation light for background rejection, and to reduce inter-event deadtime.

The quantum efficiency is defined as the ratio of photoelectrons escaping the anode to incident photons, and is wavelength dependent. Therefore, spectral matching between the quantum efficiency vs wavelength curve of the PMT and the scintillation emission spectra is essential, and is an important role of the secondary fluor. Collection efficiency is the fraction of free photoelectrons that are successfully focussed onto the first dynode, initiating the amplification process of the signal. The product of the quantum and collection efficiencies is an important objective to maximise when selecting PMTs.

Fast timing has important implications for position reconstruction, as explained in chapter 2, and is characterised by the transit time spread (TTS) of photoelectrons traversing the dynode to the anode. This spread is the result of the following effects:

1. Photoelectrons emitted from different points on the photocathode must traverse different distances to the first dynode (D_1)
2. Photoelectrons are created with different initial velocities
3. Secondary photoelectrons created at the dynodes traverse slightly different path lengths before reaching the anode

This results in a distribution of times each secondary electron arrives at the anode. If this spread is large, it is impossible to know whether an early hit is due to a close event vertex or a shorter than average transit time (or vice versa).

Additionally, the mean TTS should be short, ensuring PMTs reach threshold and trigger as soon as possible. This reduces inter-event deadtime.

Dark noise is the name given to spurious PMT signals not arising from exposure to scintillation light. These signals may be produced by a few different processes, the dominant source being thermionic emission of electrons from the photocathode. In addition, thermionic emission from the dynodes themselves and, at high operating voltages, emission from the side wall of the photocathode-D_1 system, contribute to the dark noise.

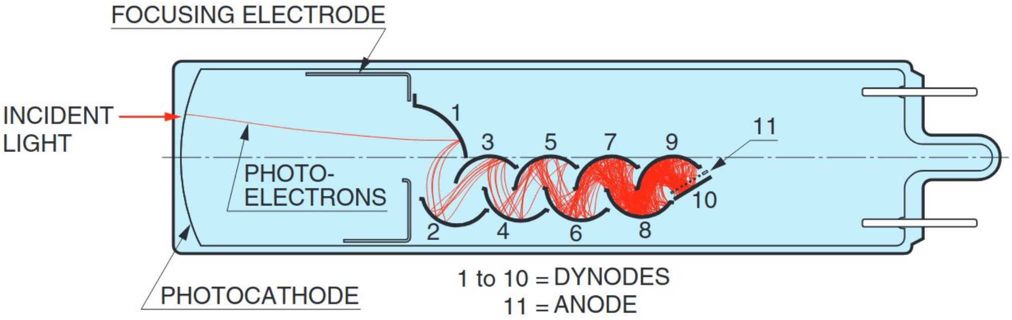

Figure 1.12: Schematic diagram of linear focussing PMT and possible photon electron trajectories through the dynode stack. Taken from [45].

A secondary component of the dark noise arises from radioactivity in the PMT materials, particularly the photocathode glass envelope. ^{40}K β-decays produce electrons that may enter the PMT and undergo charge amplification, further reducing the signal-to-noise ratio.

Typical dark noise reduction strategies include:

- Elimination of radioactive contaminants in the PMT materials
- Use of small area cathodes
- Masking of photo-emissive materials from the cathode-D_1 system walls
- Refrigeration of the cathode, thus reducing the dominant source of thermionic emission.

Balancing the timing, signal-to-noise ratio, detection efficiency, coverage and overall cost of the PMTs in a large scale liquid scintillator detector is crucial. The PMT attributes play a key role in determining the reconstruction resolution and low energy threshold of a detector, and thus broadly dictates the possible physics analyses.

Chapter 2

The SNO+ Detector

This thesis utilised data collected exclusively by the SNO+ detector during it's full-fill scintillator phase. The following chapter details the key hardware components of the detector, the primary science goals and detecting medium of each operational phase, and briefly explains the event building and reconstruction algorithms.

2.1 Detector Overview

The SNO+ experiment is the successor to the Sudbury Neutrino Observatory (SNO), as detailed in section 1.1, and makes use of a lot of the same hardware: a 6 m radius acrylic vessel (AV), shielded by 7 kilo-tonnes of ultra pure water (UPW), and instrumented with $\sim$ 9300 8" Hamamatsu R1408 photomultiplier tubes [30]. The PMTs are attached to a stainless steel PMT support structure, at an average radius of 9 m, and placed within 30 cm concentrator buckets, which increase the coverage from 31 % to 54 % [46]. The quantum efficiency vs wavelength, alongside the TTS, is shown in figure 2.2, with the typical TTS around 3.6 ns at FWHM.

The key difference between SNO+ and its predecessor is the detecting medium: where SNO used heavy water, the SNO+ experiment is filled with 780t of liquid scintillator. Figure 2.1 shows the main features of the SNO+ detector.

Located deep underground (2 km depth; 6000 m.w.e.) and exhibiting ultra-pure detector conditions, the SNO+ detector is well suited for many rare-event searches, such as low-energy solar neutrinos, reactor antineutrinos, geoneutrinos, $0\nu\beta\beta$ and invisible nucleon decays.

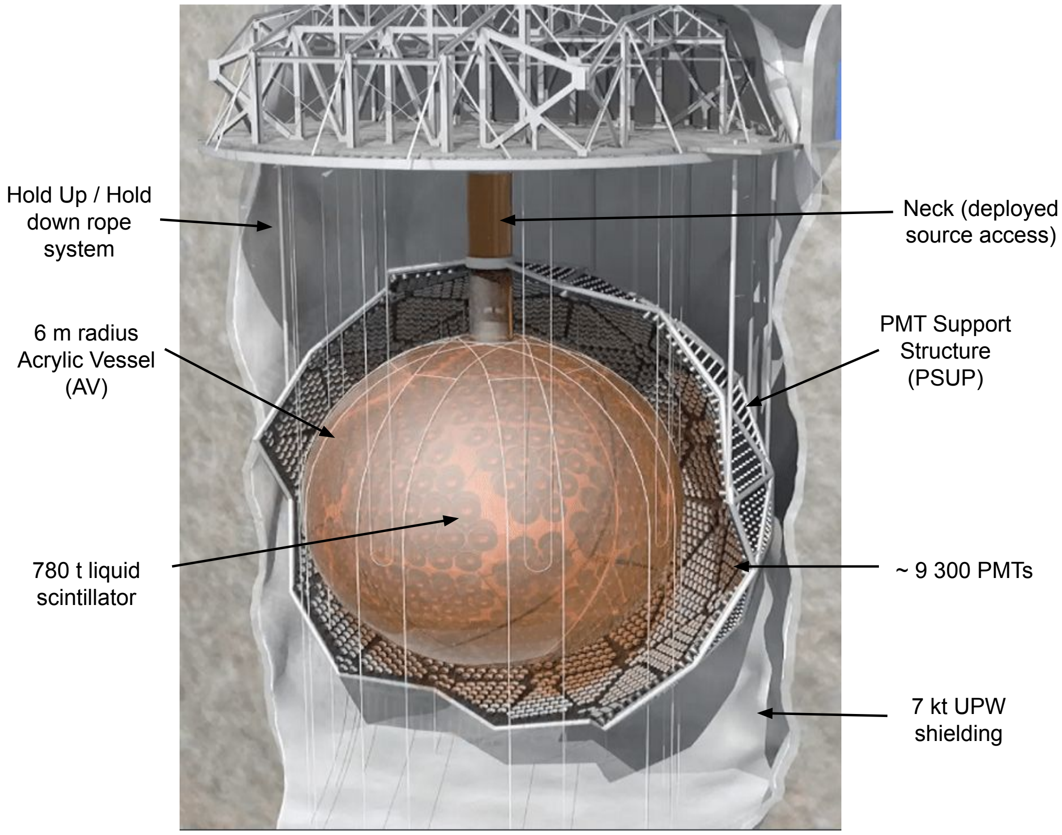

Figure 2.1: The SNO+ detector. Adapted from [47].

2.1.1 Experimental Phases

The experiment has transitioned through a number of different operational phases, defined by the composition of the detecting medium. Each phase featured unique conditions, rendering them suitable for different physics analyses.

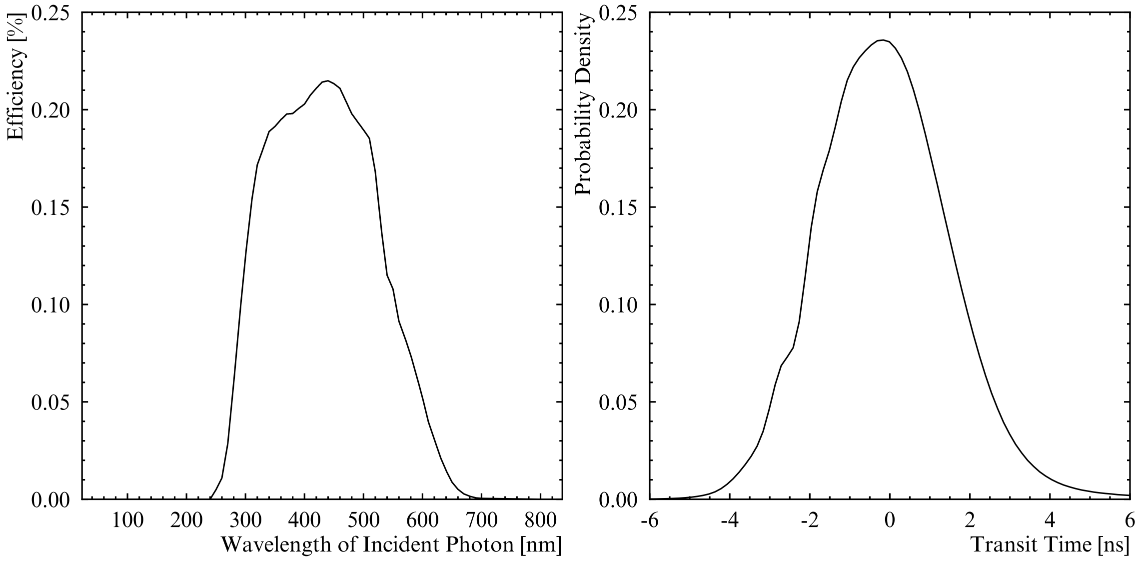

Figure 2.2: SNO+ PMT (r1408) wavelength dependent efficiency and transit time spread. Taken from RAT 7.0.15 MC software.

Water Phase

Prior to filling with liquid scintillator, the AV was filled with 905 tonnes of UPW, and between May 2017 and July 2019, SNO+ operated as a water Cerenkov detector.

The water phase was designed to allow the characterisation of the PMT response, data acquisition systems and test the detector's performance [48]. In addition, the background levels of detector hardware components (external backgrounds) were assessed, for example the external water, PSUP, hold-down ropes and the acrylic vessel itself.

In addition, SNO+ was able to undertake physics data-taking campaigns, with the main goals of improving limits on invisible nucleon decays [49, 50], observing neutron-proton captures [51], 3.5 MeV threshold observation of solar neutrinos [52] and the first observation of reactor antineutrinos with a water-Cerenkov detector [53].

Partial Fill

After water phase commissioning and data taking was complete, the deployment of the SNO+ liquid scintillator began. Liquid scintillator was injected into the AV near the top of the neck, while

UPW was removed from the bottom. This configuration was necessary to the different density and immiscibility of LAB in UPW.

From March to October 2020, the scintillator fill was interrupted by the COVID-19 pandemic. This led to an unexpected dual-phase scintillator-UPW configuration for a prolonged period of time, with 365t of scintillator added to the AV, and a PPO concentration of 0.6 g/L. An interface layer formed at $z = 0.75$ m, where $(0, 0, 0)$ is defined as the centre of the AV. This interface layer was clearly visible in photographs, as shown in figure 2.3.

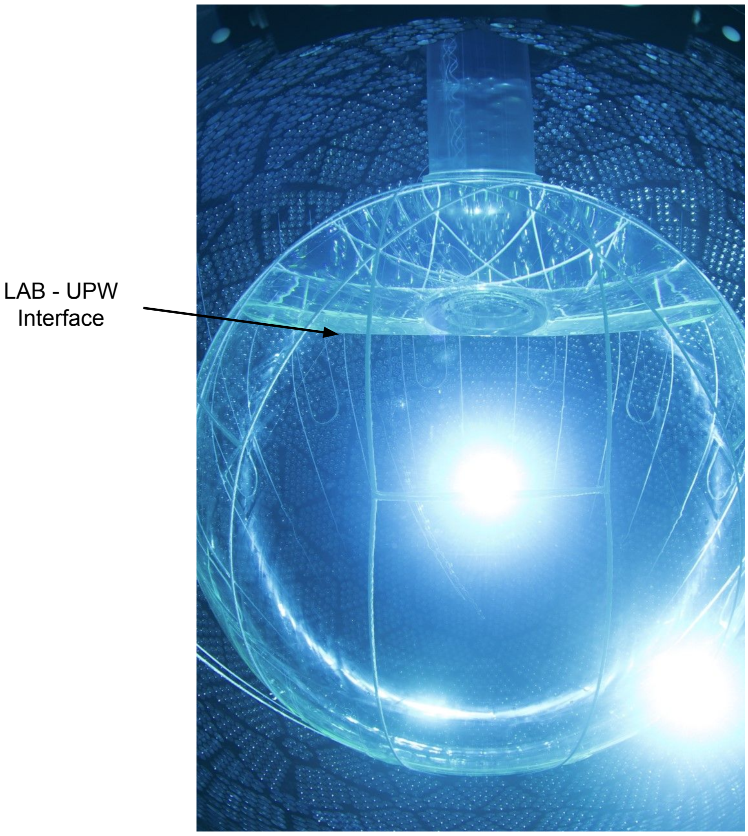

Figure 2.3: A photograph from within the PSUP during the partial fill phase. The LAB-UPW interface is clearly visible, with LAB filling from the top.

This partial fill configuration yielded approximately 130 days of physics data, from which the first SNO+ measurement of long-baseline neutrino oscillations was reported [54], alonside the first observation of ^{8}B solar neutrinos in SNO+ scintillator [55].

Full Fill (0.6 g/L PPO)

Following the lifting of COVID-related restrictions, filling with liquid scintillator was resumed. The remaining UPW was removed from the bottom and a total of 780t of LAB + 0.6 g/L PPO was loaded into the AV. Between April to June 2021, the detector was left to settle, before calibrations and physics data taking proceeded.

These data were primarily used for the first demonstration of event-by-event direction reconstruction for ^{8}B solar neutrinos [56].

Full Fill (2.2 g/L PPO)

After data-taking was completed in the full-fill 0.6 g/L phase, PPO loading was resumed. A final concentration of 2.2 g/L PPO was reached at the end of April 2022, and the detector was again left to settle until March 2023. During this time, new calibrations of the light yield, scintillator emission time profile and measurements of the internal backgrounds were performed. These data form the basis of this thesis, yielding approximately 145 days of livetime.

During the 2.2 g/L full-fill phase, the primary physics goals were: accurate calibrations using in-situ radiogenic backgrounds of the scintillator light yield and emission time profiles, characterisation of the internal ^{238}U and ^{232}Th backgrounds in preparation for the ^{130}Te loading, a first demonstration of multi-site event discrimination for use in ^{8}B solar neutrino studies, measurement of the ^{8}B solar neutrino flux, reactor and geoneutrino measurements, supernova monitoring, neutron multiplicity measurements in liquid scintillator and exotic physics searches.

Full Fill (2.2 g/L PPO + 2.2 mg/L bisMSB + 5.5 mg/L BHT)

bisMSB was first added to the scintillator cocktail in July 2023, with a final concentration of 2.2 mg/L reached in November 2023. This was done to increase the light yield of the SNO+ scintillator, leading to lower energy thresholds and more precise reconstruction.

In addition, an antioxidant, BHT, was added to stabilise the scintillator performance, helping alleviate light yield degradation via oxidisation quenching.

At the time of writing, this is the configuration of the SNO+ detector. The primary physics goals are the same as the 2.2 g/L full-fill PPO phase, with added efforts to perform target-out analyses for the $0\nu\beta\beta$ program and taking advantage of the improved position reconstruction and light yield for existing analyses.

Tellurium-Loaded Phase (2025 onwards)

The flagship physics analysis of the SNO+ experiment is the direct search for $0\nu\beta\beta$ decay. To this end, the SNO+ liquid scintillator will be doped with 0.5% by mass natural Tellurium, which contains the double-beta decaying isotope, ^{130}Te.

The selection of ^{130}Te was dependent on a number of factors [57]. The relatively long $2\nu\beta\beta$ half-life of 7.71 x10^{20} years [58] reduces the background from this process, with other candidates, such as ^{150}Nd, having 100x shorter $2\nu\beta\beta$ half-lives [59]. In addition, Te does not require enrichment, as the natural isotopic abundance of ^{130}Te is 34.08%. This has major benefits in terms of cost and ease of deployment relative to other $2\nu\beta\beta$ isotopes. Finally, the expected Q-value of ^{130}Te's $0\nu\beta\beta$ decay is 2.5 MeV, which is above the majority of the signals from natural radioactivity contaminations in the scintillator and external components of the detector. This leads to a much less congested ROI, making the signal extraction of $0\nu\beta\beta$ decay easier.

The tellurium loaded phase is initially scheduled for a 0.5% by mass deployment of ^{130}Te, with the planned increase to to 2%, or more, in the longer term.

2.2 The SNO+ Liquid Scintillator

The SNO+ liquid scintillator currently consists of linear alkyl-benzene (LAB) solvent, (2,5)diphenyhloxazole (PPO) primary fluor, (1,4)bis(2-methylstyryl)benzene (bisMSB) wavelength shifter and Butylated Hydroxytoluene (BHT) anti-oxidant stabiliser. Table 2.1 gives an overview of the current and target concentrations of each component of the SNO+ scintillator cocktail.

Previous experiments, such as KamLAND [60] and Borexino [61], deployed large scale liquid scintillator cocktails using pseudodocumene (PC) as the solvent alonside PPO. However, due to the placement of the SNO+ detector deep underground and PC's toxicity, volatility, low flash point and incompatibility with acrylic, it was deemed unfeasible for use in SNO+. In light of these issues, extensive trials of different solvents were performed, with LAB selected for use in SNO+ [62].

Chemical Species	**Role**	**Target Concentration**	**Current Concentration**
LAB	Solvent		
PPO	Primary Fluor	2.0 g L^{-1}	2.2 g L^{-1}
bisMSB	Secondary Fluor	5.0 mg L^{-1}	2.2 mg L^{-1}
BHT	Anti-Oxidant	5.5 mg L^{-1}	5.5 mg L^{-1}
Natural Te	$0\nu\beta\beta$ source	0.5 % by mass	

Table 2.1: Target and current concentrations of each component of the SNO+ liquid scintillator, relative to the upcoming tellurium phase.

LAB has been shown to have competitive light yield with PC-based cocktails, whilst being compatible with and having near identical refractive index to acrylic. In addition, LAB has a high flashpoint, is non-toxic, and, once purified, has Raleigh scattering lengths in excess of 70 m at 546 nm. Thus, it exhibits many of the positive traits discussed in section 1.4.1, whilst being much safer to handle and deploy underground than PC.

The concentration of the primary fluor, PPO, was optimised with Monte-Carlo (MC) simulations to maximise non-radiative transfer from the solvent, whilst limiting PPO self absorption. As the fluor concentration increases, the gains in light yield are expected to plateau, as shown in figure 2.4. Coupling this behaviour with the time needed to deploy the fluor underground, a 2 g/L target concentration was set, with the final concentration measured as 2.2 g/L PPO.

The addition of a secondary fluor and wavelength shifter, bisMSB, was proposed to further increase the light yield of the scintillator. This is important, as it acts to offset the quenching effect of the tellurium deployment.

Similar Monte-Carlo studies were performed to optimise the target concentration of bisMSB. Due to the $\sim$ cm scale absorption lengths of bisMSB, it was found that adding only 1-2 mg/L bisMSB would siginficantly boost the light yield of the scintillator. bisMSB achieves this light yield boost in two ways: by radiatively absorbing scintillation light emitted by the PPO and re-emitting it at a wavelength closer to the peak efficiency of the SNO+ PMTs ($\sim$ 450 nm), and reducing PPO self-absorption losses. Figure 2.5 shows the respective absorption and emission spectra of LAB, PPO and bisMSB.

BHT is a non-scintillating ‘optically neutral’ anti-oxidant, added to limit the degradation of the SNO+ scintillator’s optical clarity due to oxidation reactions. This has no impact on the light yield

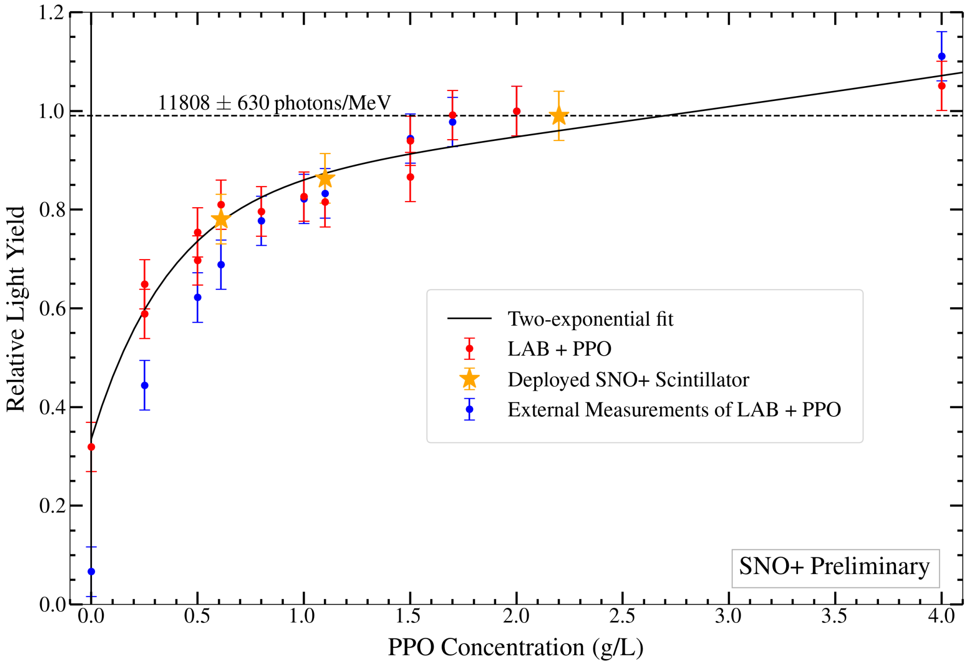

Figure 2.4: Light yield (photons / MeV deposited) relative to 2.0 g/L LAB + PPO bench top sample. Yellow points show light yield measurements of samples taken from the detector, red shows results from bench top samples and blue shows external measurements. Figure taken from [63]

or extinction lengths, and was added to maximise the longevity of the experiment.

Finally, the $0\nu\beta\beta$ isotope, ^{130}Te, will be loaded into the scintillator using a novel technique [64] to combine Te with 1,2-butane-diol, forming a soluble metal ion ligand complex. This produces a scintillator doped with the $0\nu\beta\beta$ isotope, with an initial target concentration of 0.5% by mass. The impact of tellurium loading on the scintillator is to decrease the following characteristics: extinction length, light yield, and the scintillation emission timing spread. Figure 2.6 shows the impact of tellurium-loading on the charge intensity spectrum for four LAB + 2.2 g/L PPO samples, two with 0.5% tellurium (yellow) and two without (blue). Comparison of the end points of these spectra allow quantification of the relative change in light yield expected with 0.5% loading.

However, due to the high light yield achieved in the full-fill bisMSB phase, the reduced Te-phase light yield is expected to be more than sufficient to maintain the energy resolution required to

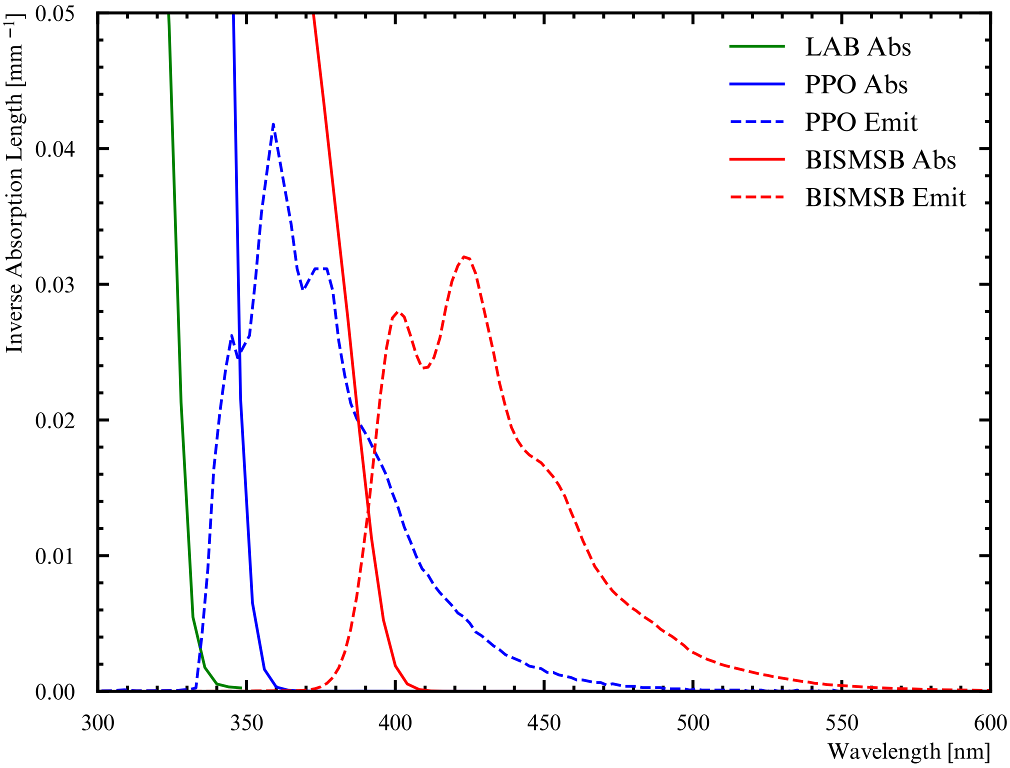

Figure 2.5: The dominant absorption and emission spectra of LAB (solvent), primary fluor (PPO) and wavelenth shifter (bisMSB), as simulated in RAT 7.0.15.

resolve the $0\nu\beta\beta$ signal. Additionally, the reduced extinction length was determined to still be in excess of 10 m at 380 nm. Finally, the faster timing, due to quenching effects observed with Te-doped scintillators, will further improve the performance of time-based classifiers and reconstruction resolution, for instance the multi-site classifier detailed in later parts of this thesis.

2.3 The SNO+ Trigger System

Following photon generation, propagation and detection by the PMTs, the detector needs a method to collect PMT hits together and package them into physics events. This is achieved by the SNO+ triggering system and event builder.

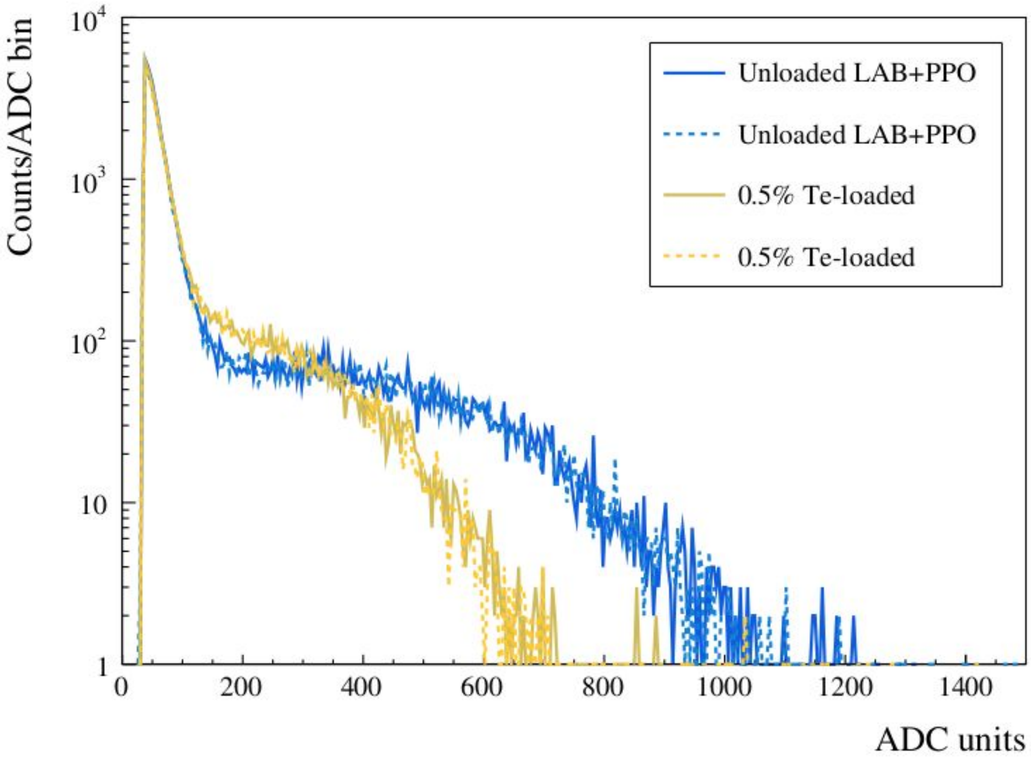

Figure 2.6: Comparison of light yield distributions for LAB + 2.2 g/L PPO scintillator samples, with (yellow) and without (blue) 0.5% tellurium loading. Two samples of each class are included (dashed and solid lines). The end points of the spectra provide an estimate of the relative change in LY expected following tellurium loading. Figure taken from [33].

Without a triggering system to discriminate between potential physics events and spurious noise, the sheer quantity of data recorded would be infeasible to save and analyse. Thus, the SNO+ triggering system utilises carefully calibrated thresholds to bring the event rate down to manageable levels whilst selecting physics of interest with low energy thresholds and acceptable signal sacrifice.

A simple schematic of the SNO+ trigger system is shown in figure 2.7. At a basic level, physics events are detected and packaged into discrete events via the summation of PMT charge pulses and their comparison to trigger thresholds. When a trigger threshold is crossed, PMT hit information is read out within a 400 ns ‘event window’. An event builder script packages this information into a discrete event object, saving it to a ZDAB file format. These ZDABs are continuously offloaded to external repositories, where data cleaning and position and energy reconstruction takes place.

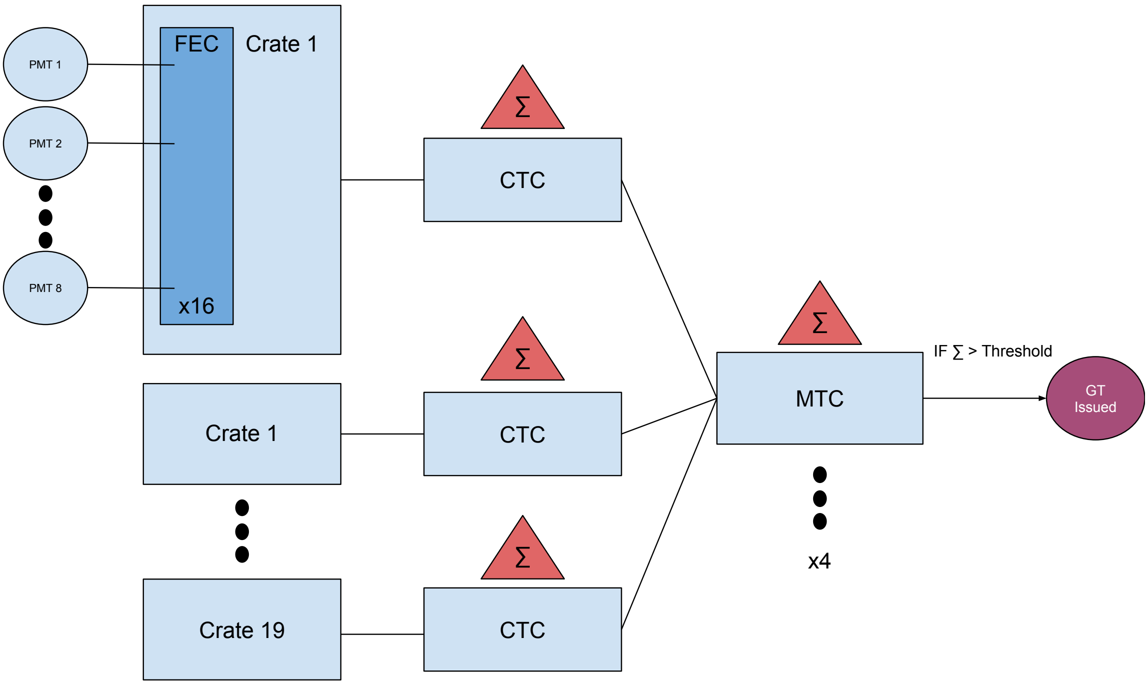

Figure 2.7: A simplified diagram of the SNO+ trigger system. The output charge of each PMT is summed per crate at the CTC, with the crate-level output subsequently summed per global trigger at the MTC. If threshold is reached, the MTC issues a global trigger (GT).

The SNO+ trigger system consists of groups of 8 PMTs connected to front end cards (FECs), with 16 FECs grouped together to form a crate. In total, SNO+ features 19 crates, each connected to a total of 512 PMTs. PMT hit level information is continuously digitised and saved for 400 ns on each FEC, awaiting a global trigger (GT) readout signal to be issued.

Each crate is in turn connected to a crate trigger card (CTC), which is responsible for summing the output pulses of every PMT in the crate. The combined pulses from each CTC are then summed at the master trigger cards (MTC), with a single MTC responsible for each of the 4 SNO+ event triggers:

1. ESumHigh, a high-gain copy of the PMT voltage pulse

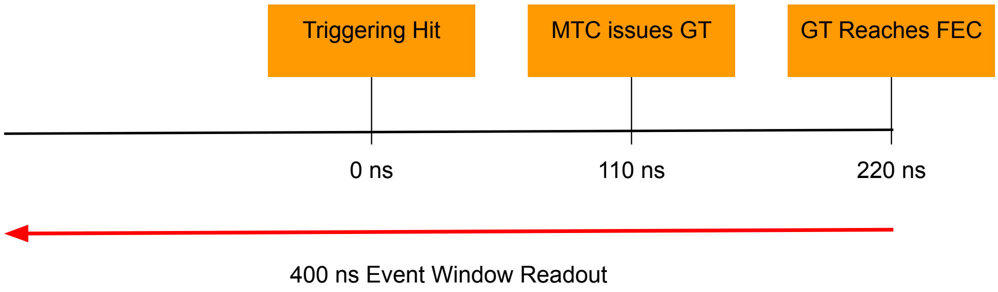

Figure 2.8: Timeline of event readout after GT is issued. Due to cable delays between the FECs and MTCs, it takes 110 ns for the triggering hit to reach the MTC, and the same again for the GT to reach the FECs. Thus, the event window includes hits 180 ns before and 220 ns after the triggering hit of the GT.

2. ESumLow, a low-gain copy of the PMT voltage pulse
3. N100, the number PMTs above threshold within a 100 ns window
4. N20, the number of PMTs above threshold within a 20 ns window.

The ESumHigh and ESumLow global triggers correspond to the total charge seen by all the PMTs, and the N100 and N20 triggers correspond to the total number of triggered PMTs within a 100 ns and 20 ns window, respectively. Of these triggers, the N100 trigger is the most commonly issued, with the current threshold set at 15 hits.

If any of the four MTC pulses cross an event trigger, a GT is issued. This sends a signal for the information digitised on each FEC to be read out and sent to the event builder and constructed into a physics event. Due to cable delays, information is saved 180 ns before and 220 ns after the PMT hit that triggered the GT, as shown in figure 2.8.

Following issuing of a GT and the subsequent packaging of PMT hit information by the builder, reconstruction of event position, time and energy is performed. These algorithms are the focus of the last section of the chapter.

2.4 Event Reconstruction in the Scintillator Phase

After data is packaged by the builder into ZDAB files, and these data have been offloaded to the grid for data processing, reconstruction of the physical variables of interest begins.

Once reconstruction is complete, event information is saved into two derived file formats: RATDS and NTUPLE. These derived file formats are commonly used in physics analysis, and differ based on the quantity of information they contain (NTUPLES contain no PMT hit-level information).

The distribution of PMT hit times and positions constitute the observables in a SNO+ event. From these data, physics quantities are reconstructed, namely: position, time and energy.

This section details the reconstruction algorithms employed during the SNO+ full-fill 2.2g/L PPO phase, with particular emphasis on position reconstruction and the use of calibrated scintillation emission time models.

Reconstruction in SNO+ assumes every event is a 2.5 MeV electron interaction, with point-like energy deposition. Practically, this makes little difference, except to note the reconstructed energy is the 'equivalent electron energy'.

2.4.1 Position Reconstruction

Position reconstruction hinges on two algorithms: QuadFit ('Quad') and MultiPDF. Quad is a simple analytical method, solving 4 simultaneous equations for cartesian coordinates and time. This solution seeds MultiPDF, which uses the method of maximum likelihood to refine the fitted position-time vertex.

QuadFitter

Given a perfect detector, with no spread in PMT timing or scintillation emission, the reconstructed position and time would be exactly determined by any 4 PMT hit times, positions and the speed of light in the detector medium using the simple time of flight equation:

$$|\vec{x}_h - \vec{x}_{ev}| = c_{ef}(t_h - t_{ev}) \tag{2.1}$$

Where $\vec{x}_{h,ev}$ and $t_{h,ev}$ are the position and time of a PMT hit and particle interaction, respectively. Four PMT hits yield four simultaneous equations, which may be solved to obtain $\vec{x}_{ev}$ and t_{ev}, without error.

In real life, there is a spread in the scintillation emission time, as well as in the transit time spread of photoelectrons traversing the PMT dynode stack. Therefore, the result of equation 2.1 depends on the choice of PMTs chosen.

To deal with this smearing, many random groups of 4 PMTs are evaluated for each event, creating a 'point-cloud' of (x_{ev}, y_{ev}, z_{ev}, t_{ev}) solutions. Straight-line light paths are assumed, with an effective velocity, c_{ef}, chosen to minimise the reconstructed radial bias of events in the inner 5 m. The median result of each point-cloud is returned as the QuadFitter position-time reconstructed vertex.

MultiPDF

The Quad vertex is used as a seed for the MultiPDF fit. For a given trial vertex, ($\vec{x}_{ev}$, t_{ev}), the log-likelihood is calculated by comparing to a 'time-residual' PDF. The specific PDF used is dependent on the radius of the Quad vertex, with PDFs created for each 1 m radial shell.

Time residuals are of central importance to SNO+, providing a position-independent quantity from which to estimate the reconstruction vertex goodness-of-fit and perform particle discrimination.

$$t_{res} = t_{hit} - t_{t.o.f} - t_{ev} \tag{2.2}$$

Where t_{hit} is the PMT hit time, $t_{t.o.f}$ is the calculated travel time, assuming straight line paths, from the reconstructed vertex to a given PMT, and t_{ev} is the reconstructed event time relative to the start of the event trigger. MultiPDF features PDFs constructed using events simulated in 1 m radial shells. Due to refraction between the scintillator, acrylic and external water, the validity of the straight line paths is reduced at higher radii. Using PDFs for each radial shell lessens the impact of these deviations with radius.

The residual time, t_{res}, providing accurate position and time reconstruction, should be equivalent to the convolution of the scintillator's emission time profile and the detector response.

In the context of position reconstruction, a trial position-time vertex is used to calculate an event's time residual distribution according to equation 2.14. Straight line paths and an effective velocity in each medium (scintillator, acrylic, and external water) are used to calculate the time residual of each hit. This residual distribution is then compared to the time residual PDF to compute the log-likelihood of the trial vertex. The log-likelihood is calculated on a hit-by-hit basis:

$$log(\mathcal{L}) = \sum_{i=0}^{N_{hits}} log(P(t^i_{res})) \tag{2.3}$$

Where $P(t^i_{res})$ is the likelihood of the time residual of the i^{th} PMT hit in the event, t^i_{res}, and N_{hits} is the total number of PMT hits.

The time residual PDF is shown in figure 2.9. It is constructed from 100,000 2.5 MeV electrons, simulated within the innermost 4 m fiducial volume (FV) of the scintillator volume. Equation 2.14 is used, however the MC-truth information for t_{ev} and $t_{t.o.f}$ are used. The primary sources of uncertainty in the position reconstruction are down to the widths of the PMT TTS and scintillation emission profiles, which together dominate the shape of the distribution. These two distributions are convolved to produce an irreducible spread in the reconstructed time residuals, depending on individual PMT performance and the physics governing the emission, propagation and detection of scintillation light.

However, as shown in figure 2.10, the reconstructed position resolution in the full-fill phase is $\sim$ 11 cm, which is sufficient for accurate definitions of fiducial volumes and external background rejection.

2.4.2 Energy Reconstruction

If a successful position-time vertex is reconstructed, the energy fitter attempts to reconstruct the equivalent electron energy of the event. At a basic level, energy is reconstructed by noting its linear relationship with the total number of PMTs hit during an event (Nhit). Thus, a simple straight line fit to Nhit vs Energy, obtained from calibrated simulations, may be used.

This linear relationship holds for events near the centre of the AV, with energies up to about 5 MeV. At larger energies, the curve deviates from the linear relationship due to saturation, where

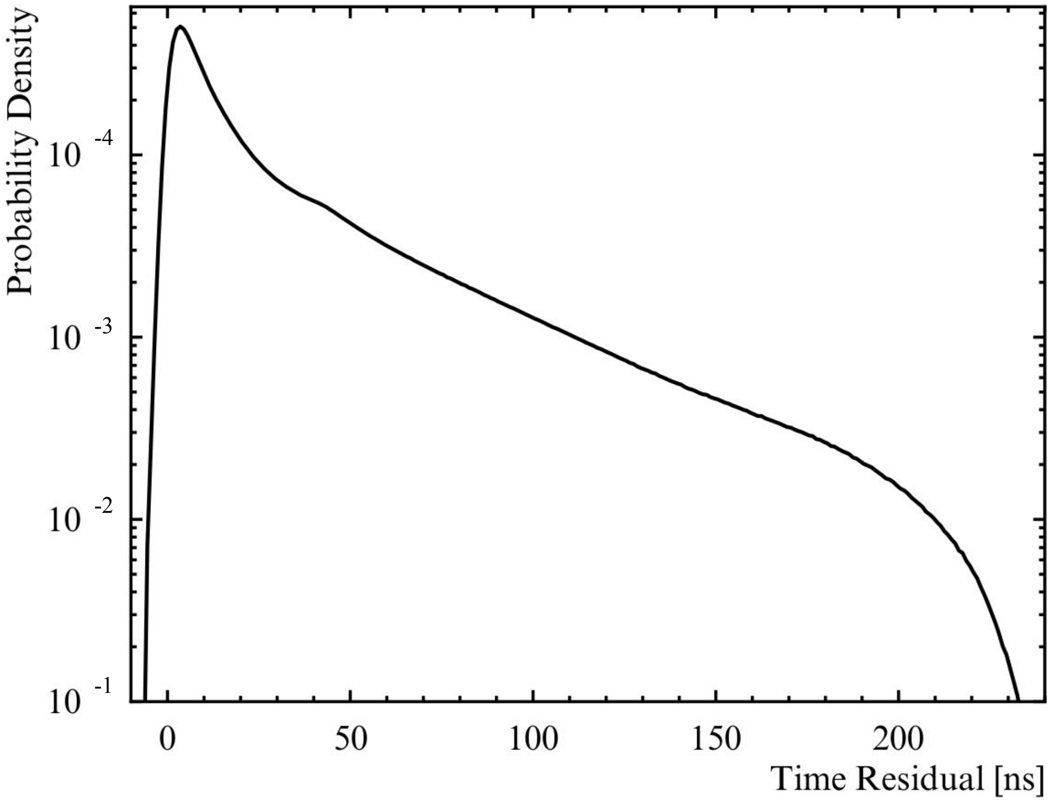

Figure 2.9: Scintillator emission time PDF used for maximum-likelihood position fit given in Equation 2.3 in 2.2 g/L scintillator phase. Constructed from 100 000 2.5 MeV electrons reconstructing within 3-4 m radial shell.

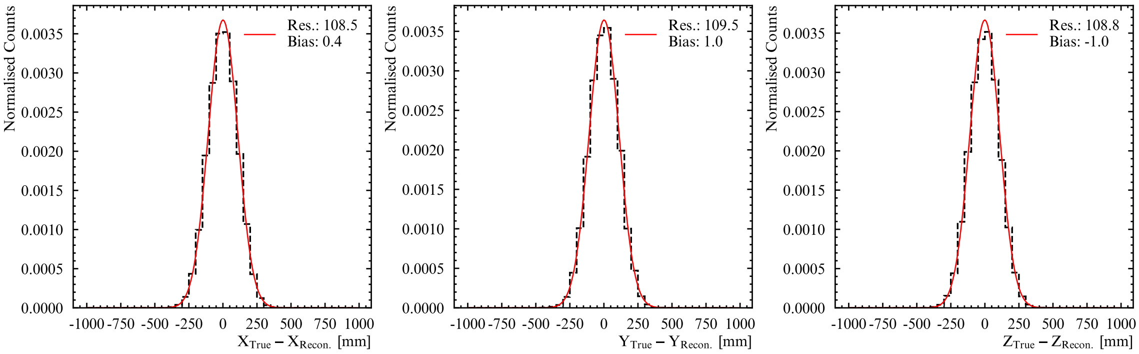

Figure 2.10: Resolution and bias of 2.2 g/L phase position fitter. Distributions obtained from 2.5 MeV electron simulations within 4 m FV.

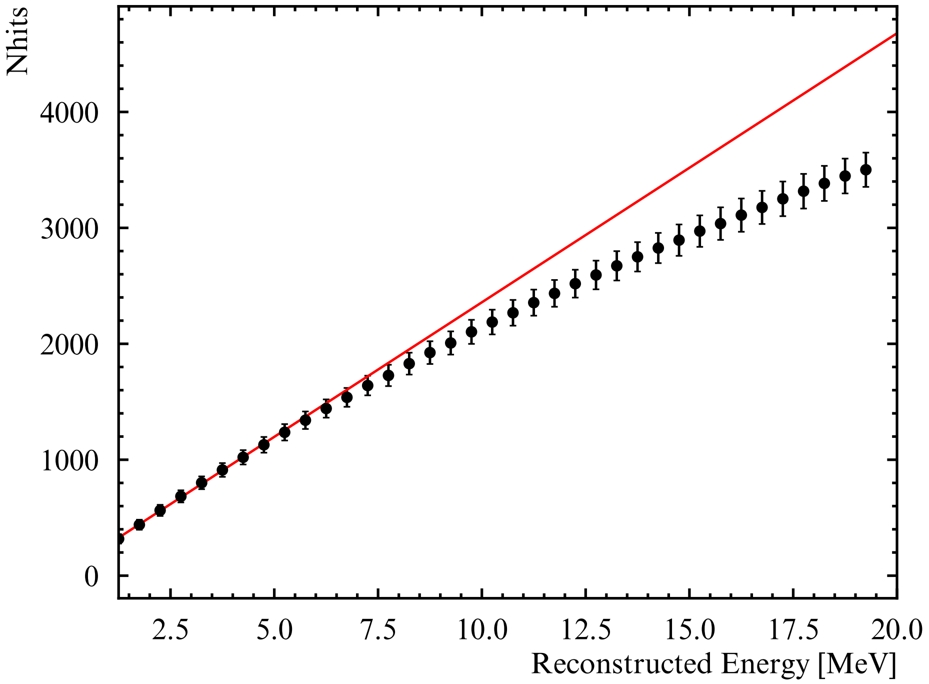

Figure 2.11: Reconstructed energy vs Nhits for electrons reconstructing within a 4 m FV. The red curve is a linear fit performed in 1-5 MeV range. Error is the standard deviation of nHits in each 0.1 MeV bin.

PMTs are hit multiple times during a single event window. As described above, PMTs may only register a single hit per event, and so Nhit is no longer linearly proportional to energy and would yield an underestimate of the true energy, as shown in figure 2.11.

Non-linearities in the Nhit-energy relationship also emerge for events occurring towards the edges of the AV. Light emitted towards the edges travels less distance before hitting PMTs on the near side, leading to many more multi-photon hits on a cluster of nearest PMTs. This effect is particularly pronounced for events reconstructing beyond 5 m a radius in the AV, as shown in figure 2.12. PMTs also feature angle of incidence dependent collection and quantum efficiencies, further distorting the Nhit-energy relationship.

The current energy fitter attempts to correct for each of these effects. The position dependence is removed by constructing a look-up table of position-dependent correction factors, using MC sim-

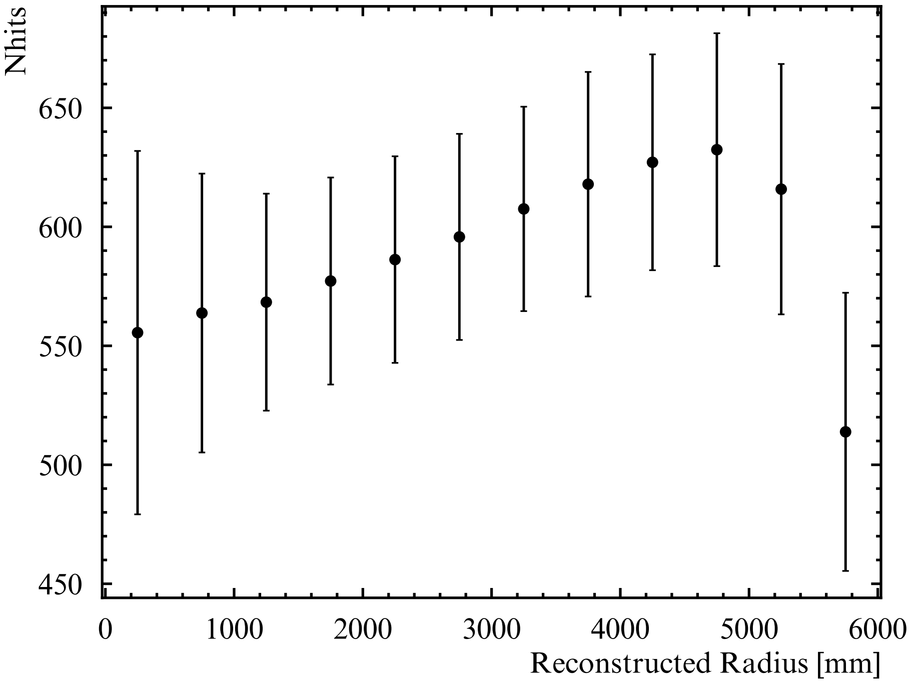

Figure 2.12: Dependence of Nhits on reconstructed radius from 2.5 MeV electrons simulated throughout AV. Error bars are standard deviation of Nhits in each 0.5 m radial bin.

ulations. The multi-hit non-linearity is handled by translating nhits to a linear estimator of the energy [33].

To first order, poisson statistics governs the recorded Nhit for an event. As such, the uncertainty in the reconstructed energy for a 2.5 MeV electron is simply proportional to the square root of the Nhit, yielding an uncertainty of $\sim 3\%$.

2.5 Simulations in SNO+

An adapted version of the Reactor Analysis Toolkit (RAT) is used to simulate events in SNO+ [65]. This library incorporates Geant4 [66] and GLG4sim [67] for detailed MC simulations of the detector geometry, generation of scintillation photons and their propagation, as well as the front-end electronics and data acquisition systems. The reconstruction algorithms explained above are directly incorporated into the framework, alongside an interface with ROOT for data handling.

Chapter 3

In-Situ Bi-Po Decays as a Background Monitoring and Calibration Tool

This chapter explains the use of internal Bi-Po coincidence decays as a background monitoring and calibration source. These Bi-Po events derive from the ^{238}U and ^{232}Th decay chains, and produce a source of α (Po) and β (Bi) events in the detector. They are of crucial importance, facilitating measurements of the internal ^{238}U and ^{232}Th levels in the SNO+ scintillator. The internal radioactive purity level determines the success or failure of the SNO+ physics program: with low purity, too many background events will be present in analysers' energy regions of interest (ROIs), and rare events, such as $0\nu\beta\beta$, will be rendered unobservable. Thus, at each stage of the detector fill, careful analysis of the ^{238}U and ^{232}Th levels was carried out and compared to the target purity levels. Bi-Po coincidences provide a relatively unambiguous signal from which indirect measurements of the purity levels are made.

In addition to their use in background characterisation and monitoring, the α and β events from Bi-Po decays find use in the calibration of the MC models governing the emission of scintillation light. Due to the coincident nature of the Bi-Po decay, it is possible to cleanly extract samples of α

and β interactions in-situ, allowing for calibration without the need for deployed sources.

This chapter first describes the source of Bi-Po events, detailing the ^{238}U and ^{232}Th decay chains. Next, the method of coincidence tagging is explained, before an analysis of the internal ^{238}U and ^{232}Th concentration for the 2.2 g/L full-fill phase is presented. In the second half, the calibration of the light yield and scintillation emission timing, using these same in-situ Bi-Po decays, is presented.

3.1 Primordial Isotopes and the Bi-Po Coincidence Decay

Isotopes with half-lives longer than the age of the earth are known as 'primordial isotopes'. Two important examples for SNO+ are ^{238}U and ^{232}Th , with half-lives of 4.5 $\times 10^9$ and 1.4 $\times 10^{10}$ years, respectively. These isotopes are the progenitors of long decay chains of daughter isotopes, which themselves constitute backgrounds to many SNO+ analyses. The decay chains of interest are shown in figures 3.2 and 3.1. Thus, it is imperative the ^{238}U and ^{232}Th concentration inside the SNO+ liquid scintillator, and surrounding construction materials, is both minimised and well characterised.

Due to their extremely long half-lives, the concentration of these primordials remains approximately constant on the time scale of a human experiment. This constant, slow rate of decay from the progenitor isotopes lead to the decay chains attaining a state of 'secular equilibrium', where the decay rate of each isotope in the chain is equal. This secular equilibrium is only possible when the half-lives of the progenitors are significantly larger than the half-lives of each daughter. In this case, each population in the chain is drip fed from the population above, and given the rate of population decay follows λN, where λ is the decay constant and N the number of isotopes, the population of each level gradually builds until the rate of input is equal to the rate of decay.

This state of secular equilibrium, attained by the internal ^{238}U and ^{232}Th chains, means a measurement of the decay rate of any isotope in the chain is equal to the decay rate of any other isotope in the chain, including the progenitor. Thus, it is possible to infer concentration of ^{238}U and ^{232}Th in the scintillator by extracting a decay rate for a single daughter in the chain.

To perform this measurement, it is necessary to isolate events from the constant background noise in the detector. For these purposes, any event that slips into our selection of daughter isotope decays is defined as an 'accidental'. This is no mean feat, and seemingly impossible: after all, the

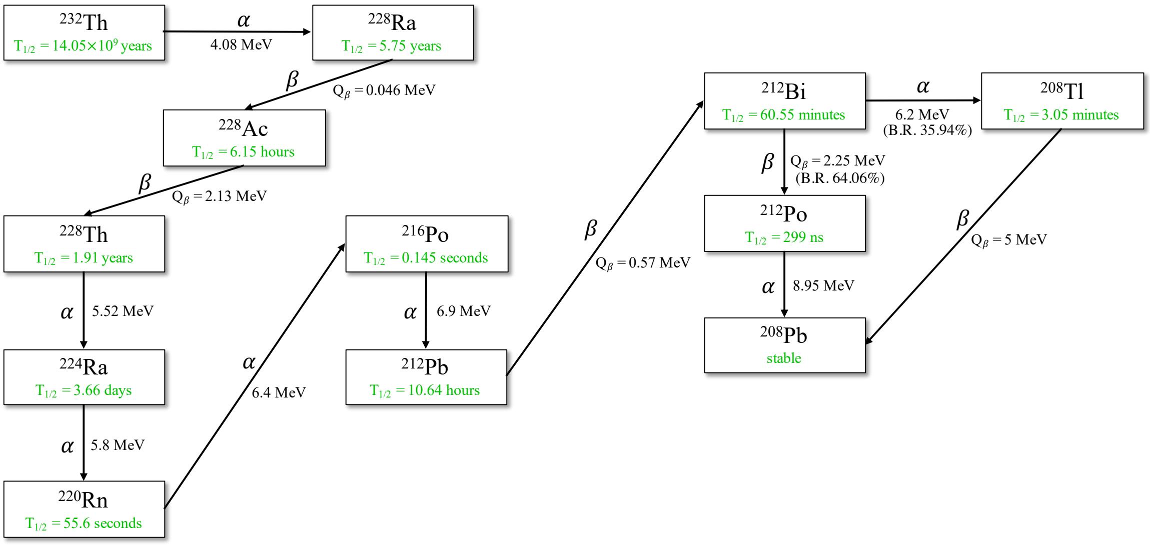

Figure 3.1: Decay chain of ^{232}Th . The rate of tagged coincidences between ^{212}Bi-^{212}Po is used to measure the effective concentration of ^{232}Th in the liquid scintillator. Figure taken from [55].

detector only sees a distribution of PMT hit times and positions resulting from α, β or γ interactions, without information of the specific interaction type.

Fortunately, both ^{238}U and ^{232}Th chains feature ^{214}Bi-^{214}Po or ^{212}Bi-^{212}Po coincidence events, as shown in figures 3.1 and 3.2. Due to the short half-lives of the Po decay, it is possible to tag Bi-Po decays by looking for pairs of events, reconstructing at approximately the same position and within a time window corresponding to a few half-lives of the Po candidate. This coincidence signal serves as an effective fingerprint to isolate Bi-Po events from the other interactions. By placing constraints on the reconstructed energy, based on the known nuclear physics of the respective decays, Bi-Po events are isolated with minimal contamination from accidentals.

By exploiting the coincidence between the Bi and Po decays, the decay rate of the respective decay chains is measured. These rate measurements, given the assumption of secular equilibrium, allow a calculation of the ^{238}U and ^{232}Th concentrations in the SNO+ liquid scintillator.

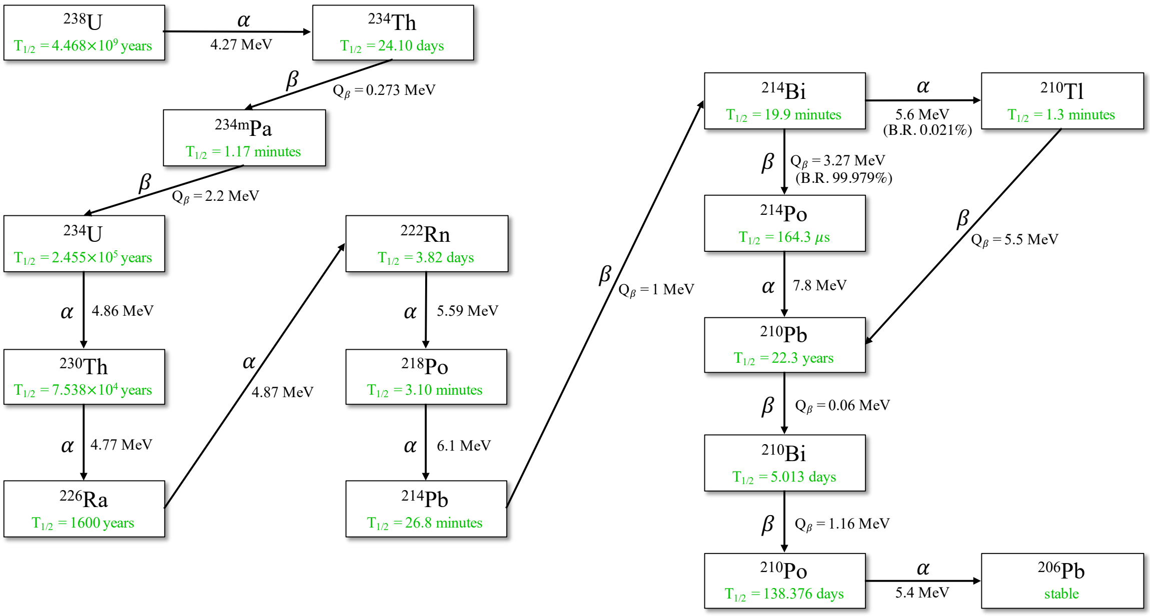

Figure 3.2: Decay chain of ^{238}U . The rate of tagged coincidences between ^{214}Bi-^{214}Po is used to measure the effective concentration of ^{238}U in the liquid scintillator. Figure taken from [55].

3.2 ^{238}U and ^{232}Th Purity Measurements in the 2.2 g/L Phase

Maintaining low radiopurity levels in the detector is essential for rare event searches. The search for $0\nu\beta\beta$ decay relies on defining an energy ROI with as few backgrounds as possible. The fewer backgrounds, the more obvious the $0\nu\beta\beta$ signal will be in an analysis. If the signal is buried beneath orders of magnitude more backgrounds, analyses lose sensitivity to the signal: statistical fluctuations in the backgrounds would more than compensate for any (slight) discrepancy between fitted models and data, regardless of the presence or otherwise of the signal[1].

Higher than expected backgrounds also increase the length of time needed to gather data before statistically significant claims about $0\nu\beta\beta$ decay may be made. In the context of SNO+ competing with other $0\nu\beta\beta$ decay searches, it is doubly important to maintain strict radiopurity requirements.

The target ^{238}U and ^{232}Th radiopurity levels for the SNO+ full-fill and tellurium phase are summarised in table 3.1. Figure 3.3 shows the background counts expected within the $0\nu\beta\beta$ ROI. It can be seen that the ^{238}U and ^{232}Th chain contributions constitute around 25% of the total

---

[1] This is a problem in the ^{8}B-ν_e analysis - see chapter 4

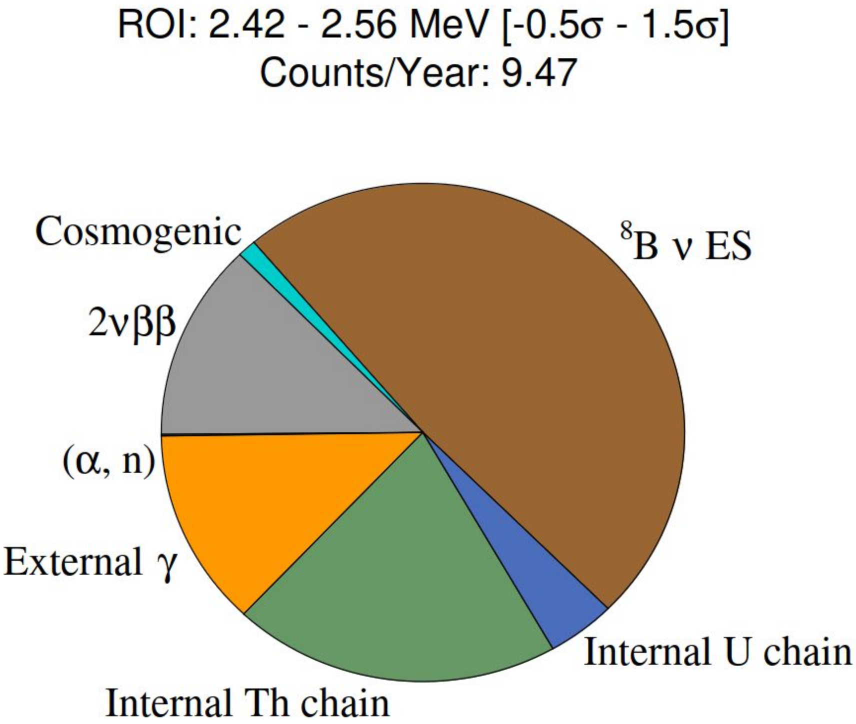

Figure 3.3: Proportions of each major background expected to contribute to the $0\nu\beta\beta$ analysis with 0.5 % ^{130}Te loading, after imposing a 2.42 to 2.56 MeV energy ROI and a 3.3 m FV cut. Taken from [68].

background budget. Thus, they are non-negligible contributors, requiring careful analysis at each stage of the experiment.

Background Type	**Full-Fill Phase Target** [g/g$_{LAB}$]	**Te Phase Target** [g/g$_{LAB}$]
Internal ^{238}U	1.60×10^{-17}	1.29×10^{-15}
Internal ^{232}Th	6.80×10^{-18}	5.57×10^{-16}

Table 3.1: The ^{238}U and ^{232}Th target concentrations for the SNO+ full-fill and tellurium phases [68].

3.2.1 ^{238}U and ^{232}Th Purity in the 2.2 g/L Full-Fill Phase

This section presents measurements of the ^{238}U and ^{232}Th scintillator purity in the full-fill 2.2 g/L PPO phase. As explained above, the measurement uses the rate of coincidence-tagged ^{214}Bi-^{214}Po and ^{212}Bi-^{212}Po , respectively, to infer the concentrations of the progenitors.

Starting with the assumption of secular equilibrium, the rate of change in the population of Bi isotope is:

$$\frac{dN}{dt} = 0 = \lambda N - \lambda_{Bi} N_{Bi} \tag{3.1}$$

Where λ and N are the decay constant and population of the progenitor, and likewise for λ_{Bi}, N_{Bi} for the bismuth. Expressing this in terms of the rate of Bi-Po decays, R_{Bi-Po}:

$$\lambda N = \lambda_{Bi} N_{Bi} = R_{Bi-Po} \tag{3.2}$$

$$N = \frac{R_{Bi-Po}}{\lambda} \tag{3.3}$$

The number of progenitor atoms is now expressed in terms of the mass concentration, σ, of progenitor atoms per gram of scintillator:

$$\frac{\sigma M_{AV} N_A}{m} = \frac{R_{Bi-Po}}{\lambda} \tag{3.4}$$

Where M_{AV} is the mass of scintillator in the AV, m is the atomic weight of the progenitor and N_A is Avogadro's number. The left hand side of equation 3.4 is constructed such that σM_{AV} yields the mass of progenitor in the AV, which is then converted to moles via the relation:

$$moles = \frac{mass}{m} \tag{3.5}$$

The number of moles is then multiplied by Avogadro's number, yielding the number of atoms, N.

The final step is to rearrange equation 3.4 for σ, and account for the imperfect measurement of R_{Bi-Po}. The coincidence tagged rate of Bi-Po events is not equal to the true rate, due to the efficiencies of the analysis cuts, branching ratio of the Bi→Po decay, non-uniform reconstruction performance as a function of radius and the detector's trigger efficiency. Therefore, it is necessary to rely on Monte-Carlo simulations of Bi-Po decays, applying identical selection cuts as in data, and calculating the tagging efficiency via 3.6:

$$\epsilon = \frac{N_{Tag}}{N_{AV}} \tag{3.6}$$

Where N_{Tag} is the number of Bi-Po coincidences tagged in the simulated dataset and N_{AV} is the

Cut	**Selection Criteria**
Prompt Energy	$1.25 \leq E \leq 3.00$ MeV
Delayed Energy	$0.8 \leq E \leq 1.1$ MeV
ΔR	$R < 1$ m
Δt	3690 ns $\leq t \leq$ 1 ms
Combined 6 m FV Efficiency (%)	77.0
Combined 4 m FV Efficiency (%)	23.5

Table 3.2: Table showing optimised ^{214}Bi-^{214}Po coincidence tagging cuts.

total number of simulated pairs in the AV.

By calculating the efficiency this way, the coincidence cuts, reconstruction non-uniformity, branching ratios and trigger efficiency is inherently accounted for. Thus, R_{Bi-Po}/ϵ yields the expected number of Bi-Po coincidences in the AV, both observed and unobserved.

Finally, the equation used to calculate the ^{238}U and ^{232}Th concentration is given by 3.7:

$$\sigma = \frac{R_{Bi-Po} m}{\epsilon \lambda M_{AV} N_A} \tag{3.7}$$

3.2.2 Bi-Po Coincidence Tagging Algorithm

The decays of bismuth and polonium occur within a characteristic time window, determined by the half life of the polonium. These decays will occur at the same true position, since scintillator drift is negligible over the time scale of the coincidence and it is the same nucleus in both events. Thus, the search for Bi-Po coincidences is defined by looking at two sequential events within a Δt and ΔR window. The coincidence tagging algorithm is tuned to identify 'out of window' coincidences, defined as those Bi-Po decays taking place in different trigger windows.

In addition, the reconstructed energy of events is used to further improve the accuracy of the selection. The expected energy spectra of the prompt bismuth β and delayed polonium α decays is used to impose energy cuts on the two candidates.

Algorithm 1 summarises the coincidence tagging algorithm logic, and tables 3.2 and 3.3 contain the applied energy, ΔR and Δt cuts imposed.

The algorithm first searches for the delayed polonium candidates, before looping back in time for the prompt bismuth candidate. This is a choice to improve the runtime of the algorithm: since

Algorithm 1 Pseudocode for Bi-Po coincidence tagging algorithm.

```
Input: A ROOT TTree dataset with many events
Output: A TTree of Bi-Po coincidence tagged events
for each event $e_i$ in the dataset do
    Apply energy cut to determine if $e_i$ is a Polonium Candidate
    if $e_i$ satisfies Polonium energy cut then
        Set Polonium Candidate $P \leftarrow e_i$
        for each preceding event $e_j$ where $e_j$ occurs before $P$ and $\Delta t \leq$ time limit do
            Apply Bismuth energy cut to $e_j$
            if $e_j$ satisfies Bismuth energy cut then
                Apply radial cut $\Delta R$ to $e_j$
                if $e_j$ satisfies radial cut $\Delta R$ then
                    Save $P$ and $e_j$ as a Bi-Po coincidence tagged event
                end if
            end if
        end for
    end if
end for
```

Cut	Selection Criteria
Prompt Energy	$0.60 \leq E \leq 3.00$ MeV
Delayed Energy	$0.85 \leq E \leq 1.3$ MeV
ΔR	$R < 1$ m
Δt	$400 \leq t \leq 800$ ns
Combined 6 m FV Efficiency (%)	8.9
Combined 4 m FV Efficiency (%)	2.8

Table 3.3: Table showing optimised ^{212}Bi-^{212}Po coincidence tagging cuts.

the polonium candidates have a narrow Gaussian α energy profile, there are fewer possible events satisfying the polonium criteria than the bismuth.

These cuts are optimised by reference to Monte-Carlo simulations of Bi-Po coincidence decays, while also taking into account sources of accidental contamination. In order to obtain accurate estimates of the scintillator purity, the Bi-Po selection must be made with low accidental contamination. Therefore, the optimised cuts have a low combined efficiency, but high purity.

Figure 3.4 shows the predicted reconstructed energy spectra of ^{214}Bi-^{214}Po , ^{212}Bi-^{212}Po and the highest rate backgrounds from ^{210}Bi-^{210}Po. This plot motivates the choice of energy selection cuts on both the prompt and delayed candidates. For the ^{214}Po, there is some overlap with the ^{212}Bi-^{212}Po and ^{210}Bi decays. However, since the ^{212}Bi-^{212}Po rate is subdominant to the ^{214}Bi-^{214}Po, and the overlap with ^{210}Bi is minimal, an energy selection cut of 0.8 to 1.1 MeV is applied, yielding high efficiencies and negligible contamination. Similarly, the majority of the ^{214}Bi decays occur in a region free of backgrounds, and thus a high efficiency energy selection cut of 1.25 to 3.0 MeV is made to capture most of the spectrum whilst avoiding the ^{212}Po and ^{219}Rn.

Figure 3.4 shows ^{212}Bi-^{212}Po decays occur in a more congested region of the energy spectrum. Despite this, the ^{212}Po selection may still be made cleanly by being mindful of the higher rate ^{214}Po and placing a lower energy cut at 0.85 MeV to limit contamination. For the ^{212}Bi selection, an energy range of 0.6 to 3.0 MeV was chosen. The lower limit was selected as a trade off between maintaining a high selection efficiency and limiting contamination from ^{210}Bi-^{210}Po. It is also worth noting that the bismuth energy selections have far less impact on the overall accidental contamination than the Δt and ΔR cuts, and so broader values to maximise selection efficiency was acceptable.

The radial cut was carried over from previous studies performed by Josh W. [69], where the limit was set to preserve 99% of simulated Bi-Po pairs. Despite the fact the Bi-Po decays technically occur at the same position, the reconstruction algorithms, due to finite resolution, 'blur' the positions of each decay away from the true position. Even in the case of two identical events occurring at the same position, the reconstruction algorithms will still return different fit results, according to the resolutions given in figure 2.10. This is a consequence of the probabilistic nature of scintillation emission and PMT response, which is dominated by the scintillator's emission time profile and the PMT TTS. There are further complications when resolving Bi-Po decays, given the fact the interaction types are not identical. The reconstruction algorithms, tuned with respect to electron simulations, will have an inherently different resolution for the polonium α decays then the bismuth β decays. In addition, the bismuth β decay is overwhelmingly accompanied by γ de-excitations, of various energies. These γs distort the energy deposition further from a point-like electron event assumed by the reconstruction. Thus, a pragmatic approach was adopted, with MC distributions of Bi-Po decays informing the choice of a 1 m ΔR cut.

Similarly, the Δt cut values were previously determined by Josh W. For the ^{214}Bi-^{214}Po, a range of 3960 ns to 1 ms was chosen, with the lower limit set to exclude Bi-Po212 with $< 2\%$ sacrifice. For ^{212}Bi-^{212}Po, the polonium has a significantly shorter half-life, meaning most Bi-Po212 decays occur within a single event window. In order to remove in-window decays, a lower limit of 400 ns was set.

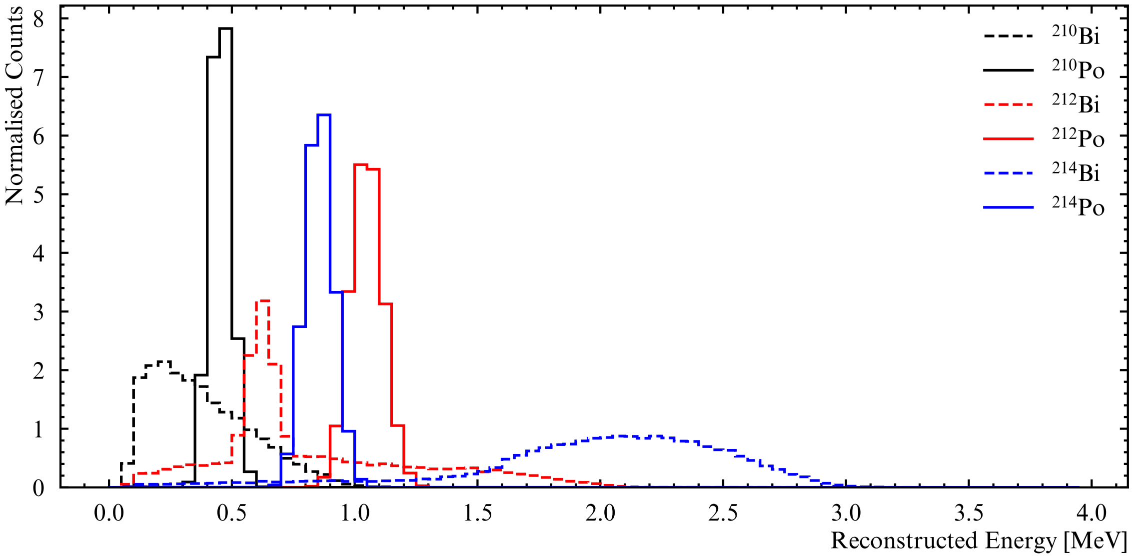

Figure 3.4: Energy spectra of Bi-Po210, Bi-Po212 and Bi-Po214 decays.

This leads to a sacrifice of $\sim$ 50%, but is unavoidable to maintain a clean out of window selection.

3.2.3 Purity Measurement Results

Finally, the out of window Bi-Po tagging algorithm detailed above was used to measure the effective ^{238}U and ^{232}Th concentration in the 2.2 g/L PPO phase. This yielded the 'effective' concentration, as it assumes secular equilibrium in the decay chain throughout the measurement period. In practice, this assumption was occasionally broken by periodic radon ingress, caused by planned maintenance on deck or changes in the external cavity water temperature. Thus, the measurements here constitute an under-estimate for the intrinsic purity of the scintillator.

Runs were carefully selected with the most stringent data quality criteria in order to maximise position and energy reconstruction performance (so-called 'gold-list' runs; see table 3.4). Furthermore, detailed consideration of the impact of muons, muon followers and other high Nhit events on the measurement livetime was performed. Despite the muon rate being as low as $\sim$ 3 hr^{-1} at the detector, when they do occur, they are capable of fully saturating the detector's PMTs for multiple event windows. In addition, their passage through the detector can lead to short lived cosmogenic backgrounds, dubbed muon follower events. These burst of followers are capable of producing many

Criteria	**Cut Values**
Run Type	Physics
Run Length	$t \geq 30$ mins
Crates On	ALL
PMT Coverage	coverage > 70%
Panel Coverage	coverage > 80%
Outward Facing PMT Status	ON

Table 3.4: Gold list run selection criteria, as determined by the SNO+ run selection group.

spurious Bi-Po coincidences, and must be removed. In addition to muons and their followers, there are other high Nhit events that may produce accidentals. Since early 2022, high Nhit event hotspots have formed in the neck of the detector. While these high Nhit neck events can be efficiently removed with a z position cut, they produce so much light in the detector that residual photons frequently produce retrigger events. These retriggers are the real danger, and primarily reconstruct around the centre of the AV. At the time of writing, their source remains unknown, however possible light leakage via the external laser calibration fiber system is being investigated.

With no muon track fitter available in the scintillator phase, the easiest method to maintain data cleanliness was to remove high Nhit events and impose a deadtime veto (20 s), within which no Bi-Po coincidences were tagged. This deadtime is a conservative cut, allowing retriggers and cosmogenic activity to decay following a neck hotspot or muon, respectively. This deadtime was subsequently subtracted from the reported run length to maintain accurate live-time calculations, from which the Bi-Po event rates were derived.

Figures 3.5 to 3.9 show the Δt, ΔR and energy distributions for the tagged populations of ^{214}Bi-^{214}Po and ^{212}Bi-^{212}Po. These distributions allow the verification of the purity of the selected Bi-Po events by comparison of the ΔR and energy spectra to pure Bi-Po simulations, as well as fitting the observed Δt distribution according to equation 3.8.

$$f(t) = Ae^{-\lambda t} + B \tag{3.8}$$

The fitted constant A is a normalisation, and B parametrises a flat accidental contribution to the distribution. Thus, the Δt distribution allows a comparison between the fitted decay constant and the accepted value of ^{214}Po, and estimate the fractional accidental counts within the sample, C, via equation 3.9.

$$C = \frac{BN^{bins}}{N^{Tags}} \tag{3.9}$$

These figures show broad agreement between the shapes of the fitted and observed ΔR distributions for both isotopes. Additionally, the fitted decay constant for ^{214}Po is 4.01 ms^{-1}, within 3σ of the accepted value of 4.22 ms^{-1}, and the accidental contamination term, B, consistent with 0.

Unfortunately, such fits to the Bi-Po212 Δt distribution are not possible. This is due to the current trigger system (see 2.3), particularly the automatic retrigger. For each N100 GT issued, the detector automatically issues another GT 460 ns after the primary trigger. By unhappy coincidence, the typical decay times of out of window ^{212}Po occurs within this automatic retrigger event window, leading to a Δt distribution with only 1 populated bin at 460 ns. Thus, no fits may be performed. Despite this, the agreement between the ΔR and energy spectra of both isotopes lend confidence to the purity of the selected populations.

Figures 3.10 and 3.11 show the effective ^{238}U and ^{232}Th concentrations, respectively, between 1st June 2022 to the 10th March 2023, with the ^{238}U plot binned by day and the lower rate ^{232}Th concentration binned by month. Each graph shows the result for two difference FVs: the entire AV (blue) and a restricted 4 m FV (black).

Both graphs show a rapid exponential decay (note log scale on y-axis) in rate following the end of the PPO fill. This is simply ^{222}Rn and ^{212}Pb ingress, brought in during the fill and subsequent recirculation and mixing procedures, decaying away. From the 1st June 2022, no further large-scale recirculation of the AV scintillator were performed and the initial equilibrium breaking ingress had decayed away. From this point, the decay chains were assumed to have reached secular equilibrium for the gold-list data selection.

Taking the data from the 1st June 2022, a weighted average concentration was calculated, with the weights given by the error on each point. Table 3.5 shows the weighted average concentrations

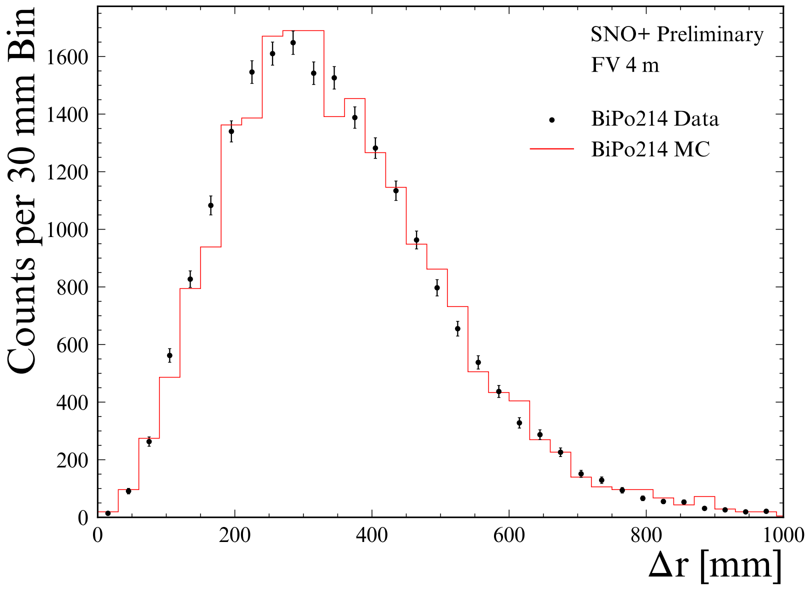

Figure 3.5: ^{214}Bi-^{214}Po ΔR

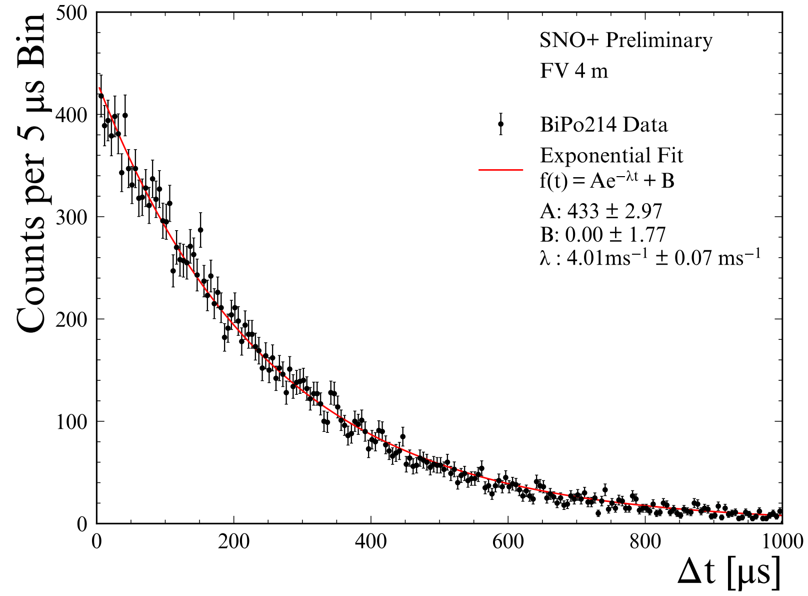

Figure 3.6: ^{214}Bi-^{214}Po Δt

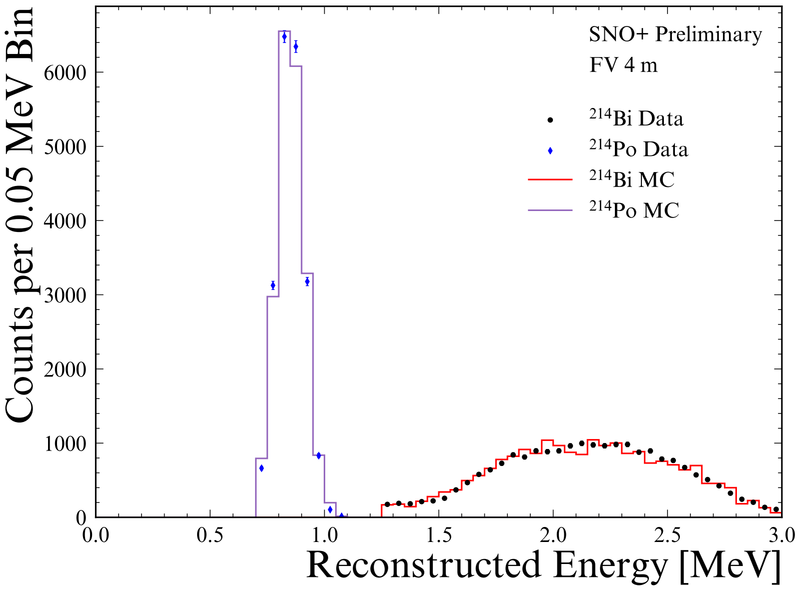

Figure 3.7: ^{214}Bi-^{214}Po Energy Spectra

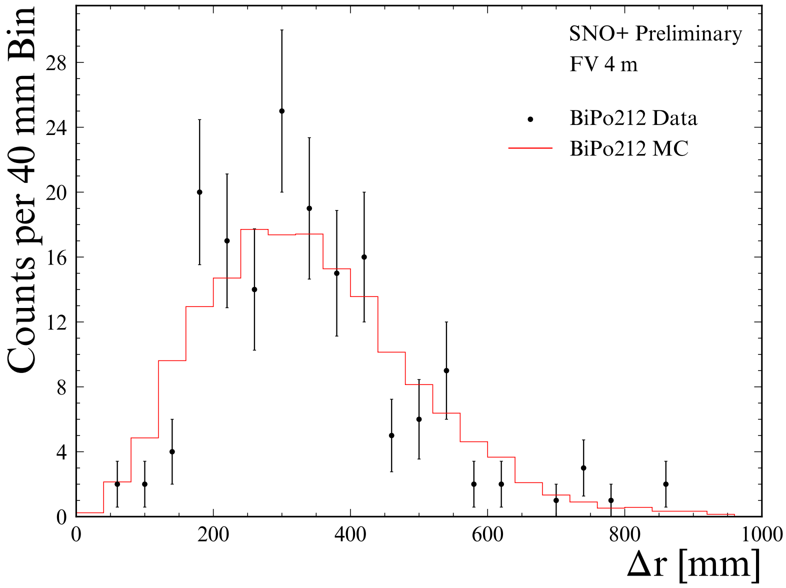

Figure 3.8: ^{212}Bi-^{212}Po ΔR

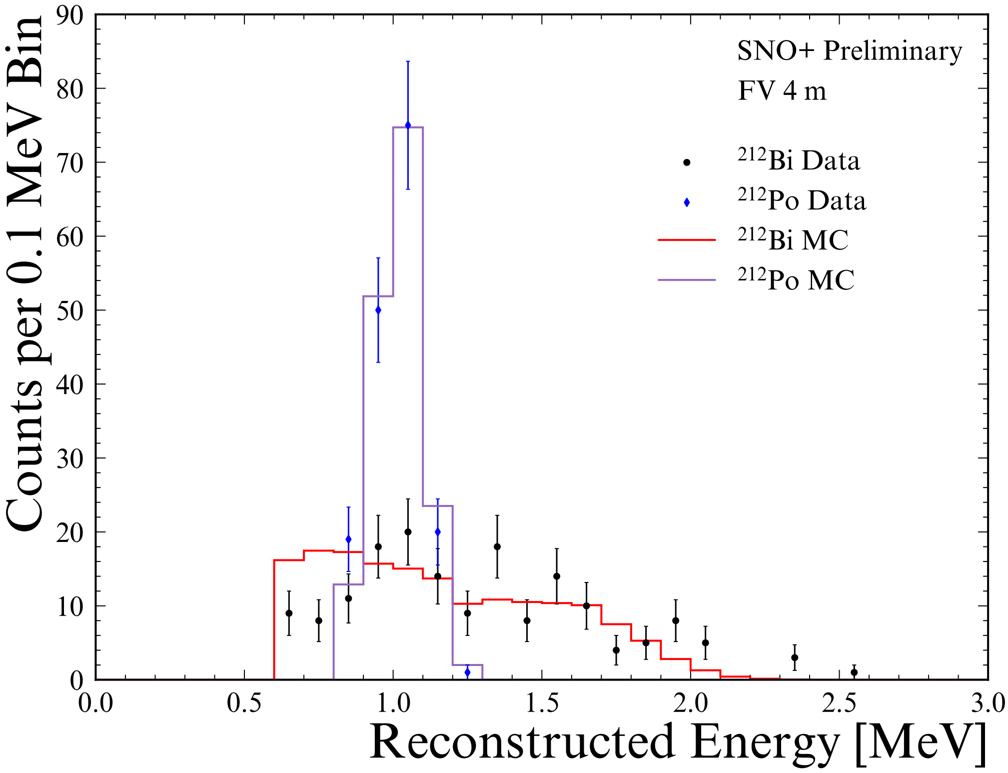

Figure 3.9: Comparison between Bi-Po212 energy spectra for MC (red) and data (black).

for both isotopes and FVs. By way of comparison to table 3.1, it can be seen that the measured concentrations for both isotopes are higher than the target concentrations for the full-fill phase. Despite this, the levels remain below the tellurium phase targets. While not ideal, the current concentrations of ^{238}U and ^{232}Th are sufficient for the $0\nu\beta\beta$ phase, providing the concentrations remain subdominant to the ^{8}B-ν_e elastic scattering rate ($\sim 10^{-15}$ level). More stringent purity targets were made on the full-fill phase, given results from experiments such as Borexino, and the relative ease of purifying the undoped liquid scintillator, relative to the other scintillator components. Specifically, the tellurium itself is expected to have significantly higher levels of ^{238}U/^{232}Th, even after purification. Therefore, the lower the purity levels achieved in the full-fill phase, the greater the tolerance to higher than expected levels in subsequent phases.

One final feature of the data to note is the inhomogeneity with volume most obvious in the ^{214}Bi-

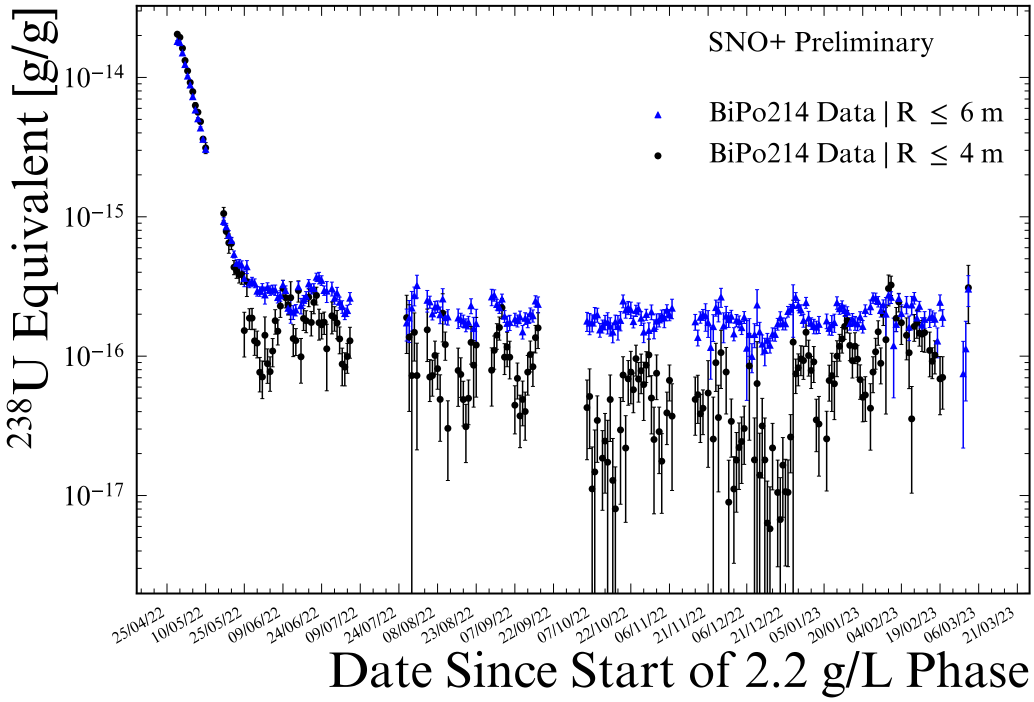

Figure 3.10: Effective ^{238}U concentration of the 2.2 g/L PPO phase from end of PPO fill campaign (29th April 2022) to 10th March 2023. Errors are statistical only.

Chain	**6 m Measured Concentration** [g / g_{LAB}]	**4 m Measured Concentration** [g / g_{LAB}]
^{238}U	$2.00 \pm 0.02 \times 10^{-16}$	$4.32 \pm 0.14 \times 10^{-17}$
^{232}Th	$1.27 \pm 0.07 \times 10^{-16}$	$5.29 \pm 0.76 \times 10^{-17}$

Table 3.5: Measured ^{238}U and ^{232}Th concentrations for 6 m and 4 m FV in the full-fill phase.

^{214}Po rate. Given a uniformly mixed scintillator in perfect thermal equilibrium, the concentration of ^{238}U would be homogenous across the full volume. Thus, figure 3.10 should show the same average for both the 4 m and 6 m FVs. The deviation from uniformity is due to temperature driven convection currents around the AV edges [70]. Slight changes (< 1 degree) in the temperature of the external water drives this convection. By examining the spatial distribution of ^{214}Bi-^{214}Po tags during periods with above and below average cavity water deviations, this convection current has

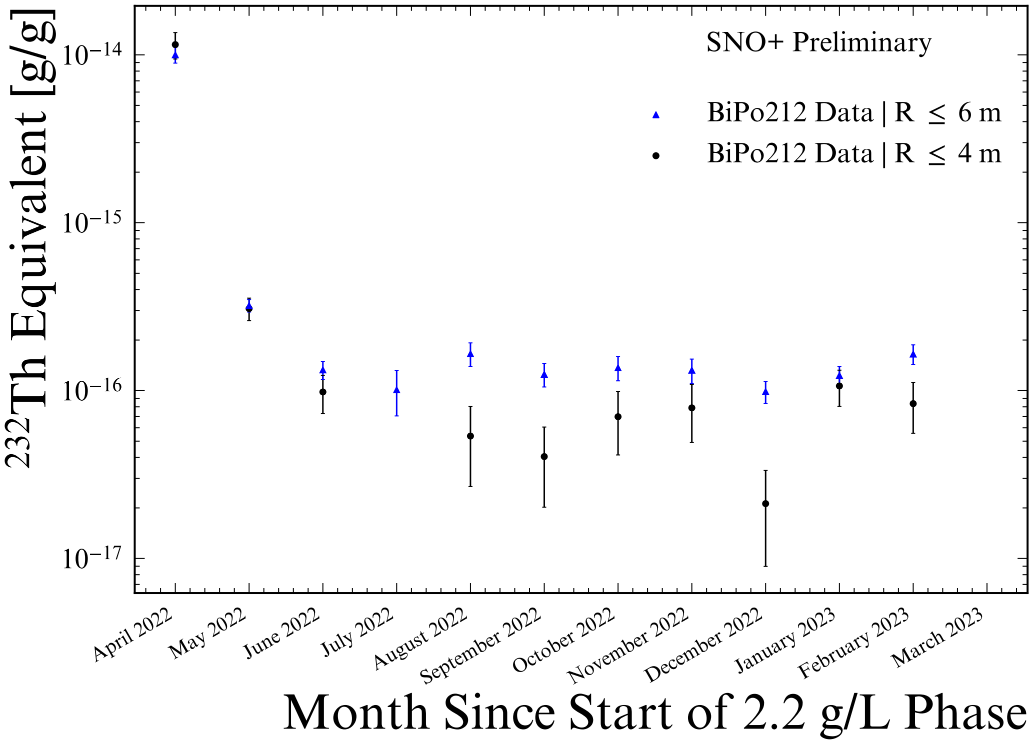

Figure 3.11: Effective ^{232}Th concentration of the 2.2 g/L PPO phase from end of PPO fill campaign (29th April 2022) to 10th March 2023. Errors are statistical only.

been understood to feature 2 different modes:

1. $T_{Cavity} < T_{AV}$: 'Normal' mode. Scintillator in the AV cools and falls along the edges, with displacing scintillator rising through the centre.
2. $T_{Cavity} > T_{AV}$: 'Diffusion' mode. Scintillator warms at the AV edges and diffuses into the bulk.

Figures 3.12a and 3.12b show a typical $\rho^2 - Z$ spatial distribution plot of Bi-Po214 events for each mode.

During normal operations, the convection current effectively insulate the inner regions of the AV from radon ingress via the neck. This would-be secular-equilibrium breaking ingress does not

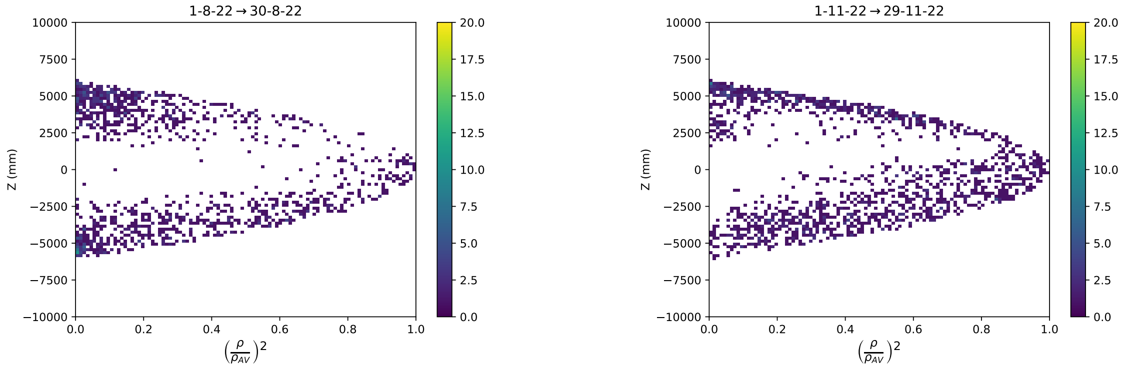

(a) ^{212}Bi-^{212}Po coincidence tags following the normal convection pattern along the AV wall.

(b) ^{214}Bi-^{214}Po coincidence tags diffusing into the bulk from the AV walls.

Figure 3.12: Comparison of normal (left) and diffusion (right) convection modes observed with ^{214}Bi-^{214}Po coincidence tagging.

diffuse throughout the bulk, but sinks along the AV walls while decaying. Thus, the internal ^{238}U concentration for the scintillator is best captured by the 4 m FV estimate, which is less influenced by chain-breaking background ingresses. The normal convection mode brings an unexpected benefit to physics analyses: the majority of analysers impose FV cuts to isolate the inner 5 m or so, in order to reject external backgrounds and benefit from better tuned reconstruction. Thus, they remove the excess internal ^{214}Bi-^{214}Po backgrounds along the AV walls.

3.3 Calibration of the Full-Fill Phase Scintillator Model Using In-Situ Bi-Po214

This section details the scintillator model calibrations using in-situ Bi-Po214, which was the only available internal calibration method for the full-fill phase. As discussed in section 2.4, accurately determined probability distribution functions (PDFs) of the observed time residual distributions are a central component to position reconstruction. Thus, these calibrations were of crucial importance to all SNO+ physics analyses, as they facilitated accurate MC simulations of the scintillator response.

The SNO+ scintillator model consists of a series of parameters governing the scintillation physics and light propagation processes detailed in chapter 1.3.3, for β^- and α interactions. The specific calibrations amenable to Bi-Po tuning are:

- Intrinsic light yield, L, in units of photons per MeV deposited
- Birks' constant, K_B, which determines the degree of ionisation quenching experienced by α particle interactions
- Scintillation emission time constants of the empirical model given by equation 3.11.
- Normalisations of wavelength dependent absorption lengths.

These parameters are tunable with in-situ ^{214}Bi-^{214}Po, since these events present a large population of separable β^- and α interactions occurring within the scintillator bulk. This calibration method has the additional advantage of being possible without deploying sources and disturbing the quiet data taking operation of the detector. Thus, physics uptime is maximised and the stringent radiopurity requirements are maintained throughout the calibration.

This thesis will focus on the initial calibration of the light yield, Birks' constant and a detailed examination of the optimisation of the scintillation emission timing model. The wavelength dependent absorption models were extensively tuned by S. Riccetto [71], using the tagged population of ^{214}Bi-^{214}Po provided by the author.

Previous phases of SNO+ were calibrated using ^{214}Bi-^{214}Po in a similar manner [69], [72]. It was found the timing model was easier to calibrate once the light yield was fixed, due to the dependence of the observed timing on the light yield. With higher light yields, the proportion of multi-photon hits recorded is increased. These multi-photon hits lead to faster threshold crossing by PMTs. This causes observed time residuals to appear more peaked, and therefore changes the best-fit emission time parameters. Thus, the full-fill model calibrations proceeded in the following order:

1. Light Yield
2. Birks' Constant
3. Absorption Model
4. Emission Timing Model

In each case, complications arising from AV reflections were excluded by performing the calibrations in a restricted 4 m FV.

3.3.1 Dataset

The calibration was carried out using tagged ^{214}Bi-^{214}Po events within a 4 m FV, using the pure Bi-Po selection cuts given in table 3.2. A total of 501 runs passing the gold list criteria of table 3.4 were selected, between 20th May and 14th June 2022. This yielded a livetime of 20.8 days and a population of 854 coincidence events.

3.3.2 Light Yield

First, the raw light yield, L_γ (photons per MeV deposited), was tuned. In RAT, the number of scintillation photons generated is obtained by sampling a Poisson distribution, with mean N_γ given by:

$$N_\gamma = \Delta E_q L_\gamma \theta_i \tag{3.10}$$

Where ΔE_q is the energy transferred for this distance step in the simulation, obtained from Birks' law (equation 1.35) and θ_i is a particle-specific additional quenching factor ($\theta_i = 1$ for electrons). These photons are uniformly spread across the length of the step, and given uniformly random direction vectors in a 4π solid angle, with wavelengths sampled from the appropriate emission spectrum.

The population of ^{214}Bi events falling within a 4 m FV was used to tune the light yield, L_γ, from $L_\gamma = 10800$ to $L_\gamma = 14000$ photons per MeV. The updated L_γ was obtained using the ratio of the median NHits between data and MC as a scaling factor. Figure 3.13 shows the distributions before and after the correction was applied. Following this update, the simulation accurately matches the observed light yield for electron interactions in the scintillator.

3.3.3 Birks' Constant

After calibrating the light yield parameter, the response of the scintillator to α interactions was tuned. As described in section 1.3.4, α-particles undergo ionisation quenching. This leads to fewer scintillation photons being emitted as compared to an equivalent energy electron interaction.

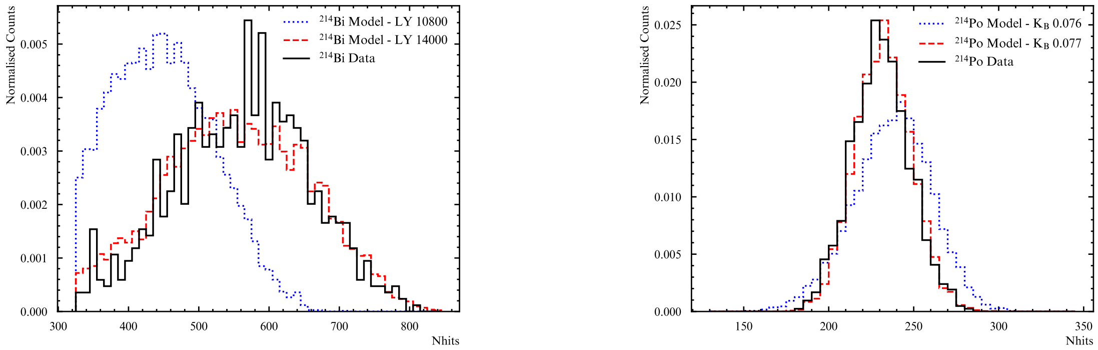

(a) ^{214}Bi MC before (blue) and after (red) light yield tuning, compared to tagged ^{214}Bi data (black), within a 4 m FV.

(b) ^{214}Po MC before (blue) and after (red) light yield tuning, compared to tagged ^{214}Po data (black), within a 4 m FV.

Figure 3.13: Comparison of nhit distributions for ^{214}Bi (left) and ^{214}Po (right) before and after light yield and Birks' constant tuning, respectively.

This quenching factor is parametrised by the simulation in terms of θ in equation 3.10. The tuning proceeded analogously to the light yield parameter: a population of ^{214}Po events within a 4 m FV were compared to a simulated population. The ratio of the mean Nhits between data and MC was used as a correction factor to the θ parameter, leading to an update from $\theta = 0.076$ to $\theta = 0.077$. Figures 3.13a and 3.13b show the impact of this calibration, which led to excellent agreement.

After this calibration of the scintillation light yield for β^- (^{214}Bi) and α (^{214}Po) interactions, the energy reconstruction algorithms were recoordinated. Thus, this work directly facilitated reliable energy reconstruction in the 2.2 g/L scintillator phase.

3.4 Emission Time Model

The remainder of this chapter concerns the tuning of the scintillation emission time model. An accurate scintillation emission time model is of paramount importance to all SNO+ physics analyses: without well calibrated PDFs, the position reconstruction (and by extension, the energy reconstruction) will not yield accurate results, as it depends wholly on knowledge of the emission time PDF for use in the maximum likelihood vertex fit. Furthermore, time-based classifiers (e.g. the multisite classifier developed in chapter 4.2 and α/β discriminators) are unable to be relied upon.

Given the calibrated light yield and α quenching factor, the time constants and weightings of the empirical emission time model was calibrated. Firstly, the method used to calibrate previous phases of SNO+ is described. Afterwards, a novel improvement to this calibration procedure, using Bayesian Optimisation with Gaussian Processes, is explained. Finally, the new method's result is compared to those obtained by the traditional technique.

The empirical scintillation emission time model is restated here in equation 3.11. The objectives of this calibration are to optimise the model decay constants, τ_i, and respective weightings, A_i.

$$f(t) = \sum_{i=1}^{4} A_i \frac{e^{-t/t_i} - e^{-t/t_r}}{t_i - t_r} \tag{3.11}$$

As described in section 1.3.4, β and α interactions lead to different emission time profiles in liquid scintillator. Thus, two different sets of decay constants and respective amplitudes needed to be tuned in the emission time model.

The traditional method of tuning the empirical emission time model is to proceed through the following steps:

1. Use benchtop measurements to extract parameter estimates and define restricted domain for fine-tuning model.
2. Use grid search to simulate ^{214}Bi-^{214}Po events with different combinations of timing parameters close to those found on the benchtop. Compare the resulting time residual distributions in simulation to those observed in tagged Bi-Po214 data.
3. Using the best-fit parameters, regenerate the position reconstruction PDFs ('recoordinating' the fitters).
4. Reprocess the tagged Bi-Po214 data with updated fitters.

If the reprocessing of the data leads to significant changes in the data-MC agreement, the process must be repeated iteratively until convergence is met. Fortunately, this was not found to be necessary for the 2.2 g/L tuning.

In each case, the time residual distributions, as defined in equation 2.2, were used to compare the agreement between the MC and data. This agreement was quantified by computing the χ^2 statistic between the two:

$$\chi^2 = \sum_{bins} \frac{(O_i - M_i)^2}{M_i} \tag{3.12}$$

Where O_i are the observed counts in bin i in the data and M_i is the same in MC. To account for differences in the number of hits between the data and MC, the MC time residuals were normalised to the total counts in the data.

Benchtop Measurement Apparatus

Using samples of scintillator taken from the AV, the scintillation emission time parameters were measured on the benchtop, using the apparatus shown in figure 3.14. While the setup was installed previously for slow scintillator studies [73], the measurements and subsequent fits are the original work of the author.

Three PMTs were installed within a sealed dark box: a r9880U 'trigger' PMT, a r6594 'charge collection' PMT and a second r9880U 'measurement' PMT. A vial of deoxygenated 2.2 g/L scintillator was placed on the plinth, with the ^{90}Sr electron source placed facing away from the measurement PMT. This facing away configuration was chosen to reduce the impact of Cerenkov light on the measured scintillation timing parameters.

^{90}Sr β-decays to ^{90}Y (Q-value 545.9 keV), which subsequently β-decays to ^{90}Zr (Q-value 2278.5 keV), which is stable. The approximate energy loss of an electron through the 1 mm glass vial wall and 2 mm diameter scintillating fibre is around 1 MeV [74], meaning only the β particle from the ^{90}Y interacts in the scintillator.

The time delay between the signals from the trigger and measurement PMTs created the timing profile of the sample. The charge collection PMT was used to estimate the energy deposited in the scintillator, allowing cuts to be placed that remove 'tail' events (events with low energy deposits that have poorly resolved time profiles). Furthermore, the measurement PMT was placed within a dark box with an adjustable iris, in order to reduce the occupancy to $\sim 10\%$. This was done to limit the impact of multi-photon hits on the measurement PMT, which at 10% occupancy were considered negligible. The masking box also ensured hits travelled directly from the scintillator, removing the possibility of scattered or reflected light triggering the measurement PMT.

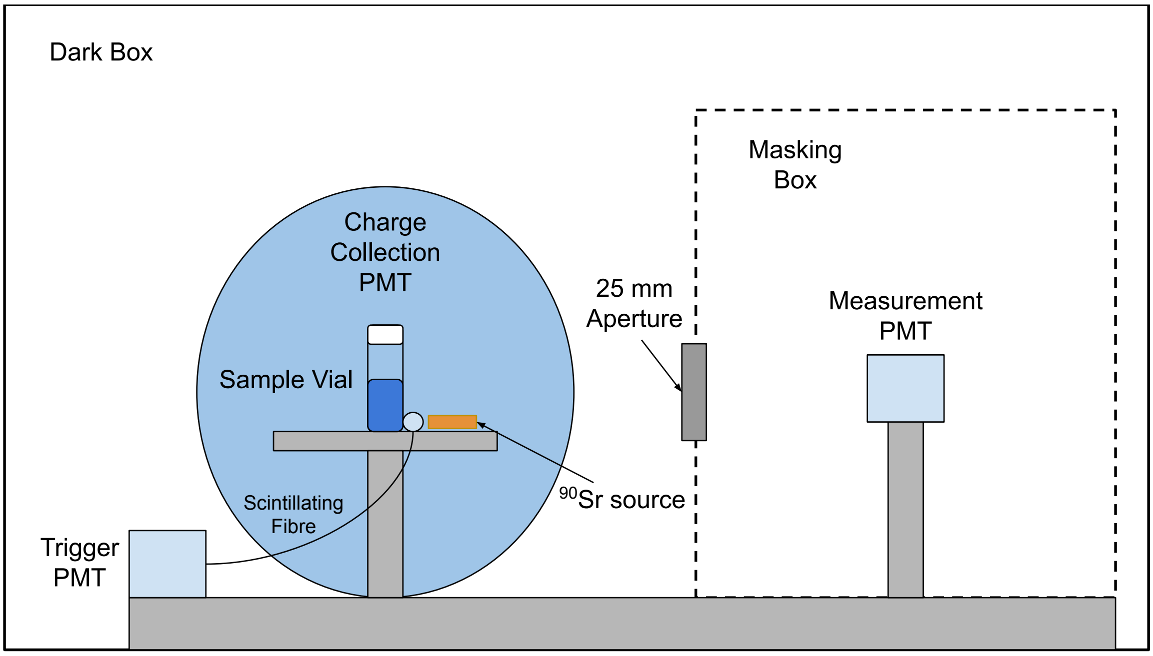

Figure 3.14: Side view of the benchtop timing setup. A ^{90}Sr source is placed facing away from the measurement PMT, removing Cerenkov light from the measurement. The Δt between the trigger and measurement PMT builds the emission time distribution, while the charge collection PMT allows a calibration of a low energy cut, removing poorly resolved events. A masking box was placed around the measurement PMT to reduce the probability of multiple photon hits.

The time resolution of the measurement apparatus was directly built into the emission time parameter fits via convolution with the impulse response function (IRF). A sample of distilled water was used in place of scintillator, and a time profile built from purely Cherenkov light was produced. The resulting time profile, shown in figure 3.15 shows the response of the apparatus to an instantaneous impulse of light.

A Gaussian fit to the IRF gave a time resolution of approximately 390 ps, however, the distribution is not perfectly Gaussian. For this reason, the IRF was directly convolved with the measured scintillator time profiles for each measurement fit.

Equation 3.14 shows the objective function fitted to the Δt time profiles obtained from the 2.2 g/L AV sample.

$$f(t) = (1 - F_{Cher}) \left(\sum_i A_i \frac{e^{-t/\tau_i} - e^{-t/\tau_r}}{\tau_i - \tau_r} \right) + F_{Cher}\delta(t) \tag{3.13}$$

$$F(t) = f(t)(t_0 - t) * IRF(t) \tag{3.14}$$

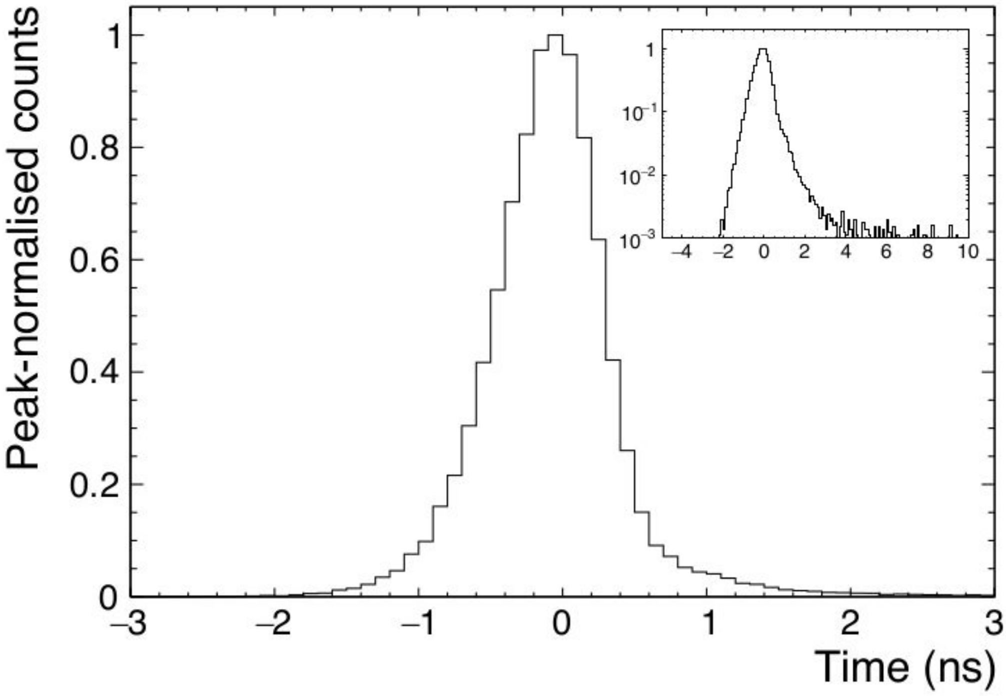

Figure 3.15: Impulse response function obtained from running the measurement on a water sample. This allows an estimate of the intrinsic resolution of the apparatus from a Gaussian fit standard deviation ($\sim$ 390 ps). Taken from [27].

Where F_{Cher} is the fraction of light produced by Cerenkov processes (consistent with zero in the facing away configuration), τ_i and A_i are the previously explained scintillation emission decay constants and weightings for the empirical model, and t_0 is a constant offset to account for the time of flight between the sample and measurement PMT.

ROOT's Minuit optimiser [75] was used to perform a fit by minimising the negative log-likelihood. The results for the β timing fits with the 2.2 g/L AV sample is shown in figure 3.16.

The best fit performance was obtained with a double exponential model, with decay constants of 5.2 and 15.7 ns, respectively, and a rise time of 1.2 ns. The Cherenkov component was on the $\sim$ 0.5 % level, as expected for a facing away source configuration.

This result is unsurprising: with a PPO concentration of 2.2 g/L, non-radiative coupling to the fluor dominates and LAB emission is negligible. Additionally, due to the short path lengths of

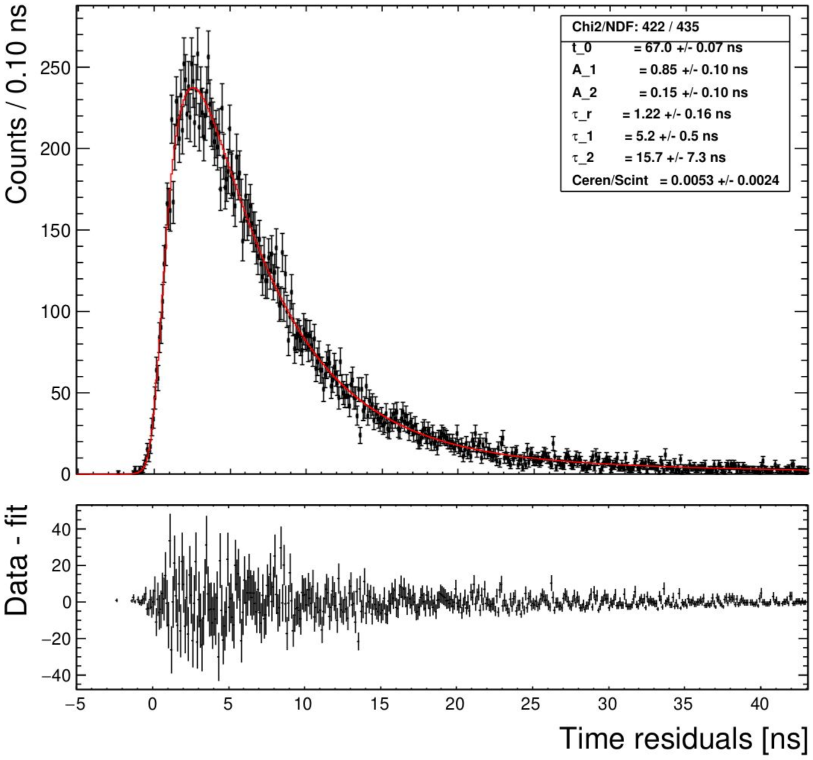

Figure 3.16: Benchtop β timing fit result using a double exponential model. Decay constants and amplitudes were used as seeds for the grid search tuning of the 2.2 g/L optics model.

scintillation light through the sample vial ($\sim$ mm scale), the impact of absorption, re-emission and scattering is not visible. Thus, the simple model is able to capture the characteristic emission time profile of the PPO well.

The β benchtop timing result was used as a seed for the calibration of the SNO+ emission time model. As no α sources were on hand, similar measurements for the α timing could not be performed. However, since the α/β timing differences are not extreme, the same seed was used for the α timing calibrations.

Timing Calibration Results

Using the process outlined in section 3.2.3, the full-fill emission time model was calibrated. The final tuned parameters for α and β interactions are summarised in table 3.6. Figures 3.17 and 3.18 show a comparison between the time residuals obtained from simulation and ^{214}Bi-^{214}Po data within a 4

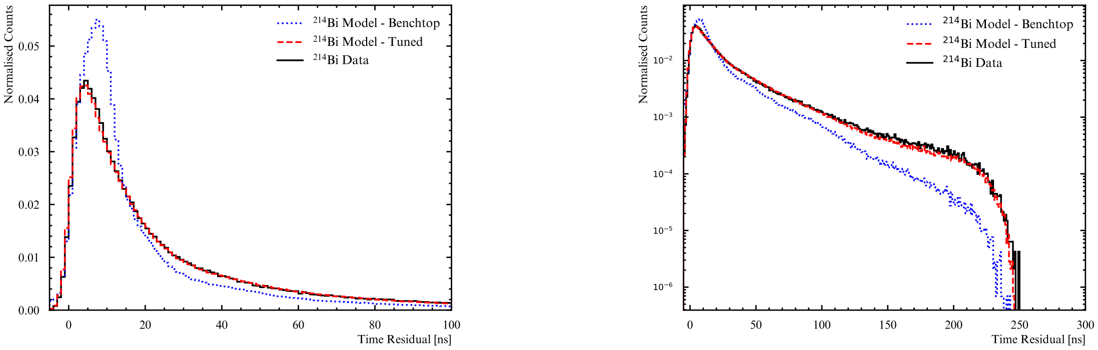

(a) ^{214}Bi MC before (blue) and after (red) emission time tuning, compared to tagged ^{214}Bi data (black), within a 4 m FV.

(b) ^{214}Po MC before (blue) and after (red) emission time tuning, compared to tagged ^{214}Po data (black), within a 4 m FV.

Figure 3.17: Comparison of time residual distributions for ^{214}Bi for peak (left) and tail (right) regions, before and after emission time tuning.

m FV, with the MC using the benchtop measurements (dashed) or the tuned constants (solid).

As figures 3.17 and 3.18 show, there is a significant discrepancy between the data-MC agreement with the benchtop parameters and those following the tuning process. This owes to the more complicated environment surrounding events in the detector vs those in the benchtop sample vial. Whereas the benchtop measurement was performed to suppress as many other optical processes as possible, the detector data time residuals are the result of the convolution of scattering, absorption, re-emission, quenching, AV reflection and the PMT/electronics resolutions. At typical path lengths,

Parameter	**Benchtop Result**	**β Grid Search Result**	**Bayesian Optimiser β Result**	**α Grid Search Result**
t_1	5.2 ns	5.0 ns	5.5	4.1
t_2	15.7 ns	24.5 ns	41.0	21.0
t_3		399.0 ns	60.0	84.0
t_4			500.0	197.0
t_r	1.22 ns	0.85 ns	1.22	0.85
A_1	0.850	0.656	0.755	0.523
A_2	0.150	0.252	0.145	0.303
A_3		0.092	0.05	0.070
A_4			0.05	0.104

Table 3.6: Comparison of benchtop and grid search scintillation emission time model parameters for the full-fill phase.

the light passes through ∼ meters of scintillator and cavity water, as well as 5.5 cm of acrylic. No longer are these other optical processes negligible, and greatly impact the observed time residual distributions.

The failure of the benchtop measurements to produce accurate time residual distributions compared to the data is due to mismodelling in the extinction and scattering lengths, as well as discrepancies in the electronics/detector response simulation between MC and data.

Given this fact, there are two options: a) fix the emission timing to the measured benchtop values and tune the other parameters in the model or b) change the emission timing constants to account for the convolved impacts of other optical processes. Since there was no reliable method to independently measure the scintillator's extinction lengths, a pragmatic choice was made to fine tune the emission timing model to compensate for this mismodelling, until such a time as the SNO+ optical fibre calibration system is fully operational.

Despite the success of this tuning in accurately reconstructing the observed time residual distributions, the time necessary to produce calibrated PDFs such as in figure 3.17 was extreme. The presented tuning took months to complete, with literally billions of events simulated. Since this calibration is necessary whenever the composition of the scintillator is altered, and given no reliable analysis may be performed until the timing is optimised, it was mission critical to develop a better method. The next section describes this improved, and fully automated, emission tuning method.

3.5 Calibration using Bayesian Optimisation with Gaussian Processes

This next section details an attempt to improve the tuning process for the scintillation emission timing model. As shown above, the grid search method over each parameter is extremely costly, both in terms of wall time and the quantity of Monte-Carlo produced. It is therefore desirable to develop methods that more efficiently tune the model, whilst preserving the same strong agreement between data and Monte-Carlo.

To this end, a Bayesian optimisation algorithm was developed [76]. Bayesian optimisation is a technique widely used for tuning the hyperparameters in neural networks, which suffer many of

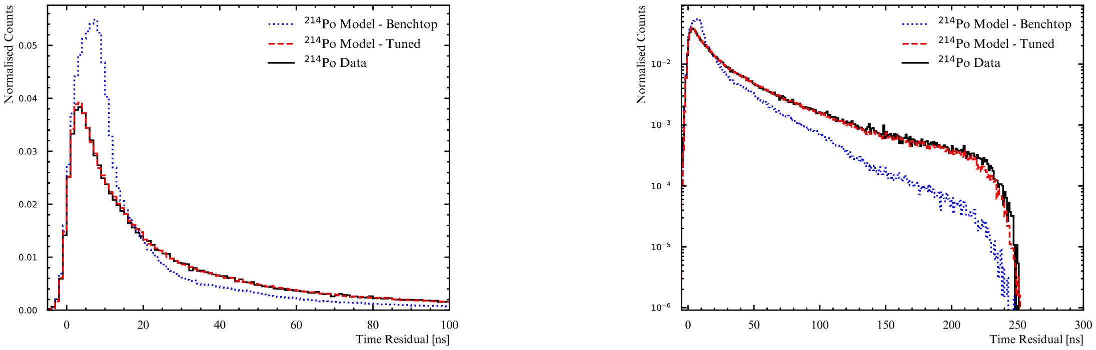

(a) ^{214}Po MC before (blue) and after (red) emission time tuning, compared to tagged ^{214}Po data (black), within a 4 m FV.

(b) ^{214}Po MC before (blue) and after (red) emission time tuning, compared to tagged ^{214}Po data (black), within a 4 m FV.

Figure 3.18: Comparison of time residual distributions for ^{214}Po for peak (left) and tail (right) regions, before and after emission time tuning.

the same problems as found in the timing tuning [77]. Hyperparameter tuning requires a machine learning model to be re-trained on input data for each set of parameters; there is no analytic function describing how the model output relates to the hyperparameters, and so such tunings are expensive and time consuming.

In this case, the convolution of the emission timing model in equation 3.11 with the detector response, along with the other optical processes, means an analytic relationship between the model parameters and the data-MC agreement is unknown. The scintillator model may also be considered an expensive to evaluate black-box function, with a large number of hyperparameters to tune. Thus, it was determined Bayesian Optimisation would be well suited to solve these problems.

Bayesian optimisation seeks to minimise the number of evaluations of the model (i.e. the number of times we must simulate Monte-Carlo with a set of parameters) by reformulating the problem in terms of surrogate, acquisition and the objective functions. The objective function is the χ^2 loss between the data and Monte Carlo, and is unknown. Conversely, the surrogate is a known, easy to evaluate prediction of the objective at an array of unmeasured points, where the predictions are conditioned on the measured samples from the objective. This surrogate is therefore the prior and posterior (depending on the step in the process) of the Bayesian optimisation. Finally, the acquisition function is a simple function that uses the measured values of the objective and the predictions made by the surrogate to select the next point in the domain to sample the objective. There are many choices of acquisition function. In this work we present the results from three common choices: expected improvement, most probable improvement and upper confidence bound.

3.5.1 The Objective Function

First, the objective function to be minimised must be defined. In this case, it is simply the χ^2 between the time residual distributions of the simulated model and the measured data:

$$\chi^2 = \sum_{bins} \frac{(O_i - M_i)^2}{M_i} \tag{3.15}$$

Where O_i and M_i are the observed and measured frequencies in bin i, respectively.

The objective function is non-trivially dependent on the parameters in the emission time model of equation 3.11, as it is the result of the full MC simulations and convolves many optical and detector effects.

3.5.2 The Surrogate Function

In order to make a prediction of the objective function across the domain, a cheap to evaluate prior, or surrogate, is defined. This surrogate function is chosen to be a Gaussian Process (GP), which is an infinite-dimension analogue to a multi-variate Gaussian distribution [78]. A Gaussian Process is a probability distribution over an infinite family of functions, defined by a mean, $\mu(\vec{X})$, and a covariance function, $\Sigma(\vec{X}, \vec{X}')$:

$$f(\hat{X}) \sim \mathcal{GP}(\mu(\vec{X}), \hat{\Sigma}(\vec{X}, \vec{X}')) \tag{3.16}$$

Where $\vec{X}$ is a 9-dimensional input vector:

$$\vec{X} = \begin{pmatrix} t1 \\ t2 \\ t3 \\ t4 \\ t_r \\ A_1 \\ A_2 \\ A_3 \\ A_4 \end{pmatrix} \tag{3.17}$$

There is a constraint on the weights, A_i, such that:

$$\sum_{i=1}^{4} A_i = 1 \tag{3.18}$$

Gaussian Processes

For this work, we take the convenient choice of a zero mean prior, $\mu(\vec{X}) = 0$. The covariance function, $\hat{\Sigma}$, is determined by the choice of kernel.

The choice of kernel, K, determines the family of functions over which the GP is defined. This choice is informed either by intuition or trial-and-error; for example, if you expect periodicity in your objective function with respect to the inputs, you would use a kernel that encodes for this feature. At each predicted point, the surrogate is defined as the GP's mean function with uncertainty given by the covariance function. This mean is the weighted average of all the infinite functions evaluated at that point.

As an example, consider data with a single feature, e.g. the time residual data-MC agreement as a function of t_1. In order to understand the GP as a probability distribution over functions, we define a set of 100 evenly spaced points in the domain of t_1 (see figure 3.19). The GP ensures the value of the objective at each point follows a normal distribution, with the set of sampled points constituting a multivariate-normal distribution. This joint distribution is defined by the mean and covariance function. By sampling from this multivariate distribution, we instantiate a single function from the infinite family of functions defined by the GP. The central tendency of the multivariate normal distribution at each point is determined by the mean function, and the type of functions sampled are governed by the covariance function. For example, figure 3.19 shows functions sampled using the RBF [79] and Matern [80] covariance functions and both with a zero mean. The RBF covariance function enforces smoothness in the sampled functions, whereas the Matern kernel encodes more local variability.

The choice of kernel function, K, is of central importance to the predictions of the surrogate. Depending on the kernel, the family of functions over which the GP is defined is different. Two common choices of kernel are the Radial Basis Kernel (RBF), which provides a measure of similarity between points in the domain, and the Matern Kernel:

$$K_{RBF} = exp\left(-\frac{||\vec{X} - \vec{X}'||^2}{2l^2}\right) \tag{3.19}$$

$$K_{Matern} = \frac{1}{\Gamma(\nu)2^{\nu-1}}\left(\frac{\sqrt{2\nu}}{l}||\vec{X} - \vec{X}'||^2\right)^{\nu} B_{\nu}\left(\frac{\sqrt{2\nu}}{l}||\vec{X} - \vec{X}'||^2\right) \tag{3.20}$$

Where $\vec{X}$ and $\vec{X}'$ are two input parameter vectors and l is a tunable free parameter, governing the length scale over which the points are seen as similar. The Matern kernel includes additional parameters, such as the gamma function, Γ, modified Bessel function, B, and ν, which controls the smoothness of the output.

This work uses the RBF kernel, assuming that the objective function scales smoothly with the 'distance' between input parameters.

After observing data, the mean and covariance functions are updated via equations 3.21 and 3.22.

$$\mu(\vec{X}) = \hat{K}_{m,p}(\hat{K}_{m,m} + \sigma^2\mathcal{I})^{-1}\vec{F} \tag{3.21}$$

$$\hat{\Sigma} = \hat{K}_{p,p} - \hat{K}_{p,m}(\hat{K}_{m,m} + \sigma^2\mathcal{I})^{-1}\hat{K}_{m,p} \tag{3.22}$$

Where $\vec{X}$ is an input parameter vector, $\hat{K}$ is a kernel function that returns the covariance between two points, $\mathcal{I}$ is the identity matrix, $\vec{F}$ is a vector of measured objective function values and σ^2 is a

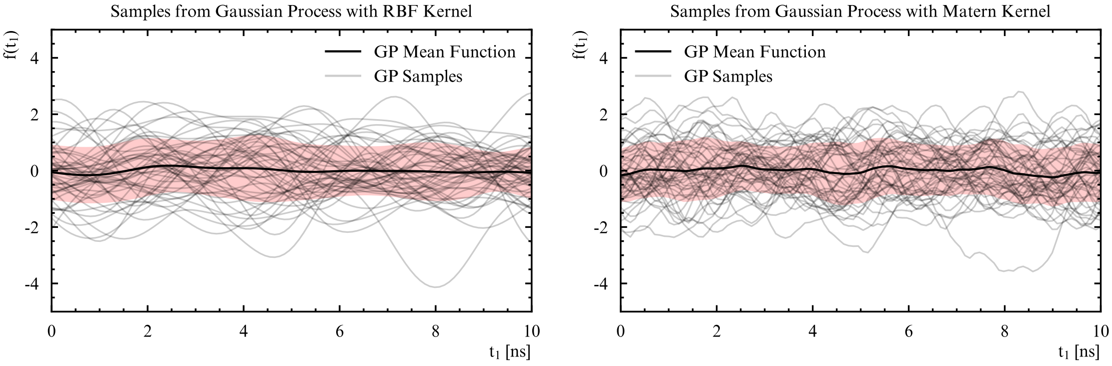

Figure 3.19: Comparison of function samples from a GP defined with zero mean and RBF (left) or Matern (right) covariance function. Functions evaluated at 100 predicted points in both cases, according to equations 3.19 and 3.20, respectively.

user-defined (or model learned) measurement uncertainty. In the above, subscript m and p refers to measured and predicted input parameter vector-points, respectively.

For any finite set of n objective evaluations, $[F(\vec{X_1}), F(\vec{X_2}), \cdots, F(\vec{X_n})]$, the GP assumption ensures the set form a jointly Gaussian multi-variate distribution. This is important, as it allows closed-form solutions for the covariance and mean functions given in equations 3.21 and 3.22, since the conditional probability distribution, given new data, is also jointly Gaussian.

As the objective is sampled at more points, the updated surrogate function is constrained to pass through each measured point. This effectively re-weights the GP distribution over functions, making some in the family more or less likely. The updated mean and covariance functions therefore progressively make better informed predictions of the true objective as more measurements are made.

3.5.3 Acquisition Functions

An acquisition function is the part of the optimiser that determines the next point to sample the objective function, given the predictions of the surrogate across the domain. As with other optimisation techniques, there is a tradeoff to be made bewtween exploration and convergence. Different acquisition functions prioritise different strategies and therefore are more or less prone to being trapped in local minima. The choice of acquisition function is therefore sensitive to the true objective surface over which the optimisation is performed, and thus different functions are good for different cases.

Lower Confidence Bound

The acquisition function considered in this work is the Lower Confidence Bound (LCB)[2] [81]:

$$LCB(\vec{X}) = \kappa\sigma(\vec{X}) - \mu(\vec{X}) \tag{3.23}$$

Where κ is a tunable parameter.

The LCB acquisition formula may be thought of as a simpler version of the EI function: a balance is sought between selecting sample points with high uncertainty and low mean function, with the trade-off between exploring new points and exploiting good solutions explicitly controlled by the parameter κ, rather than 'baked in' using the uncertainty of each predicted point.

3.5.4 Algorithm Implementation

As stated before, the main goal of this work was to develop a fully automated and much more time-efficient method of performing timing calibrations. To this end, the algorithm was designed to use HTCondor's direct acyclic graphs (DAGs) [82], which execute a series of discrete scripts (nodes) in sequence. The calibration process was broken down into the following steps:

- Propose a set of time constants to evaluate
- Simulate
- Calculate simulated time residuals
- Evaluate agreement between simulation and data residuals

At the conclusion of each iteration, a post script runs to determine if the optimisation has converged (i.e. has the last 5 samples been within 5% of each other in the feature space), or if the maximum iterations have been reached.

---

[2] Note that this formulation is the negative of the usual version, to transform the problem into an acquisition function to be maximised

Practical Implementation Details and Constraints

Despite the apparent advantages over the laborious grid search method, Bayesian optimisation is not without drawbacks. In particular, there is a bottleneck resulting from the matrix inversion necessary for the mean and covariance function update rules, as shown in equations 3.21 and 3.22. This bottleneck scales according to $\mathcal{O}(n^3)$, where n is the number of measured points, or iterations of the algorithm.

Despite this, the algorithm is not expected to require more than the order of 10s of function evaluations to converge on a good solution, and if so, the matrix inverse step will remain subdominant to the MC simulation runtime.

A second drawback of Bayesian optimisation is the 'curse of dimensionality'. Like many algorithms, as the dimensionality of the feature space increases, the problem rapidly becomes intractable. In this case, the specific problem is defining a mesh of predicted points sufficiently fine to capture local and global minima in the non-trivial feature-space domain, upon which to evaluate the surrogate. In fact, with a 9-dimensional feature space, the number of points necessary for sufficient coverage was prohibitive. Upon identifying this, it was decided to run the optimisation for a subset of parameters only, and iteratively tune each subset until convergence or maximum iterations was reached.

With this in mind, the algorithm performs sequential tunings of (t_i, A_i) pairs, before running a final 1D optimisation of the rise time parameter.

Something to note for all these algorithms is the RBF kernel's length scale parameter, l. As the RBF kernel is used to return a similarity measure between point-vectors, it is necessary to first transform the features such that they possess length scales related to the impact each parameter has on the objective. It is known that t_1 is dominant over t_4, however t_1 varies over ~ 5 ns scale and t_4 over a ~ 100 ns scale. In order to treat point-vectors with similar t_1 but different t_4 values as 'close', we must weight the length scales appropriately.

Automatic Relevance Determination

The standard methodology is to first fit the GP to the data for a given set of parameters, and then maximise the log-likelihood of the model fitting the data by adjusting the length scales of the method. This is known as 'Automated Relevance Determination' (ARD) [83], and allows the GP to 'learn' the correct scales for each feature as part of the optimisation process. Equation 3.24 shows the log-likelihood to be optimised.

$$log(\mathcal{L}) = -\frac{1}{2}\vec{y}^T\hat{K}^{-1}\vec{y} - \frac{1}{2}log(|\hat{K}|) - \frac{n}{2}log(2\pi) \tag{3.24}$$

Where $\vec{y}$ is a vector of measured objective values and $\hat{K}$ is the covariance matrix between each measured point. This equation may be understood by considering each term:

- $\frac{1}{2}\vec{y}^T\hat{K}^{-1}\vec{y}$ is the data-model fit term. It decreases with length scale, meaning the model gets progressively less flexible.
- $\frac{1}{2}log(|\hat{K}|)$ is a complexity penalty term. As the length scale increases, the penalty decreases, since more smoothly varying functions are less complex.
- $\frac{n}{2}log(2\pi)$ is a normalisation offset to the likelihood and does not impact results.

The partial derivatives of equation 3.24 with respect to the length scales make it amenable to gradient descent optimisation. However, since we are only tuning two parameters at once, a simple grid search over length scales is simpler to implement and still does not present a bottleneck in computation time.

Each time the algorithm updates the surrogate with a new measurement, the ARD script evaluates 3.24 to find the length scales which allow the surrogate to best fit the available data. With progressively more samples, these learned scales converge to stable solutions. Figure 3.20 shows the log-marginal likelihood space for a later iteration of the (t_2, A_2) tuning loop.

3.5.5 Results

The automated Bayesian optimiser result is compared to the grid search and benchtop solution in figure 3.23. The optimiser is able to rapidly learn the form of the best fit parameter space for (t_1, t_2), as shown in figure 3.21. This figure shows the surrogate learns the dominance of t_1 over t_2 in less than 20 iterations, before converging to an initial best fit of $(4.5, 41.0$ ns). Next, the amplitudes (A_1,

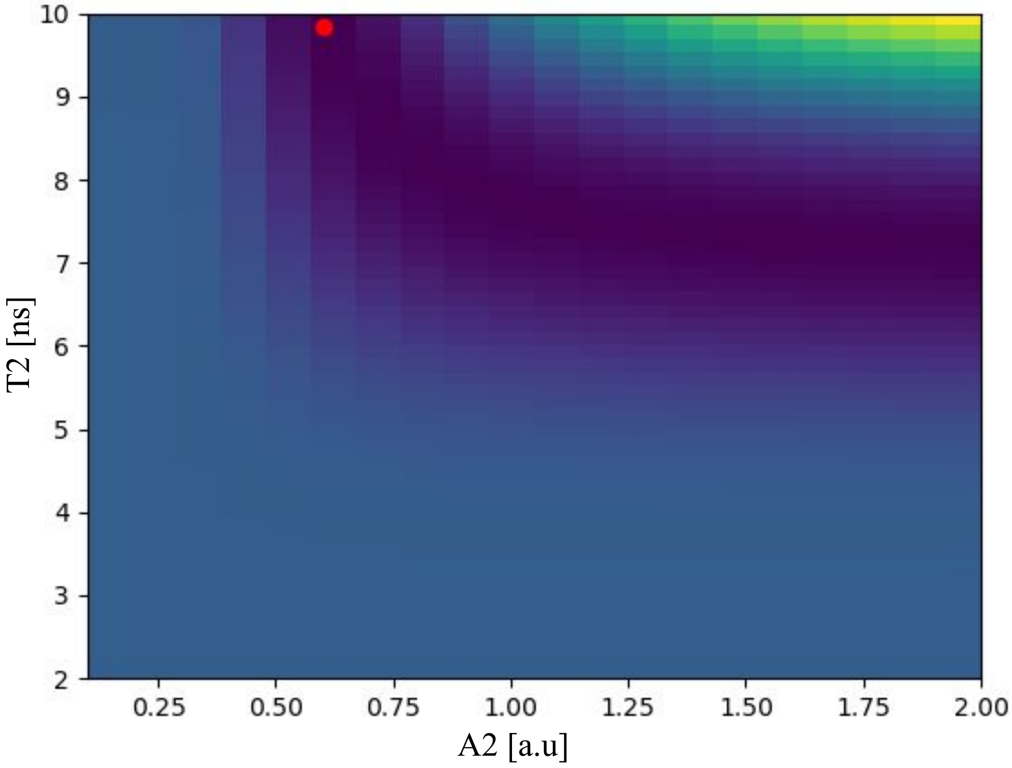

Figure 3.20: Log-marginal likelihood grid search over the RBF kernel length scale parameters for (t_2, A_2). This result implies the surrogate varies over $\sim$ 10 ns scale of t_2.

A_2) were optimised, while keeping the time constants fixed. Figure 3.22 shows the convergence of A_1 to around 0.8, producing the solution in 3.22.

The process is repeated for the third and fourth time constants, with the final result given in 3.23. This demonstrates the power of the Bayesian optimiser: in less than 6 hours, a tuning with comparable accuracy to the exhaustive grid search method was obtained.

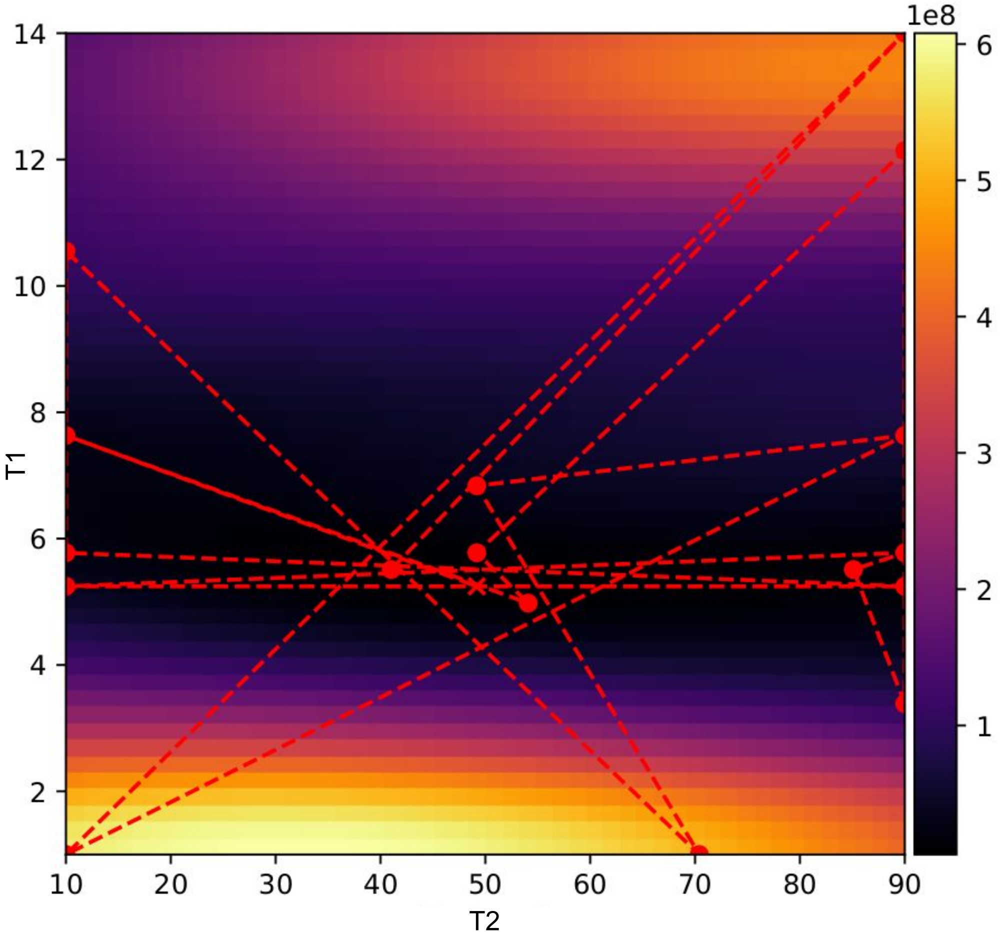

Figure 3.21: 2D Parameter space over t_1 and t_2. The red points show the sampled points, and the colour is mapped to the χ^2 between the data and MC time residuals. The optimiser rapidly learns the dominance fo t_1 over t_2, as shown by the lower predicted χ^2 valley between 4 to 6 ns in t_1. The converged solution is (5.51, 41.02).

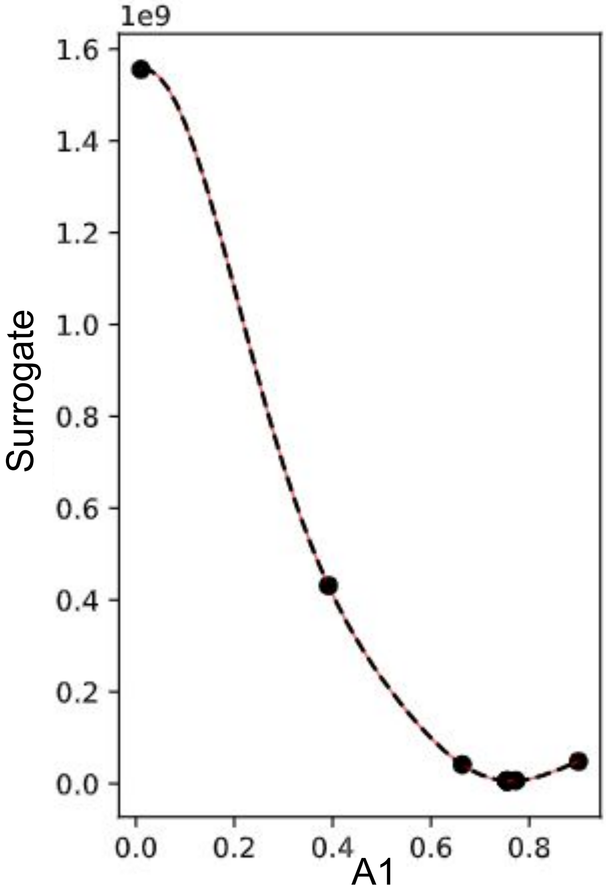

Figure 3.22: Tuning over A_1. The optimiser rapidly converges to a best fit solution of 0.755.

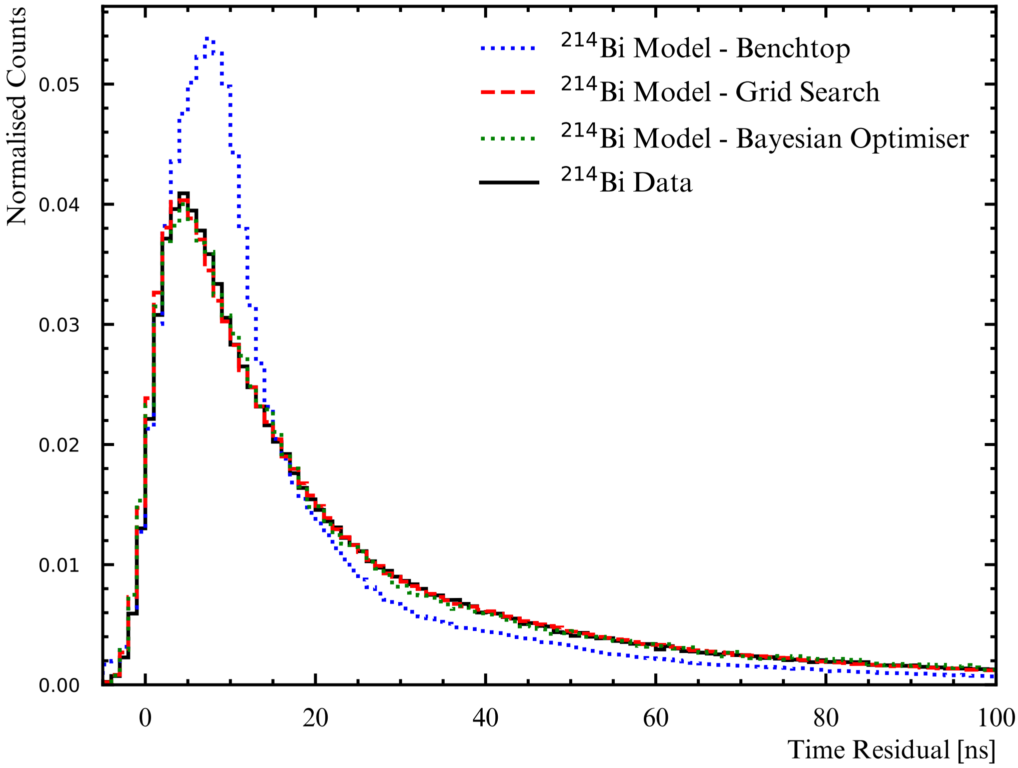

Figure 3.23: Comparison of the benchtop (blue), grid search (red) and Bayesian optimiser solution (green) to tagged ^{214}Bi data (black).

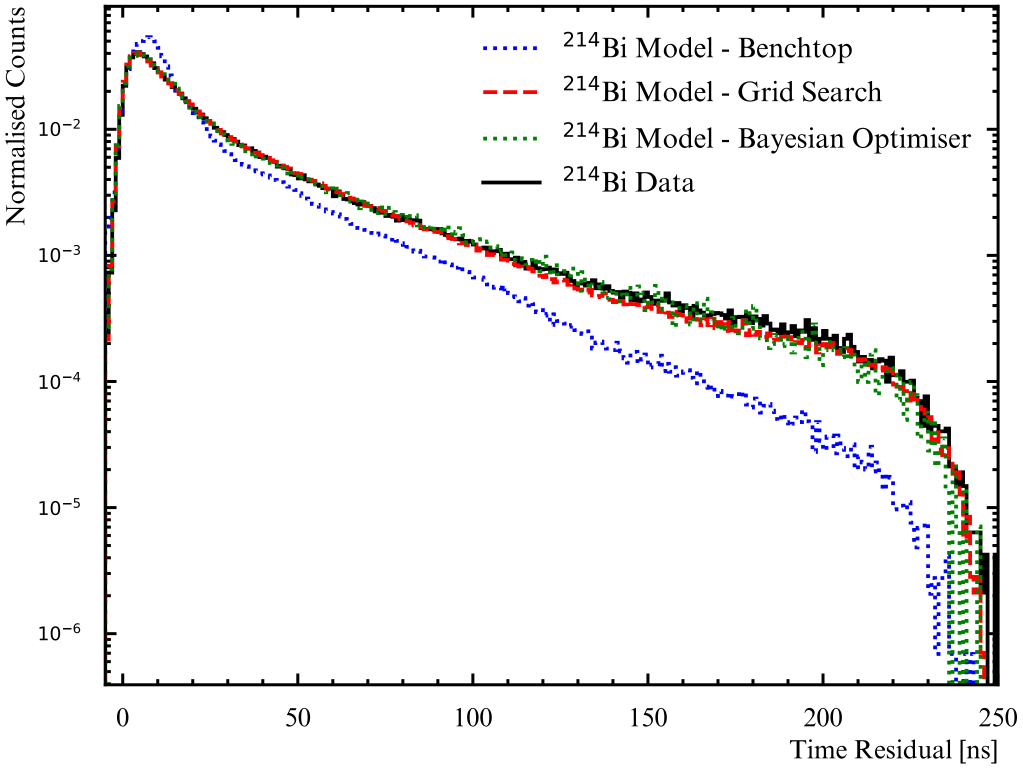

Figure 3.24: Comparison of the tails of the benchtop (blue), grid search (red) and Bayesian optimiser solution (green) to tagged ^{214}Bi data (black).

Chapter 4

Multisite Event Discrimination for ^{8}B Solar Neutrino Analysis in a Background Dominated Domain

This chapter showcases the first application of 'multisite' event discrimination in a liquid scintillator detector, within the context of a ^{8}B solar neutrino flux measurement. The energy ROI is defined between 2.5 to 5.0 MeV, within a 4.5 m FV.

In liquid scintillator, interactions may be classified into two broad classes: events which create single, point-like energy depositions, and those which deposit energy over a broad area at a series of discrete vertices. The former are electron-like events, such as those produced by ^{8}B-ν_e elastic scattering or β-decay. The latter class, multisite, frequently arise from radioactive decays involving γ particles, such as internal ^{208}Tl β-γ decays, and it is the Compton scattering of these γs that produces multiple energy deposition vertices for a single event.

The γ emissions arise following the β-decay of the parent nucleus to an excited state of the daughter, which relaxes via the emission of a γ photon to the ground state.

The observable difference between single-site and multisite events is a broadening in the time residual distributions of multisite events. The discrete energy depositions at each Compton scattering vertex does not conform to the position reconstruction's assumption of events containing a single interaction vertex. Thus, when equation 2.2 is used to calculate the time residuals for a given fit-vertex, the resulting distribution is broader than expected. It is this broadening of multisite event residual distributions, compared to single-site events, that forms the basis of discriminating single-site and multisite interactions.

As discussed in section 1.2.3, a measurement of the ^{8}B-ν_e flux in the 2.0 to 5.0 MeV transition region is of great interest to the community, due to its sensitivity to BSM physics. It is also a challenging ROI to perform a measurement, owing to high rate of internal radioactive backgrounds. To date, no ^{8}B flux measurements have been reported with the necessary precision to probe new physics in this energy range.

In liquid scintillator, the dominant background in this ROI is internal ^{208}Tl β-γ decays. These decays feature a prominent γ emission at 2.6 MeV as the daughter ^{208}Pb relaxes to the ground state, as shown in equation 4.1.

$$^{208}\mathrm{Tl} \rightarrow^{208}\mathrm{Pb}^{*} + e^{-} \rightarrow^{208}\mathrm{Pb} + e^{-} + \gamma(2.6\ \mathrm{MeV}) \tag{4.1}$$

Where * denotes an excited state.

Due to their multisite character, and the single-site nature of the ^{8}B-ν_e signal, multisite event discrimination provides a method to improve the precision of solar neutrino flux measurements. This chapter begins by explaining the estimated ^{8}B-ν interactions are expected in the SNO+ detector. Next, details of the multisite classifier, developed using the calibrated scintillator model described in chapter 3, are provided. Indicative results from an Asimov study are presented, showing the probable scale of improvement multisite event discrimination can provide. Finally, the multisite classifier is applied to real detector data and used to measure the ^{8}B solar neutrino flux in the 3.5 to 5.0 MeV region.

4.1 ^{8}B Solar Neutrino Interactions in SNO+

^{8}B solar neutrino interactions are detected via elastic scattering reactions with atomic electrons. Interactions with nuclei, while possible, deposit too little energy into the scintillator to be detected.

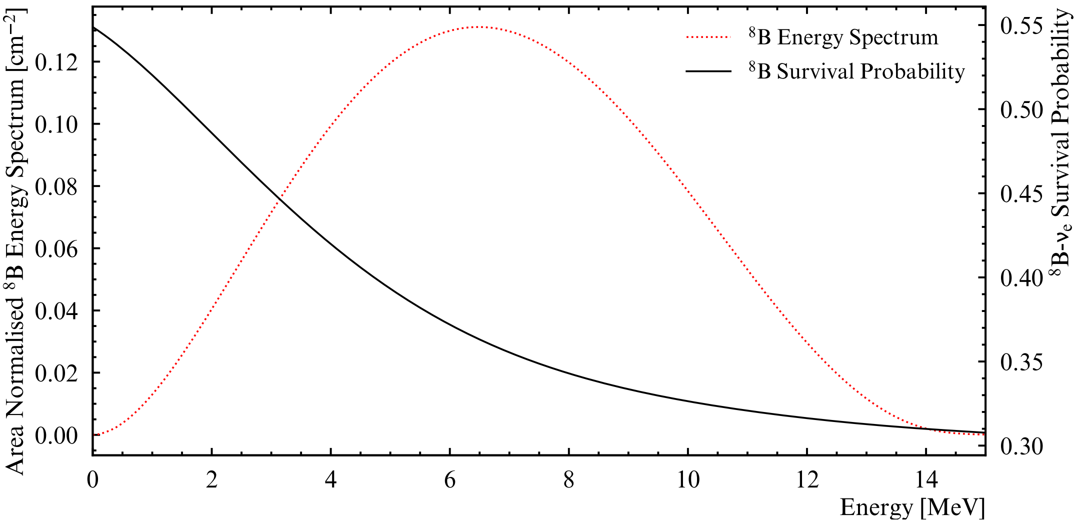

Figure 4.1: Solar ν_e energy spectrum (red) and survival probability (black). Energy spectrum created using values from [84], and survival probability curve taken from [55], which calculated P_{ee} for the B16 GS98 Standard Solar Model [85] using mixing parameters from [86].

Elastic scattering reactions may proceed via both charge current and neutral current interactions. However, since the energy scale of the solar spectrum is far below that necessary to produce muons or tau particles, the CC pathway is only available to ν_e. Considering a solar neutrino of incident energy E_ν and producing a recoil electron with kinetic energy T, the differential cross-section is, to first order:

$$\frac{d\sigma}{dT} = \frac{\sigma_0}{m_e c^2}\left[g_L^2 + g_R^2(1 - \frac{T}{E_\nu})^2 - g_L g_R \frac{T}{E_\nu}\frac{m_e c^2}{E_\nu}\right] \tag{4.2}$$

Where m_e is the electron mass, $g_L = sin^2(\theta_W) \sim 0.23$, $g_R = g_L \pm \frac{1}{2}$ (+ for ν_e, $-$ for other flavours) and $\sigma_0 \sim 8.81 \times 10^{-45}$ cm^2. θ_W is the weak mixing angle. The recoiling electron's kinetic energy is limited by relativistic kinematics to an upper limit, T_{max}, given by equation 4.3:

$$T_{max} = \frac{2E_\nu^2}{m_e + 2E_\nu} \tag{4.3}$$

Additional radiative corrections must be made to the first order result of 4.2 when measuring ^{8}B above 5 MeV, changing the first order result by around 5%.

Component	**fraction**	**Electrons**
LAB	0.99747	131
PPO	0.00253	116

Table 4.1: Mass fractions and number of electrons for LAB and PPO in the full-fill scintillator [87].

Given an elastic scattering interaction produces a recoil electron, this electron excites scintillator molecules, which eventually release this energy as scintillation light, through the processes detailed in section 1.3.3. The expected rate of ^{8}B-ν_e elastic scattering events may be calculated by integrating the cross-section in equation 4.2 with respect to the neutrino energy, while taking account of the ν_e survival probability, $P_{ee}(E)$. Given a predicted flux, $\Phi_{^8B}$, and knowledge of the number of electron targets in the scintillator, n_e, the expected interaction rate is calculated according to equation 4.4:

$$R = \Phi_{^8B} n_e \int P(E) \left[P_{ee}(E)\sigma_e(E) + (1 - P_{ee}(E))\sigma_{\mu,\tau}(E) \right] dE \tag{4.4}$$

Where $P(E)$ is the normalised ν_e energy spectrum shown in figure 4.1. The number of electron targets in the scintillator may be obtained from equation 4.5:

$$n_e = \frac{N_A M_0 \sum_i f_i n_i}{m_0} \tag{4.5}$$

Where f_i is the fraction of the scintillator cocktail made up of component i, n_i is the number of electrons in component i, M_0 is the total mass of the solvent, N_A is Avogadro's constant and m_0 is the molecular weight of the solvent. Using the values in table 4.1, 235 g/mol as the molecular weight of LAB, and 784 tonnes as the total mass of scintillator, the number of electron targets in the full-fill phase was calculated to be 2.63 $\times 10^{32}$.

The predicted ^{8}B flux, $\Phi_{^8\mathrm{B}}$, was taken from [88] and given by:

$$\Phi_{^8\mathrm{B}} = 5.16^{+2.5\%}_{-1.7\%} \times 10^6 \ \ \mathrm{cm}^{-2}\mathrm{s}^{-1} \tag{4.6}$$

Finally, the expected rate of ^{8}B events in the full-fill phase is given by equation 4.7, with the integral evaluated numerically over the range of possible neutrino energies, given by the spectrum in figure 4.1. For each energy, the cross section in 4.2 was first integrated up to the end-point of the recoil electron spectrum, given by 4.3. The neutrino survival probabilities were similarly evaluated using the curve given in figure 4.1.

$$R = 0.136 \text{ hr}^{-1} \tag{4.7}$$

In order to obtain an estimate of the expected number of signal events in the ROI, the predicted rate in 4.4 needs to be scaled by the efficiency of the analysis cuts. These include fiducial volume, energy, reconstruction efficiency and any data cleaning cuts imposed on the data. The event selection cuts used in this analysis are given in table 4.2.

4.2 The Multisite Classifier

This work developed a classifier to leverage subtle differences in timing between single site physics signals (e.g. solar neutrino interactions) and multisite backgrounds (e.g. internal ^{208}Tl). The classification method itself has been studied in MC (see [89], [33]) in the context of rejecting multisite cosmogenics (e.g. ^{60}Co, ^{22}Na) for $0\nu\beta\beta$ decay searches, but has not been used for full-fill scintillator backgrounds, and never applied to data.

The classifier uses the time residuals of a given event, along with PDFs produced of the time residual distributions of a single-site and multisite event, to return the log-likelihood ratio:

$$\Delta log(\mathcal{L}) = \frac{1}{Nhits} \sum_{i}^{Nhits} log \left(\frac{P_S(t^i_{res})}{P_B(t^i_{res})} \right) \tag{4.8}$$

Where $Nhits$ is the total number of PMT hits in the event, P_S, P_B are the respective time residual PDFs of the signal and background, and t^i_{res} is the time residual calculated with equation 2.2 for the i^{th} PMT hit.

An example of the time residual PDFs used in this work are shown in figure 4.2a. Each PDF was created from 50 000 simulated internal ^{8}B-ν_e or ^{208}Tl events. Selection cuts of 4.5 m FV and a 2.5 to 5.0 MeV ROI were applied to events building the PDFs.

As shown in figure 4.2a, the differences in the time residual distributions themselves appears small. However, due to the careful calibration of the time profiles explained in chapter 3, alongside the high statistics used to create the PDFs, these small discrepancies are significant, with the bin-

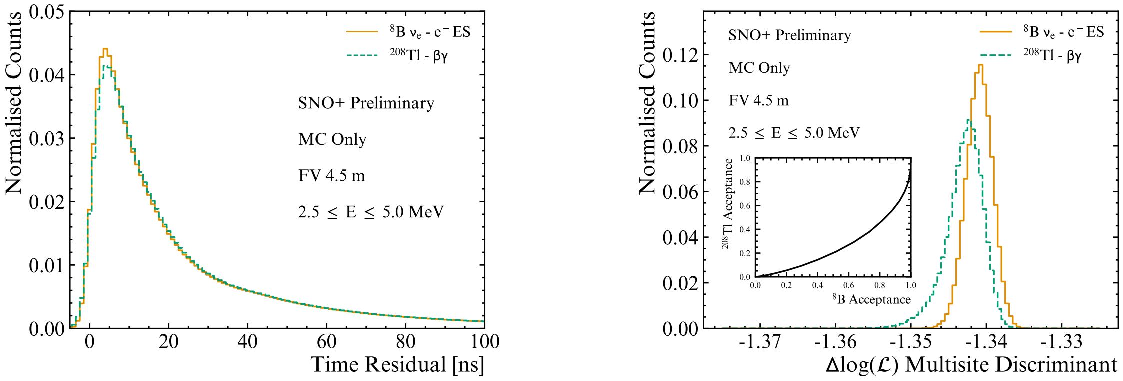

(a) Comparison of time residual PDFs used in multisite classifier for ^{208}Tl $\beta\gamma$ and ^{8}B-ν_e ES.

(b) Multisite discriminant distributions for ^{208}Tl $\beta\gamma$ and ^{8}B-ν_e ES, showing statistical separation power for single-site versus multsite events.

Figure 4.2: Underlying multisite classifier PDFs and statistical separation for ^{208}Tl $\beta\gamma$ and ^{8}B-ν_e ES.

by-bin statistical errors too small to be seen. Thus, the $\Delta log(\mathcal{L})$ multisite classifier is able to achieve a good statistical separation between the two event types, as shown in figure 4.2b.

The analysis process proceeds as follows: first, generate PDFs in energy and multisite discriminant for each process within the domain; second, use these PDFs in an extended-binned maximum log-likelihood fit, to obtain the normalisations of the dominant multisite background (^{208}Tl) and single-site signal (^{8}B-ν_e ES). The fitting formula is given in equation 4.9. The key metric of success is the width of the profile log-likelihood curves for the extracted ^{8}B normalisation with and without multisite information included in the fits. To this end, the fits were performed thrice, using a) energy only, b) multisite only and c) a combination of energy and multisite.

If the multisite classifier encodes enough *different* information from the energy PDFs, the combination of these PDFs in the fit should yield a more precise ^{8}B flux than each alone would obtain.

A bespoke likelihood-based fitting framework was developed to recover the solar neutrino and ^{208}Tl normalisations in the dataset. The analysis is a binned-extended log-likelihood fit, where the log-likelihood is calculated according to:

$$log(\mathcal{L}) = \sum_k n_k log\left(\sum_i N_i P_{ik}\right) - \sum_i N_i \tag{4.9}$$

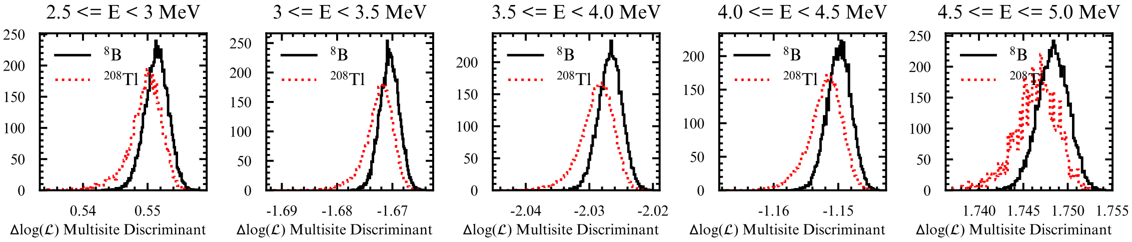

Figure 4.3: Comparison of discrimination power as a function of energy. From left to right: $2.5 < E \leq 3.0$ MeV, $3.0 < E \leq 3.5$ MeV, $3.5 < E \leq 4.0$ MeV, $4.0 < E \leq 4.5$ MeV, $4.5 < E \leq 5.0$ MeV. Separation between event classes remains approximately the same over the 2.5 to 5.0 MeV energy ROI.

Where n_k is the number of events in bin k of the dataset, N_i is the normalisation of event type i and P_{ik} is the counts in the energy or multisite PDF for event type i. In each case, the $-2log(\mathcal{L})$ is minimised with respect to the normalisations, N_i.

4.2.1 Energy Dependence of the Multisite Classifier

The underlying time residual PDFs, used in the multisite classifier calculation, exhibit an irreducible energy dependence. At high energies ($\sim$ E $>$ 8 MeV), this is due to the effect of multiple photon hits on the PMTs. These multi-photon hits generate an excess of photoelectrons in the PMTs, leading to a more rapid crossing of the PMT trigger thresholds. Since each PMT may only register a single hit per event window, this biases the observed time residuals towards earlier scintillation emission times.

Before applying the classifier to data, it was important to verify the extent to which it was sensitive to changes in energy across the ROI. To this end, the ROI was split into 0.5 MeV bins, and the multisite classifier time residual PDFs regenerated within these energy bins. The output discriminant distributions for ^{208}Tl and ^{8}B-ν_e ES events reconstructing within each respective bin are shown in figure 4.3.

As can be seen in figure 4.3, the energy dependence of the classifier is negligible within the analysis ROI. While care should be taken when applying the classifier more generally to higher energy events, it was deemed acceptable to use a single set of time residual PDFs, as given in figure 4.2a.

4.3 Asimov Dataset Analysis

To get a handle on the impact of including multisite information in the maximum likelihood extraction code, the fit was run on an Asimov dataset [90]. This dataset was created from simulated events within a 4.5 m FV and 2.5 to 5.0 MeV ROI. PDFs were created in both energy and multisite for each event class, and then scaled to their relative normalisations given by a background model. These normalisations were then scaled up to represent the event counts for 1 year of livetime. For each event, the underlying time residual PDFs were created from ^{208}Tl and ^{8}B-ν_e events, using equivalent energy and FV cuts.

4.3.1 Background Model: 2.5 to 5.0 MeV

The Asimov model was constructed from PDFs in multisite and energy space, scaled to their respective predicted normalisations assuming 1 year of livetime. The contributing event classes considered are summarised in table 4.9, with the dominant contributions coming from ^{212}Bi-^{212}Po in-window pile-ups, ^{214}Bi, ^{208}Tl β-γ decays and ^{8}B-ν elastic scattering. ^{210}Tl β-γ and (α, n) coincidences were also considered, but contribute negligible rates to the ROI.

The respective contributions of these event classes to the ROI were estimated by assuming the ^{8}B interaction rate given in equation 4.4, and using the ^{214}Bi-^{214}Po and ^{212}Bi-^{212}Po measured rates described in chapter 3. Since the ^{208}Tl and ^{210}Tl are members of the ^{232}Th and ^{238}U chains, respectively, the predicted number of events within a given ROI and livetime are directly calculable, given the number of tagged $^{214(212)}$Bi-$^{214(212)}$Po events, the Bi-Po tagging efficiencies and the respective Bi $\rightarrow$ Tl branching ratios.

In order to obtain realistic normalisations for the Asimov dataset, the background rates in a real sample of detector data were analysed. Gold list runs, as defined by table 3.4, were selected between 1st June 2022 to 10th March 2023, yielding 145.7 days of livetime. Events were selected using the analysis cuts summarised in table 4.2.

After data selection, a set of generalised coincidence cuts were applied to the dataset. This was done in order to tag and remove coincident backgrounds, such as Bi-Po and (α, n) decays. Table 4.3 summarises the coincidence cuts applied. These cut values were tuned to remove $> 95\%$ ^{214}Bi-^{214}Po,

Cut Name	**Acceptance Criteria**	**Comment**
Energy	$2.5 \leq \text{E} \leq 5.0$ MeV	Energy ROI providing clean selection of ^{8}B-ν single-site and ^{208}Tl multisite interactions.
FV	4.5 m	Trade-off between sufficient single-site statistics and excluding external backgrounds.
ITR	$0.22 \leq \text{ITR} \leq 0.3$	Applied same ITR cut as used in section 4.4. Important for removing neck follower events.
Retrigger Cut	$\Delta t > 460$ ns	Inter-event time calculated from 50 MHz clock. Ensures no retrigger events enter selection (e.g. neck followers, muon followers, electronic noise).
Reconstruction Valid	True	Ensures both position and energy reconstruction converged.
High Nhit	Nhit ≤ 5000	High nhit veto trigger. If triggered, 20 s deadtime implemented to remove neck hotspot and muon followers. Livetime adjusted by number of veto windows triggered.
Data Cleaning Mask	0x2100000042C2	Automated data cleaning checks applied during pre-processing. Flags events due to instrumental noise.

Table 4.2: Event selection cuts applied to the dataset. Identical cuts were applied when creating the PDFs used in this analysis.

with < 5% signal sacrifice in the energy and FV ROI, with the final efficiencies for each event class listed in table 4.4.

Prompt Energy	$E \leq 10$ MeV
Delayed Energy	$0.2 \leq E \leq 10$ MeV
ΔT	$\Delta T \leq 4$ ms
ΔR	$0 \leq \Delta R \leq 2$ m

Table 4.3: General coincidence tagging cuts.

While the out-of-window cuts are effective at removing ^{214}Bi-^{214}Po and (α, n), the ^{212}Bi-^{212}Po is almost unaffected. This is because in-window ^{212}Bi-^{212}Pos are overwhelmingly more frequent than out-of-window ^{214}Bi-^{214}Pos, due to the relatively short half life of ^{212}Po compared to the detector's trigger window. Initially, in-window classifiers were used to remove these events. These classifiers use MC simulations of in-window and out-of-window decays to create characteristic time residual PDFs, which are subsequently used to assign a likelihood for a given event being an in-window decay. However, the multisite and in-window classifiers, both being time-based, are correlated. In order to

Isotope	**Efficiency**
^{214}Bi-^{214}Po	0.953
^{212}Bi-^{212}Po	0.048
(α, n)	0.960
^{208}Tl	0.053
^{210}Tl	0.025
$^{8}B_{\nu e}$	0.036
$^{8}B_{\nu\mu}$	0.036

Table 4.4: General coincidence cut efficiencies in a 4.5 m FV and 2.5 to 5.0 MeV ROI.

not bias this analysis with data cleaning dependent cuts, in-window rejection was not implemented.

^{214}Bi-^{214}Po and ^{212}Bi-^{212}Po Expectations

The first step in measuring the relative rates of each background was to calculate the number of Bi-Po coincidences falling within the ROI. BiPo tagging was applied to every run in the 145.7 day livetime dataset, using the coincidence cuts summarised in table 3.2. The number of coincidence tags was scaled up by the combined efficiency, providing an estimate of the total number of prompt events occurring in the 6 m FV, both observed and unobserved.

Next, this number was multiplied by the fraction of $^{214(212)}$Bi events that fell within the 4.5 m FV and energy ROI. Finally, this quantity was multiplied by the fraction of events remaining following the generalised coincidence tagging, given in table 4.4.

The total number of $^{214(212)}$Bi-$^{214(212)}$Po events expected in the FV and ROI is given by:

$$N_{BiPo} = \frac{N^{Tag}_{clean}}{\epsilon^{FV}_{clean}} \epsilon^{FV}_{ROI}(1 - \epsilon_{Tag}) \tag{4.10}$$

This may be simplified to:

$$N_{BiPo} = N^{AV}_{BiPo} \epsilon^{FV}_{ROI}(1 - \epsilon_{Tag}) \tag{4.11}$$

Where N_{BiPo} is the number of Bi-Po events expected to contribute to the FV and ROI, N^{AV}_{BiPo} is the predicted number of Bi-Po events in the full 6 m volume, N^{Tag}_{clean} is the number of Bi-Po events tagged in the dataset with the clean selection cuts, ϵ_{clean} is the tagging efficiency of the clean selection cuts, ϵ^{FV}_{ROI} is the fraction of Bi-Po events that fall into the FV and ROI and ϵ_{Tag} is the tagging efficiency of the ROI and FV coincidence tagging.

Table 4.5 shows the expected contributions of ^{214}Bi-^{214}Po and ^{212}Bi-^{212}Po to the dataset.

Isotope	$\boldsymbol{N_{clean}^{Tag}}$	$\boldsymbol{N_{BiPo}^{AV}}$	$\boldsymbol{\epsilon_{ROI}^{FV}}$	$\boldsymbol{N_{BiPo}}$
^{214}Bi-^{214}Po	4583.0	12946	0.062979	38.528
^{212}Bi-^{212}Po	138.00	3201.2	0.019424	59.227

Table 4.5: Expected number of Bi-Po events contributing to the ROI and FV. N_{BiPo} was obtained via equation 4.10 and the ROI and FV tagging efficiencies summarised in table 4.4, for a 145.7 day livetime.

^{208}Tl and ^{210}Tl Expectations

The ^{208}Tl(^{210}Tl) expectation value was obtained from the respective number of $^{212(214)}$Bi-$^{212(214)}$Po events in the AV, and the respective branching ratio:

$$N_{Tl} = N_{BiPo}^{AV}\Gamma_{Bi\rightarrow Tl}\epsilon_{ROI}^{FV}(1-\epsilon_{Tag}) \tag{4.12}$$

Where N_{Tl} is the expected contribution of ^{208}Tl(^{210}Tl) to the ROI and FV, and $\Gamma_{Bi\rightarrow Tl}$ is the branching ratio of ^{212}Bi(^{214}Bi) to ^{208}Tl(^{210}Tl).

Table 4.6 summarises the branching ratios, efficiencies and background model expectation values for each Tl isotope.

Isotope	$\boldsymbol{\Gamma_{Bi\rightarrow Tl}}$	$\boldsymbol{\epsilon_{ROI}^{FV}}$	$\boldsymbol{N_{Tl}}$
^{208}Tl	0.3594	0.4305	468.9
^{210}Tl	0.0002100	0.3425	0.9081

Table 4.6: The expected number of Tl events contributing to the ROI and FV for the 145.7 day livetime. N_{Tl} is obtained via equation 4.12, N_{BiPo}^{AV} from table 4.5 and the ROI and FV tagging efficiencies in table 4.4.

Table 4.6 demonstrates the negligible contribution of ^{210}Tl to the ROI, yielding less than 1 count for this 145.7 day livetime.

(α, n) Expectation

α particles, primarily created by ^{210}Po decays, are able to interact with atoms in the scintillator, producing free neutrons and a prompt γ. The most common reactions involve carbon and oxygen,

Pre-Cuts $N_{(\alpha,n)}$	ϵ_p	ϵ_d	$N_{(\alpha,n)}$
24.3	0.118	0.309	0.4

Table 4.7: Pre-cut (α, n) estimate obtained from the product of neutron conversion efficiency, ^{210}Po rate and 145.7 day livetime. $\epsilon_{p,d}$ give the fraction of prompt and delayed events reconstructing within the ROI and FV. $N_{(\alpha,n)}$ is the final estimate, accounting for the generalised tagging efficiency.

as shown in equations 4.13 and 4.14:

$$\alpha +^{13} \mathrm{C} \rightarrow^{16} \mathrm{O} + n \tag{4.13}$$

$$\alpha +^{18} \mathrm{O} \rightarrow^{21} \mathrm{Ne} + n \tag{4.14}$$

These neutrons are captured with a time constant of $\sim 200\ \mu s$, producing a delayed 2.2 MeV γ. Because of the coincident nature of (α, n) interactions, they may be efficiently removed using the generalised coincidence tagging cuts.

The predicted rate of (α, n) interactions is given by equation 4.15.

$$N_{(\alpha,n)} = 6.1 \times 10^{-8} R_{210} t_{live} (\epsilon_p + \epsilon_d) \tag{4.15}$$

Where 6.1 $\times 10^{-8}$ is the α to neutron conversion factor, R_{210} is the rate of ^{210}Po decays, t_{live} is the livetime and $\epsilon_{p,d}$ is the ROI and FV efficiency for prompt and delayed events, respectively.

The two efficiencies are important, as both prompt and delayed events may contribute to the events in the ROI. Thus, the efficiencies of prompt (α, n) and delayed 2.2 MeV γ events were estimated individually from MC.

Table 4.7 shows the predicted contribution of (α, n) events to the ROI and FV. The average ^{210}Po rate of 35 mBq m^{-3} was taken from [91].

^{8}B Solar Neutrino Expectations

The expected number of solar neutrino events in the dataset, was calculated using equation **??**, for ν_e or ν_μ:

$$N_{\nu_{e,\mu}} = \epsilon_{e,\mu} t_{live} \epsilon_{ROI}^{FV} (1 - \epsilon_{Tag}) R \tag{4.16}$$

Where R is the total expected solar neutrino interaction rate in the detector, $\epsilon_{e,\mu}$ is the fraction of R arising from ν_e or ν_μ, t_{live} is the livetime, ϵ^{FV}_{ROI} is the fraction of events reconstructing within the ROI and FV, and ϵ_{Tag} is the generalised tagging efficiency, given in table 4.4.

The expected contributions of each flavour of solar neutrino are summarised in table 4.8.

Flavour	$\boldsymbol{R}$	$\boldsymbol{\epsilon_{e,\mu}}$	$\boldsymbol{\epsilon^{FV}_{ROI}}$	$\boldsymbol{N_{\nu_{e,\mu}}}$
$\boldsymbol{\nu_e}$	475.5	0.7541	0.114	41.0
$\boldsymbol{\nu_\mu}$	475.5	0.2459	0.109	12.7

Table 4.8: Expected number of ^{8}B solar neutrinos in the ROI and FV.

Thus, the expected total number of ^{8}B-ν elastic scattering events in the dataset is 53.7.

4.3.2 Data-Model Agreement

The dataset and background model is displayed in figure 4.4, with a summary of the pre and post analysis cuts expectations given in table 4.9. As can be seen by the plot, the background estimates are in good agreement with the dataset.

Interaction	**Pre-Cuts Estimated Counts** [livetime^{-1}]	**Post-Cuts Estimated Counts** [livetime^{-1}]	**Comment**
^{8}B-ν Elastic Scattering	475.6	53.7	Single-site β interaction.
^{208}Tl β-γ Decay	1150.5	468.9	Dominant Multisite background
^{210}Tl β-γ Decay	2.7	0.9	Negligible multisite background
^{214}Bi β-γ Decay	12 946.0	38.5	Prompt multisite background event in out-of-window coincidence. A small fraction of in-window Bi-Po events may also contribute.
^{212}Bi-^{212}Po Pileup	3201.2	59.2	In-window multisite background
(α, n) Coincidence	24.3	0.4	Negligible out-of-window coincidence

Table 4.9: Summary of event classes considered by Asimov model, contributing to 2.5 to 5.0 MeV ROI, before and after analysis cuts were considered. A livetime of 145.7 days was used.

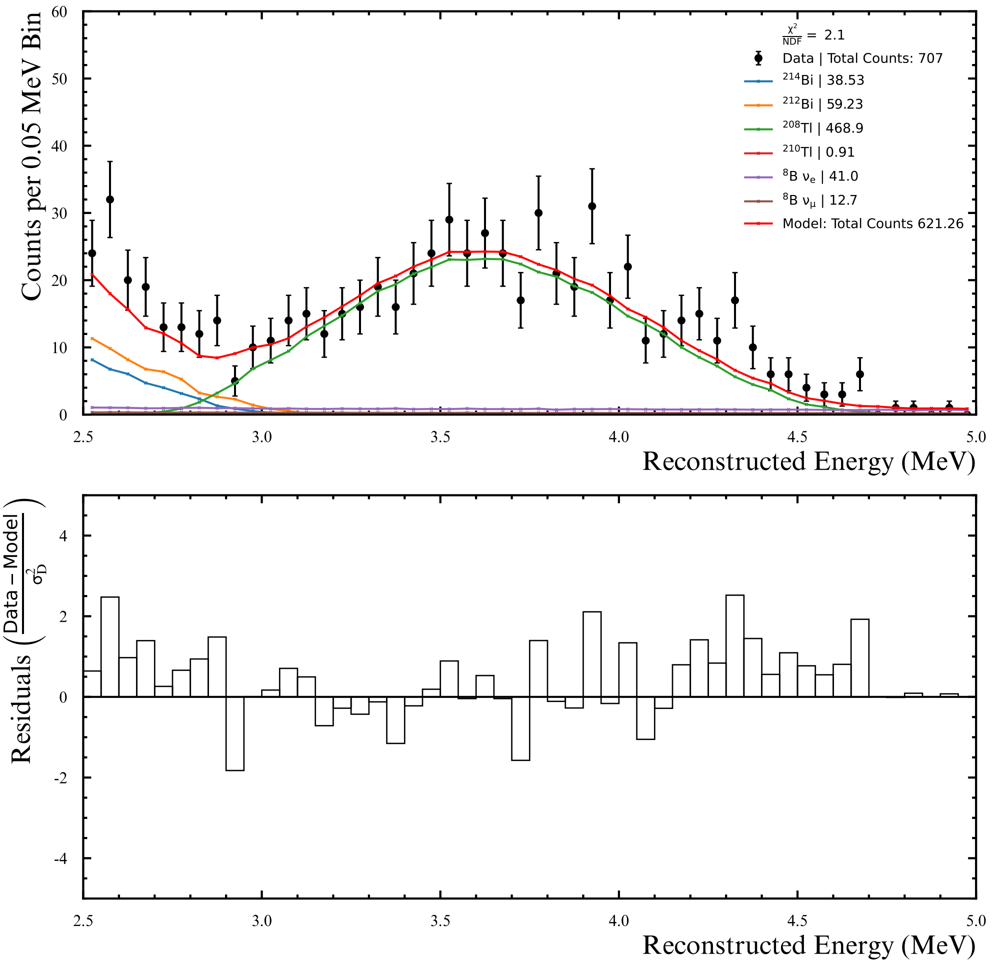

Figure 4.4: Background model showing predicted background and signal event numbers compared to the dataset.

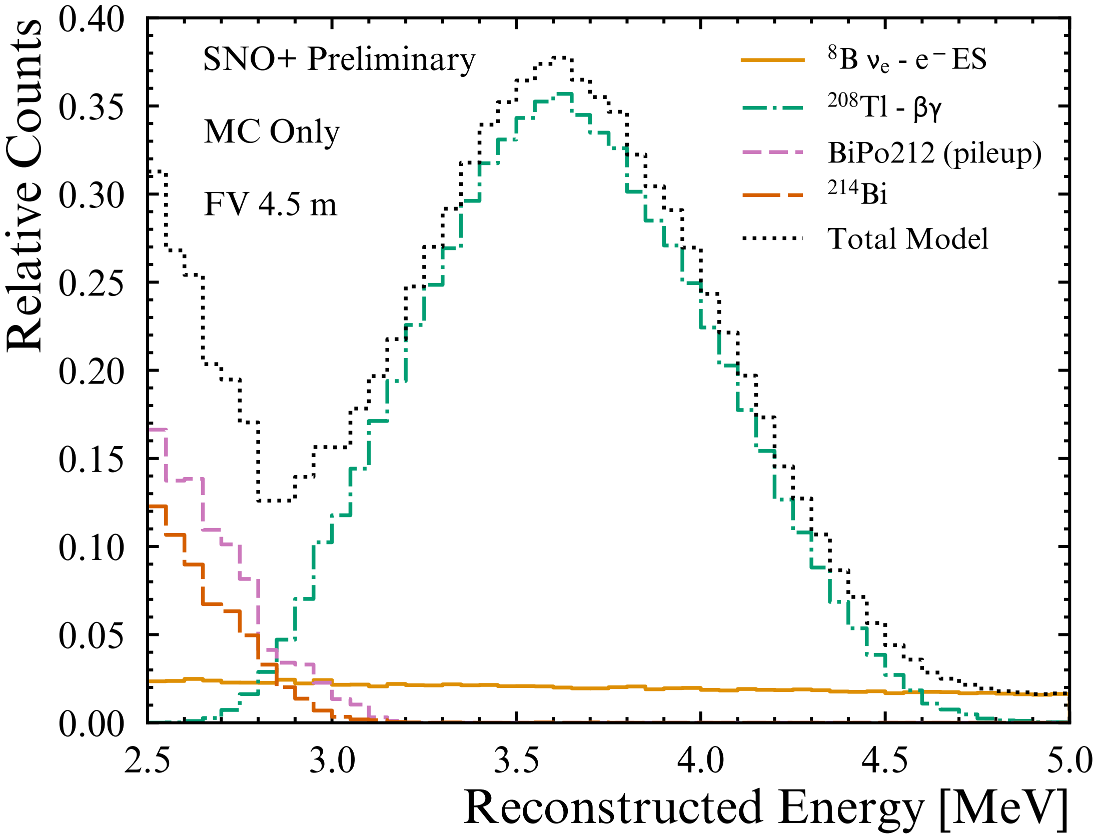

Figure 4.5: Asimov dataset energy distributions with total model dotted black.

4.3.3 Asimov Dataset

Given the background normalisations given in table 4.9, each rate was scaled to represent 1 year of livetime. Monte-Carlo simulations were used to produce PDFs in both energy and multisite discriminant space, according to equation 4.8, with the underlying time residual PDFs shown in figure 4.2a. The energy and multisite PDFs for each event were scaled to the 1 year normalisations, producing figures 4.5 and 4.6. Figures 4.5 and 4.6 show the Asimov dataset in energy and multisite space. Of note is the multisite distributions: there are clear differences between each event type, clearly demonstrating the possibility of using multisite information for particle ID. ^{214}Bi and ^{212}Bi-^{212}Po show differences in multisite, owing to the prevalence of ^{212}Bi-^{212}Po pileup events. Since the polonium decay often occurs within the same trigger window as the ^{212}Bi, their resulting multisite

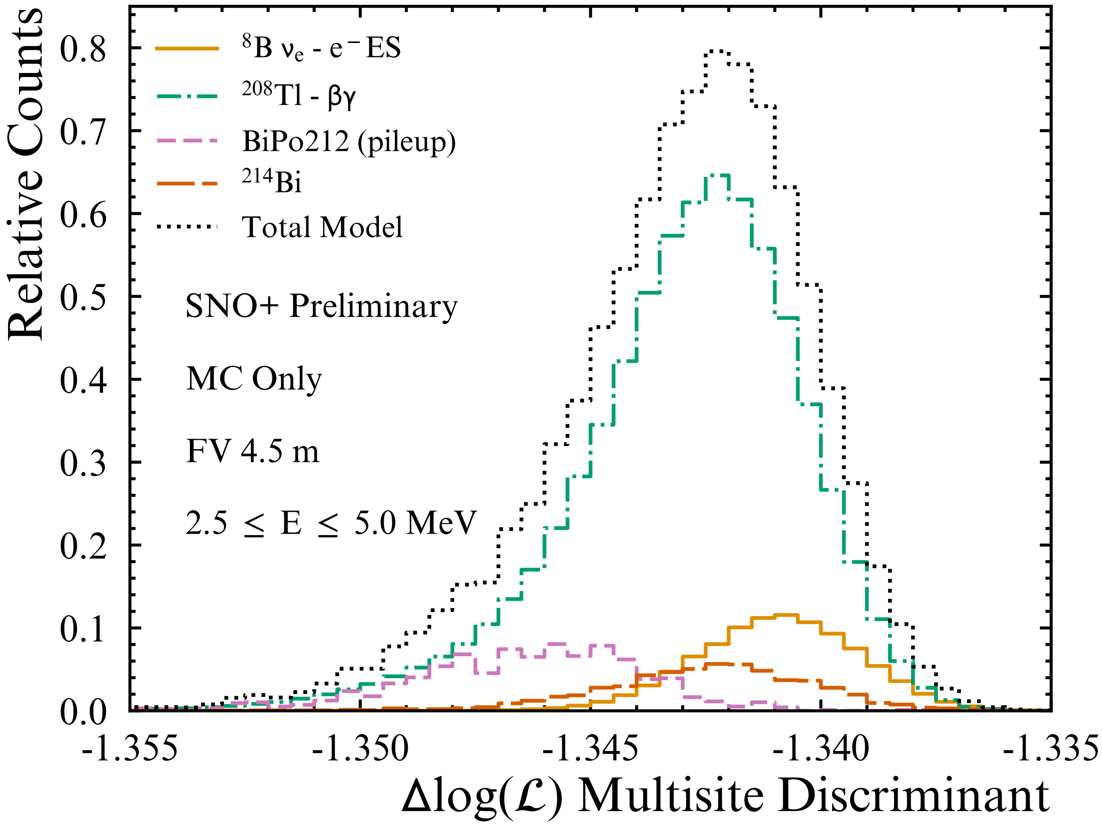

Figure 4.6: Asimov dataset multisite discriminant distributions with total model dotted black.

discriminant is significantly different to ^{214}Bi.

After creating the Asimov dataset, the same PDFs were used in the binned-extended likelihood fit given in equation 4.9. For simplicity, only the ^{208}Tl and ^{8}B-ν PDFs were floated in the fit, with the ^{214}Bi, ^{212}Bi-^{212}Po pileup and ^{210}Tl normalisations fixed to their expected values.

Given only 2 normalisations to fit, a simple grid-search was implemented. The 2D log-likelihood space, as a function of $N^{^{208}\text{Tl}}$ and $N^{^8\text{B}}$ are given in figure 4.7. Of more interest is the profile log-likelihood curves given in figure 4.8. This figure directly compares the sensitivity of the fits using only energy PDF information (orange), only multisite (green), and the combined energy and multisite fit (black). It is clear that the multisite encodes approximately the same quantity of information as the energy PDFs, and since the combined fit is ∼20% more precise, the information is not degenerate.

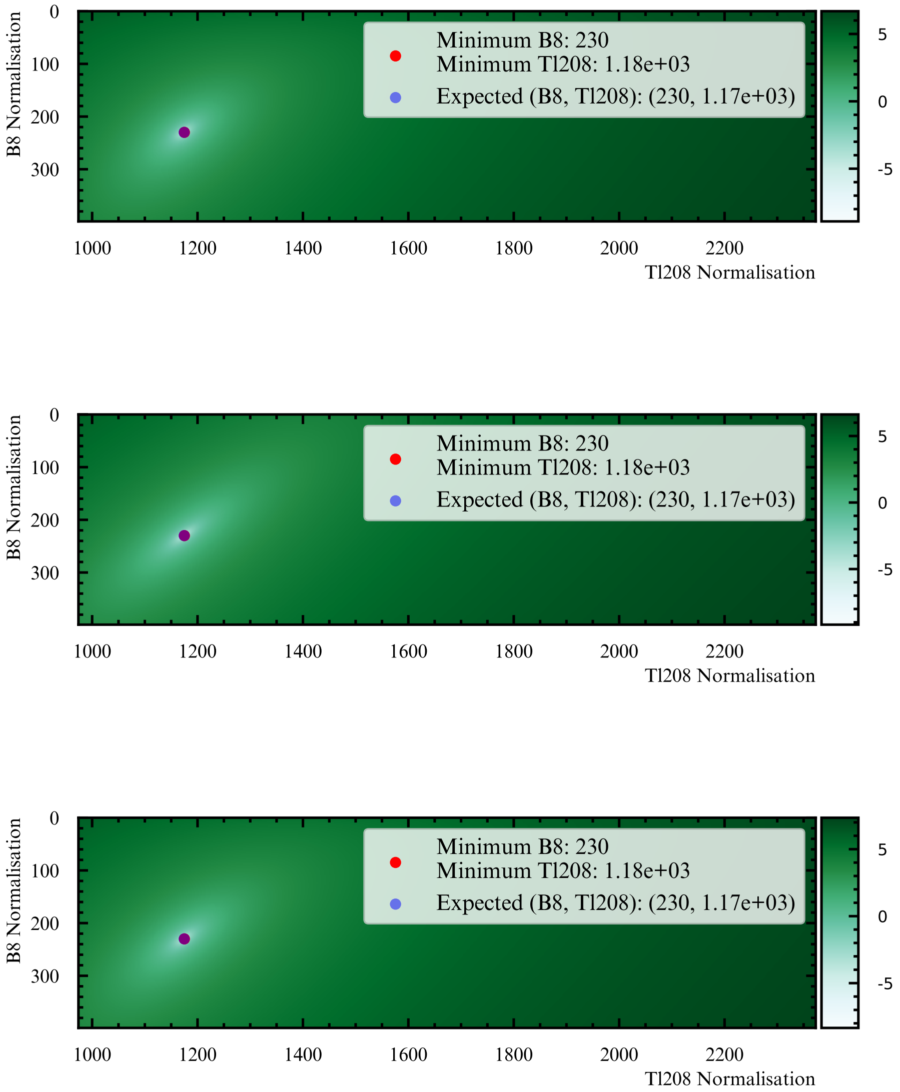

Figure 4.7: $\Delta log(\mathcal{L})$ space obtained from Asimov fits. Top: energy; Middle: multisite; Bottom: combined fits.

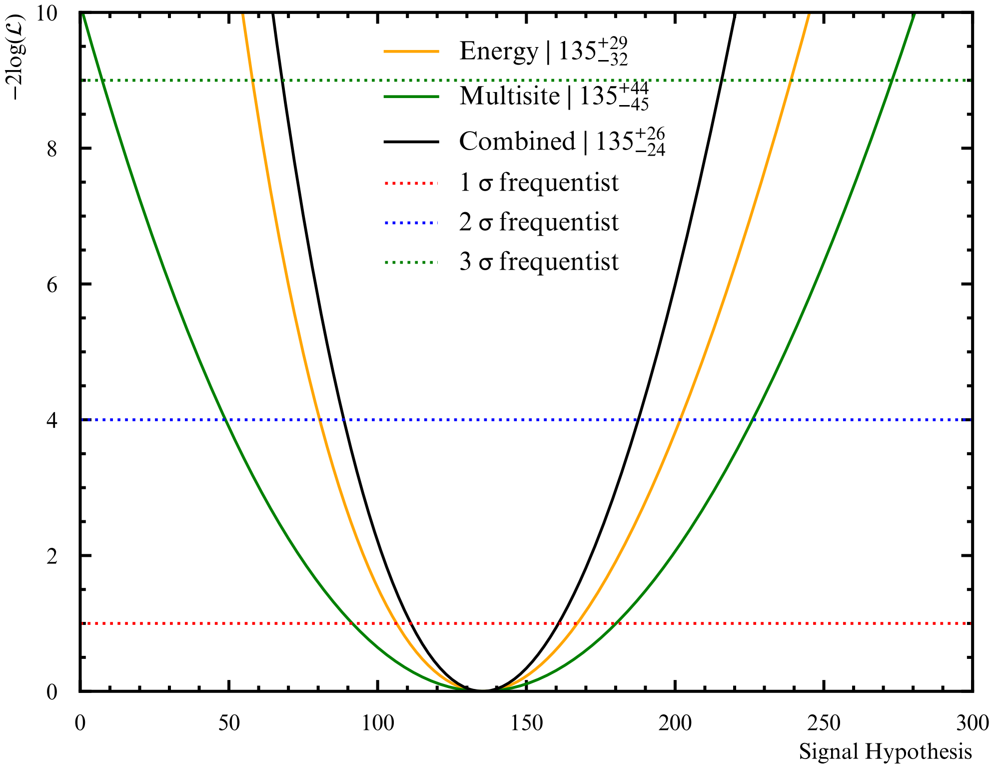

Figure 4.8: Profile log-likelihood curves for ^{8}B normalisations, obtained from Asimov dataset shown in figures 4.5 and 4.6. Fit results from energy (dashed orange), multisite (dashed green) and a combination (solid black) show greater precision when multisite and energy information is combined. Frequentist 1 σ confidence intervals shown as horizontal dotted line (black).

This result is encouraging, and demonstrated a clear improvement over the traditional energy-only flux extraction codes.

4.3.4 Fluctuated Datasets

A second use of simulated data was to create many fluctuated datasets. These were produced by using the expected number of events for each isotope as the mean of a Poisson distribution, and returning a set of Poisson-fluctuated normalisations. These normalisations were then used to scale the energy and multisite PDFs, creating a more realistic fluctuated dataset.

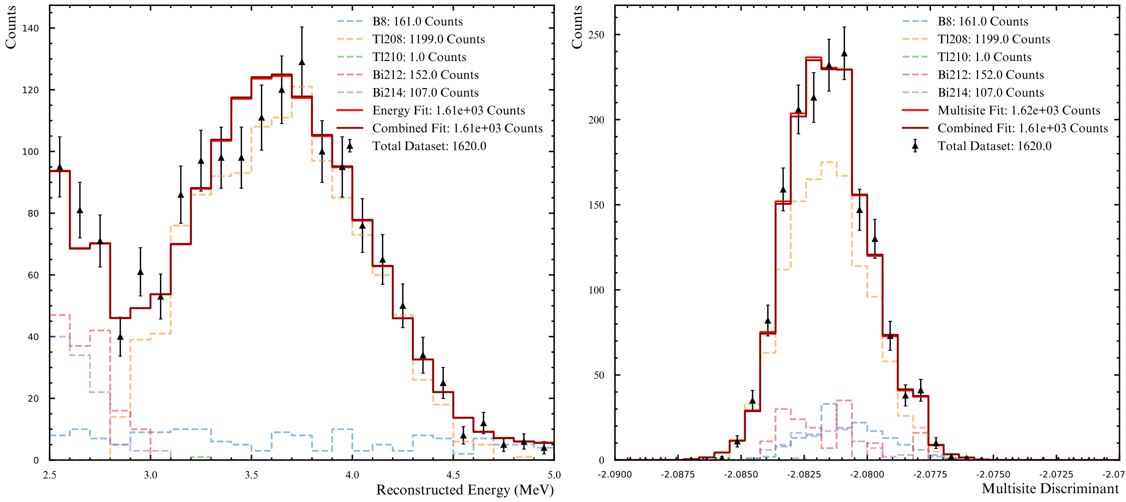

Figure 4.9: Example fluctuated dataset using 1 year Asimov counts as mean of poisson distribution for each isotope. Right: dataset in energy space; left: dataset in multisite space.

These simulated spectrums were then used to analyse the bias and pull of the maximum likelihood signal extraction fits. A total of 10 000 fluctuated datasets were produced, and the signal extraction framework was applied to each. Given the true normalisation of the ^{8}B-ν_e interactions in each, the bias and pull of the fits were evaluated according to equations 4.17 and 4.18.

$$B = \frac{N_{Fit} - N_{True}}{N_{True}} \tag{4.17}$$

$$P = \begin{cases} \frac{N_{True} - N_{Fit}}{\sigma_{+}^{Fit}}, & \text{if } N_{True} > N_{Fit} \\ \frac{N_{Fit} - N_{True}}{\sigma_{-}^{Fit}}, & \text{if } N_{Fit} > N_{True} \end{cases} \tag{4.18}$$

Where B is bias, P is pull and $\sigma_{+/-}$ is the profile log-likelihood 1-sigma uncertainties.

These two quantities allow us to quantify the quality of fits we expect on real data. The bias yields a Gaussian distribution, with the offset of the mean from zero giving an estimate of the accuracy and its width giving a measure of the resolution. The pull distribution should be compared to a standard Gaussian, with mean zero and unit standard deviation. This allows us to check that the uncertainties given by the frequentist intervals are correct: an underestimate would yield a broader pull distribution than the standard Gaussian, and vice versa.

Figure 4.9 shows an example fluctuated dataset, and figures 4.10 show the results of passing all 10 000 datasets through equations 4.17 and 4.18. These plots show that, for our signal to noise ratio, we can expect large biases. This is unsurprising, as the expected ^{208}Tl normalisation is an order of magnitude higher than the signal. In this high background region, the multisite fit has particularly poor sensitivity.

In this case, the pull distributions suggest that the frequentist confidence intervals are providing a good uncertainty estimation in each case, however the combined fit may suffer from additional correlations not fully accounted for.

4.4 Verification of the Multisite Classifier on Data

Prior to running a full likelihood analysis on the data, it was necessary to verify the multisite classifier results obtained from MC studies matched those from data. To do this, a clean sample of multisite

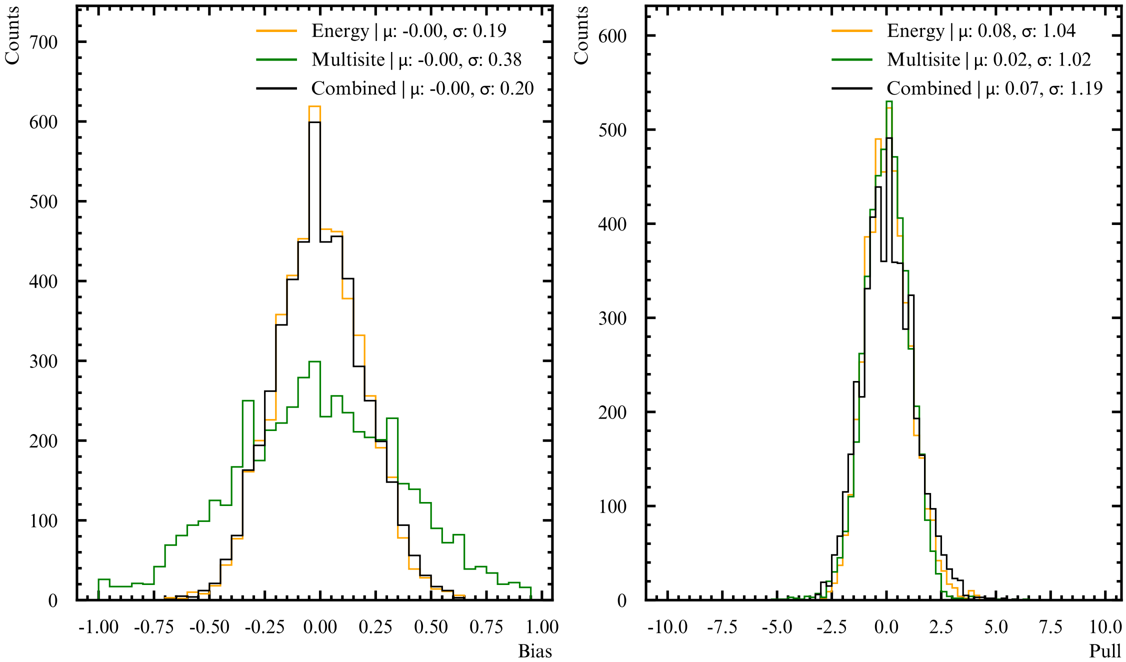

Figure 4.10: Bias (left) and pull (right) for 10 000 fake datasets. In each case the fit is performed with each of the three methods: energy only (orange), multisite only (green) and combined (black).

and single site events in data needed to be selected and their classifier results compared to those obtained from equivalent simulated events.

A natural choice of easily isolated multisite event in data was ^{214}Bi, which decays via β-γ. The emitted γ has a broad range of energies, which Compton scatters through the scintillator and deposits energy at each scatter vertex. ^{214}Bi were tagged using the coincidence tagging algorithm described in section 3.2.2.

For single-site, a high energy selection of ^{8}B-ν interactions was chosen. Beyond 6 MeV, there are no backgrounds, and so the single site events were selected by applying a simple energy cut.

Table 4.10 shows the set of cuts used to extract each sample, with identical cuts applied to the MC. During the data selection, anomalous events were identified according to their in-time-ratio (ITR) classifier result. For a given event, the ITR classifier calculates the ratio of prompt (-2.5 to 5 ns) time residuals to the total PMT hits. This ratio is useful in identifying non-physics retrigger events, such as the neck hotspot followers discussed in section 3.2.3. Upon inspection of the ITR distributions for the ^{214}Bi candidates, shown in figure 4.11, a cut was imposed on both multisite and single-site events. By way of comparison to the MC, the data ITR distribution included additional events, distorting the otherwise Gaussian spectrum. These events may be attributed to neck hotspot followers, which feature unusually flat time residual distributions. Thus, a relatively strong ITR cut was imposed to ensure the purity of the selected events.

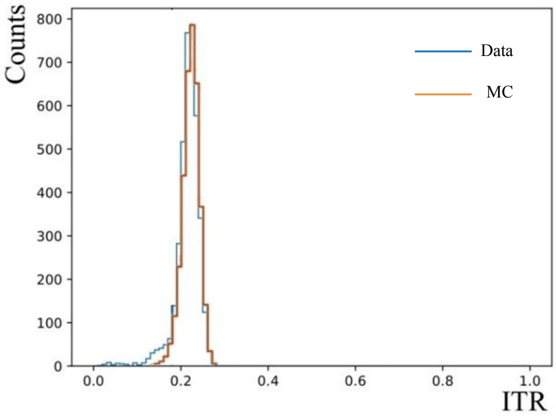

Figure 4.11: Comparison of ^{214}Bi data and MC ITR distributions.

The time residual PDFs used in the multisite classifier were the same as shown in figure 4.2a.

Cut	**Pass Criteria**
^{214}Bi Energy	$1.25 \leq E \leq 3.00$ MeV
^{8}B-ν Energy	$E \geq 6$ MeV
FV	$R_{recon} \leq 4.5$ m
ITR	$0.22 \leq ITR \leq 0.30$

Table 4.10: Selection cuts for multisite ^{214}Bi and single-site ^{8}B-ν events, applied to data and MC. ^{214}Bi were tagged using the cuts given in 3.2.

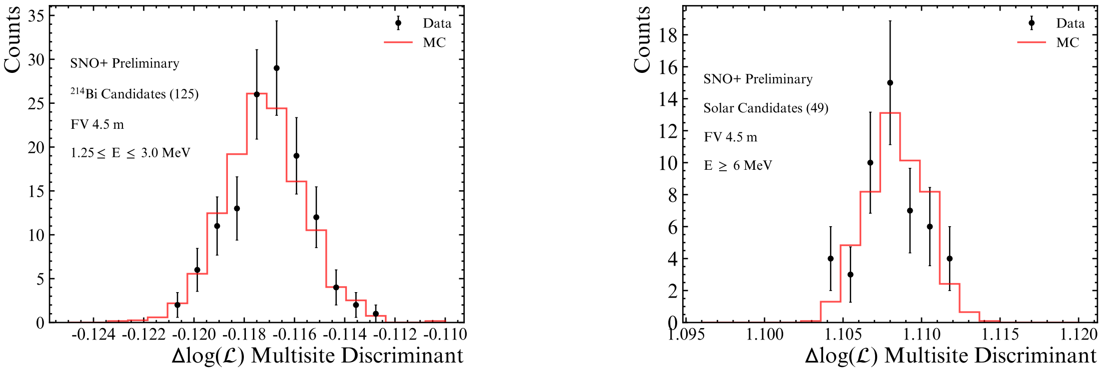

(a) $\Delta log(\mathcal{L})$ multisite discriminant comparison between ^{214}Bi data and MC.

(b) $\Delta log(\mathcal{L})$ multisite discriminant comparison between ^{8}B-ν_e ES data and MC.

Figure 4.12: Comparison between data and MC $\Delta log(\mathcal{L})$ multisite discriminant distributions for multisite (^{214}Bi) and single-site (^{8}B-ν_e ES) events.

Despite differences in the PDF generated event types, FVs and energy, the main purpose of this cross-check was to ensure the MC and data classifier results agree. Therefore, we are really checking if the time residual distributions inputted to the classifier are equivalent. As only the inputs may change, and the machinery of the classifier remains constant, this is still a fair test.

The results of this cross-check are given in figure 4.12. The multisite and single site $\Delta log(\mathcal{L})$ distributions show agreement for both samples, demonstrating the classifier results from MC are reliable. This was an important check, as the ultimate likelihood fit to recover the ^{8}B normalisation is sensitive to the shape of the multisite classifier distributions for each event type. If the distributions in this check did not conform between MC and data, the likelihood fit would lose sensitivity, as the shape information encoded in the likelihood would not be present in the data.

4.5 First Demonstration of Multisite Event Discrimination for ^{8}B-ν Flux Measurement in 3.5 to 5.0 MeV ROI

This analysis proceeded by floating the dominant multisite background from internal ^{208}Tl alongside the single-site signal arising from ^{8}B-ν elastic scattering, employing a focused energy range of 3.5 to 5.0 MeV. This energy range efficiently removes the residual Bi-Po backgrounds, leaving a clean selection of ^{208}Tl, ^{8}B and a negligible ^{210}Tl contribution. This choice was made to simplify the fitting code, reducing the problem to a 2D grid search over the single-site and multisite normalisations. Table 4.11 shows the updated ^{208}Tl and ^{8}B-ν event normalisations for the restricted ROI. The dataset was selected by applying the cuts summarised in table 4.2.

Event Type	**Updated ROI Efficiency**	**Predicted Normalisation**
^{208}Tl	0.273	297.4
^{8}B-ν_e	0.06	21.5
^{8}B-ν_μ	0.06	7.0

Table 4.11: Updated normalisations for ^{208}Tl multisite and ^{8}B-ν single-site events in the restricted 3.5 to 5.0 MeV ROI. Obtained by in the same manner as predictions given in tables 4.6 and 4.8, but using a restricted 3.5 to 5.0 MeV ROI.

Using equation 4.9, the normalisations of the ^{208}Tl and ^{8}B were grid searched over, creating a map of the $-2log(\mathcal{L})$ space. The minimum value was returned to yield the best fit normalisations for each interaction type, and used to scale the energy and multisite PDFs to create the data-model comparison plots shown in figure 4.13. Finally, the profile log-likelihood curves, shown in figure 4.14, were obtained from the 2D likelihood space by selecting the minimum value for each ^{8}B normalisation.

Fit	**Predicted ^{208}Tl**	**Fitted ^{208}Tl**	**Predicted ^{8}B-ν ES**	**Fitted ^{8}B-ν ES**
Energy	297.4 ± 29.7	174^{+17}_{-17}	28.5	27.0^{+13}_{-11}
Multisite	297.4 ± 29.7	201^{+15}_{-15}	28.5	0.0^{+9}_{-9}
Combined	297.4 ± 29.7	184^{+13}_{-12}	28.5	17.0^{+9}_{-8}

Table 4.12: Fitted Normalisations compared to background model predictions in 3.5 to 5.0 MeV ROI and 4.5 m FV. A 10% statistical uncertainty from the ^{212}Bi-^{212}Po tagging is propagated to the ^{208}Tl prediction. Fit uncertainties come from the 1σ frequentist intervals on the respective profile-likelihood curves. The predicted ^{8}B flux is presented without error, as it is used only to cross-check the scale of the fitted ^{8}B normalisations.

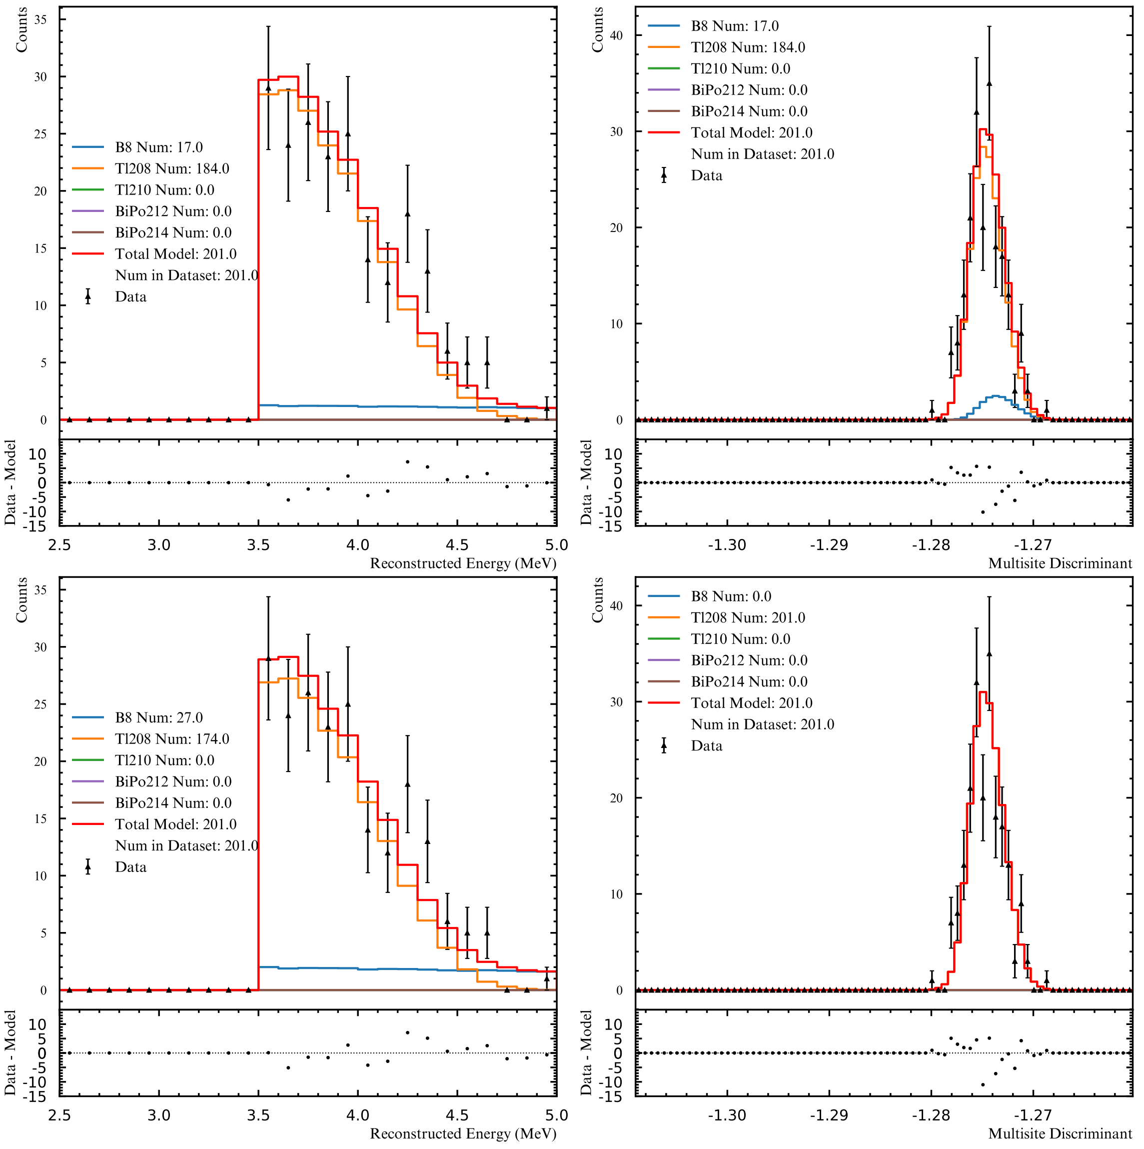

Figure 4.13: Fitted model and data using an unconstrained fit between 3.5 to 5.0 MeV. Upper row: energy and multisite model and data using combined likelihood; Bottom row: model and data using individual energy and multisite fits.

By comparing the widths and minima of the profile likelihood curves, the impact of including multisite information was quantified. As shown in figure 4.13, the multisite and energy fit are both consistent with the combined fit result at the 1σ level, and consistent with each other within 2σ. The energy fit prefers higher values for the ^{8}B normalisation, and is consistent with the background

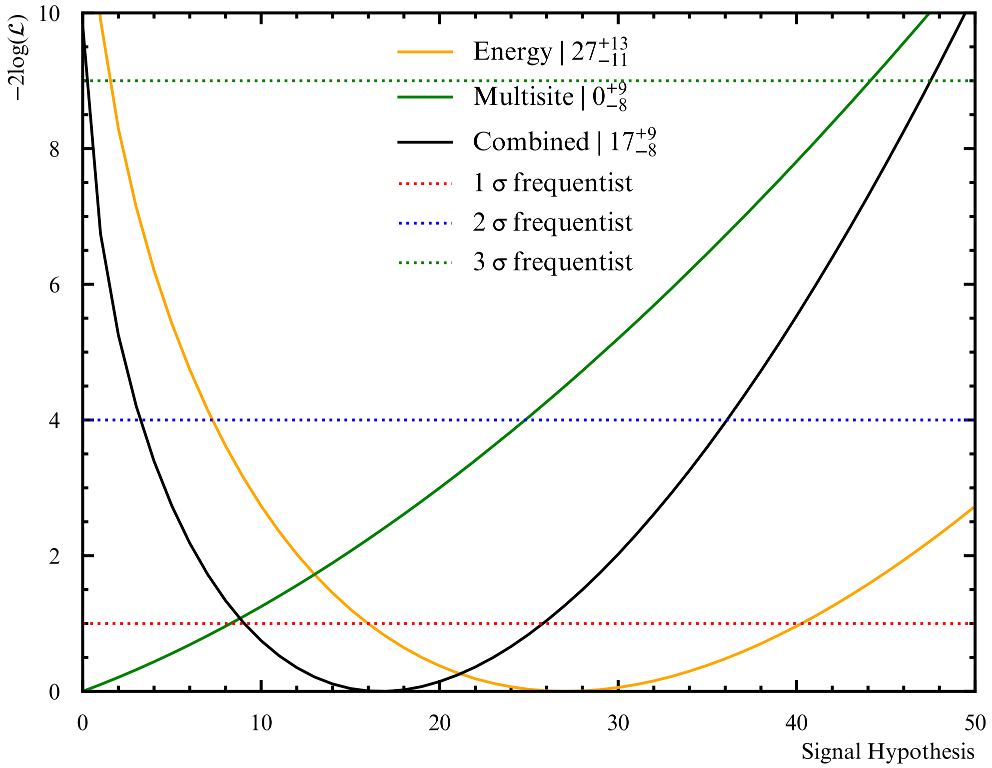

Figure 4.14: Profile $-2log(\mathcal{L})$ curve as a function of ^{8}B-ν_e event normalisation, obtained from energy (orange), multisite (green) and combined (black) fits on dataset.

model prediction of 28.5 ^{8}B-ν_e counts within 1σ (see table 4.12), and the combined is consistent with the prediction well within 2σ. The multisite fit minimises to 0 ^{8}B counts, which is a consequence of the low signal statistics and modest separation power of the multisite discriminant. Monte-Carlo studies, carried out on fluctuated datasets with the mean event rates fixed to the background model normalisations, consistently showed multisite fits minimising to 0 counts, with the picture improving with additional statistics (compare bias distributions of figure 4.15 and 4.10). Thus, it is expected that the multisite fits will gain sensitivity as more data is gathered.

Based on the results in figure 4.14, including multisite information in the combined likelihood significantly reduces the relative uncertainty of the predictions, with the width of the combined curve ~ 30 % smaller than the energy curve. However, the absolute error on the combined measurement is

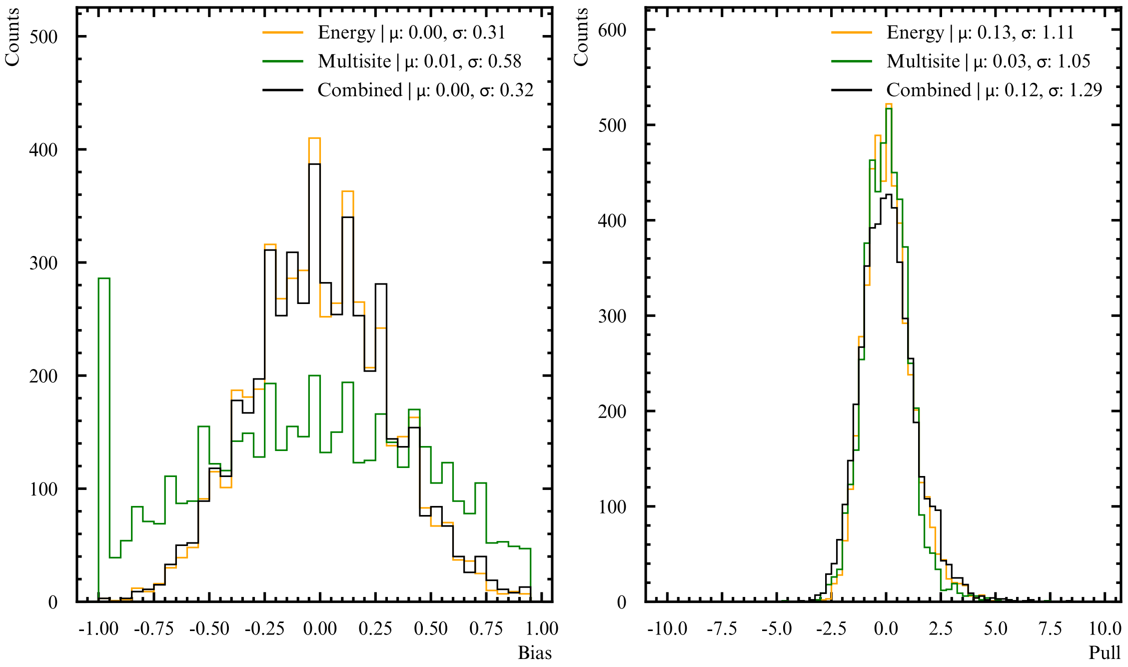

Figure 4.15: Bias and pull distributions for 5000 fake datasets, with mean rates fixed to the 145.7 day livetime. Large proportion of multisite fits minimise to zero, shown as a spike at bias = -1, mimicking the data result.

slightly higher than on the energy fit (48 % vs 52 %). Thus, background suppression techniques in addition to multisite discrimination will be necessary to perform this measurement. These additional techniques are considered in the final section below.

4.6 Conclusions and Future Work

This chapter has shown, by virtue of having well calibrated models of the SNO+ liquid scintillator emission timing, it is possible to exploit the differences in the time residual distributions between single-site β^- and multi-site $\beta^-\gamma$ interactions to perform multisite background discrimination. This was shown for the first time in a liquid scintillator detector, improving the precision of the unconstrained ^{8}B flux extraction by 30%, as compared to the uncertainty of an energy fit alone.

While a significant milestone in its own right, this successful demonstration is the first step in a longer journey towards a precise measurement of the solar neutrino transition region. As shown by figure 4.14, multisite alone is insufficient to resolve the fine details necessary to determine the ν_e survival probabilities with energy. This owes to its modest discrimination power and small S/B ratio. Therefore, future analyses will need to incorporate additional, complimentary techniques. One such method is ^{212}Bi-^{208}Tl coincidence tagging, which is currently under active development by the collaboration [92]. By adding additional and distinct information from ^{212}Bi-^{208}Tl tagging to the likelihood in equation 4.9, the sensitivity of subsequent solar analyses will be much increased.

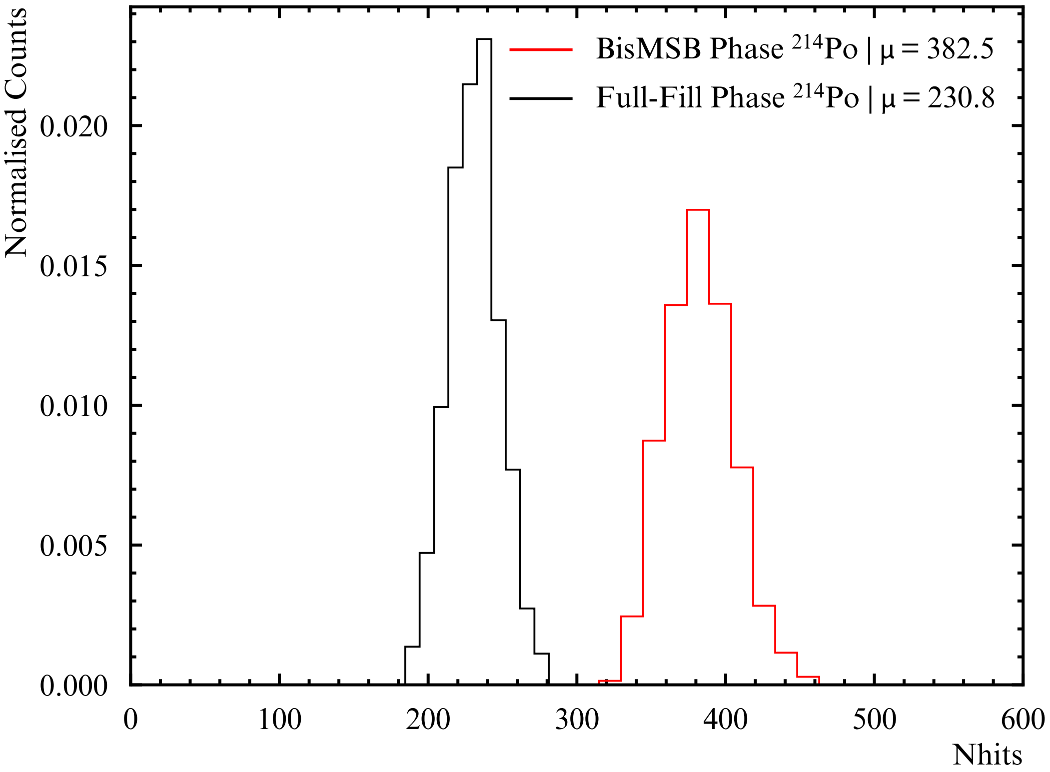

Figure 4.16: Comparison of tagged ^{214}Po distributions in the full-fill (black) and bismsb phases (red). The average of the distributions are given as μ, showing an average increase in nhit of $\sim 66\%$ following bisMSB addition. The bisMSB ^{214}Po gratefully received from Po-Wei Huang.

It is worth considering the impacts of including data taken during the current bisMSB and imminent tellurium phase. Addition of bisMSB, as described in section 2.1.1, will improve the performance of the energy and position reconstruction, and by extension the multisite classifier.

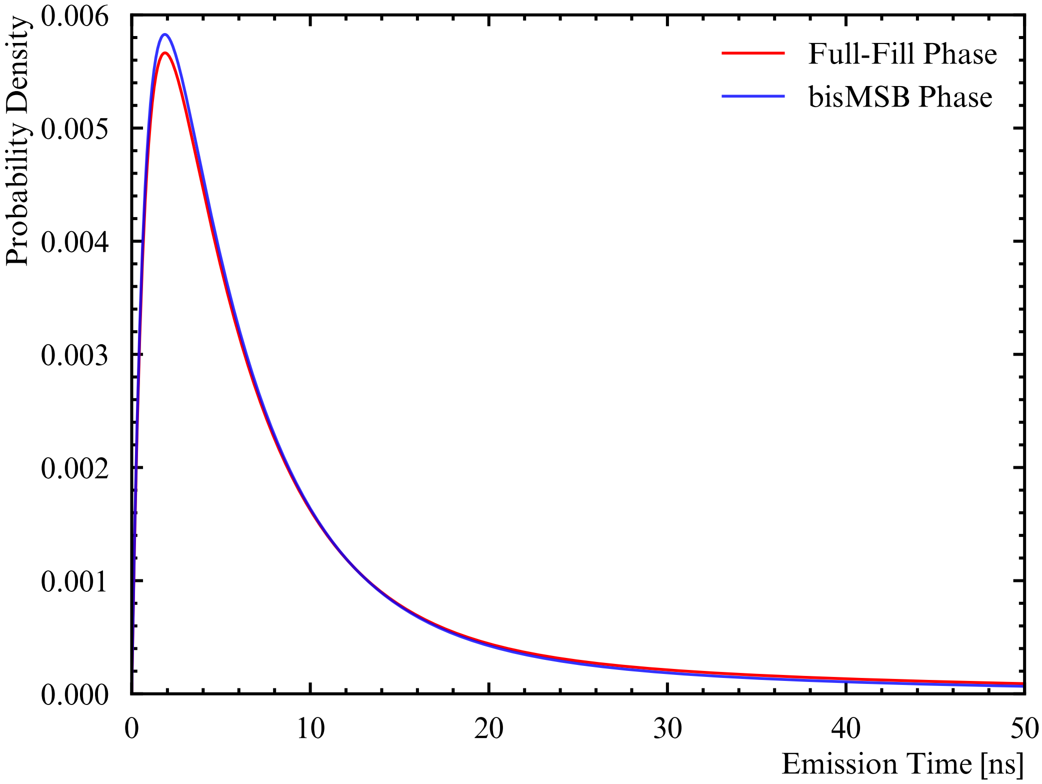

Figure 4.17: Comparison of the tuned scintillation emission models for the full-fill (red) and bisMSB phase (blue). Since the bisMSB does not directly compete with PPO for LAB excitation energy, the emission timing models are similar. Figure created using results from 3.4 and [93].

Figure 4.16 shows the increase in nhits obtained after loading 2.2 mg/L bisMSB. The additional light yield obtained from the bisMSB provides a boost to the energy reconstruction. This increased light yield is owed to a significant reduction in energy losses from PPO self-absorption [94], as well as better matching between the bisMSB emission spectrum and the peak efficiencies of the PMTs. The reduction in self-absorption is also the source of faster observed time residuals, as shown in figure 4.18. Since the energy transfer between PPO and bisMSB is radiative, the scintillation emission time model in the bisMSB phase is largely similar to that tuned for the full-fill phase. Figure 4.17 compares the two emission time models.

However, by efficiently absorbing the PPO emission on the $\sim$ cm scale and re-emitting it at

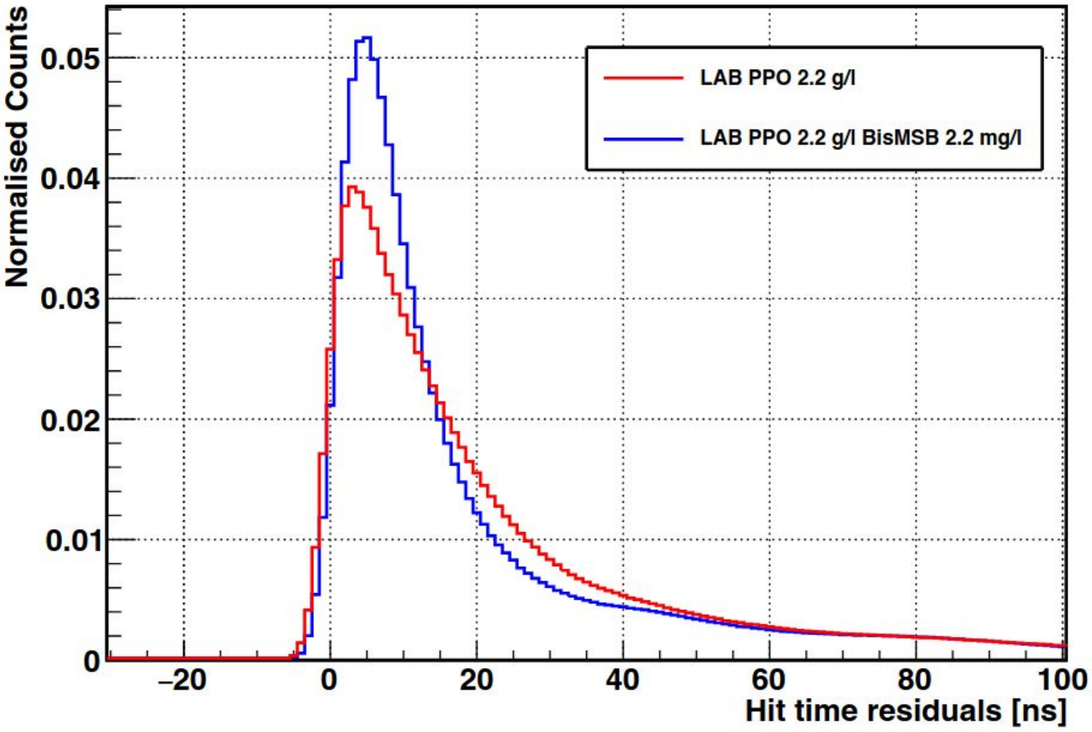

Figure 4.18: Time residual PDFs used by the position reconstruction in the full fill (red) and bisMSB phase (blue). The faster timing with bisMSB is a result of reduced PPO absorption-reemission. This led to improvements in resolution from 11 cm to 6 cm. Taken from [95].

longer wavelengths away from the peak PPO absorption, bisMSB addition leads to more peaked time residual distributions. This results in a significantly improved position resolution, going from $\sim$ 11 cm in full fill to $\sim$ 6 cm in the bisMSB phase [95].

The multisite discrimination power is better in bisMSB, due to faster observed time residual distributions after its addition. With faster timing, the emission spectrum narrows, reducing the dominant source of uncertainty in the position reconstruction algorithms. Therefore, single-site events appear more point-like, and their time residuals are rendered more distinct from multisite backgrounds.

A comparison of the light yield before and after bisMSB addition, along with a comparison of the preliminary scintillation emission model, is given in figure 4.16 and 4.17. An Asimov dataset was constructed using this preliminary bisMSB optics model, with normalisations scaled to the previous

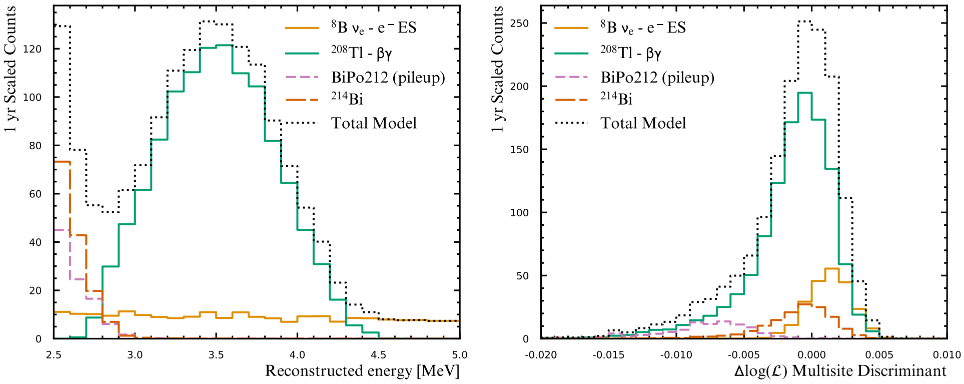

Figure 4.19: Asimov dataset constructed from 60 000 events per class, simulated with tuned bisMSB optics model.

Simulated Phase	Energy	Multisite	Combined
Full-fill	230^{+37}_{-35}	230^{+48}_{-47}	230^{+28}_{-29}
bisMSB	230^{+29}_{-31}	230^{+36}_{-37}	230^{+23}_{-24}

Table 4.13: Comparison of the Asimov 1σ frequentist errors obtained from simulating the full-fill and bisMSB phase.

1 year expected rates and shown in figure 4.19. Running the binned-extended log-likelihood fits exactly as before, but with updated bisMSB time residual and energy PDFs, yields the profile likelihood curves in figure 4.20.

Table 4.13 summarises the fit uncertainties for the Asimov results in the full-fill and bisMSB phase. Figure 4.20 suggests the multisite performance is improved in bisMSB by approximately 17.5%. Both the multisite and energy fits show improvement with bisMSB, due to the increased light yield and improved timing, respectively. It can be seen that the multisite precision in the bisMSB phase performs approximately as well as the energy fit in the full-fill result, constituting a 23% improvement in precision. The energy fit shows a slightly more modest 17% increase.

The inclusion of bisMSB data, alongside complimentary ^{208}Tl tagging information, provides an exciting prospect for a high precision measurement of the ^{8}B flux in the transition region for the first time. This will finally allow theoretical predictions of non-standard physics to be compared to experiments and possibly open the door for unexpected new physics in the neutrino sector.

As a final note, it is worth mentioning multisite event discrimination in the context of SNO+'s

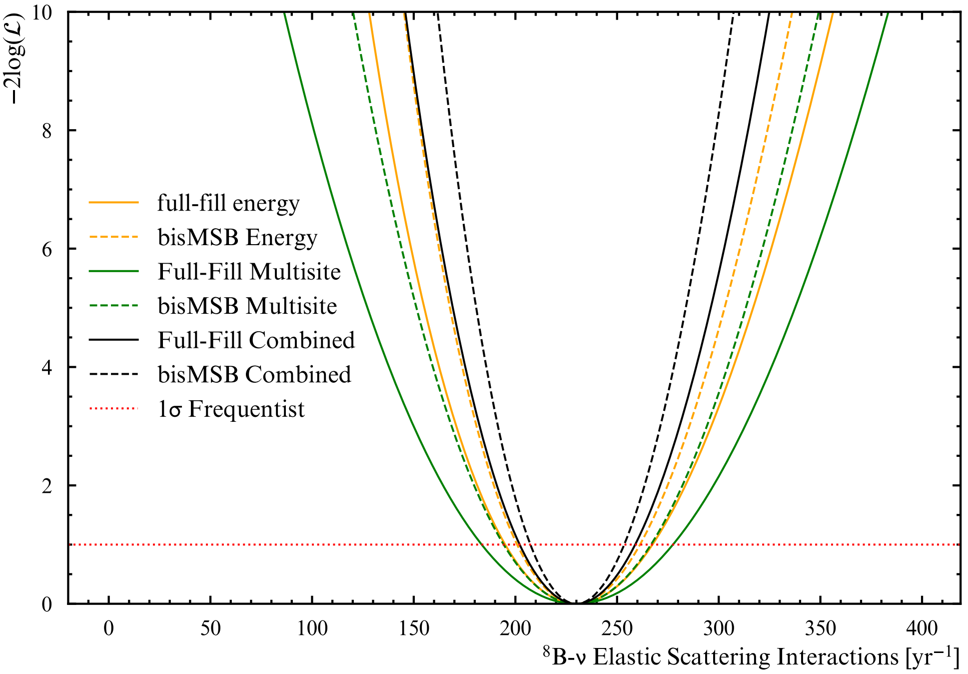

Figure 4.20: Profile log-likelihood curves comparing fit precision on Asimov datasets constructed from full-fill (solid lines) and bisMSB phase (dashed lines). bisMSB leads to significantly improved precision for all fits performed.

flagship physics goal: the detection of neutrinoless double-beta decay. As shown in [89] and [33], multisite discrimination is crucial for the rejection of cosmogenic backgrounds within the $0\nu\beta\beta$ ROI. Following this successful first demonstration of multisite event discrimination in data, work has begun on recoordinating the classifier for use in $0\nu\beta\beta$ studies.

Chapter 5

Conclusions

The solar neutrino spectrum between 2.0 to 6.0 MeV, known as the 'transition region,' remains a largely unexplored domain, offering a unique window into the intermediate oscillation behavior of neutrinos—between vacuum oscillations and the MSW effect. This region holds profound significance, as the ν_e survival probabilities are highly sensitive to physics beyond the Standard Model, including non-standard neutrino interactions and additional neutrino flavors. A precise measurement of the solar neutrino flux within this energy range could yield groundbreaking insights into these phenomena.

Recognizing this immense potential, this thesis has demonstrated that large liquid scintillator detectors are exceptionally well-suited for achieving high-precision measurements in this region. Their superior light yields, combined with strong position and energy resolution capabilities, position ^{8}B solar neutrinos well above detection thresholds. However, the primary challenge arises from the high rate of internal ^{208}Tl β-γ decays, which compromises sensitivity in this energy domain.

To address this, this work has pioneered the application of multisite event discrimination in ^{8}B solar neutrino flux measurements within the transition region. By exploiting subtle differences in the topologies between single-site ν elastic scattering events and multisite ^{208}Tl decays, it was shown that time residual distributions can reveal critical distinctions. The construction of a multisite log-likelihood ratio discriminant, grounded in well-calibrated scintillator emission time PDFs, encodes information on par with that of energy PDFs, significantly enhancing the analysis.

Central to this achievement was the development of a novel scintillation model calibration method, leveraging naturally occurring ^{214}Bi-^{214}Po coincidences. This allowed for the extraction of clean β and α interaction samples, tuning of light yield, α-quenching factors, and scintillation emission time constants for the full-fill phase. A major innovation of this thesis was the implementation of Bayesian optimization, which drastically reduced the computational effort required for tuning, yielding results comparable to previous grid-search methods. This streamlined approach represents a significant leap forward for future analyses, providing an efficient framework for prompt tunings in the upcoming tellurium phase.

Furthermore, precise measurements of internal ^{238}U and ^{232}Th concentrations were conducted using coincident Bi-Po tagging. These concentrations, both below the tellurium-phase targets, were crucial in constructing the background model for the multisite analysis, allowing for accurate predictions of ^{208}Tl, ^{210}Tl, and Bi-Po rates within the analysis ROI. In addition, these measurements provided important constraints for the background models of all full-fill SNO+ analyses.

For the first time, this multisite discriminant was successfully applied to liquid scintillator data in the transition region, improving the combined fit precision by 30% relative to the energy-only fit. This result underscores the power of multisite discrimination in significantly enhancing sensitivity to ^{8}B solar neutrinos.

Looking ahead, future SNO+ analyses will greatly benefit from the bisMSB phase, where position and energy resolutions are set to improve dramatically. With the addition of 2.2 mg/L bisMSB, light yield is expected to increase by around 40%, with position resolution improving from 11 cm to an impressive 6 cm. These advancements are driven by bisMSB's ability to reduce self-absorption of scintillation light by the primary fluor, PPO, significantly reducing the spread of the observed time residual distributions and improving position reconstruction. The projected enhancement in ^{8}B solar neutrino measurement precision in the bisMSB phase is estimated at 17.5%, based on an Asimov dataset, further solidifying the potential for future breakthroughs.

In the near term, the integration of multisite event discrimination with ^{212}Bi-^{208}Tl coincidence tagging promises to unlock even greater precision. By incorporating this additional discriminant into the likelihood analysis, the sensitivity of ^{8}B solar neutrino flux measurements will be further refined. As we continue to leverage these advanced methodologies, future investigations into the solar neutrino flux in the transition region will draw on a comprehensive combination of energy, multisite discrimination, and ^{208}Tl tagging. These innovations will bring us closer than ever to a precise measurement of ^{8}B neutrinos, offering critical insights into both neutrino oscillations and potential BSM physics.

Bibliography

[1] W. Pauli. Dear radioactive ladies and gentlemen. *Phys. Today*, 31N9:27, 1978.

[2] F. Reines. The neutrino: from poltergeist to particle. *Rev. Mod. Phys.*, 68:317–327, Apr 1996.

[3] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956.

[4] John N. Bahcall, Walter F. Huebner, Stephen H. Lubow, Peter D. Parker, and Roger K. Ulrich. Standard solar models and the uncertainties in predicted capture rates of solar neutrinos. *Rev. Mod. Phys.*, 54:767–799, Jul 1982.

[5] Bruce T. Cleveland, Timothy Daily, Jr. Raymond Davis, James R. Distel, Kenneth Lande, C. K. Lee, Paul S. Wildenhain, and Jack Ullman. Measurement of the solar electron neutrino flux with the homestake chlorine detector. *The Astrophysical Journal*, 496(1):505, mar 1998.

[6] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, May 1968.

[7] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten. Reanalysis of the gallex solar neutrino flux and source experiments. *Physics Letters B*, 685(1):47–54, February 2010.

[8] The GNO Collaboration. Complete results for five years of gno solar neutrino observations. *Physics Letters B*, 616(3–4):174–190, June 2005.

[9] The SAGE Collaboration. Measurement of the solar neutrino capture rate with gallium metal. iii. results for the 2002–2007 data-taking period. *Phys. Rev. C*, 80:015807, Jul 2009.

[10] John N. Bahcall and R. L. Sears. Solar neutrinos. *Annual Review of Astronomy and Astrophysics*, 10(Volume 10, 1972):25–44, 1972.

[11] The SNO Collaboration. Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Phys. Rev. Lett.*, 89:011301, Jun 2002.

[12] A. Bellerive, J.R. Klein, A.B. McDonald, A.J. Noble, and A.W.P. Poon. The sudbury neutrino observatory. *Nuclear Physics B*, 908:30–51, July 2016.

[13] M. Schwartz. Feasibility of using high-energy neutrinos to study the weak interactions. *Phys. Rev. Lett.*, 4:306–307, Mar 1960.

[14] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28(5):870–880, 11 1962.

[15] The Particle Data Group. Review of particle physics. *Phys. Rev. D*, 110:030001, Aug 2024.

[16] The KM3Net Collaboration. Letter of intent for km3net 2.0. *Journal of Physics G: Nuclear and Particle Physics*, 43(8):084001, jun 2016.

[17] The Borexino Collaboration. Comprehensive measurement of pp-chain solar neutrinos. *Nature*, 562(7728):505–510, Oct 2018.

[18] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, May 1978.

[19] The Borexino Collaboration. Precision measurement of the ^{7}Be solar neutrino interaction rate in borexino. *Phys. Rev. Lett.*, 107:141302, Sep 2011.

[20] The Borexino Collaboration. First evidence of *pep* solar neutrinos by direct detection in borexino. *Phys. Rev. Lett.*, 108:051302, Feb 2012.

[21] The Borexino Collaboration. Measurement of the solar ^{8}B neutrino rate with a liquid scintillator target and 3 mev energy threshold in the borexino detector. *Phys. Rev. D*, 82:033006, Aug 2010.

[22] The Borexino Collaboration. Final results of borexino on cno solar neutrinos. *Phys. Rev. D*, 108:102005, Nov 2023.

[23] The KATRIN Collaboration. Direct neutrino-mass measurement based on 259 days of katrin data, 2024.

[24] Ettore Majorana. Teoria simmetrica dell'elettrone e del positrone. *Il Nuovo Cimento (1924-1942)*, 14(4):171–184, Apr 1937.

[25] Tsutomu Yanagida. Horizontal Symmetry and Masses of Neutrinos. *Progress of Theoretical Physics*, 64(3):1103–1105, 09 1980.

[26] M. Goeppert-Mayer. Double beta-disintegration. *Phys. Rev.*, 48:512–516, Sep 1935.

[27] Josephine Paton. *Directional Reconstruction in Liquid Scintillator Neutrino Detectors Using Time Based Separation of Cerenkov Light.* PhD thesis, University of Oxford, 2023.

[28] Andrea Giuliani and Alfredo Poves. Neutrinoless double-beta decay. *Advances in High Energy Physics*, 2012(1):857016, 2012.

[29] J. Schechter and J. W. F. Valle. Neutrinoless double-β decay in su(2)$\times$u(1) theories. *Phys. Rev. D*, 25:2951–2954, Jun 1982.

[30] The SNO+ Collaboration. The sno+ experiment. *Journal of Instrumentation*, 16(08):P08059, August 2021.

[31] Stefano Dell'Oro, Simone Marcocci, Matteo Viel, and Francesco Vissani. Neutrinoless double beta decay: 2015 review. *Advances in High Energy Physics*, 2016:1–37, 2016.

[32] The KamLAND-Zen Collaboration. Search for majorana neutrinos near the inverted mass hierarchy region with kamland-zen. *Phys. Rev. Lett.*, 117:082503, Aug 2016.

[33] Tereza Kroupova. *Improving the Sensitivity to Neutrinoless Double Beta Decay in SNO+.* PhD thesis, University of Oxford, 2020.

[34] The Borexino Collaboration. Constraints on flavor-diagonal non-standard neutrino interactions from borexino phase-ii. *Journal of High Energy Physics*, 2020(2):38, Feb 2020.

[35] *Progress in Particle and Nuclear Physics*, 123:103927, March 2022.

[36] J B Birks. Chapter 1: Introduction. In *The Theory and Practice of Scintillation Counting*, International Series of Monographs on Electronics and Instrumentation, pages 1–13. Pergamon Press, 1967.

[37] John Betteley Birks. *The theory and practice of scintillation counting: International series of monographs in electronics and instrumentation*, volume 27, chapter 3, pages 39–67. 1967.

[38] J B Birks. Chapter 8: Organic liquid scintillators. In *The Theory and Practice of Scintillation Counting*, International Series of Monographs on Electronics and Instrumentation, pages 269–351. Pergamon Press, 1967.

[39] John Betteley Birks. *The theory and practice of scintillation counting: International series of monographs in electronics and instrumentation*, volume 27, chapter 8, pages 269–320. 1967.

[40] 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.

[41] Lord Rayleigh. Xxxiv. on the transmission of light through an atmosphere containing small particles in suspension, and on the origin of the blue of the sky. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 47(287):375–384, 1899.

[42] Cabannes, Jean. Relation entre le degré de polarisation et l'intensité de la lumière diffusée par des molécules anisotropes. nouvelle détermination de la constante d'avogadro. *J. Phys. Radium*, 1(5):129–142, 1920.

[43] Xiang Zhou, Qian Liu, Michael Wurm, Qingmin Zhang, Yayun Ding, Zhenyu Zhang, Yangheng Zheng, Li Zhou, Jun Cao, and Yifang Wang. Rayleigh scattering of linear alkylbenzene in large liquid scintillator detectors. *Review of Scientific Instruments*, 86(7), July 2015.

[44] J B Birks. Chapter 5: The detection of scintillations. In *The Theory and Practice of Scintillation Counting*, International Series of Monographs on Electronics and Instrumentation, pages 96–180. Pergamon Press, 1967.

[45] Hamamatsu. *Photomultiplier Tubes: Basics and Applications*, 2017. Accessed: 2024-10-05.

[46] The SNO Collaboration. The sudbury neutrino observatory. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 449(1-2):172–207, 2000.

[47] The SNO+ Collaboration. The sno+ experiment. *Journal of Instrumentation*, 16(08):P08059, 2021.

[48] The SNO+ Collaboration. Optical calibration of the sno+ detector in the water phase with deployed sources. *Journal of Instrumentation*, 16(10):P10021, 2021.

[49] The SNO+ Collaboration. Search for invisible modes of nucleon decay in water with the sno+ detector. *Physical Review D*, 99(3):032008, 2019.

[50] The SNO+ Collaboration. Improved search for invisible modes of nucleon decay in water with the sno+ detector. *Physical Review D*, 105(11):112012, 2022.

[51] The SNO+ Collaboration. Measurement of neutron-proton capture in the sno+ water phase. *Physical Review C*, 102(1):014002, 2020.

[52] The SNO+ Collaboration. Measurement of the ^{8}B solar neutrino flux in SNO+ with very low backgrounds. *Phys. Rev. D*, 99:012012, Jan 2019.

[53] The SNO+ Collaboration. Evidence of antineutrinos from distant reactors using pure water at sno+. *Physical Review Letters*, 130(9):091801, 2023.

[54] The SNO+ Collaboration. Initial measurement of reactor antineutrino oscillation at SNO+. *arXiv preprint arXiv:2405.19700*, 2024.

[55] Ana Sofia Carpinteiro Inácio. *Data Analysis of the Water and Scintillator Phases of SNO+: from Solar Neutrino Measurements to Double Beta Decay Sensitivity Studies.* PhD thesis, UNIVERSIDADE DE LISBOA FACULDADE DE CIÊNCIAS, 2022.

[56] The SNO+ Collaboration. Event-by-event direction reconstruction of solar neutrinos in a high light-yield liquid scintillator. *Physical Review D*, 109(7):072002, 2024.

[57] The SNO+ Collaboration. Current status and future prospects of the sno+ experiment. *Advances in High Energy Physics*, 2016(1):6194250, 2016.

[58] The CUORE Collaboration. Measurement of the $2\nu\beta\beta$ decay half-life of ^{130}Te with cuore. *Phys. Rev. Lett.*, 126:171801, Apr 2021.

[59] The NEMO Collaboration. Measurement of the double-β decay half-life of ^{150}Nd and search for neutrinoless decay modes with the NEMO-3 detector. *Physical Review C—Nuclear Physics*, 80(3):032501, 2009.

[60] The KamLAND Collaboration. Light output response of kamland liquid scintillator for protons and 12c nuclei. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 622(3):574–582, 2010.

[61] The Borexino Collaboration. Measurements of liquid scintillator properties for the borexino detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 400(1):53–68, 1997.

[62] The SNO+ Collaboration. Development, characterisation, and deployment of the sno+ liquid scintillator. *Journal of Instrumentation*, 16(05):P05009, May 2021.

[63] Benjamin Tamm Amanda Bacon. Sno+ scintillator (lab + 2.2 g/l ppo) light yield plot for approval - docdb 7406. *Internal SNO+ Document*, 2022.

[64] DJ Auty, D Bartlett, SD Biller, D Chauhan, M Chen, O Chkvorets, S Connolly, X Dai, E Fletcher, K Frankiewicz, et al. A method to load tellurium in liquid scintillator for the study of neutrinoless double beta decay. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1051:168204, 2023.

[65] RAT Collaboration. Reactor analysis toolkit documentation. `https://rat.readthedocs.io/en/latest/overview.html`, 2024. Accessed: 2024-10-08.

[66] S. Agostinelli et al. GEANT4–a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003.

[67] Glenn Horton-Smith. Glg4sim: Generic liquid-scintillator anti-neutrino detector, 2007. Accessed: 2024-10-08.

[68] Helen O'Keeffe and Valentina Lozza. Expected radioactive backgrounds in sno+. Technical Report 507, SNO+, 2010. Internal SNO+ document.

[69] Jia-Shian Wang. *Supernova Neutrinos and Measurement of Liquid Scintillator Backgrounds in SNO+*. PhD thesis, University of Oxford, 2022.

[70] J. D. Wilson. Thermally-driven scintillator flow in the sno+ neutrino detector, 2022.

[71] Serena Riccetto. Rat optics 2.2g/l summary cm mar2023. *Internal SNO+ Document 7728*, 2023.

[72] Iwan Morton-Blake. *First Measurement of Reactor Antineutrinos in Scintillator at SNO+ and Study of Alternative Designs for Large-Scale Liquid Scintillator Detectors*. PhD thesis, University of Oxford, 2021.

[73] Steven D. Biller, Edward J. Leming, and Josephine L. Paton. Slow fluors for effective separation of cherenkov light in liquid scintillators. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 972:164106, August 2020.

[74] Martin Berger, J Coursey, and M Zucker. Estar, pstar, and astar: Computer programs for calculating stopping-power and range tables for electrons, protons, and helium ions (version 1.21), 1999-01-01 1999.

[75] F. James and M. Roos. Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations. *Comput. Phys. Commun.*, 10:343–367, 1975.

[76] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

[77] James Bergstra, Daniel Yamins, and David Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine*

Learning, volume 28 of *Proceedings of Machine Learning Research*, pages 115–123, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.

[78] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, Dec 1998.

[79] David Lowe and D Broomhead. Multivariable functional interpolation and adaptive networks. *Complex systems*, 2(3):321–355, 1988.

[80] Bertil Matérn. *Spatial variation*, volume 36. Springer Science & Business Media, 2013.

[81] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.

[82] University of Wisconsin-Madison Center for High Throughput Computing, Computer Sciences Department. Htcondor manual. `https://htcondor.readthedocs.io/en/latest/index.html`, 2024. Accessed: 2024-10-06.

[83] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.

[84] W. T. Winter, S. J. Freedman, K. E. Rehm, and J. P. Schiffer. The ^{8}B neutrino spectrum. *Phys. Rev. C*, 73:025503, Feb 2006.

[85] Núria Vinyoles, Aldo M. Serenelli, Francesco L. Villante, Sarbani Basu, Johannes Bergström, M. C. Gonzalez-Garcia, Michele Maltoni, Carlos Peña-Garay, and Ningqiang Song. A new generation of standard solar models. *The Astrophysical Journal*, 835(2):202, January 2017.

[86] The Particle Data Group. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8):083C01, 08 2020.

[87] Benjamin Tam. Lab certificate of analysis- docdb 7572. *Internal SNO+ Document*, 2022.

[88] Edoardo Vitagliano, Irene Tamborra, and Georg Raffelt. Grand unified neutrino spectrum at earth: Sources and spectral components. *Rev. Mod. Phys.*, 92:045006, Dec 2020.

[89] Jack Dunger and Steven D. Biller. Multi-site event discrimination in large liquid scintillation detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 943:162420, November 2019.

[90] Isaac Asimov. Franchise. In *If: Worlds of Science Fiction*. Fantasy House, 1955.

[91] Valentina Lozza. Radioactive backgrounds in sno+: A starting guide - docdb 8091. *Internal SNO+ Document*, 2023.

[92] Gulliver Milton. Tl208 pre collaboration meeting background call. Technical Report 8316, SNO+, 2024. Internal SNO+ document.

[93] Po-Wei Huang. Recordination of optical model in batch4 bismsb detector phase -docdb 8266. *Internal SNO+ Document*, 2024.

[94] Rafael Hunt-Stokes and Po-Wei Huang. bismsb timing study - e&o collaboration meeting - march 2024 - docdb 8157. *Internal SNO+ Document*, 2024.

[95] William Parker. Position recoordination for labppo_2p2_bismsb_2p2_scintillator - docdb 8326. *Internal SNO+ Document*, 2024.